

# LEAKAGE AND SECOND-ORDER DYNAMICS IMPROVE HIPPOCAMPAL RNN REPLAY

Anonymous authors

Paper under double-blind review

## ABSTRACT

Biological neural networks (like the hippocampus) can internally generate “replay” resembling stimulus-driven activity. Recent computational models of replay use noisy recurrent neural networks (RNNs) trained to path-integrate. Replay in these networks has been described as Langevin sampling, but new modifiers of noisy RNN replay have surpassed this description. We re-examine noisy RNN replay as sampling to understand or improve it in three ways: (1) Under simple assumptions, we prove that the gradients replay activity should follow are time-varying and difficult to estimate, but readily motivate the use of hidden state leakage in RNNs for replay. (2) We confirm that hidden state adaptation (negative feedback) encourages exploration in replay, but show that it incurs non-Markov sampling that also slows replay. (3) We propose the first model of temporally compressed replay in noisy path-integrating RNNs through hidden state momentum, connect it to underdamped Langevin sampling, and show that, when combined with adaptation, it counters slowness while maintaining exploration. We verify our findings via path-integration of 2D paths in T-maze and triangular environments and of high-dimensional paths of synthetic rat place cell activity.

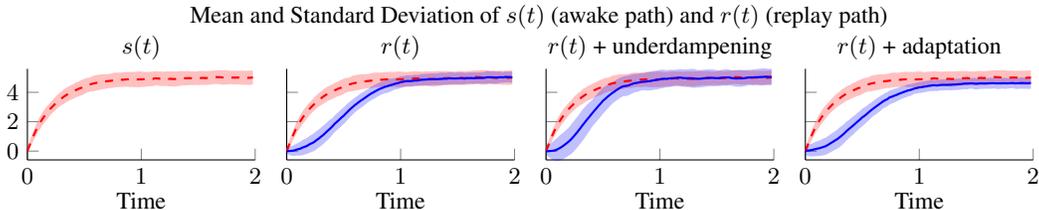


Figure 1: **Underdamped dynamics accelerate offline replay, adaptation slows it.** Here we simulate a noisy RNN  $r(t)$  that optimally path-integrates an Ornstein-Uhlenbeck process  $s(t)$  from its velocity  $s'(t)$ . We assume  $r(t)$  minimizes the loss in Equation (8) and thus evolves according to its *score function*  $\nabla_{r(t)} \log p(r(t))$  (Equations 12 and 16), performing a variant of Langevin sampling when no input is given. Above, we compare three modifiers of RNN activity: the default (no modification, a.k.a. *overdamped*), our proposed *underdamped* (momentum), and *adaptation* (negative feedback) dynamics. Each modifier affects the replay distribution  $p(r(t))$  in different ways: underdamped sampling accelerates  $p(r(t))$  towards  $p(s(t))$ , decreasing the distance between them, while adaptation slows convergence of  $p(r(t))$  towards  $p(s(t))$ , increasing this distance.

## 1 INTRODUCTION

During quiescent periods such as sleep or wakeful resting, some neural circuits internally generate activity resembling that of active periods (Tingley & Peyrache, 2020). Such “replay” phenomena have been observed in the prefrontal (Euston et al., 2007; Peyrache et al., 2009), sensory (Kenet et al., 2003; Xu et al., 2012), motor (Hoffman & McNaughton, 2002), and entorhinal cortices (Gardner et al., 2022); the anterior thalamus (Peyrache et al., 2015); and the hippocampus (Buzsáki, 1986; Skaggs & McNaughton, 1996; Nádasdy et al., 1999; Lee & Wilson, 2002; Foster, 2017). Of these circuits, the hippocampus is particularly interesting because its robustness in tasks like navigation (O’Keefe & Nadel, 1978; Burgess et al., 1994; McNaughton et al., 1996) and planning (Pfeiffer &

Foster, 2013; Miller et al., 2017) during active states seems crucially tied to its spontaneous activity during quiescent states (Buzsáki, 1989; 2015; Tononi & Cirelli, 2014; Ólafsdóttir et al., 2015; 2018).

While some works have produced replay using supervised generative models (Deperrois et al., 2022), most existing models of hippocampal activity treat replay as an emergent byproduct of careful network design. Relevant network parameters include connectivity structures (Shen & McNaughton, 1996; Milstein et al., 2023), local plasticity mechanisms (Hopfield, 2010; Litwin-Kumar & Doiron, 2014; Theodoni et al., 2018; Haga & Fukai, 2018; Asabuki & Fukai, 2025), firing rate adaptation (Chu et al., 2024; Azizi et al., 2013; Itskov et al., 2011; Dong et al., 2021; Li et al., 2024), and input modulation (Kang & DeWeese, 2019). While these models reproduce aspects of replay, they are typically motivated by empirical findings and lack rigorous theoretical justification.

A more principled model of hippocampal function with emergent replay is sequential predictive learning (Krishna et al., 2024; Levenstein et al., 2024), wherein neural circuits predict dynamic environment or task variables from imperfect observations thereof, e.g., path-integrating velocity measurements to track a position. This normative description of the hippocampus as a sequence predictor (Levy, 1989; Stachenfeld et al., 2017) matches hippocampal encodings of upcoming stimuli (Davachi & DuBrow, 2015) and prediction errors (Aitken & Kok, 2022; Miller et al., 2023), and neural activity sweeps that represent possible future trajectories (Kay et al., 2020; Johnson & Redish, 2007). Unlike traditional, hand-crafted models of hippocampal circuits (i.e., continuous attractor networks (Samsonovich & McNaughton, 1997; Battaglia & Treves, 1998)), sequential predictive learning models are trained from data. Nonetheless, they account for the emergence of place cells (Recanatesi et al., 2021; Levenstein et al., 2024; Chen et al., 2024), grid cells (Cueva & Wei, 2018; Sorscher et al., 2019), and head direction cells (Cueva et al., 2020; Uria et al., 2022); can incorporate phenomena like theta oscillations (Levenstein et al., 2024); and exhibit quiescent replay activity (Krishna et al., 2024; Levenstein et al., 2024; Chen et al., 2024).

Krishna et al. (2024) provided the first theoretical foundation for replay in sequential predictive learning networks, showing analytically that they generate diffusive replay (Stella et al., 2019) during quiescent activity (i.e., noise-driven activity in the absence of inputs) by Langevin sampling (Besag, 1994) from the waking activity distribution using its score function. Subsequent empirical work (Levenstein et al., 2024) introduced new mechanisms to induce *exploration* in replay through negative feedback, i.e., neural adaptation. Exploration is the notion that replay expresses a variety of behavioral sequences (Davidson et al., 2009; Pfeiffer, 2020), and is associated with long trajectories in neural space, visitation of multiple attractor basins, and transitions that were not present in awake activity. Adaptation can destabilize attractors and induce sudden transitions in replay activity (Itskov et al., 2011; Dong et al., 2021; Li et al., 2024; Levenstein et al., 2024), thereby facilitating exploration, and is thought to play a key role in the dynamics of replay *in vivo* (Levenstein et al., 2019). However, existing theory on sequential predictive learning cannot account for these mechanisms. Furthermore, sequential predictive learning models do not currently account for the *temporal compression* of replay sequences relative to awake sequences of activity (Nádasdy et al., 1999; Buzsáki, 2015; Michelmann et al., 2019; Farooq & Dragoi, 2019). This phenomenon, which could be caused by short-term facilitation (Leibold et al., 2008; Jaramillo & Kempter, 2017), is not currently captured in replay from any trained RNN model to our knowledge.

Overall, while sequential predictive learning is a promising model of hippocampal function and replay, its recent empirical advances have outpaced its theoretical foundations. Moreover, it is unclear how to incorporate phenomena like temporal compression in these models, or what inductive biases in these RNNs’ design affect replay. We remedy these shortcomings by characterizing how RNN design and Langevin sampling statistics affect each other. Some results describe how RNN design affects the speed and variance of replay activity, while others start from the Langevin sampling formulation of replay and either explain existing architectural choices as useful inductive biases, or propose new mechanisms to again modulate sampling. In summary, we answer three questions:

1. Optimal path-integration in the presence of noise requires RNNs to learn the score function of the noisy activity distribution. What is this function, might it inform RNN design? *The score function is time-variant and difficult to estimate, even for simple distributions, but our expression of it motivates the addition of leakage (linear dynamics) in RNNs.*

- 108 2. Adaptation (negative feedback) empirically induces exploration in replay. How does it affect  
 109 Langevin sampling? *Adaptation induces non-Markov second-order Langevin sampling that*  
 110 *destabilizes attractors, which can both help (diversify) and hurt (slow) replay.*  
 111  
 112 3. Traditional generative models benefit from a wide array of sampling techniques. Could new sam-  
 113 pling methods improve replay in noisy RNNs? *Underdamped Langevin sampling via momentum*  
 114 *quickens neural replay, like temporal compression induced by short-term facilitation in vivo, and*  
 115 *mitigates slowing from adaptation while maintaining exploration.*

## 116 2 BACKGROUND

117  
 118  
 119 In this section, we provide background and summarize results from prior work (Krishna et al., 2024)  
 120 that has described replay in path-integrating neural circuits as the sampling during quiescence (i.e.,  
 121 the absence of any inputs) of neural activity states from the distribution of waking, task-like activity.  
 122 In short, the dynamics of recurrent networks that learn to optimally path-integrate noisy inputs cause  
 123 the network’s states even in the absence of inputs to resemble those attained during actual task  
 124 performance. We also provide details on mechanisms used in the literature to improve or bias replay,  
 125 which we explore the effects of in more detail in this work. For an overview of this section, see  
 126 Appendix B.

### 127 2.1 LANGEVIN DYNAMICS

128  
 129 Langevin sampling from an unknown distribution  $p(x)$  entails stochastic gradient ascent of an iterate  
 130  $x(t)$  along the log-likelihood of  $p(x)$  via its *score function*  $\nabla \log p(x)$  or an estimate thereof:

$$131 \quad 132 \quad x'(t) = x(t) + \nabla \log p(x) + \sqrt{2}\eta(t), \quad (1)$$

133 where  $\eta(t)$  is Gaussian white noise. While Equation 1 describes *overdamped* dynamics, there also  
 134 exist *underdamped* Langevin dynamics<sup>1</sup> (Equations 2 or 3, see Chapter 6 of Pavliotis (2014)) that  
 135 converge faster to the target distribution  $p(x)$  and better utilize noisy gradients (Cheng et al., 2018):

$$136 \quad 137 \quad x''(t) = \nabla \log p(x) - \gamma x'(t) + \sqrt{2\gamma}\eta(t), \quad \text{or} \quad (2)$$

$$138 \quad 139 \quad x'(t) = v(t), \quad v'(t) = \nabla \log p(x) - \gamma v(t) + \sqrt{2\gamma}\eta(t) \quad (3)$$

140  
 141 In this work, we consider noisy RNNs that have implicitly learned to perform Langevin sampling  
 142 of their own, fixed distributions of activity during task performance, when driven by just intrinsic  
 143 noise and in the absence of inputs. That is, the activity of the RNNs at each unrolled timestep in  
 144 the absence of inputs represents a plausible and likely vector of network activity during actual task  
 145 performance in the presence of inputs. In particular, we view such networks in the context of replay,  
 146 where neural circuits recapitulate task-like activity even during sleep.

### 147 2.2 OFFLINE REPLAY IN RNNs

148  
 149 This work focuses on RNNs that implicitly learn to act as generative models over a fixed distribution  
 150 of input sequences. Krishna et al. (2024) have shown how noisy RNNs trained to *path-integrate*  
 151 their inputs implicitly learn statistics that produce Langevin sampling of *their own task-relevant*  
 152 *activity distribution* when no input is given, demonstrating statistically faithful replay sequences  
 153 during quiescence. That is, the RNNs’ activity in the absence of inputs “replays” states from the  
 154 same distribution as RNN activity with inputs during actual task performance. This leads to the  
 155 generation or “replay” of output sequences resembling those that the network sees during training or  
 156 task performance with inputs. Path-integration is particularly relevant to neuroscience: animals can  
 157 leverage motion cues, observations, or prior experiences to accurately estimate positions (Seelig &  
 158 Jayaraman, 2015; Chrastil, 2025), and neural circuits like the entorhinal cortex (Sorscher et al., 2019)  
 159 have been identified to perform such computations. Here we summarize the finding that noisy RNNs  
 160 trained to path-integrate input time-series learn the *score function* of the input distribution (Krishna  
 161 et al., 2024).

<sup>1</sup>Discretized underdamped Langevin dynamics are a form of Hamiltonian MCMC (Cheng et al., 2018).

**Definition 1.** A noisy recurrent neural network (RNN) has hidden states  $\mathbf{r}(t)$  that evolve at each timestep  $t$  via some (nonlinear) function of its previous hidden states, an input signal  $\mathbf{u}(t)$ , and noise:

$$\mathbf{r}(t + \Delta t) = f(\mathbf{r}(t), \mathbf{u}(t), \sigma_r \boldsymbol{\eta}(t)) \quad (4)$$

**Definition 2.** A path-integration objective  $\mathcal{L}(t)$  penalizes the difference between a state variable  $\mathbf{s}(t)$  and a learnable linear projection of the RNN hidden state  $\mathbf{r}(t)$  at every timestep  $t$ :

$$\mathcal{L}(t) = \mathbb{E}_\eta \|\mathbf{s}(t) - \mathbf{D}\mathbf{r}(t)\|_2 \quad (5)$$

**Assumption 1.** *Krishna et al. (2024)* assume the hidden state dynamics of a noisy RNN can be decomposed as the sum of two separate functions<sup>2</sup> and a white noise term  $\sigma_r \boldsymbol{\eta}(t) \sim \mathcal{N}(0, \sigma_r^2 \Delta t)$ .

$$\mathbf{r}(t + \Delta t) - \mathbf{r}(t) = \Delta \mathbf{r}(t) \approx \Delta \mathbf{r}_1(t) + \Delta \mathbf{r}_2(t) + \sigma_r \boldsymbol{\eta}(t) \quad (6)$$

**Assumption 2.** The optimal  $\mathbf{r}(t)$  minimizes  $\mathcal{L}(t)$  such that  $p(\mathbf{r}^*(t))$  is normal around  $\mathbf{D}^\dagger \mathbf{s}(t)$ :

$$p(\mathbf{r}^*(t) | \mathbf{s}(t)) \sim \mathcal{N}(\mathbf{D}^\dagger \mathbf{s}(t), \mathbf{I} \sigma_r^2 \Delta t) \quad (7)$$

**Lemma 1.** With Assumptions 1 and 2,  $\mathcal{L}(t + \Delta t)$  is upper bounded by tracking and denoising terms:

$$\mathcal{L}(t + \Delta t) \leq \mathcal{L}_{upper}(t + \Delta t) = \|\mathbf{s}'(t) - \mathbf{D}\Delta \mathbf{r}_2(t)\| + \|\mathbf{D}\|_F \mathbb{E}_\eta \|\Delta \mathbf{r}_1(t) + \sigma_r \boldsymbol{\eta}(t - \Delta t)\|_2 + \mathbb{E}_\eta \|\mathbf{D}\sigma_r \boldsymbol{\eta}(t)\|_2 \quad (8)$$

**Assumption 3.** The optimal  $\mathbf{r}(t)$  greedily minimizes  $\mathcal{L}_{upper}$  at each time  $t$  only.

$$\{\mathbf{r}^*(t)\}_{t=0}^T = \arg \min_{\{\mathbf{r}(t)\}_{t=0}^T} \int_{t=0}^T \mathcal{L}(t + \Delta t) \Delta t \approx \left\{ \arg \min_{\mathbf{r}(t)} \mathcal{L}_{upper}(t + \Delta t) \right\}_{t=0}^T \quad (9)$$

**Theorem 1.** Given Assumption 3, the optimal  $\Delta \mathbf{r}(t)$  follows  $\mathbf{s}'(t)$  and the score function of  $p(\mathbf{r}(t))$ .

$$\Delta \mathbf{r}^*(t) = \arg \min_{\Delta \mathbf{r}_2(t)} \mathcal{L}_{upper}(t + \Delta t) + \arg \min_{\Delta \mathbf{r}_1(t)} \mathcal{L}_{upper}(t + \Delta t) + \sigma_r \boldsymbol{\eta}(t) \quad (10)$$

$$= \mathbf{D}^\dagger \mathbf{s}'(t) \Delta t + \sigma_r^2 \Delta t \nabla_{\mathbf{r}(t)} \log p(\mathbf{r}(t)) + \sigma_r \boldsymbol{\eta}(t) \quad (11)$$

During training, noisy RNNs are (unless otherwise stated, see Section 2.3) provided  $\mathbf{u}(t) = \mathbf{s}'(t)$  (or a nonlinear observation thereof) to perform path-integration and focus on denoising.

