




Dimensionality of population-level latent mechanisms encoding spatial representations

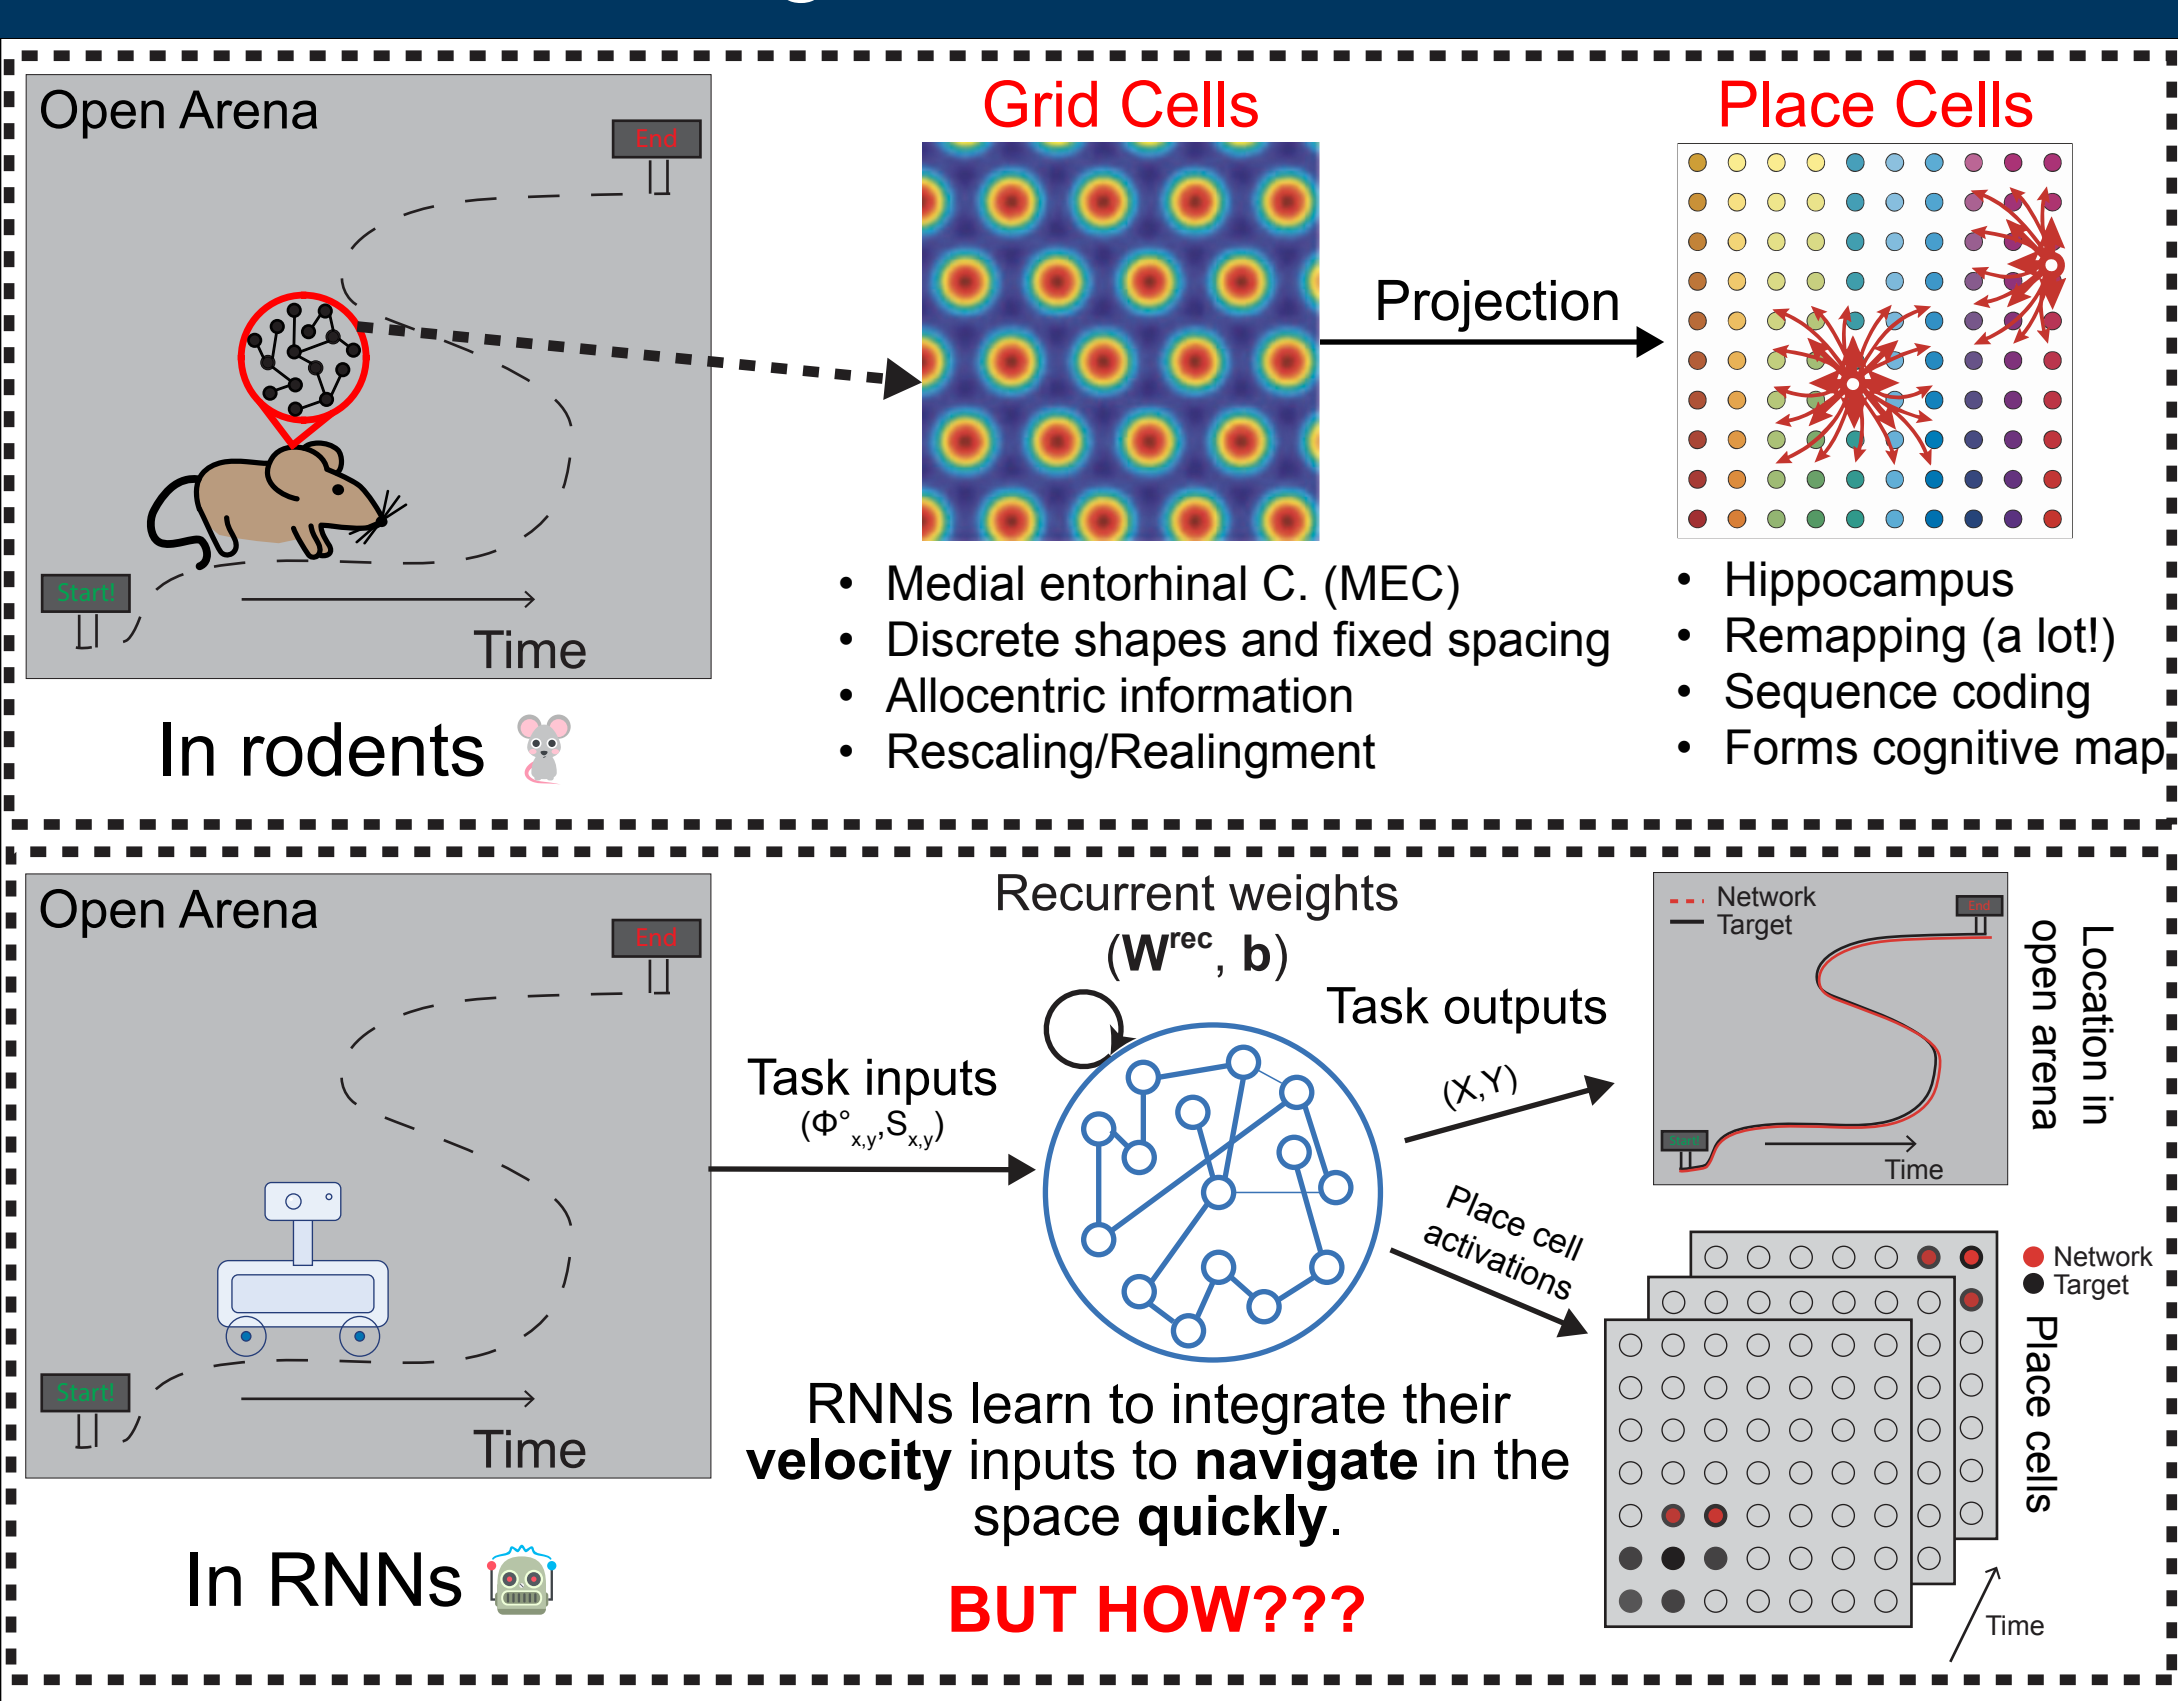
Nur E. Delikkaya^{1,2,4}, Alperen Cimen^{3,4,5}, Francisco Acosta⁴, Adele Myers⁴, Andy Alexander^{5,6}, Fatih Dinc^{4,7}, Nina Miolane⁴