**Theorem 2.** In the absence of input (quiescence), a noisy RNN already trained to path-integrate  $\mathbf{s}(t)$  from  $\mathbf{s}'(t)$  will perform gradient ascent along the score function of  $\mathbf{r}(t)$ . If  $p(\mathbf{r}(t))$  and  $p(\mathbf{s}(t))$  are stationary, then this ascent is Langevin sampling (Equation 1). If the variance of  $\sigma_r \boldsymbol{\eta}(t)$  is scaled by a factor of 2, then  $p(\mathbf{r}(t))$  is guaranteed to converge to the steady-state distribution  $p(\mathbf{r}) = p(\mathbf{D}^\dagger \mathbf{s})$ .

$$\text{If } \mathbf{r}(t + \Delta t) = \mathbf{r}(t) + \sigma_r^2 \Delta t \nabla_{\mathbf{r}(t)} \log p(\mathbf{r}(t)) + \sqrt{2} \sigma_r \boldsymbol{\eta}(t), \text{ then } \lim_{t \rightarrow \infty} p(\mathbf{r}(t)) = p(\mathbf{D}^\dagger \mathbf{s}) \quad (12)$$

### 2.3 EXISTING METHODS OF BIASING REPLAY IN RNNs

**Neural adaptation.** Biological neurons can mitigate prolonged or low-frequency activity via negative feedback, or *adaptation* (Benda, 2021; Gutkin & Zeldenrust, 2014). This feedback has proven important for describing *in vivo* hippocampal activity and replay (Itskov et al., 2011; Levenstein et al., 2019), and in computational models of replay has been shown to encourage long replay trajectories (exploration) by preventing neural activations from getting stuck in attractor basins (Dong et al., 2021; Li et al., 2024; Levenstein et al., 2024). Like Levenstein et al. (2024), we define adaptation as negative feedback  $\mathbf{c}(t)$  added to RNN activity  $\mathbf{r}(t)$  (Equation 4) after training<sup>3</sup>:

$$\mathbf{r}(t + \Delta t) = f(\mathbf{r}(t), \mathbf{u}(t), \sigma_r \boldsymbol{\eta}(t)) - \mathbf{c}(t), \quad \Delta \mathbf{c}(t) = \frac{1}{\tau_a} (-\mathbf{c}(t) + b_a \mathbf{r}(t)) \quad (13)$$

<sup>2</sup>Each function  $\Delta \mathbf{r}_1(t)$ ,  $\Delta \mathbf{r}_2(t)$  can depend on variables beyond  $t$ , but they are omitted for concision.

<sup>3</sup>Subtraction of the moving average  $\mathbf{c}(t)$  also arises naturally from greedy minimization of  $\mathcal{L}(t) + \frac{1}{2} \|\mathbf{c}(t)\|_2^2$ .

**Masked training.** Denoisers and autoencoders benefit from *masked training*, wherein some regions of input data are set to zero before model processing (Zhang et al., 2023). Levenstein et al. (2024) introduce masked training for path-integration by periodically masking the input  $\mathbf{u}(t)$  (the observation of  $\mathbf{s}'(t)$ ): only at every  $k$ -th timestep does the RNN observe a nonzero input (Equation 14).

$$\mathbf{u}(t) = \begin{cases} \mathbf{s}'(t), & t \bmod k = 0 \\ \mathbf{0}, & \text{otherwise} \end{cases} \quad (14)$$

Levenstein et al. (2024) found that masked training makes replay sequences more coherent and makes manifolds of neural activity more similar to the spatial layout of the environment. We found that masked training improves replay stability, so we use it (with  $k \geq 3$ ) in training all RNNs.

### 3 ESTIMATING THE SCORE FUNCTION OF NOISY RNN ACTIVITY

Noisy RNNs trained to path-integrate implicitly learn the score function of their activity. Previous works have not examined the score function; in fact, they assume the distribution of RNN activity  $p(\mathbf{r}(t))$  is stationary (Theorem 2), and thus the score function  $\nabla_{\mathbf{r}(t)} \log p(\mathbf{r}(t))$  depends only on  $\mathbf{r}(t)$  (Krishna et al., 2024). However, we refute this assumption, and in doing so reveal the role of leakage in path-integration: even if the RNN path-integrates a simple Gaussian process, the score function requires information beyond  $\mathbf{r}(t)$ , which it employs through linear leakage (decay) of  $\mathbf{r}(t)$ . This linearity suggests that leakage is useful for path-integration, which we confirm experimentally.

#### 3.1 CHALLENGES IN SIMPLE DISTRIBUTIONS

First, we examine how the score function of optimal path-integrating RNN activity  $\mathbf{r}(t)$  has nonstationarities that are challenging to perfectly estimate, even for simple Gaussian processes.

**Assumption 4.** *The observed states  $\mathbf{s}(t)$  form some Gaussian process:  $p(\mathbf{s}(t)) \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{s}(t)}, \boldsymbol{\Sigma}_{\mathbf{s}(t)})$ .*

**Theorem 3.** *With Assumption 4, the score function of trained activity  $\mathbf{r}(t)$  has a closed form which, while linear in  $\mathbf{r}(t)$ , is nonlinear with respect to the parameters of  $p(\mathbf{s}(t))$  (see Appendix C.2):*

$$\sigma_r^2 \Delta t \nabla_{\mathbf{r}(t)} \log p(\mathbf{r}(t)) = - \underbrace{\sigma_r^2 \Delta t (\mathbf{I} \sigma_r^2 \Delta t + \mathbf{D}^\dagger \boldsymbol{\Sigma}_{\mathbf{s}(t)} (\mathbf{D}^\dagger)^T)^{-1}}_{\boldsymbol{\Lambda}(t)} (\mathbf{r}(t) - \mathbf{D}^\dagger \boldsymbol{\mu}_{\mathbf{s}(t)}) \quad (15)$$

With  $\boldsymbol{\Lambda}(t)$  as the *leakage matrix* of  $\mathbf{r}(t)$ , we can already gain some insight from Equation 15.

**Remark 1.** If  $p(\mathbf{s}(t))$  is Gaussian, then the score of  $p(\mathbf{r}(t))$  is simply a linear function of  $\mathbf{r}(t)$ , but its parameters are nonlinear functions of time, and are only as stationary as  $p(\mathbf{s}(t))$ .

**Remark 2.** The eigenvalues of the leakage matrix  $\boldsymbol{\Lambda}(t)$  are always between 0 and 1 (see Appendix C.2). Moreover, the eigenvalues of  $\boldsymbol{\Lambda}(t)$  and  $\mathbf{D}^\dagger \boldsymbol{\Sigma}_{\mathbf{s}(t)} (\mathbf{D}^\dagger)^T$  inversely correlate: strong decay of  $\mathbf{r}(t)$  implies weak noise  $\boldsymbol{\Sigma}_{\mathbf{r}(t)}$ , and weak decay of  $\mathbf{r}(t)$  implies strong noise  $\boldsymbol{\Sigma}_{\mathbf{r}(t)}$ .

To further illustrate how estimating the score function of  $\mathbf{r}(t)$  (Equation 15) can be challenging, we now examine a scalar RNN  $r_{ou}(t)$  trained to path-integrate Ornstein-Uhlenbeck processes.

**Assumption 5.** *The observed states follow a scalar Ornstein-Uhlenbeck process parameterized by  $\theta, \mu, \sigma_s$ :  $s'_{ou}(t) = \theta(\mu - s_{ou}(t)) + \sigma_s \eta(t)$ , where  $p(s(0)) \sim \mathcal{N}(0, \sigma_0^2)$ . In other words,  $p(s_{ou}(t))$  parameterizes a directed random walk from  $s_{ou}(0)$  to  $\mu$ .*

**Remark 3.** The score function of the optimal  $r_{ou}(t)$  under Assumption 5 follows from Equation (27) (see Appendix D):

$$\sigma_r^2 \Delta t \nabla_{r_{ou}(t)} \log p(r_{ou}(t)) = \frac{-\sigma_r^2 \Delta t (r_{ou}(t) - \mu (1 - e^{-\theta t}))}{\sigma_r^2 \Delta t + \frac{\sigma_s^2}{2\theta} (1 - e^{-2\theta t}) + \sigma_0^2 e^{-\theta t}} \quad (16)$$

The score function, and thus the optimal quiescent activity, is evidently complex and nonstationary:  $\lim_{t \rightarrow 0} \sigma_r^2 \Delta t \nabla_{r_{ou}(t)} \log p(r_{ou}(t)) = -\frac{\sigma_r^2 \Delta t}{\sigma_r^2 \Delta t + \sigma_0^2} r_{ou}(t)$ , while  $\lim_{t \rightarrow \infty} \sigma_r^2 \Delta t \nabla_{r_{ou}(t)} \log p(r_{ou}(t)) = -\frac{\sigma_r^2 \Delta t}{\sigma_r^2 \Delta t + \sigma_s^2 / 2\theta} (r_{ou}(t) - \mu)$ . While one could force

stationarity by implicitly assuming the process starts at  $s_{ou}(0) = \mu$ , or assuming the steady-state dynamics ( $t \rightarrow \infty$ ) are the most important, we argue that any such approach would miss a fundamentally relevant aspect of the Ornstein-Uhlenbeck process for navigation: *intention*. Unlike the Wiener process (Appendix D), the Ornstein-Uhlenbeck process can describe a random walk that *intentionally* navigates from  $s_u(0)$  to  $\mu$ , rather than one that simply wanders around  $\mu$ . Thus, for navigation, the non-stationary, or “early”, dynamics of the Ornstein-Uhlenbeck process are the most salient. Given the relevance of non-stationary dynamics for navigation, our analyses focus on the entire course of replay dynamics (rather than steady-state distributions), examining properties like speed and path diversity (exploration).

### 3.2 THE ADVANTAGE OF LEAKAGE

Theorem 1 and Remark 1 suggest that linear leakage may be a useful inductive bias for RNNs learning to path-integrate: the score function for a Gaussian process is linear with respect to  $\mathbf{r}(t)$ , although the parameters of said linearity are nonlinear functions of time. We examine the utility of leakage by comparing two RNNs trained to path-integrate:  $\mathbf{r}(t + \Delta t) = \kappa\mathbf{r}(t) + f_1(\mathbf{r}(t), \mathbf{u}(t), \sigma_r\boldsymbol{\eta}(t))$  (RNN with leakage  $0 < \kappa < 1$ ) and  $\mathbf{r}(t + \Delta t) = f_2(\mathbf{r}(t), \mathbf{u}(t), \sigma_r\boldsymbol{\eta}(t))$  (RNN without leakage), where  $f_1, f_2$  are shallow ReLU layers trained separately (see Appendix A.5). The second is more reminiscent of traditional RNNs in machine learning (e.g., ReLU or gated RNNs), which do not typically employ leakage. In Figure 2, we show that leakage is useful for path-integration, especially with masked training ( $k > 1$ ). Our results suggest that leakage is a useful component of path-integrating RNNs, although other architectures may be able to successfully learn without it (e.g., the layer-norm RNN of Levenstein et al. (2024)).

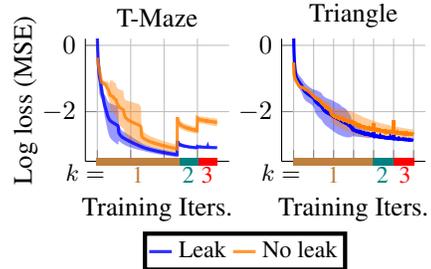


Figure 2: **Leakage helps path-integration.** Here we train RNNs on two tasks, ablating the leakage term. Leakage helps training, especially when losses increase with the masking difficulty  $k$ . Means are solid, standard deviations are faint.

## 4 SECOND-ORDER LANGEVIN SAMPLING FOR NEURAL REPLAY

In the previous section, we examined the score function of trained path-integrating RNN activity. Now, we are assuming that trained RNNs have well-estimated the score function, and follow it during quiescent (internally-driven) activity to generate replay via Langevin dynamics. Here we ask how modulating the dynamics of such RNNs affects the distribution of replay—in other words, how changing the RNN dynamics biases the distribution of replay. We see that adaptation (negative feedback) incurs a non-ideal form of second-order Langevin sampling, so we propose a complementary alternative that explicitly performs underdamped second-order sampling.

### 4.1 ADAPTATION AS UNDERDAMPED LANGEVIN DYNAMICS

First we examine adaptation (negative feedback), as defined in Section 2.3.

**Proposition 1.** *Adding adaptation (Equation (13)) to a trained path-integrating RNN during quiescence (Equation (12)) incurs Langevin sampling with negative feedback:*

$$\Delta\mathbf{r}(t) = \sigma_r^2\Delta t \nabla_{\mathbf{r}(t)} \log p(\mathbf{r}(t)) + \sqrt{2}\sigma_r\boldsymbol{\eta}(t) - \mathbf{c}(t), \quad \Delta\mathbf{c}(t) = \frac{1}{\tau_a}(-\mathbf{c}(t) + b_a\mathbf{r}(t)) \quad (17)$$

For the clearest illustration of the effects of adaptation, let us examine a stationary  $p(\mathbf{r}(t))$ —a simplification which we argued in Section 3.1 is not realistic, but is nonetheless intuitive.

**Assumption 6.** *The observed states are normal (Assumption 4) and stationary:  $p(\mathbf{s}(t)) \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ .*

**Theorem 4.** *Adding adaptation to an RNN trained to path-integrate states drawn from a stationary Gaussian distribution (Assumption 6 and Equation (17)) produces the following second-order*

stochastic dynamics during quiescence (see Appendix E):

$$\begin{aligned} \mathbf{r}''(t) = & \left( \frac{b_a}{\tau_a} \mathbf{I} + \sigma_r^2 \Delta t \frac{d^2}{d\mathbf{r}(t)^2} \log p(\mathbf{r}(t)) \right) \mathbf{r}'(t) - \frac{b_a}{\tau_a} \sigma_r^2 \Delta t \nabla_{\mathbf{r}(t)} \log p(\mathbf{r}(t)) + \frac{1}{\tau_a} \mathbf{r}(t) \\ & - \sigma \frac{b_a}{\tau_a} \boldsymbol{\eta}(t) + \sigma \boldsymbol{\eta}'(t) \end{aligned} \quad (18)$$

Comparing Equation 18 with Equation 2, the two indeed resemble each other: adaptation seems to induce a form of underdamped Langevin dynamics. This may help explain the observed utility of adaptation for generating replay (Itskov et al., 2011; Levenstein et al., 2019; 2024). Moreover, since we established in Section 3 that the score function of even a basic stochastic process is difficult to estimate, the effectiveness of underdamped Langevin sampling for working with noisy gradients (Cheng et al., 2018) may be useful in realistic settings where  $\nabla_{\mathbf{r}(t)} \log p(\mathbf{r}(t))$  is poorly estimated.

However, interpreting Equation (18) as underdamped Langevin sampling reveals some shortcomings thereof as a sampling method.

**Remark 4.** The coefficient of  $\mathbf{r}'(t)$  is usually constant, and should be positive to ensure convergence (Pavliotis (2014), pg. 183), but  $\frac{b_a}{\tau_a} \mathbf{I} + \nabla_{\mathbf{r}(t)} \log p(\mathbf{r}(t))$  is not constant and could be negative.

**Remark 5.** Underdamped Langevin sampling from  $\mathbf{r}(t)$  should not have a negative sign in front of  $\nabla_{\mathbf{r}(t)} \log p(\mathbf{r}(t))$  if the intention is to maximize  $p(\mathbf{r}(t))$ .

## 4.2 REPLAY VIA EXPLICIT UNDERDAMPED LANGEVIN DYNAMICS

While adaptation is a biologically plausible mechanism of performing a variant of underdamped Langevin dynamics in RNNs, we propose an alternative method, following Equation 3, to more clearly and directly perform underdamped Langevin sampling.

**Definition 3.** We implement *explicitly underdamped* dynamics via a momentum term  $\mathbf{v}(t)$  controlled by friction  $\lambda_v$  (when  $\lambda_v = 1$ , dynamics are identical to the overdamped case from Equation (4)):

$$\mathbf{v}(t + \Delta t) = (1 - \lambda_v) \mathbf{v}(t) + \underbrace{f(\mathbf{r}(t), \mathbf{u}(t), \sigma_r \boldsymbol{\eta}(t)) - \mathbf{r}(t))}_{\Delta \mathbf{r}(t) \text{ if } \lambda_v = 1}, \quad \Delta \mathbf{r}(t) = \mathbf{v}(t + \Delta t) \quad (19)$$

Our proposed sampling method is conceptually similar to RNNs with momentum (Nguyen et al., 2020):  $\mathbf{v}(t)$  (i.e., the velocity of  $\mathbf{r}(t)$ ) accumulates previous values when friction  $\lambda_v \in [0, 1]$  is below 1<sup>4</sup>. Moreover, our sampling mechanism shares conceptual ties with synaptic facilitation, a phenomenon associated with temporal compression *in vivo* wherein inputs that successfully trigger output activity have a temporarily increased influence on output activity (Leibold et al., 2008; Thurley et al., 2008); we leave a possible implementation of underdampening via synaptic facilitation to future work. Lastly, our mechanism could relate to momentum observed in replay trajectories *in vivo* (Krause & Drugowitsch, 2022).

## 5 NUMERICAL RESULTS

Now we examine how adaptation and underdampening (Equations 13 and 19) bias replay distributions in trained path-integrating RNNs. We first see they counter each other: adaptation slows replay, while underdampening quickens it. Then we confirm that adaptation induces exploration, and show that underdampening does not prevent exploration, but rather complements it by increasing path lengths. We note here that replay trajectories are obtained by first randomly initializing the hidden state of the RNN, following which we collect activations for several unrolled timesteps in the absence of inputs but in the presence of noise. Further details are provided in Appendix A.1. Here we use ReLU and leaky ReLU activations in trained RNNs, as justified in Appendix A.2. In Appendix G we reproduce key findings with tanh RNNs.

**Experiments.** We have five tasks on which we train RNNs and then examine how incorporating sampling mechanisms post-training (adaptation strength  $b_a$  and friction  $\lambda_v$ ) affect replay (quiescent activity) statistics. For a detailed discussion of implementation, see Appendix A.

<sup>4</sup>For a comparison of  $\lambda_v$  and the  $\gamma$  term from Equations 3 and 2, see Appendix A.6.

1. 1D Ornstein-Uhlenbeck (OU) process (Figure 1): the optimal (with respect to Equation (8)) path-integrating RNN has a closed form (Equation (16)), which we use in lieu of a trained RNN.
2. 2D T-maze and triangle (Figure 3): we simulate each direction (two in T-maze, six in triangle) as a 2D OU process, and train a ReLU RNN to integrate 2D paths from every direction.
3. Rat trajectories (Appendix F, Figure 8): like Krishna et al. (2024), we simulate 512 place cells that encode 2D directed (biased) or undirected (unbiased) random walks from RatInABox (George et al., 2024), train a ReLU RNN to path-integrate place cell activity, and decode activity in 2D.

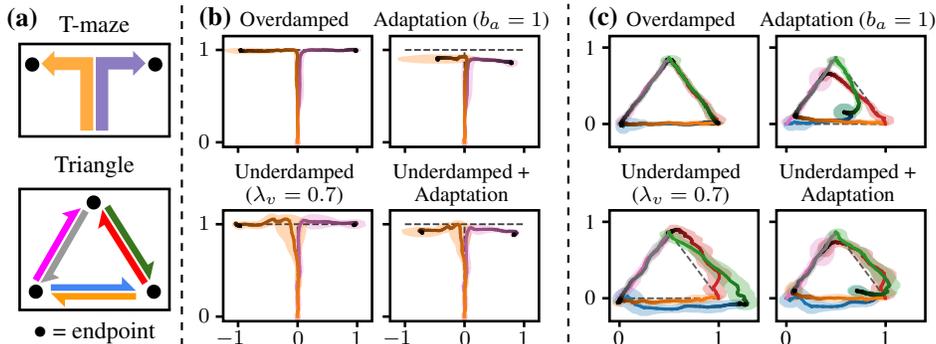


Figure 3: **Underdamping and adaptation counter each other.** Here we show replay from RNNs trained to path-integrate in T-maze or triangular environments. (a) Awake paths in each task form a mixture of Ornstein-Uhlenbeck processes, one for each direction of travel. Awake paths reach their endpoints and stay there. In (b) and (c) are mean replay paths (darkening over time) from T-maze and triangle tasks, simulated for the same time as awake paths. Standard deviations are faint, ideal path means are dashed. Like in Figure 1, adaptation slows convergence towards endpoints, underdamping quickens it. Also, the two mechanisms induce deviations that negate each other.

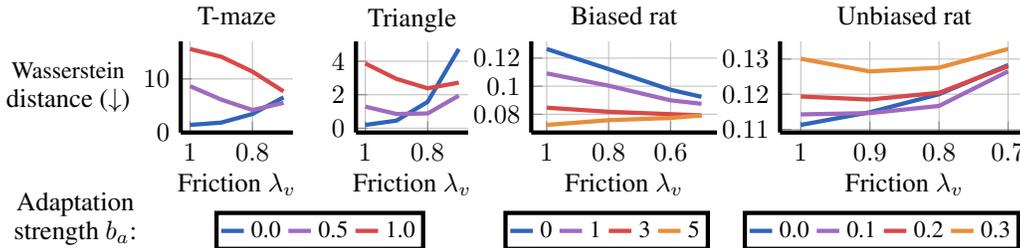


Figure 4: **Underdamping improves replay fidelity in the presence of adaptation.** We compute the Wasserstein distance (dissimilarity) between awake and replay path distributions ( $p(\{s(t)\}_{t=0}^T)$  and  $p(\{r(t)\}_{t=0}^T)$ ), varying friction and adaptation strength (see Appendix A for details). While the two mechanisms both generally increase this distance, underdamping ( $\lambda_v < 1$ ) decreases it if adaptation is nonzero. Like in Figure 3, underdamping counters adaptation-induced deviations.

**Underdamping and adaptation seem to counter each other.** We initially observe that the two modifiers counteract each other: in Figures 1 and 3, adaptation repels trajectories away from attractors (such as endpoints), while underdamping accelerates trajectories towards them. This makes sense given Remark 5 and the nature of underdamped sampling. Then we measure how the two modifiers affect the similarity between replay (internally-driven) and awake (observed) trajectories. Figure 4 shows they both generally decrease the similarity to awake trajectories, but they do counter each other insofar as underdamping, in the presence of adaptation, increases similarity to awake trajectories. Thus far, underdamping counters adaptation qualitatively and statistically.

**Adaptation slows replay, underdamping accelerates it.** Next, we examine a key component of replay: speed. In Figures 4 and 5, we generally see that adaptation slows replay or increases the

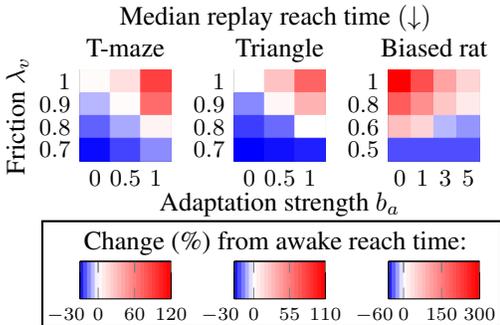


Figure 5: **Underdamping temporally compresses replay.** We calculate how long it takes awake and replay trajectories to reach their endpoints. Underdamping ( $\lambda_v < 1$ ) not only shortens this reach time, but makes it smaller than that of awake paths, *temporally compressing* awake activity. See Appendix F Figure 9 for mean reach times. We do not include unbiased rat trajectories because they do not have defined endpoints, but we do confirm in Appendix F Figure 10 that underdamping quickens them.

dissimilarity between replay and awake trajectories. The exception is the biased rat trajectories, since they are the only task where all trajectories reach the same common goal. Such path distributions resemble Ornstein-Uhlenbeck processes with steady-state mean  $\mu = 0$ : in such cases, adaptation actually accelerates convergence towards the steady-state. In Figure 5, we see that underdamping increases replay speed, performing temporal compression relative to awake activity.

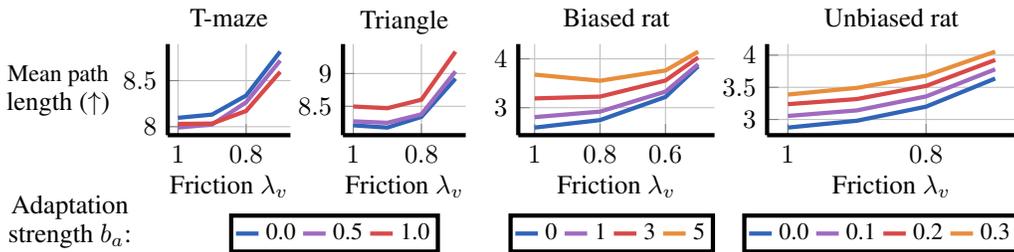


Figure 6: **On average, underdamping increases exploration via path length.** In here and Figure 7 we simulate replay paths for more time than awake paths, varying friction and adaptation strength. Shown above, underdamping ( $\lambda_v < 1$ ) increases the average length of replay paths. It may do so via increased replay speeds (Appendix F Figure 10) or via additional path transitions (Figure 7); the latter is how adaptation increases path length.

**Underdamping complements exploration from adaptation.** We have established that underdamping counters adaptation, both qualitatively and in terms of speed. However, we do not want to merely propose a mechanism that undoes adaptation. Adaptation can induce exploration, i.e., prevent replay trajectories from getting stuck in attractors. We find that underdamping does not on average prevent adaptation-induced exploration. In fact, underdamping complements adaptation for exploration: underdamped paths travel farther (Figures 6 and 7), have more variance (Appendix F Figure 11), and generally exhibit the same, if not more, exploratory behavior as they did with only adaptation (Figure 7). Underdamping maintains exploration while counteracting the slowness from adaptation.

## 6 CONCLUSIONS AND FUTURE WORK

We have re-applied Langevin sampling theory to replay in sequential predictive learning networks, producing theoretically and confirming empirically three key insights: (1) estimating the per-timestep score function of RNN activity is challenging, but does benefit from linear leakage; (2) adaptation (negative feedback) is a variant of underdamped Langevin sampling that encourages exploration (as shown in prior works) but also slows replay; (3) our new underdamping mechanism (momentum) temporally compresses replay while also increasing (on average) exploration from adaptation. These findings improve our understanding of biological neural networks that produce replay, like the hippocampus. Our proposed underdamping mechanism via momentum could be tied to short-term facilitation (which can be probed experimentally), connected to specific subregions in the hippocampus (like Chen et al. (2024)), or refined through insights from existing RNNs with momentum (Nguyen et al., 2020). Future efforts could confirm our findings in more complex environments Wood

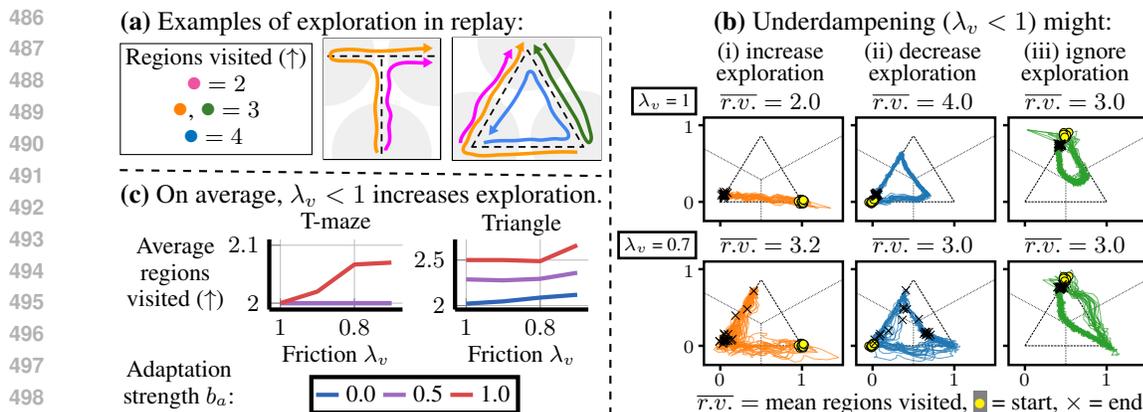


Figure 7: **On average, underdamping increases exploration via regions visited.** Another relevant aspect of exploration is transitions not present in awake activity. **(a)** Here we quantify these transitions via *regions visited*. Each region is an area where every point is closest to the same endpoint. In our T-maze and triangle tasks, awake paths go from starting points to endpoints and stay there, visiting two regions total (starting and ending). Meanwhile, adaptation ( $b_a = 1$ ) can make replay paths visit multiple endpoints (regions visited  $> 2$ ). **(b)** When we use adaptation and underdamping ( $\lambda_v < 1$ ), we see that replay paths might visit (i) more, (ii) fewer, or (iii) the same regions (see Appendix F Figure 12 for more details). **(c)** However, on average, across several different trained models, underdamping increases regions visited, thus increasing adaptation.

et al. (2018); Levenstein et al. (2024). Like preceding works, our network models are rate-based, but extending our work to spike-based models of replay and sequential predictive learning (Saponati & Vinck, 2023; Asabuki & Fukai, 2025; Bono et al., 2023) would be interesting (this may involve extending Langevin sampling to inhomogeneous Poisson processes). There still remain several avenues for connecting Langevin sampling to neural replay.

## REPRODUCIBILITY STATEMENT

For our theoretical contributions, i.e., Theorems 3 and 4, we provide proofs in Appendix C and E (proofs for Theorems 1 and 2 may be found in Section 2 of Krishna et al. (2024)). Details on our numerical experiments, including associated hyperparameters and implementation details are provided in Appendix A. Our code has been submitted as Supplementary Material (see README.md for instructions) and will be made publicly available upon publication.

## REFERENCES

- Fraser Aitken and Peter Kok. Hippocampal representations switch from errors to predictions during acquisition of predictive associations. *Nature Communications*, 13(1), 2022.
- Sina Alemohammad, Josue Casco-Rodriguez, Lorenzo Luzi, Ahmed Imtiaz Humayun, Hossein Babaei, Daniel LeJeune, Ali Siahkoochi, and Richard G Baraniuk. Self-consuming generative models go MAD. In *International Conference on Learning Representations (ICLR)*, 2024.
- Abdullahi Ali, Nasir Ahmad, Elgar de Groot, Marcel Antonius Johannes van Gerven, and Tim Christian Kietzmann. Predictive coding is a consequence of energy efficiency in recurrent neural networks. *Patterns*, 3(12), 2022.
- Toshitake Asabuki and Tomoki Fukai. Predictive learning rules generate a cortical-like replay of probabilistic sensory experiences. *eLife*, 13, 2025.
- Amir H Azizi, Laurenz Wiskott, and Sen Cheng. A computational model for preplay in the hippocampus. *Frontiers in computational neuroscience*, 7:161, 2013.

- 540 Francesco P Battaglia and Alessandro Treves. Attractor neural networks storing multiple space  
541 representations: a model for hippocampal place fields. *Physical Review E*, 58(6), 1998.  
542
- 543 Jan Benda. Neural adaptation. *Current Biology*, 31(3), 2021.
- 544 Julian E Besag. Discussion of “Representations of knowledge in complex systems” by Ulf Grenander  
545 and Michael I Miller. *Journal of the Royal Statistics Society B*, 56(4):591–592, 1994.  
546
- 547 Nicolas Bonneel, Julien Rabin, Gabriel Peyré, and Hanspeter Pfister. Sliced and radon wasserstein  
548 barycenters of measures. *Journal of Mathematical Imaging and Vision*, 51(1), 2015.  
549
- 550 Jacopo Bono, Sara Zannone, Victor Pedrosa, and Claudia Clopath. Learning predictive cognitive  
551 maps with spiking neurons during behavior and replays. *Elife*, 12, 2023.
- 552 Neil Burgess, Michael Recce, and John O’Keefe. A model of hippocampal function. *Neural networks*,  
553 7(6-7), 1994.
- 554 György Buzsáki. Hippocampal sharp waves: their origin and significance. *Brain research*, 398(2),  
555 1986.  
556
- 557 György Buzsáki. Two-stage model of memory trace formation: a role for “noisy” brain states.  
558 *Neuroscience*, 31(3), 1989.  
559
- 560 György Buzsáki. Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and  
561 planning. *Hippocampus*, 25(10), 2015.
- 562 Yusi Chen, Huanqiu Zhang, Mia Cameron, and Terrence Sejnowski. Predictive sequence learning in  
563 the hippocampal formation. *Neuron*, 112(15), 2024.  
564
- 565 Xiang Cheng, Niladri S. Chatterji, Peter L. Bartlett, and Michael I. Jordan. Underdamped Langevin  
566 MCMC: A non-asymptotic analysis. In *Proceedings of the 31st Conference On Learning Theory*,  
567 volume 75, 2018.
- 568 Selmaan N Chettih, Emily L Mackevicius, Stephanie Hale, and Dmitriy Aronov. Barcoding of  
569 episodic memories in the hippocampus of a food-caching bird. *Cell*, 187(8), 2024.  
570
- 571 Elizabeth R. Chrastil. Human path integration and the neural underpinnings. In *Encyclopedia of*  
572 *the Human Brain (Second Edition)*, pp. 157–170. Elsevier, 2025. ISBN 978-0-12-820481-8. doi:  
573 <https://doi.org/10.1016/B978-0-12-820480-1.00016-4>.
- 574 Tianhao Chu, Zilong Ji, Junfeng Zuo, Yuanyuan Mi, Wen-hao Zhang, Tiejun Huang, Daniel Bush,  
575 Neil Burgess, and Si Wu. Firing rate adaptation affords place cell theta sweeps, phase precession  
576 and procession. *bioRxiv*, 2024.  
577
- 578 Mark M Churchland and Krishna V Shenoy. Preparatory activity and the expansive null-space. *Nature*  
579 *Reviews Neuroscience*, 25(4), 2024.
- 580 Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in  
581 vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9), 2023.  
582
- 583 Christopher J Cueva and Xue-Xin Wei. Emergence of grid-like representations by training recur-  
584 rent neural networks to perform spatial localization. In *International Conference on Learning*  
585 *Representations*, 2018.
- 586 Christopher J Cueva, Peter Y Wang, Matthew Chin, and Xue-Xin Wei. Emergence of functional and  
587 structural properties of the head direction system by optimization of recurrent neural networks. In  
588 *International Conference on Learning Representations*, 2020.  
589
- 590 Lila Davachi and Sarah DuBrow. How the hippocampus preserves order: the role of prediction and  
591 context. *Trends in cognitive sciences*, 19(2), 2015.  
592
- 593 Thomas J Davidson, Fabian Kloosterman, and Matthew A Wilson. Hippocampal replay of extended  
experience. *Neuron*, 63(4), 2009.

- 594 Nicolas Deperrois, Mihai A Petrovici, Walter Senn, and Jakob Jordan. Learning cortical representa-  
595 tions through perturbed and adversarial dreaming. *eLife*, 11:e76384, 2022. ISSN 2050-084X.  
596
- 597 Xingsi Dong, Tianhao Chu, Tiejun Huang, Zilong Ji, and Si Wu. Noisy adaptation generates Lévy  
598 flights in attractor neural networks. *Advances in Neural Information Processing Systems (NeurIPS)*,  
599 34, 2021.
- 600 David R Euston, Masami Tatsuno, and Bruce L McNaughton. Fast-forward playback of recent  
601 memory sequences in prefrontal cortex during sleep. *Science*, 318(5853):1147–1150, 2007.  
602
- 603 Usman Farooq and George Dragoi. Emergence of preconfigured and plastic time-compressed  
604 sequences in early postnatal development. *Science*, 363(6423), 2019.
- 605 David J Foster. Replay comes of age. *Annual Review of Neuroscience*, 40:581–602, 2017. ISSN  
606 1545-4126.
- 607 Richard J Gardner, Erik Hermansen, Marius Pachitariu, Yoram Burak, Nils A Baas, Benjamin A  
608 Dunn, May-Britt Moser, and Edvard I Moser. Toroidal topology of population activity in grid cells.  
609 *Nature*, 602(7895):123–128, 2022. ISSN 1476-4687.  
610
- 611 Tom M George, Mehul Rastogi, William de Cothi, Claudia Clopath, Kimberly Stachenfeld, and  
612 Caswell Barry. RatInABox, a toolkit for modelling locomotion and neuronal activity in continuous  
613 environments. *eLife*, 13:e85274, 2024. ISSN 2050-084X.
- 614 Roberto Gozalo-Brizuela and Eduardo C Garrido-Merchan. ChatGPT is not all you need. a state of  
615 the art review of large generative AI models. *arXiv preprint arXiv:2301.04655*, 2023.  
616
- 617 Boris Gutkin and Fleur Zeldenrust. Spike frequency adaptation. *Scholarpedia*, 9(2), 2014.  
618
- 619 Tatsuya Haga and Tomoki Fukai. Recurrent network model for learning goal-directed sequences  
620 through reverse replay. *eLife*, 7:e34171, 2018. ISSN 2050-084X.
- 621 Kari L Hoffman and Bruce L McNaughton. Coordinated reactivation of distributed memory traces in  
622 primate neocortex. *Science*, 297(5589):2070–2073, 2002.
- 623 John J Hopfield. Neurodynamics of mental exploration. *Proceedings of the National Academy of  
624 Sciences*, 107(4):1648–1653, 2010.  
625
- 626 Vladimir Itskov, Carina Curto, Eva Pastalkova, and György Buzsáki. Cell assembly sequences arising  
627 from spike threshold adaptation keep track of time in the hippocampus. *Journal of Neuroscience*,  
628 31(8), 2011.
- 629 Jorge Jaramillo and Richard Kempter. Phase precession: a neural code underlying episodic memory?  
630 *Current Opinion in Neurobiology*, 43, 2017.  
631
- 632 Adam Johnson and A David Redish. Neural ensembles in CA3 transiently encode paths forward of  
633 the animal at a decision point. *Journal of Neuroscience*, 27(45), 2007.
- 634 Louis Kang and Michael R DeWeese. Replay as wavefronts and theta sequences as bump oscillations  
635 in a grid cell attractor network. *eLife*, 8:e46351, 2019. ISSN 2050-084X.  
636
- 637 Kenneth Kay, Jason E Chung, Marielena Sosa, Jonathan S Schor, Mattias P Karlsson, Margaret C  
638 Larkin, Daniel F Liu, and Loren M Frank. Constant sub-second cycling between representations  
639 of possible futures in the hippocampus. *Cell*, 180(3), 2020.
- 640 Tal Kenet, Dmitri Bibitchkov, Misha Tsodyks, Amiram Grinvald, and Amos Arieli. Spontaneously  
641 emerging cortical representations of visual attributes. *Nature*, 425(6961):954–956, 2003. ISSN  
642 1476-4687.
- 643 Emma L Krause and Jan Drugowitsch. A large majority of awake hippocampal sharp-wave ripples  
644 feature spatial trajectories with momentum. *Neuron*, 110(4), 2022.  
645
- 646 Nanda H Krishna, Colin Bredenber, Daniel Levenstein, Blake Aaron Richards, and Guillaume  
647 Lajoie. Sufficient conditions for offline reactivation in recurrent neural networks. In *International  
Conference on Learning Representations (ICLR)*, 2024.

- 648 Albert K Lee and Matthew A Wilson. Memory of sequential experience in the hippocampus during  
649 slow wave sleep. *Neuron*, 36(6):1183–1194, 2002. ISSN 0896-6273.
- 650
- 651 Christian Leibold, Anja Gundlfinger, Robert Schmidt, Kay Thurley, Dietmar Schmitz, and Richard  
652 Kempter. Temporal compression mediated by short-term synaptic plasticity. *Proceedings of the*  
653 *National Academy of Sciences*, 105(11), 2008.
- 654 Daniel Levenstein, György Buzsáki, and John Rinzel. NREM sleep in the rodent neocortex and  
655 hippocampus reflects excitable dynamics. *Nature Communications*, 10(1), 2019.
- 656
- 657 Daniel Levenstein, Aleksei Efremov, Roy Henha Eyono, Adrien Peyrache, and Blake Richards.  
658 Sequential predictive learning is a unifying theory for hippocampal representation and replay.  
659 *bioRxiv*, 2024.
- 660 William B Levy. A computational approach to hippocampal function. In *Psychology of learning and*  
661 *motivation*, volume 23. Elsevier, 1989.
- 662
- 663 Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. Diffusion-  
664 LM improves controllable text generation. *Advances in Neural Information Processing Systems*  
665 (*NeurIPS*), 35, 2022.
- 666
- 667 Yujun Li, Tianhao Chu, and Si Wu. Dynamics of adaptive continuous attractor neural networks.  
668 *arXiv preprint arXiv:2410.06517*, 2024.
- 669 Ashok Litwin-Kumar and Brent Doiron. Formation and maintenance of neuronal assemblies through  
670 synaptic plasticity. *Nature Communications*, 5(1):5319, 2014. ISSN 2041-1723.
- 671
- 672 Lorenzo Luzi, Paul M Mayer, Josue Casco-Rodriguez, Ali Siahkoohi, and Richard Baraniuk.  
673 Boomerang: Local sampling on image manifolds using diffusion models. *Transactions on Machine*  
674 *Learning Research (TMLR)*, 2024.
- 675 Daniel C McNamee, Kimberly L Stachenfeld, Matthew M Botvinick, and Samuel J Gershman.  
676 Flexible modulation of sequence generation in the entorhinal–hippocampal system. *Nature*  
677 *Neuroscience*, 24(6), 2021.
- 678
- 679 Bruce L McNaughton, Carol A Barnes, Jason L Gerrard, Katalin Gothard, Min W Jung, James J  
680 Knierim, H Kudrimoti, Y Qin, WE Skaggs, M Suster, et al. Deciphering the hippocampal polyglot:  
681 the hippocampus as a path integration system. *Journal of Experimental Biology*, 199(1), 1996.
- 682
- 683 Sebastian Michelmann, Bernhard P Staresina, Howard Bowman, and Simon Hanslmayr. Speed of  
684 time-compressed forward replay flexibly changes in human episodic memory. *Nature Human*  
*Behaviour*, 3(2), 2019.
- 685
- 686 Adam MP Miller, Alex D Jacob, Adam I Ramsaran, Mitchell L De Snoo, Sheena A Josselyn, and  
687 Paul W Frankland. Emergence of a predictive model in the hippocampus. *Neuron*, 111(12), 2023.
- 688
- 689 Kevin J Miller, Matthew M Botvinick, and Carlos D Brody. Dorsal hippocampus contributes to  
690 model-based planning. *Nature neuroscience*, 20(9), 2017.
- 691
- 692 Aaron D Milstein, Sarah Tran, Grace Ng, and Ivan Soltesz. Offline memory replay in recurrent  
693 neuronal networks emerges from constraints on online dynamics. *The Journal of Physiology*, 601  
694 (15):3241–3264, 2023.
- 695
- 696 Koichi Miyasawa. An empirical Bayes estimator of the mean of a normal population. *Bulletin of the*  
697 *International Statistical Institute*, 38(4):181–188, 1961.
- 698
- 699 Zoltán Nádasdy, Hajime Hirase, András Czurkó, Jozsef Csicsvari, and György Buzsáki. Replay and  
700 time compression of recurring spike sequences in the hippocampus. *Journal of Neuroscience*, 19  
701 (21):9497–9507, 1999. ISSN 0270-6474.
- 702
- 703 Tan Nguyen, Richard Baraniuk, Andrea Bertozzi, Stanley Osher, and Bao Wang. MomentumRNN:  
704 Integrating momentum into recurrent neural networks. *Advances in Neural Information Processing*  
*Systems (NeurIPS)*, 33, 2020.

- 702 John O’Keefe and Lynn Nadel. *The hippocampus as a cognitive map*. Oxford University Press, 1978.  
703
- 704 Victor M Panaretos and Yoav Zemel. Statistical aspects of wasserstein distances. *Annual Review of*  
705 *Statistics and its Application*, 6(1), 2019.
- 706 Grigorios A Pavliotis. Stochastic processes and applications. *Texts in Applied Mathematics*, 60, 2014.  
707
- 708 Adrien Peyrache, Mehdi Khamassi, Karim Benchenane, Sidney I Wiener, and Francesco P Battaglia.  
709 Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nature*  
710 *Neuroscience*, 12(7):919–926, 2009. ISSN 1546-1726.
- 711 Adrien Peyrache, Marie M Lacroix, Peter C Petersen, and György Buzsáki. Internally organized  
712 mechanisms of the head direction sense. *Nature Neuroscience*, 18(4):569–575, 2015. ISSN  
713 1546-1726.
- 714 Brad E Pfeiffer. The content of hippocampal “replay”. *Hippocampus*, 30(1), 2020.
- 715 Brad E Pfeiffer and David J Foster. Hippocampal place-cell sequences depict future paths to  
716 remembered goals. *Nature*, 497(7447), 2013.
- 717 Stefano Recanatesi, Matthew Farrell, Guillaume Lajoie, Sophie Deneve, Mattia Rigotti, and Eric  
718 Shea-Brown. Predictive learning as a network mechanism for extracting low-dimensional latent  
719 space representations. *Nature Communications*, 12(1), 2021.
- 720 Alexei Samsonovich and Bruce L McNaughton. Path integration and cognitive mapping in a  
721 continuous attractor neural network model. *Journal of Neuroscience*, 17(15), 1997.
- 722 Matteo Saponati and Martin Vinck. Sequence anticipation and spike-timing-dependent plasticity  
723 emerge from a predictive learning rule. *Nature Communications*, 14(1), 2023.
- 724 Johannes D Seelig and Vivek Jayaraman. Neural dynamics for landmark orientation and angular path  
725 integration. *Nature*, 521(7551), 2015.
- 726 Bin Shen and Bruce L McNaughton. Modeling the spontaneous reactivation of experience-specific  
727 hippocampal cell assemblies during sleep. *Hippocampus*, 6(6):685–692, 1996.
- 728 William E Skaggs and Bruce L McNaughton. Replay of neuronal firing sequences in rat hippocampus  
729 during sleep following spatial experience. *Science*, 271(5257), 1996.
- 730 Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution.  
731 In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, 2019.
- 732 Ben Sorscher, Gabriel Mel, Surya Ganguli, and Samuel Ocko. A unified theory for the origin of  
733 grid cells through the lens of pattern formation. In Hanna M Wallach, Hugo Larochelle, Alina  
734 Beygelzimer, Florence d’Alché Buc, Emily B Fox, and Roman Garnett (eds.), *Advances in Neural*  
735 *Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- 740 Kimberly L Stachenfeld, Matthew M Botvinick, and Samuel J Gershman. The hippocampus as a  
741 predictive map. *Nature neuroscience*, 20(11), 2017.
- 742 Federico Stella, Peter Baracs, Joseph O’Neill, and Jozsef Csicsvari. Hippocampal reactivation of  
743 random trajectories resembling brownian diffusion. *Neuron*, 102(2), 2019.
- 744 Mufeng Tang, Helen Barron, and Rafal Bogacz. Sequential memory with temporal predictive coding.  
745 In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, 2023.
- 746 Mufeng Tang, Helen Barron, and Rafal Bogacz. Learning grid cells by predictive coding. *arXiv*  
747 *preprint arXiv:2410.01022*, 2024.
- 748 Panagiota Theodoni, Bernat Rovira, Yingxue Wang, and Alex Roxin. Theta-modulation drives the  
749 emergence of connectivity patterns underlying replay in a network model of place cells. *eLife*, 7:  
750 e37388, 2018. ISSN 2050-084X.
- 751 Kay Thurley, Christian Leibold, Anja Gundlfinger, Dietmar Schmitz, and Richard Kempster. Phase  
752 precession through synaptic facilitation. *Neural Computation*, 20(5), 2008.

- 756 David Tingley and Adrien Peyrache. On the methods for reactivation and replay analysis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 375(1799):20190231, 2020.
- 757  
758
- 759 Giulio Tononi and Chiara Cirelli. Sleep and the price of plasticity: from synaptic and cellular  
760 homeostasis to memory consolidation and integration. *Neuron*, 81(1), 2014.
- 761 Benigno Uribe, Borja Ibarz, Andrea Banino, Vinicius Zambaldi, Dharshan Kumaran, Demis Hassabis,  
762 Caswell Barry, and Charles Blundell. A model of egocentric to allocentric understanding in  
763 mammalian brains. *bioRxiv*, 2022.
- 764 Ruth A Wood, Marius Bauza, Julija Krupic, Stephen Burton, Andrea Delekate, Dennis Chan, and  
765 John O’Keefe. The honeycomb maze provides a novel test to study hippocampal-dependent spatial  
766 navigation. *Nature*, 554(7690), 2018.
- 767 Dehong Xu, Ruiqi Gao, Wenhao Zhang, Xue-Xin Wei, and Ying Nian Wu. On conformal isometry  
768 of grid cells: Learning distance-preserving position embedding. In *International Conference on  
769 Learning Representations (ICLR)*, 2025.
- 770 Shengjin Xu, Wanchen Jiang, Mu-ming Poo, and Yang Dan. Activity recall in a visual cortical  
771 ensemble. *Nature Neuroscience*, 15(3):449–455, 2012. ISSN 1546-1726.
- 772 Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang,  
773 Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and  
774 applications. *ACM Computing Surveys*, 56(4), 2023.
- 775 Chaoning Zhang, Chenshuang Zhang, Junha Song, John Seon Keun Yi, and In So Kweon. A survey  
776 on masked autoencoder for visual self-supervised learning. In *IJCAI*, 2023.
- 777 Xiaohan Zhang, Xiaoyang Long, Sheng-Jia Zhang, and Zhe Sage Chen. Excitatory-inhibitory  
778 recurrent dynamics produce robust visual grids and stable attractors. *Cell Reports*, 41(11), 2022.
- 779 H Freyja Ólafsdóttir, Caswell Barry, Aman B Saleem, Demis Hassabis, and Hugo J Spiers. Hippocam-  
780 pal place cells construct reward related sequences through unexplored space. *eLife*, 4:e06063,  
781 2015. ISSN 2050-084X.
- 782 H Freyja Ólafsdóttir, Daniel Bush, and Caswell Barry. The role of hippocampal replay in memory  
783 and planning. *Current Biology*, 28(1):R37–R50, 2018. ISSN 0960-9822.
- 784  
785  
786  
787  
788  
789  
790  
791  
792  
793  
794  
795  
796  
797  
798  
799  
800  
801  
802  
803  
804  
805  
806  
807  
808  
809

810 SUPPLEMENTARY MATERIAL

811  
812 A METHODS

813  
814 A.1 EXPERIMENTS

815  
816 A.1.1 ORNSTEIN-UHLENBECK

817  
818 Awake trajectories  $s(t)$  are simulated with  $\Delta t = 0.02, \sigma_s = 0.1, \sigma_0 = 0.2, \theta = 2, \mu = 5$  for  
819  $T = 100$  iterations. We used Equation (16) as the deterministic component of  $\Delta \mathbf{r}(t)$  since a  
820 1D Ornstein-Uhlenbeck process admits a closed-form expression for  $\frac{d}{dr(t)} \log p(r(t))$ . In Figure 1,  
821 underdamping corresponds to  $\lambda_v = 0.5$ , while adaptation corresponds to  $b_a = 1, \tau_a = 100$ .

822  
823 A.1.2 T-MAZE AND TRIANGLE

824  
825 **Task description.** In the 2D T-maze and triangle tasks, we simulate directed random walks  
826 (Ornstein-Uhlenbeck processes) along several directions, and train the RNN to path-integrate these  
827 walks from their velocities. In the T-maze task, there are two directions of travel: both start from  
828 the origin and go up, then one goes left and the other goes right (orange and purple, respectively,  
829 in Figure 3). Meanwhile, in the triangle task there are six directions: denoting the three corners  
830 of the equilateral triangle as  $A = (0, 0), B = (1, 0)$ , and  $C = (1/2, \sqrt{3}/2)$ , respectively, these  
831 directions are  $\overrightarrow{AB}, \overrightarrow{BC}, \overrightarrow{CA}, \overrightarrow{AC}, \overrightarrow{CB}, \overrightarrow{BA}$ , shown in Figure 3 as blue, red, gray, pink, green, and  
832 orange, respectively. Awake and replay paths last 100 timesteps, unless exploration is being measured,  
833 in which replay paths are allowed to last 400 timesteps.

834  
835 **Architecture and training.** For both tasks, we used shallow leaky-ReLU RNNs with linear output  
836 projections, as described in Appendix A.5. In the T-maze task, we used 20 hidden neurons, while  
837 in the triangle task, we used 40. We use masking difficulty  $k = 3$  in a progressive curriculum.  
838 Triangle-task RNNs are trained with  $k = 1$  for 20,000 epochs, then at  $k = 2$  and  $k = 3$  for 5,000  
839 epochs each; T-maze-task RNNs are trained likewise, but with fewer epochs (12,000 at  $k = 1$  and  
840 5,000 at  $k = 2$  and  $k = 3$  each). In each task, we train 5 RNNs with different seeds.

841  
842 **Hidden state initialization.** In each task, multiple directions start from the same point in space (for  
843 example,  $\overrightarrow{AB}$  and  $\overrightarrow{AC}$  in the triangle task), so we add onto initial hidden states some random vectors,  
844 orthogonal to the 2D output projection, specific to each direction. For example, hidden states for  $\overrightarrow{AB}$   
845 and  $\overrightarrow{AC}$  paths are both initialized to start near  $A$ , but  $\overrightarrow{AB}$  paths also start with a fixed random vector  
846  $\boldsymbol{\eta}_{AB}$  added to their initial hidden states, whereas  $\overrightarrow{AC}$  have a different fixed random vector  $\boldsymbol{\eta}_{AC}$  added  
847 to their initial hidden states. This notion of initializing hidden states in directions orthogonal to output  
848 projections has been previously discussed in computational neuroscientific contexts (Churchland &  
849 Shenoy, 2024), and recent evidence suggests that memories (which replay is similar to) have uniquely  
850 identifiable, output-orthogonal patterns of neural activation (Chettih et al., 2024).

851 A.1.3 RAT PLACE CELL TRAJECTORIES

852  
853 Our unbiased (undirected random walks within a 2D box) and biased (directed random walks that  
854 head towards the center of a 2D box) rat place cell trajectory experiments are identical to those of  
855 Krishna et al. (2024), with 5 different ReLU RNN seeds per task, but with two modifications:

- 856  
857 1. We add masked training, with difficulties  $k = 3$  in the biased task and  $k = 6$  in the unbiased task.  
858 We use a higher masking difficulty in the unbiased task to encourage longer replay trajectories; at  
859 lower values of  $k$ , unbiased replay paths are much shorter than awake trajectories. In Langevin  
860 sampling terms, we conjecture this is because the score function  $\nabla_{\mathbf{r}(t)} \log p(\mathbf{r}(t))$  is fairly weak  
861 since unbiased trajectories are uniformly distributed<sup>5</sup>.

862  
863 <sup>5</sup>This relative weakness the unbiased activity score function (especially compared to the biased activity score  
function, where trajectories all go towards the center of a 2D box) is also why we use relatively small adaptation  
strengths  $b_a$  in simulating replay from unbiased-task RNNs.

- 864 2. We use a slower, but more detailed, decoder. Krishna et al. (2024) decode position at a given  
 865 timestep as the average position of the top 3 most active place cells, which is fast but effectively  
 866 quantizes decoded positions. We use this method to initialize our positions, but then we optimize  
 867 our positions via gradient descent to minimize the mean-squared error between observed place  
 868 cell activity and the place cell activity that would correspond to these optimized positions (this  
 869 is made possible by our knowledge of the exact place cell activity function). Our procedure  
 870 takes much longer, but results in truly continuous replay trajectories, avoiding any significant  
 871 quantization.

## 872 A.2 HYPERPARAMETERS

### 873 A.2.1 ACTIVATION FUNCTION

874 In our work, we have mostly used RNNs with ReLU or leaky ReLU activation functions (except in  
 875 Appendix G); our reasons for doing so are threefold. The first is that ReLU or ReLU-like activations,  
 876 unlike tanh activations, do not always saturate, which can mitigate vanishing gradients in RNNs  
 877 (although care must be taken to avoid exploding gradients). The second is that, compared to RNNs  
 878 with multiplicative gating interactions (e.g., GRUs and LSTMs), ReLU RNNs are more biologically  
 879 plausible insofar as they can be easily interpreted as the nonnegative, sparse firing rates of spiking  
 880 networks. The third is that most other previous works in sequential predictive coding tend to use  
 881 ReLU activations, and have shown no significant differences in results when using more complex  
 882 nonlinearities:  
 883

- 884 • Ali et al. (2022), some of the first authors to report the emergence of predictive representations  
 885 emerging in RNNs, use ReLU activations.
- 886 • Krishna et al. (2024) primarily use ReLU networks, and report similar results with GRU networks.
- 887 • Levenstein et al. (2024) use ReLU activations in conjunction with layer normalization, but report  
 888 similar results when layer normalization is removed.
- 889 • Chen et al. (2024) use tanh for most results, but when they care about the emergence of sparse,  
 890 localized, nonnegative activations (i.e., place cells), they use ReLU activations.
- 891 • Sorscher et al. (2019); Xu et al. (2025), and Zhang et al. (2022) were all interested in the emergence  
 892 of grid cells in predictive representations and primarily used ReLU or ReLU-like activations;  
 893 those that tried using other nonlinearities did not observe notable differences in results.
- 894 • Tang et al. (2024), who were also interested in grid cells, use both ReLU and tanh activations, and  
 895 also report no notable differences between the two.
- 896 • Tang et al. (2023) use linear networks, and notice no significant difference when using tanh  
 897 activations.
- 898 • Tang et al. (2023) use linear networks, and notice no significant difference when using tanh  
 899 activations.

## 900 A.3 METRICS

### 901 A.3.1 WASSERSTEIN DISTANCE

902 We seek to compare distributions of trajectories. In most of our tasks, these are objects of dimension  
 903  $T \times 2$ , where  $T$  is the number of timesteps (recall that T-maze and triangle tasks are already in  
 904 2D, while rat trajectories are analyzed after being projected into 2D space). To compare such  
 905 high-dimensional distributions, we use Wasserstein distances (Panaretos & Zemel, 2019):  
 906

- 907 • In the T-maze and triangle tasks, we first calculate the Wasserstein distance between awake and  
 908 replay paths belonging to the same direction (e.g.,  $\overrightarrow{AB}$  in the triangle task). Since paths along the  
 909 same direction should resemble, if not obey, Gaussian processes, these paths should approximately  
 910 be normally distributed, and so we can use the closed-form equation for the Wasserstein distance  
 911 between two normal distributions. After computing Wasserstein distances for each direction, we  
 912 take the average as our final distance. KL divergence might have also worked in this task, but in  
 913 practice the covariance matrices of the  $(T \times 2)$ -dimensional distributions were singular.
- 914 • In the rat experiments, decoded 2D trajectories are not easily decomposed into groups of Gaussian  
 915 processes, so instead we apply sliced Wasserstein distances (Bonneel et al., 2015) to compare  
 916 the  $(T \times 2)$ -dimensional distributions. Sliced Wasserstein distances are essentially calculated by  
 917

918 taking many random projections of two distributions onto 1D, where Wasserstein distances have  
 919 a closed form. Computing KL divergences instead would have been challenging, since rat path  
 920 distributions are high-dimensional and do not admit straightforward estimations of probability  
 921 density.

### 922 A.3.2 REACH TIMES

923 In the T-maze, triangle, and biased rat tasks, awake trajectories have clearly defined endpoints  
 924 that they reach. Replay trajectories also aim for these endpoints, although second-order dynamics  
 925 modifiers like adaptation and underdampening affect how quickly or how closely replay paths reach  
 926 them. We quantify these changes by measuring the timesteps required for replay paths to get within  
 927 10% of their endpoints (for example, the timesteps required for  $\overrightarrow{AB}$  paths to get within a  $0.1|\overrightarrow{AB}|$  of  
 928  $B$ ).

### 931 A.3.3 PATH LENGTHS

932 One way we quantify exploration in replay paths is through path length. We calculate this simply as  
 933 the sum of velocity magnitudes.

### 935 A.4 REGIONS VISITED

936 Another way we quantify exploration in the T-maze and triangle tasks is through *regions visited*. As  
 937 explained in Figure 7, *regions* are portions of input space where all points are closest to the same  
 938 endpoint. In the triangle task, there are 3 endpoints (which all also act as initial points), while in the  
 939 T-maze there are 2 endpoints and 1 shared starting point, which for the purposes of region assignment  
 940 we also consider an “endpoint”. Thus, in each task there are 3 endpoints, and thus 3 regions. Awake  
 941 trajectories go from one region to another and stay there, but replay trajectories might proceed to  
 942 visit more regions. We define a “visit” as a contiguous presence within a single region for at least 10  
 943 timesteps.

### 944 A.5 DISCRETIZATION AND IMPLEMENTATION OF NOISY RNNs

945 While biological neural networks are continuous-time systems, RNNs are in practice implemented  
 946 discretely. In order to incorporate second-order processes like adaptation and underdamped sampling  
 947 in RNNs, we must discretize them, which we do as follows:

$$951 \Delta \tilde{\mathbf{r}}(t) = f(\mathbf{r}(t), \mathbf{u}(t), \sigma_r \boldsymbol{\eta}(t)) - \mathbf{r}(t) \quad (20)$$

$$952 \mathbf{c}(t + \Delta t) = \mathbf{c}(t) + \frac{1}{\tau_a} (-\mathbf{c}(t) + b_a \mathbf{r}(t)) \quad (21)$$

$$953 \mathbf{v}(t + \Delta t) = (1 - \lambda_v) \mathbf{v}(t) + \Delta \tilde{\mathbf{r}}(t) \quad (22)$$

$$954 \mathbf{r}(t + \Delta t) = \mathbf{r}(t) - \mathbf{c}(t) + \mathbf{v}(t + \Delta t) \quad (23)$$

955 We train and sample from ReLU RNNs:  $f(\mathbf{r}(t), \mathbf{u}(t), \sigma_r \boldsymbol{\eta}(t)) = \kappa \mathbf{r}(t) + \text{ReLU}(\mathbf{W}_r \mathbf{r}(t) +$   
 956  $\mathbf{W}_{in} \mathbf{u}(t) + \sigma_r \boldsymbol{\eta}(t))$  (where  $\kappa \in [0, 1]$  is also learnable) unless otherwise stated (a choice justi-  
 957 fied in Appendix A.2). In the absence of directed inputs  $\mathbf{u}(t)$ ,  $f(\mathbf{r}(t), \mathbf{u}(t), \sigma_r \boldsymbol{\eta}(t))$  should be  
 958 roughly equivalent to the noisy score function from Equation 12.

959 If the adaptation strength  $b_a = 0$ , then there is no adaptation ( $\mathbf{c}(t) = \mathbf{0}$ ). As for our *friction* term  
 960  $\lambda_v \in [0, 1]$ , if  $\lambda_v = 1$ , then there is no underdamped sampling since  $\mathbf{v}(t) = \Delta \mathbf{r}(t) + \sigma_r \boldsymbol{\eta}(t)$  no  
 961 longer accumulates previous values of  $\mathbf{v}(t)$ . In other words, the friction term  $\lambda_v$  allows for smooth  
 962 interpolation between overdamped and underdamped sampling.

963 Note that, unlike  $\mathbf{c}(t + \Delta t)$  and  $\mathbf{v}(t + \Delta t)$ ,  $\mathbf{r}(t + \Delta t)$  depends on another term calculated at  $t + \Delta t$ .  
 964 This is a symplectic Euler discretization of the second-order dynamics, which we employ for to  
 965 ensure the stability of interactions between  $\mathbf{r}(t)$  and its momentum  $\mathbf{v}(t)$ . The negative feedback  $\mathbf{c}(t)$   
 966 is much more stable, so its discretization goes unchanged.

967 Across all trained RNN experiments we use  $\tau_a = 0$ ,  $\lambda_v = 1$  for awake activity (i.e., training). When  
 968 generating replay, we always use  $\tau_a = 100$ , and either  $T = 100, 400$ , or  $500$  timesteps, depending  
 969  
 970  
 971

on whether we are measuring fidelity and speed ( $T = 100$ , like in Figures 3, 4, 5) or exploration ( $T = 400$  for T-maze and triangle,  $T = 500$  for rat tasks, like in Figures 6, 7).

## A.6 DEVIATIONS FROM TRADITIONAL LANGEVIN SAMPLING

We must note a few minor distinctions between our replay RNNs and traditional Langevin sampling.

**Stationarity.** We treat neural replay as a sequence of events, and are thus interested in the joint distribution of replay activity at all timesteps  $p(\{\mathbf{r}(t)\}_{t=1}^T)$ . This distribution, however, is not necessarily stationary across  $t$ . Stationarity would imply that  $p(\mathbf{r}(t))$  has no need or intention of traversing along any meaningful path. While a path-integrating RNN does perform gradient ascent along  $\log p(\mathbf{r}(t))$ , it does so in a piecewise manner along  $t$  rather than jointly along all  $t$  simultaneously. This variation on gradient ascent, in combination with the non-stationarity of  $p(\mathbf{r}(t))$ , means that Langevin sampling guarantees do not hold. For this reason, we make modifications to our RNNs that under stationary Langevin sampling theory might seem unprincipled.

**Underdamping.** The friction term  $\lambda_v$  applied to the velocity  $\mathbf{v}(t)$  is subtly different from  $\gamma$  in Equations 3 and 2.  $\lambda_v \in [0, 1]$  can be interpreted as trying to map  $\gamma \in [0, \infty)$  onto  $[0, 1]$ . Like  $\gamma \rightarrow \infty, \lambda_v \rightarrow 1$  removes any dependency of  $\mathbf{v}(t + \Delta t)$  on  $\mathbf{v}(t)$ .

**Noise scaling.** RNN replay, unlike Langevin sampling, is sensitive to the variance of noise used. We found that omitting the  $\sqrt{2}$  factor in front of  $\sigma_r \boldsymbol{\eta}(t)$  worked best. For this reason, we also omitted  $\lambda_v$  as a scaling term on  $\sigma_r \boldsymbol{\eta}(t)$ . If  $\lambda_v$  is allowed to modulate  $\sqrt{2} \sigma_r \boldsymbol{\eta}(t)$ , we found that it is easy to show improvements in replay fidelity when  $\lambda_v < 1$ , but we found these to come from noise scaling rather than from underdamping or momentum as mechanisms.

**Relation to diffusion models.** Throughout this work we treat RNNs as generative models using a variant of Langevin sampling. Expressive generative models, in particular *diffusion models* whose activity resembles Langevin sampling, have received much attention from the machine learning community (Alemohammad et al., 2024; Luzi et al., 2024; Yang et al., 2023; Croitoru et al., 2023; Li et al., 2022; Gozalo-Brizuela & Garrido-Merchan, 2023; Song & Ermon, 2019). The fundamental differences between our RNNs and diffusion models are twofold. The first is that each timestep of a generated replay trajectory is generated sequentially and with only one effective step along the gradient of log-likelihood, whereas a diffusion model would generate all timesteps of a path simultaneously and with many steps along the gradient of log-likelihood. The second is that diffusion models use time-varying noise levels (annealed dynamics) and are explicitly conditioned on time as an input, whereas our RNNs use the same noise level across time and are never explicitly conditioned on time.

## B OVERVIEW OF BACKGROUND

Throughout the text we make use of notions including sampling, Langevin dynamics, replay, path-integration, awake and quiescent activity, adaptation, and masked training. Here we gather these notions, which have been defined in various previous work, as follows:

1. Our work examines how trained noisy RNNs can act as generative models in the absence of inputs, which prior work including Krishna et al. (2024) consider to be models of replay in neuronal circuits. Accordingly, we define Langevin dynamics, with (Equation (1)) and without (Equations 2 and 3) momentum, as an iterative process by which an agent could attempt to sample from an unknown distribution  $p(x)$  using only knowledge of its *score function*  $\nabla \log p(x)$ , which can be learned from noisy observations drawn from the distribution. The following points provide more detail on the sampling process and why Langevin dynamics arise:
  - Biological networks such as those underlying navigation must accurately estimate environmental state variables such as position from observations of self-motion, even in the presence of intrinsic noise.
  - This noise induces a distribution over network states, i.e., activity, however, it is unknown exactly what the true noise distribution is.

- 1026
- 1027
- 1028
- 1029
- 1030
- 1031
- 1032
- 1033
- 1034
- 1035
- 1036
- 1037
- 1038
- 1039
- 1040
- 1041
- 1042
- 1043
- 1044
- 1045
- 1046
- 1047
- 1048
- 1049
- 1050
- 1051
- 1052
- 1053
- 1054
- 1055
- 1056
- 1057
- 1058
- 1059
- 1060
- 1061
- 1062
- 1063
- 1064
- 1065
- 1066
- 1067
- 1068
- 1069
- 1070
- 1071
- 1072
- 1073
- 1074
- 1075
- 1076
- 1077
- 1078
- 1079
- Accurate state estimation would thus require the network to optimally remove intrinsic noise without access to the exact parameters of the noise distribution, but given noisy network states.
  - This optimal solution can be accomplished by learning the score function of noisy network states and using it to denoise network states over time as additional inputs and noise are presented (Miyasawa, 1961).
  - When the network receives no inputs, its noise-driven dynamics still attempt to remove intrinsic noise, but the learned score function is for the distribution of waking task-like activity. This causes the network states to resemble actual samples drawn from the distribution over network states in the presence of inputs. The theoretical results show that the network dynamics during this quiescent state carry out sampling by Langevin dynamics, i.e., using the score function of the network’s activity distribution during task performance to yield quiescent network activity that resembles waking activity.
2. Biological neural circuits like the hippocampus can navigate through environments when awake, and recall navigatory episodes during rest. Accordingly, we summarize the proof from Krishna et al. (2024) that a noisy RNN trained on a biologically relevant navigation task like path-integration can act as a generative model.
- Definitions 1 and 2 simply express that our RNNs have nonlinear dynamics and internal noise, and are trained to estimate a signal  $s(t)$  by integrating observations of  $s'(t)$ . This phenomenon, known as path-integration, underlies navigation and is used to estimate one’s position from self-motion cues.
  - Assumption 1 asserts that RNN dynamics can be decomposed into additive components. Assumption 2 simply states that the optimal RNN dynamics would lead to accurate path-integration, such that the conditional probability distribution over network states (activity) given environmental states (position) would be Gaussian and identical to the additive noise in our networks. That is, the variability over network states would only be due to intrinsic noise in the system and have the same statistics. Finally, Assumption 3 assumes that the RNN dynamics are greedily optimal for path-integration. Greedy optimization is a sensible way of partitioning effort across time in path-integration: the network does the best that it can at each timestep, assuming that at each previous timestep the best possible job has been done.
  - Theorem 1 states that path-integrative RNN dynamics will use the score function of waking activity to remove intrinsic noise from the system and use inputs to update state estimates. This is a two-step greedily optimal solution for path integration in the presence of intrinsic noise: first, intrinsic noise must be removed, following which the inputs must be used to update the estimate of the current environmental state. Theorem 2 shows that these dynamics will result in Langevin sampling from the distribution of waking activity in the absence of any inputs, i.e., analogous to sleep. This means that the quiescent network dynamics sample neural states from the same distribution as waking, task-like activity, and can thus represent sequences like those during awake task performance, i.e., leading to replay.
3. Finally, we define existing methods of modulating RNN activity or training that affect RNN replay distributions: adaptation (negative feedback), which encourages diversity in replay paths, and masked training, which encourages coherence in replay.

## C SCORE FUNCTIONS OF GAUSSIAN DISTRIBUTIONS

For any matrix calculus involved, we use denominator layout.

### C.1 MULTIVARIATE GAUSSIAN DISTRIBUTION

Let's suppose  $\mathbf{r} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ . If  $\mathbf{r} \in \mathbb{R}^d$ , then:

$$p(\mathbf{r}) = \frac{1}{\sqrt{(2\pi)^d |\boldsymbol{\Sigma}|}} \exp\left(-\frac{1}{2}(\mathbf{r} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{r} - \boldsymbol{\mu})\right) \quad (24)$$

$$\log p(\mathbf{r}) \propto -\frac{1}{2}(\mathbf{r} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{r} - \boldsymbol{\mu}) \quad (25)$$

$$\nabla_{\mathbf{r}} \log p(\mathbf{r}) = -\frac{1}{2}((\boldsymbol{\Sigma}^{-1})^T + \boldsymbol{\Sigma}^{-1})(\mathbf{r} - \boldsymbol{\mu}) \quad (26)$$

$$= -\boldsymbol{\Sigma}^{-1}(\mathbf{r} - \boldsymbol{\mu}) \quad (27)$$

$$= -\sigma^{-2}(\mathbf{r} - \boldsymbol{\mu}) \text{ if } \mathbf{r} \in \mathbb{R} \quad (28)$$

### C.2 SCORE FUNCTION OF $r(t)$ FOR GAUSSIAN $s(t)$

Recall that  $p(\mathbf{r}(t)|\mathbf{s}(t)) \sim \mathcal{N}(\mathbf{D}^\dagger \mathbf{s}(t), \mathbf{I}\sigma_r^2 \Delta t)$  from Equation 7. If we suppose that  $\mathbf{s}(t)$  is normally distributed with mean and covariance  $\boldsymbol{\mu}_{\mathbf{s}(t)}, \boldsymbol{\Sigma}_{\mathbf{s}(t)}$ , then we can obtain  $p(\mathbf{r}(t))$ :

$$p(\mathbf{r}(t)) \sim \mathcal{N}(\mathbf{D}^\dagger \boldsymbol{\mu}_{\mathbf{s}(t)}, \mathbf{I}\sigma_r^2 \Delta t + \mathbf{D}^\dagger \boldsymbol{\Sigma}_{\mathbf{s}(t)} (\mathbf{D}^\dagger)^T), \quad (29)$$

which we can plug into Equation 27 to get  $\nabla_{\mathbf{r}(t)} \log p(\mathbf{r}(t))$ :

$$\nabla_{\mathbf{r}(t)} \log p(\mathbf{r}(t)) = -(\mathbf{I}\sigma_r^2 \Delta t + \mathbf{D}^\dagger \boldsymbol{\Sigma}_{\mathbf{s}(t)} (\mathbf{D}^\dagger)^T)^{-1} (\mathbf{r}(t) - \mathbf{D}^\dagger \boldsymbol{\mu}_{\mathbf{s}(t)}) \quad (30)$$

Moreover, we can use the above score function to calculate the optimal  $\Delta \mathbf{r}(t + \Delta t)$  in Equation 11:

$$\begin{aligned} \Delta \mathbf{r}^*(t + \Delta t) &= \sigma_r^2 \Delta t (\mathbf{I}\sigma_r^2 \Delta t + \mathbf{D}^\dagger \boldsymbol{\Sigma}_{\mathbf{s}(t)} (\mathbf{D}^\dagger)^T)^{-1} (-\mathbf{r}(t) + \mathbf{D}^\dagger \boldsymbol{\mu}_{\mathbf{s}(t)}) \\ &\quad + \mathbf{D}^\dagger \mathbf{s}'(t) \Delta t + \sigma_r \boldsymbol{\eta}(t) \end{aligned} \quad (31)$$

Some properties of the leakage matrix  $\sigma_r^2 \Delta t (\mathbf{I}\sigma_r^2 \Delta t + \mathbf{D}^\dagger \boldsymbol{\Sigma}_{\mathbf{s}(t)} (\mathbf{D}^\dagger)^T)$  include:

1. The covariance matrix  $\boldsymbol{\Sigma}_{\mathbf{s}(t)}$  is positive semidefinite (PSD): all its eigenvalues are  $\geq 0$ .
2.  $\mathbf{D}^\dagger \boldsymbol{\Sigma}_{\mathbf{s}(t)} (\mathbf{D}^\dagger)^T$  is also PSD<sup>6</sup>, symmetric, and therefore diagonalizable.
3. The eigenvalues of  $(\mathbf{I}\sigma_r^2 \Delta t + \mathbf{D}^\dagger \boldsymbol{\Sigma}_{\mathbf{s}(t)} (\mathbf{D}^\dagger)^T)^{-1}$  are thus all  $\leq (\sigma_r^2 \Delta t)^{-1}$ <sup>7</sup>.
4. The eigenvalues of  $\sigma_r^2 \Delta t (\mathbf{I}\sigma_r^2 \Delta t + \mathbf{D}^\dagger \boldsymbol{\Sigma}_{\mathbf{s}(t)} (\mathbf{D}^\dagger)^T)^{-1}$  are thus all  $\leq 1$ .
5. If the off-diagonal entries of the leakage matrix above are sufficiently small in magnitude, then all the diagonal entries should be less than 1, as justified by the Gershgorin Circle Theorem. In fact, if the leakage matrix is diagonal, then it must have all values less than 1 (which could be achieved via sigmoid functions or perhaps spectral normalization).
6. As for interpretation, smaller leakage eigenvalues means higher eigenvalues of  $\mathbf{D}^\dagger \boldsymbol{\Sigma}_{\mathbf{s}(t)} (\mathbf{D}^\dagger)^T$ , or essentially, more noise. The maximum determinant of the leakage matrix is 1, when there is essentially no noise in  $\mathbf{s}(t)$ .

## D ADDITIONAL SCORE FUNCTION RESULTS

**Wiener Processes.** One simple stochastic process is the Wiener process, which in terms of navigation represents an undirected random walk ( $\theta = 0$ ). Assuming  $s_w(0) = 0$ , then  $s_w(t) \sim \mathcal{N}(0, \sigma_s^2 t)$ , and therefore  $p(r_w(t)) \sim \mathcal{N}(0, \sigma_s^2 t + \sigma_r^2 \Delta t)$  from Equation 7, producing the following score:

$$\sigma_r^2 \Delta t \nabla_{r_w(t)} \log p(r_w(t)) = \sigma_r^2 \Delta t \frac{-r_w(t)}{\sigma_s^2 t + \sigma_r^2 \Delta t} \quad (32)$$

Even from a simple Wiener process, we observe that the per-timestep optimal score is not constant with respect to  $t$ : at  $t = 0$ , it equals  $-r_w(t)$ , while as  $t$  approaches  $\infty$ , it approaches 0.

<sup>6</sup>Proof: If  $\mathbf{B}$  is PSD, then  $\mathbf{x}^T \mathbf{A} \mathbf{B} \mathbf{A}^T \mathbf{x} = (\mathbf{A}^T \mathbf{x})^T \mathbf{B} (\mathbf{A}^T \mathbf{x}) = \mathbf{v}^T \mathbf{B} \mathbf{v} \geq 0$ .

<sup>7</sup>Proof: For diagonalizable  $\mathbf{A}$ , the  $i$ -th eigenvalue of  $(\lambda \mathbf{I} + \mathbf{A})^{-1} = (\mathbf{Q}(\lambda \mathbf{I} + \boldsymbol{\Lambda})\mathbf{Q}^{-1})^{-1}$  is equal to  $(\lambda + \Lambda_{ii})^{-1}$ , which can be no larger than  $\lambda^{-1}$  if  $\mathbf{A}$  is PSD.

**Ornstein-Uhlenbeck Processes.** Now we incorporate non-zero leakage ( $\theta > 0$ ) to describe a directed random walk navigating from an arbitrary starting point  $s_{ou}(0)$  towards a mean destination  $\mu$ . If  $p(s_{ou}(0)) \sim \mathcal{N}(0, \sigma_0^2)$ , then  $p(s_{ou}(t)) \sim \mathcal{N}\left(\mu(1 - e^{-\theta t}), \frac{\sigma_r^2}{2\theta}(1 - e^{-2\theta t}) + \sigma_0^2 e^{-\theta t}\right)$ , and the score function is:

$$\sigma_r^2 \Delta t \nabla_{r_{ou}(t)} \log p(r_{ou}(t)) = \sigma_r^2 \Delta t \frac{-(r_{ou}(t) - \mu(1 - e^{-\theta t}))}{\frac{\sigma_r^2}{2\theta}(1 - e^{-2\theta t}) + \sigma_0^2 e^{-\theta t} + \sigma_r^2 \Delta t} \quad (33)$$

## E ADAPTATION AS A SECOND-ORDER STOCHASTIC DIFFERENTIAL EQUATION

Let us first combine the following two coupled linear stochastic differential equations into one second-order equation:

$$dX_t = (AX_t + BY_t + M)dt + \sigma dB_t, \quad dY_t = (CX_t + DY_t)dt \quad (34)$$

$$d^2 X_t = AdX_t + BdY_t + \sigma d(dB_t) \quad (35)$$

$$= AdX_t + BdY_t + \sigma d^2 B_t \quad (36)$$

$$= AdX_t + B(CX_t + DY_t)dt + \sigma d^2 B_t, \quad Y_t = \frac{1}{Bdt}(dX_t - \sigma dB_t - AX_t dt - Mdt) \quad (37)$$

$$= AdX_t + BCX_t dt + BD \frac{1}{Bdt}(dX_t - \sigma dB_t - AX_t dt - Mdt)dt + \sigma d^2 B_t \quad (38)$$

$$= AdX_t + BCX_t dt + D(dX_t - \sigma dB_t - AX_t dt - Mdt) + \sigma d^2 B_t \quad (39)$$

$$= (A + D)dX_t + (BC - AD)X_t dt - DMdt - \sigma DdB_t + \sigma d^2 B_t \quad (40)$$

Replacing all variables involved (except  $dt, \sigma$ ) with matrices and vectors yields the same equation as long as  $\mathbf{B}$  is invertible:

$$d^2 \mathbf{x}_t = (\mathbf{A} + \mathbf{D})d\mathbf{x}_t + (\mathbf{BC} - \mathbf{AD})\mathbf{x}_t dt - \mathbf{D}\mathbf{m}dt - \sigma \mathbf{D}d\mathbf{B}_t + \sigma d^2 \mathbf{B}_t \quad (41)$$

For consistency with the notation used throughout the paper, the equation above can be written as:

$$\mathbf{x}''(t) = (\mathbf{A} + \mathbf{D})\mathbf{x}'(t) + (\mathbf{BC} - \mathbf{AD})\mathbf{x}(t) - \mathbf{D}\mathbf{m} - \sigma \mathbf{D}\boldsymbol{\eta}(t) + \sigma \boldsymbol{\eta}'(t) \quad (42)$$

If we apply the following substitutions from Equations 27 and 17:

- $\mathbf{x}(t) \leftarrow \mathbf{r}(t)$ ,
- $\mathbf{A} \leftarrow -\sigma_r^2 \Delta t \boldsymbol{\Sigma}^{-1}$ ,
- $\mathbf{B} \leftarrow -\mathbf{I}$ ,
- $\mathbf{m} \leftarrow \sigma_r^2 \Delta t \boldsymbol{\Sigma}_t^{-1} \boldsymbol{\mu}$ ,
- $\mathbf{C} \leftarrow -\frac{1}{\tau_a} \mathbf{I}$ ,
- $\mathbf{D} \leftarrow \frac{b_a}{\tau_a} \mathbf{I}$ ,

then  $\mathbf{r}''(t)$  is:

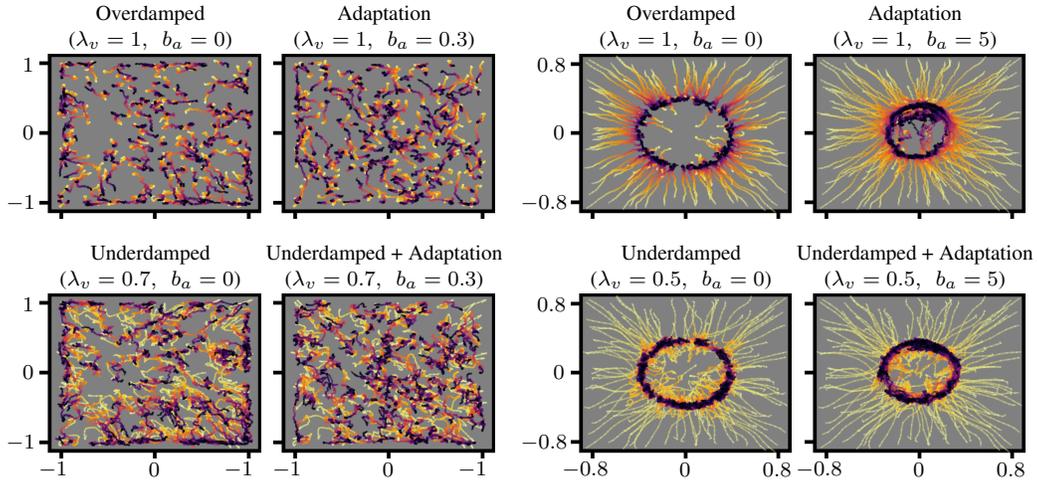
$$\begin{aligned} \mathbf{r}''(t) &= \left( \frac{b_a}{\tau_a} \mathbf{I} - \sigma_r^2 \Delta t \boldsymbol{\Sigma}^{-1} \right) \mathbf{r}'(t) \\ &\quad + \left( \frac{1}{\tau_a} \mathbf{I} + \frac{b_a}{\tau_a} \sigma_r^2 \Delta t \boldsymbol{\Sigma}^{-1} \right) \mathbf{r}(t) \\ &\quad - \frac{b_a}{\tau_a} \sigma_r^2 \Delta t \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu} - \sigma \frac{b_a}{\tau_a} \boldsymbol{\eta}(t) + \sigma \boldsymbol{\eta}'(t) \end{aligned} \quad (43)$$

1188 Recall that, if  $\mathbf{r}(t)$  follows a stationary Gaussian distribution, then  $\nabla_{\mathbf{r}(t)} \log p(\mathbf{r}(t)) = \Sigma^{-1}(-\mathbf{r}(t) +$   
 1189  $\boldsymbol{\mu})$  (Equation 27), and therefore  $\frac{d^2}{d\mathbf{r}(t)^2} \log p(\mathbf{r}(t)) = -\Sigma^{-1}$ . Then,

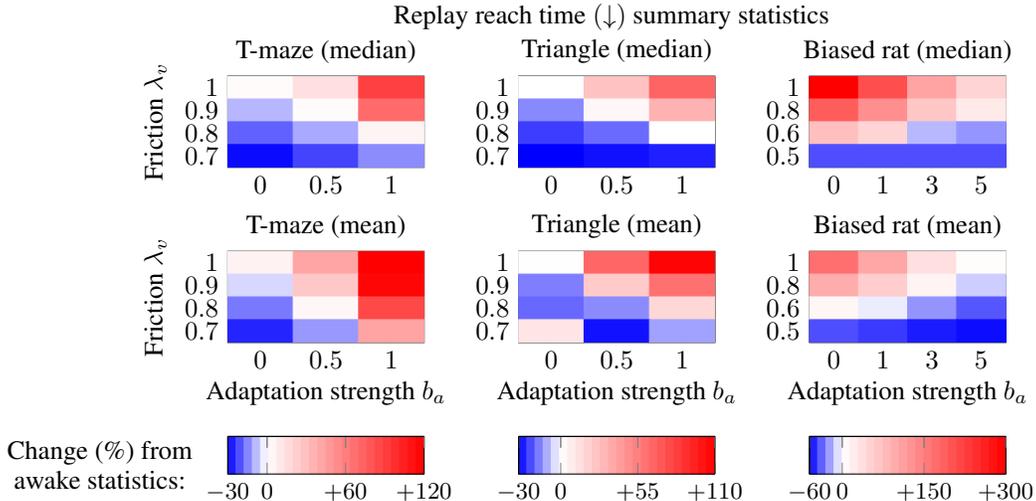
$$1191 \mathbf{r}''(t) = \left( \frac{b_a}{\tau_a} \mathbf{I} + \sigma_r^2 \Delta t \frac{d^2}{d\mathbf{r}(t)^2} \log p(\mathbf{r}(t)) \right) \mathbf{r}'(t) \\
 1192 - \frac{b_a}{\tau_a} \sigma_r^2 \Delta t \nabla_{\mathbf{r}(t)} \log p(\mathbf{r}(t)) + \frac{1}{\tau_a} \mathbf{r}(t) \\
 1193 \\
 1194 - \sigma \frac{b_a}{\tau_a} \boldsymbol{\eta}(t) + \sigma \boldsymbol{\eta}'(t) \quad (44)$$

## 1199 F ADDITIONAL RESULTS

1200 Here we present additional results or figures that supplement those of the main text.

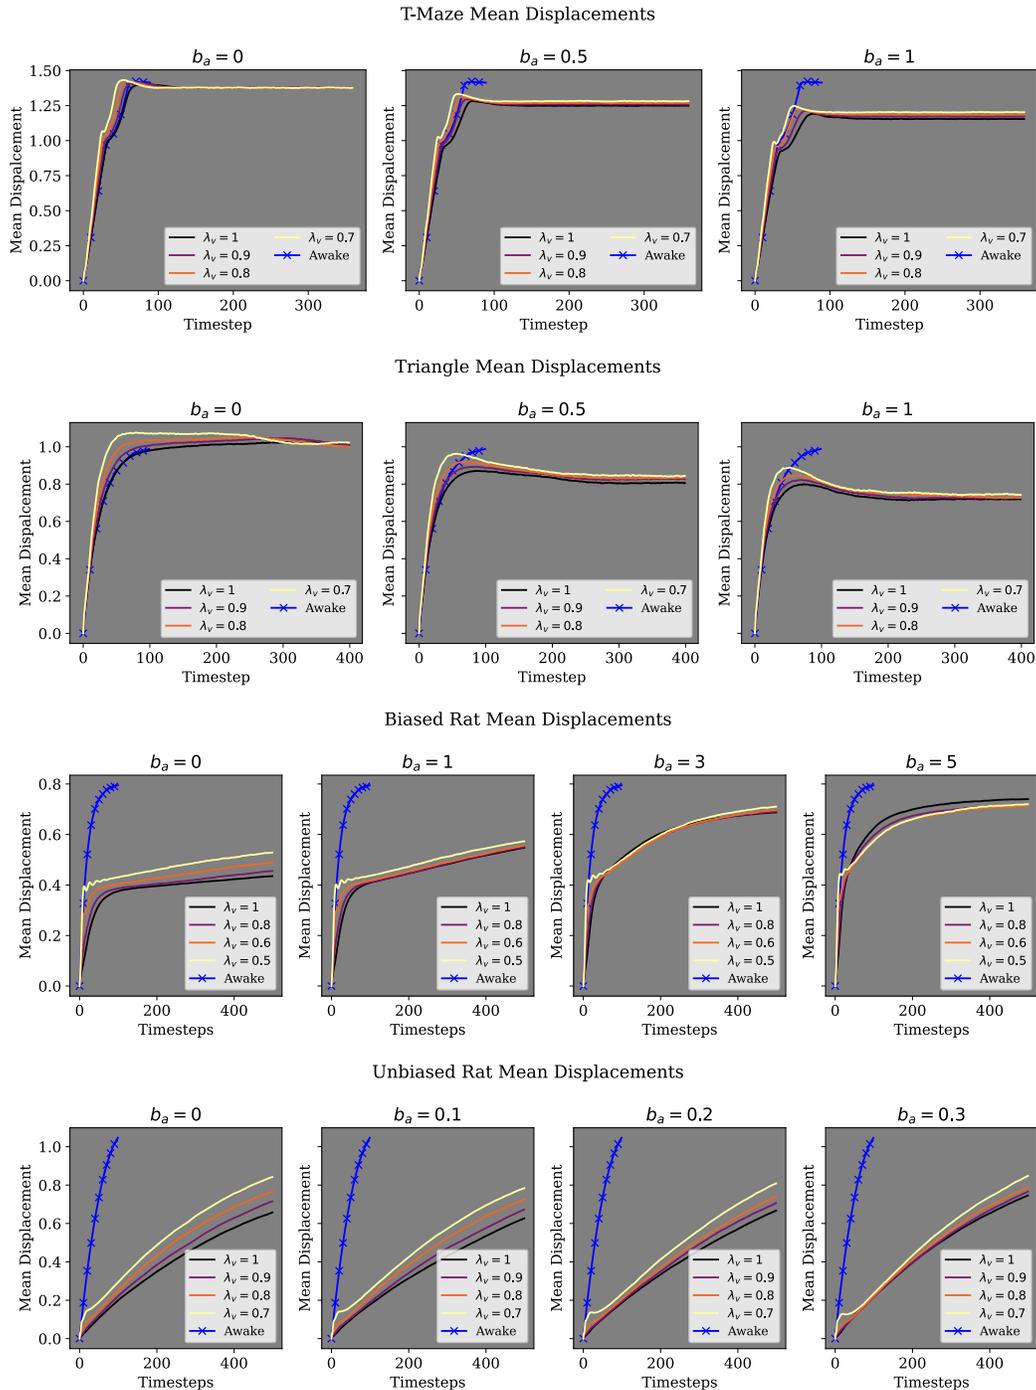


1219 Figure 8: Example decoded 2D replay trajectories in unbiased (left) and biased (right) rat tasks. Like  
 1220 in Figure 3, trajectories get darker over time.



1239 Figure 9: This is another version of Figure 5, but now with mean reach times also shown.

1240  
1241



1288  
1289  
1290  
1291  
1292  
1293  
1294  
1295

Figure 10: **Underdamping ( $\lambda_v < 1$ ) increases mean displacement of replay trajectories, especially at early timesteps.** Similarly to [McNamee et al. \(2021\)](#), here we analyze the *mean displacement*, or distance, of replay trajectories from their starting points as a function of time. For reference, mean displacement over time is also plotted for awake trajectories. In exploration experiments, we simulate replay for  $4\times$  the duration of awake trajectories. Higher mean displacements over time generally correlate with increased path length and exploration. Note that awake activity is always calculated with  $b_a = 0$ ,  $\lambda_v = 1$ , hence why awake statistics are the same across plots within the same task.

1296  
 1297  
 1298  
 1299  
 1300  
 1301  
 1302  
 1303  
 1304  
 1305  
 1306  
 1307  
 1308  
 1309  
 1310  
 1311  
 1312  
 1313  
 1314  
 1315  
 1316  
 1317  
 1318  
 1319  
 1320  
 1321  
 1322  
 1323  
 1324  
 1325  
 1326  
 1327  
 1328  
 1329  
 1330  
 1331  
 1332  
 1333  
 1334  
 1335  
 1336  
 1337  
 1338  
 1339  
 1340  
 1341  
 1342  
 1343  
 1344  
 1345  
 1346  
 1347  
 1348  
 1349

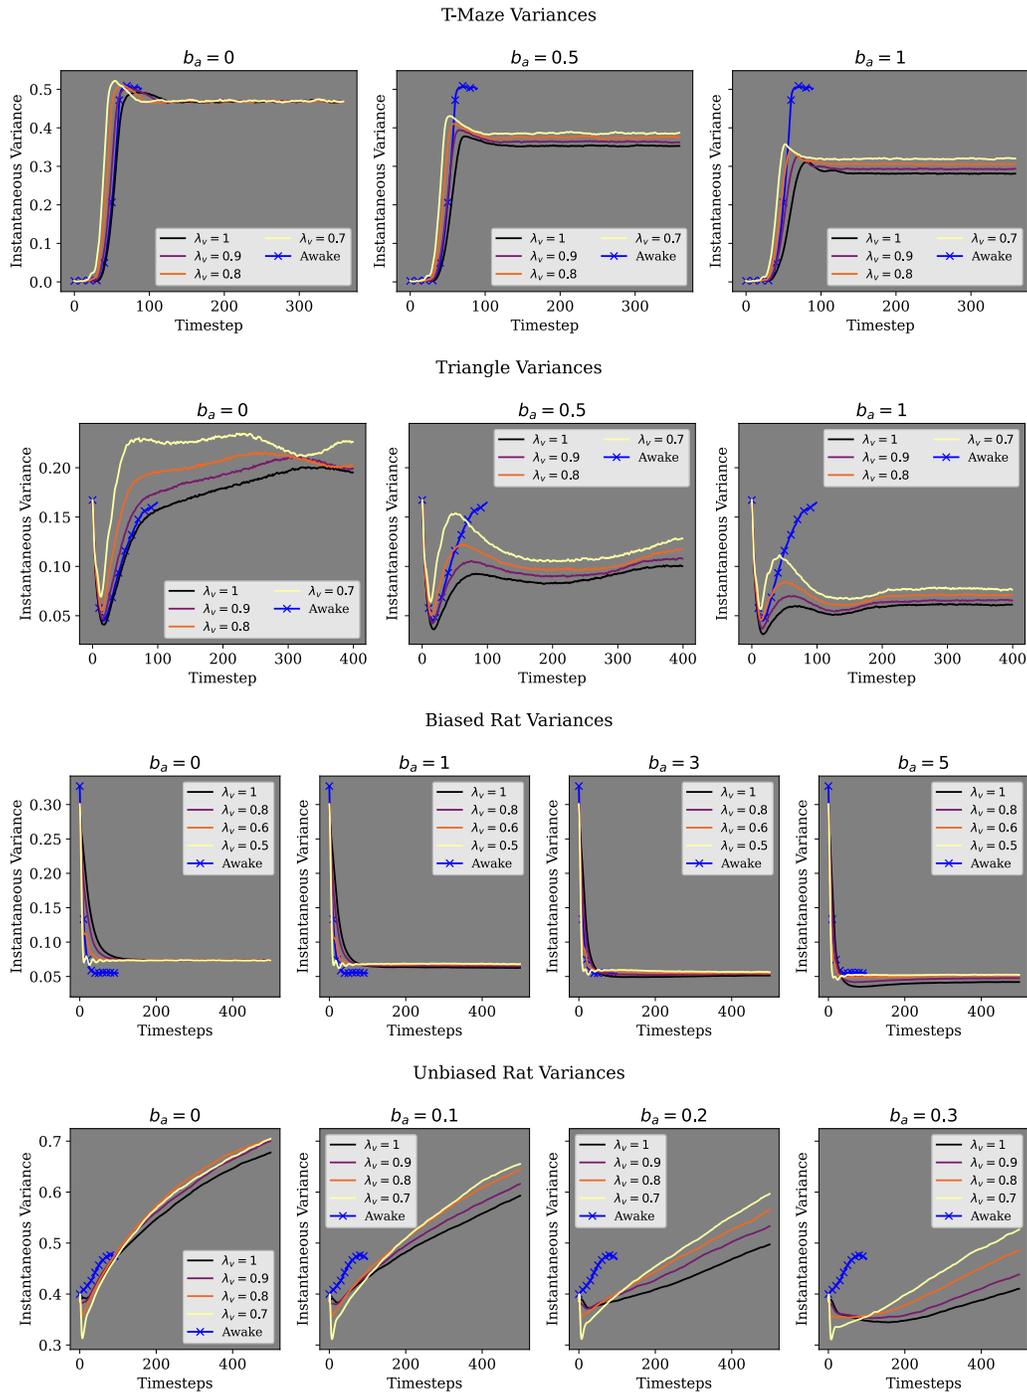


Figure 11: **Underdamping ( $\lambda_v < 1$ ) increases replay trajectory variances.** Like in Figure 10, here we analyze the variance of trajectories at each timestep. Trajectories are 2D, so the variances plotted above are simply the averages of the variances along each coordinate. Underdamping increasing variance is another confirmation that underdamping complements adaptation-induced exploration.

1350  
 1351  
 1352  
 1353  
 1354  
 1355  
 1356  
 1357  
 1358  
 1359  
 1360  
 1361  
 1362  
 1363  
 1364  
 1365  
 1366  
 1367  
 1368  
 1369  
 1370  
 1371  
 1372  
 1373  
 1374  
 1375  
 1376  
 1377  
 1378  
 1379  
 1380  
 1381  
 1382  
 1383  
 1384  
 1385  
 1386  
 1387  
 1388  
 1389  
 1390  
 1391  
 1392  
 1393  
 1394  
 1395  
 1396  
 1397  
 1398  
 1399  
 1400  
 1401  
 1402  
 1403

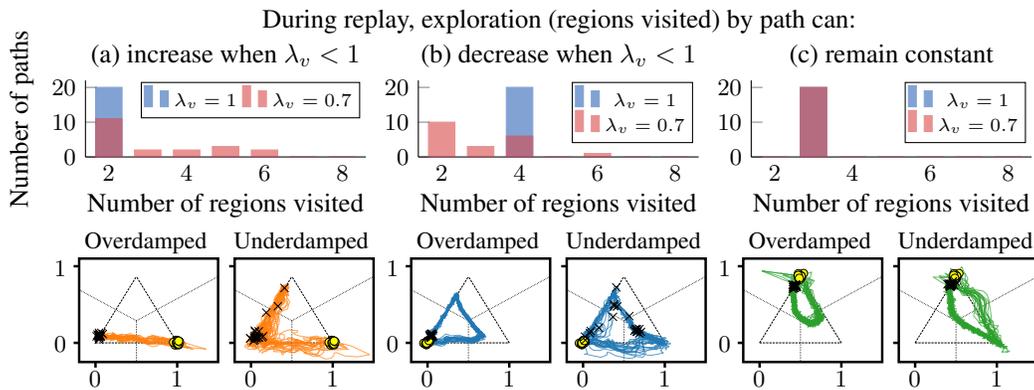


Figure 12: This is another version of Figure 7, but this time with the specific distributions of regions visited shown as bar plots above pairs of replay trajectory sets. Recall that replay paths start at yellow circles and end at  $\times$  symbols. In (a), underdamping clearly increases regions visited for several trajectories: instead of just going from  $(1, 0)$  to  $(0, 0)$ , replay paths also continue onwards to traverse between  $(0, 0)$  and  $(1/2, \sqrt{3}/2)$ . Meanwhile in (b), underdamping decreases regions visited for many replay trajectories: instead of traversing the whole triangle (4 regions visited), many replay paths stop in the region defined by  $(1, 0)$ . As for (c), underdamping did not affect regions visited since all paths started at  $(1/2, \sqrt{3}/2)$ , visited  $(1, 0)$ , and then returned to  $(1/2, \sqrt{3}/2)$ .

## G RESULTS WITH TANH ACTIVATIONS

Throughout all other parts of the text, we have used piecewise-linear (linear or leaky/non-leaky ReLU) activation functions, as defended in Appendix A.2. For completeness, here we reproduce some key findings of our work with all hyperparameters kept the same, except that RNNs now use tanh activation functions instead of piecewise-linear ones. Replay from tanh RNNs visited far fewer regions than replay from their piecewise-linear counterparts, so a reproduction of Figure 7 is omitted.

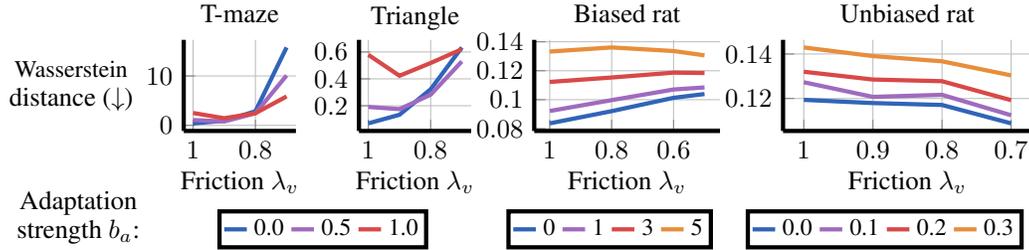


Figure 13: **Mild underdampening can improve replay fidelity in the presence of adaptation in tanh RNNs.** This is another version of Figure 4, for T-maze and triangle tasks, but with tanh RNNs.

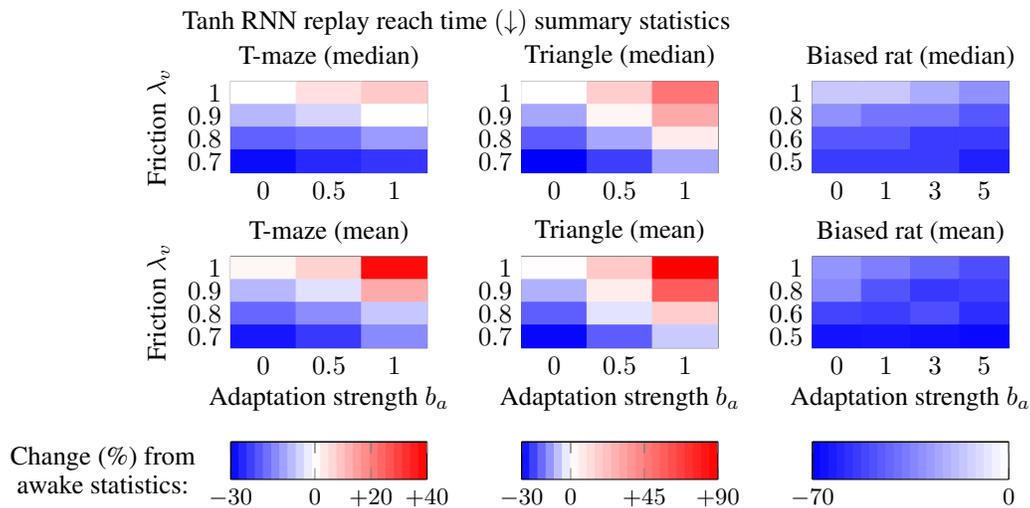


Figure 14: **Underdampening increases speed in tanh RNN replay.** This is another version of Figure 9 for T-maze and triangle tasks, but now with tanh RNNs.

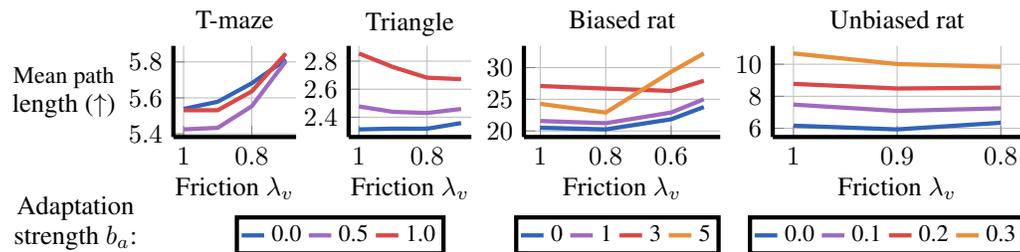


Figure 15: **In tanh RNNs, underdampening may increase or decrease exploration via path length.** This is an alternate version of Figure 6 but with tanh RNNs. Underdampening has mixed effects on tanh RNN replay path length because of tanh saturation.