¹Computer Science, ²Psychology, Sabanci Univ. Istanbul, Turkey; ³Electrical Engineering, Ozyegin Univ. Istanbul, Turkey; ⁴Geometric Intelligence Lab;

⁵Psychology and Brain Sciences, Neuroscience Research Institute; ⁶Dynamical Neuroscience; ⁷Kavli Theoretical Physics Institute UC Santa Barbara, Santa Barbara, CA 93117

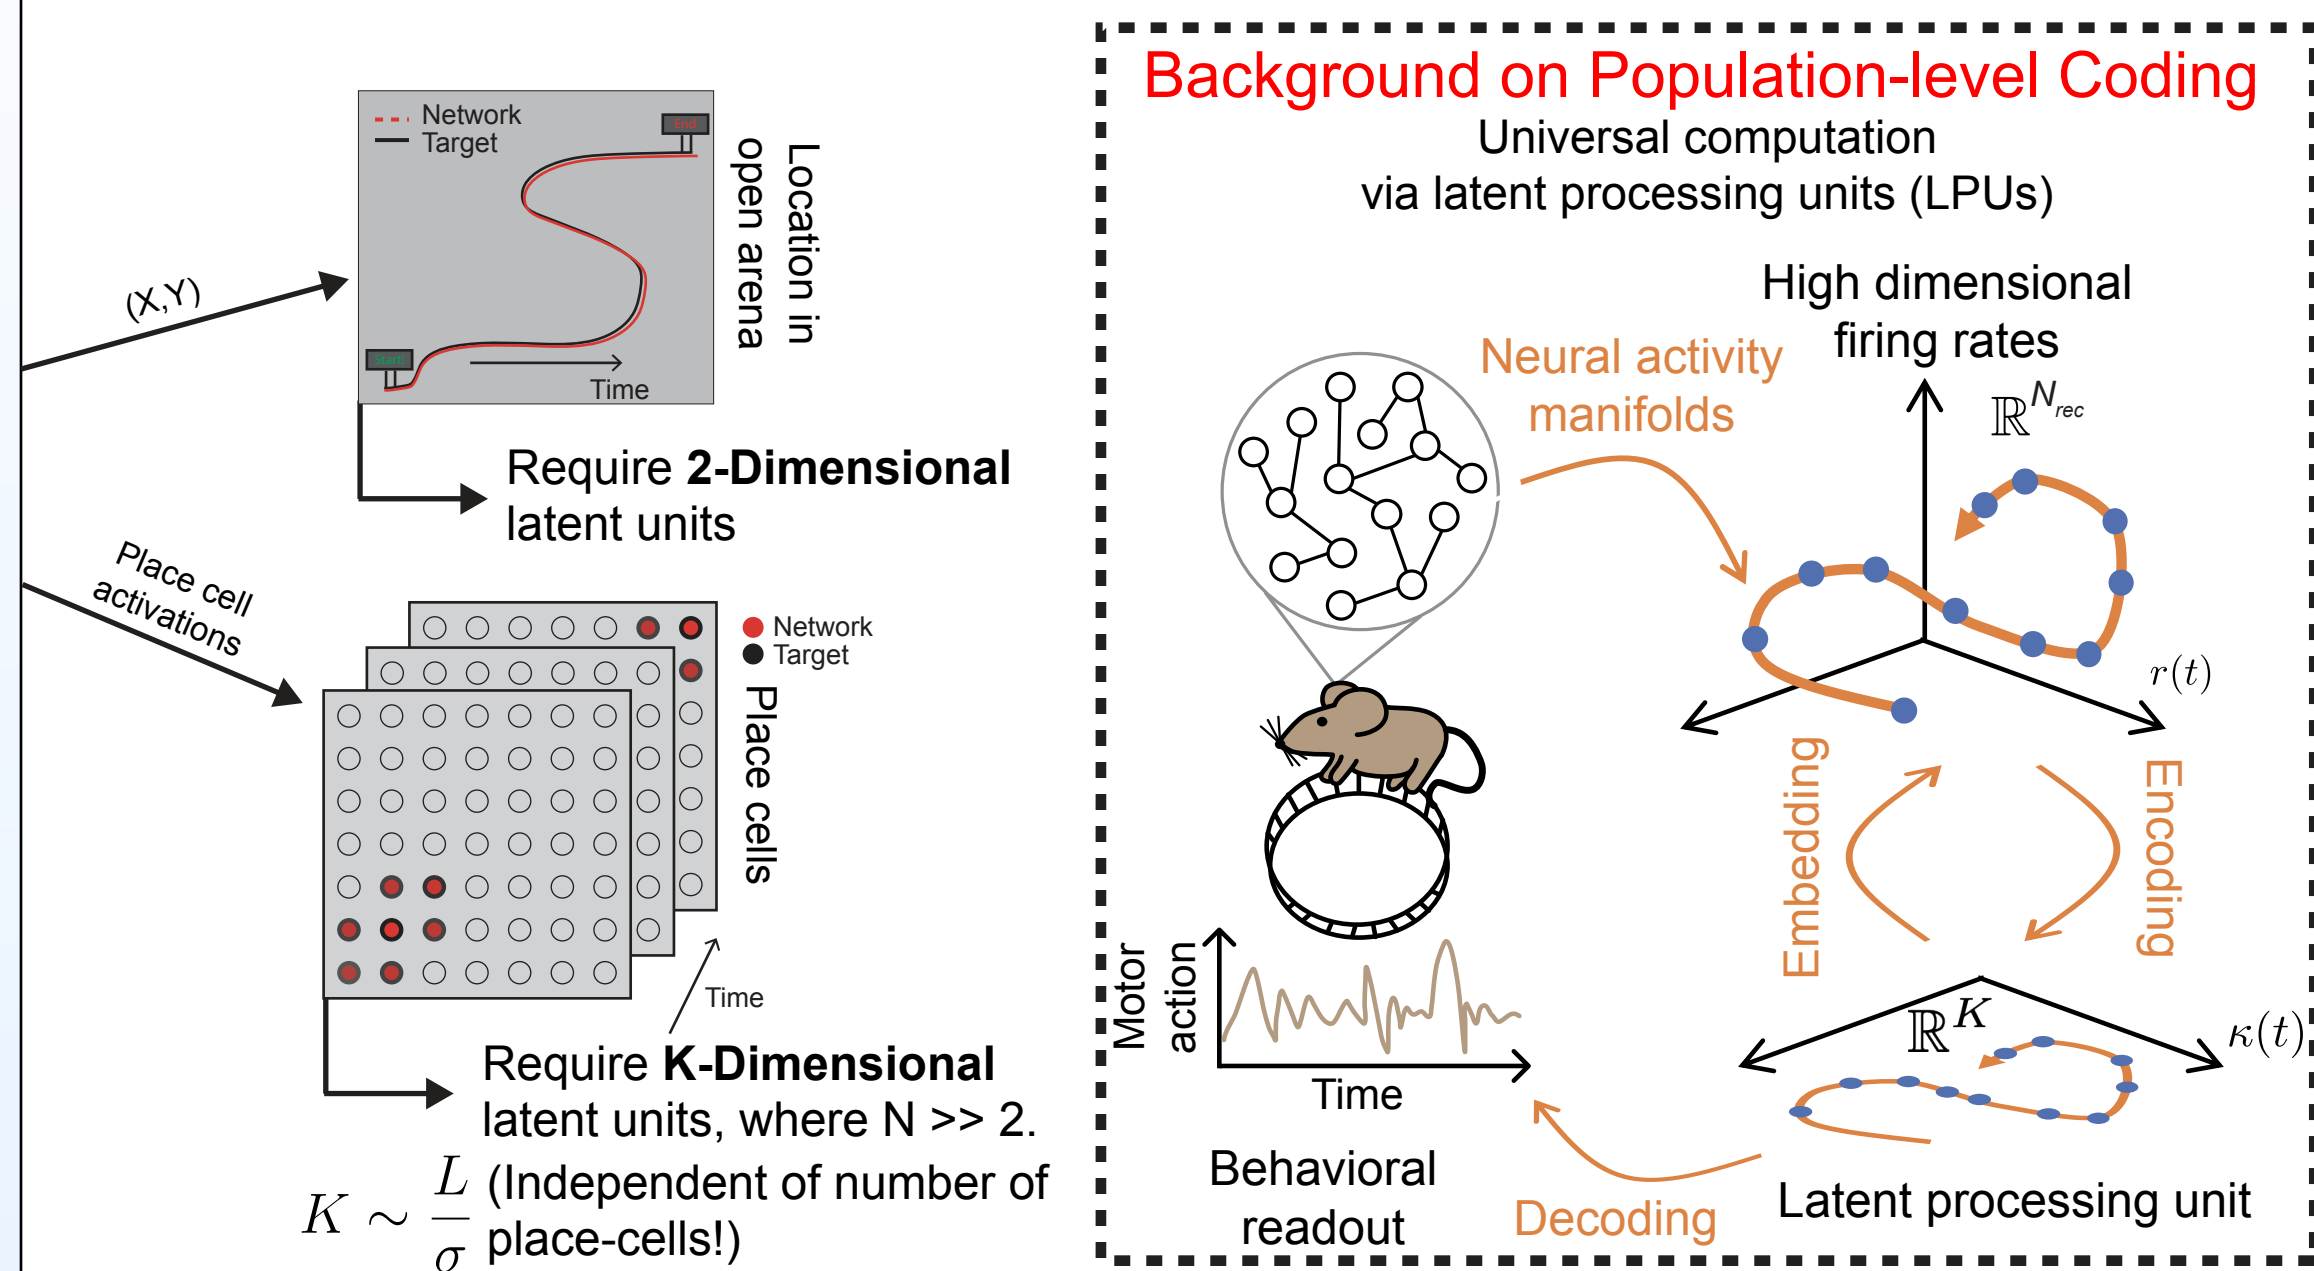


Navigation Mechanisms



Dimensionality of Spatial Navigation

The dimensionality of an RNN depends critically on the biological computation it is meant to model. When an RNN is used to represent the medial entorhinal cortex (MEC)—for example, by generating place-cell-like outputs—it requires a high-dimensional latent mechanism to reproduce the rich population codes observed in MEC. In contrast, when the RNN is trained to output continuous 2-D position directly, its latent space only requires a low-dimensional (2D) continuous attractor underlying navigation dynamics in the networks.



Theory

Building on Sorscher et al. (2023), we formulate spatial encoding as an optimization problem and extend this framework to examine the dimensionality and properties of the basis functions when outputs are either discrete

Let space is covered by place cells whose tuning curves: $f_\mu(x) = \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$. μ : Cell index, x : Position, σ : Spatial resolution

Then, our aim is to reconstruct the space with orthonormal basis functions: $\{\phi_k(x)\}_{k=1}^K$

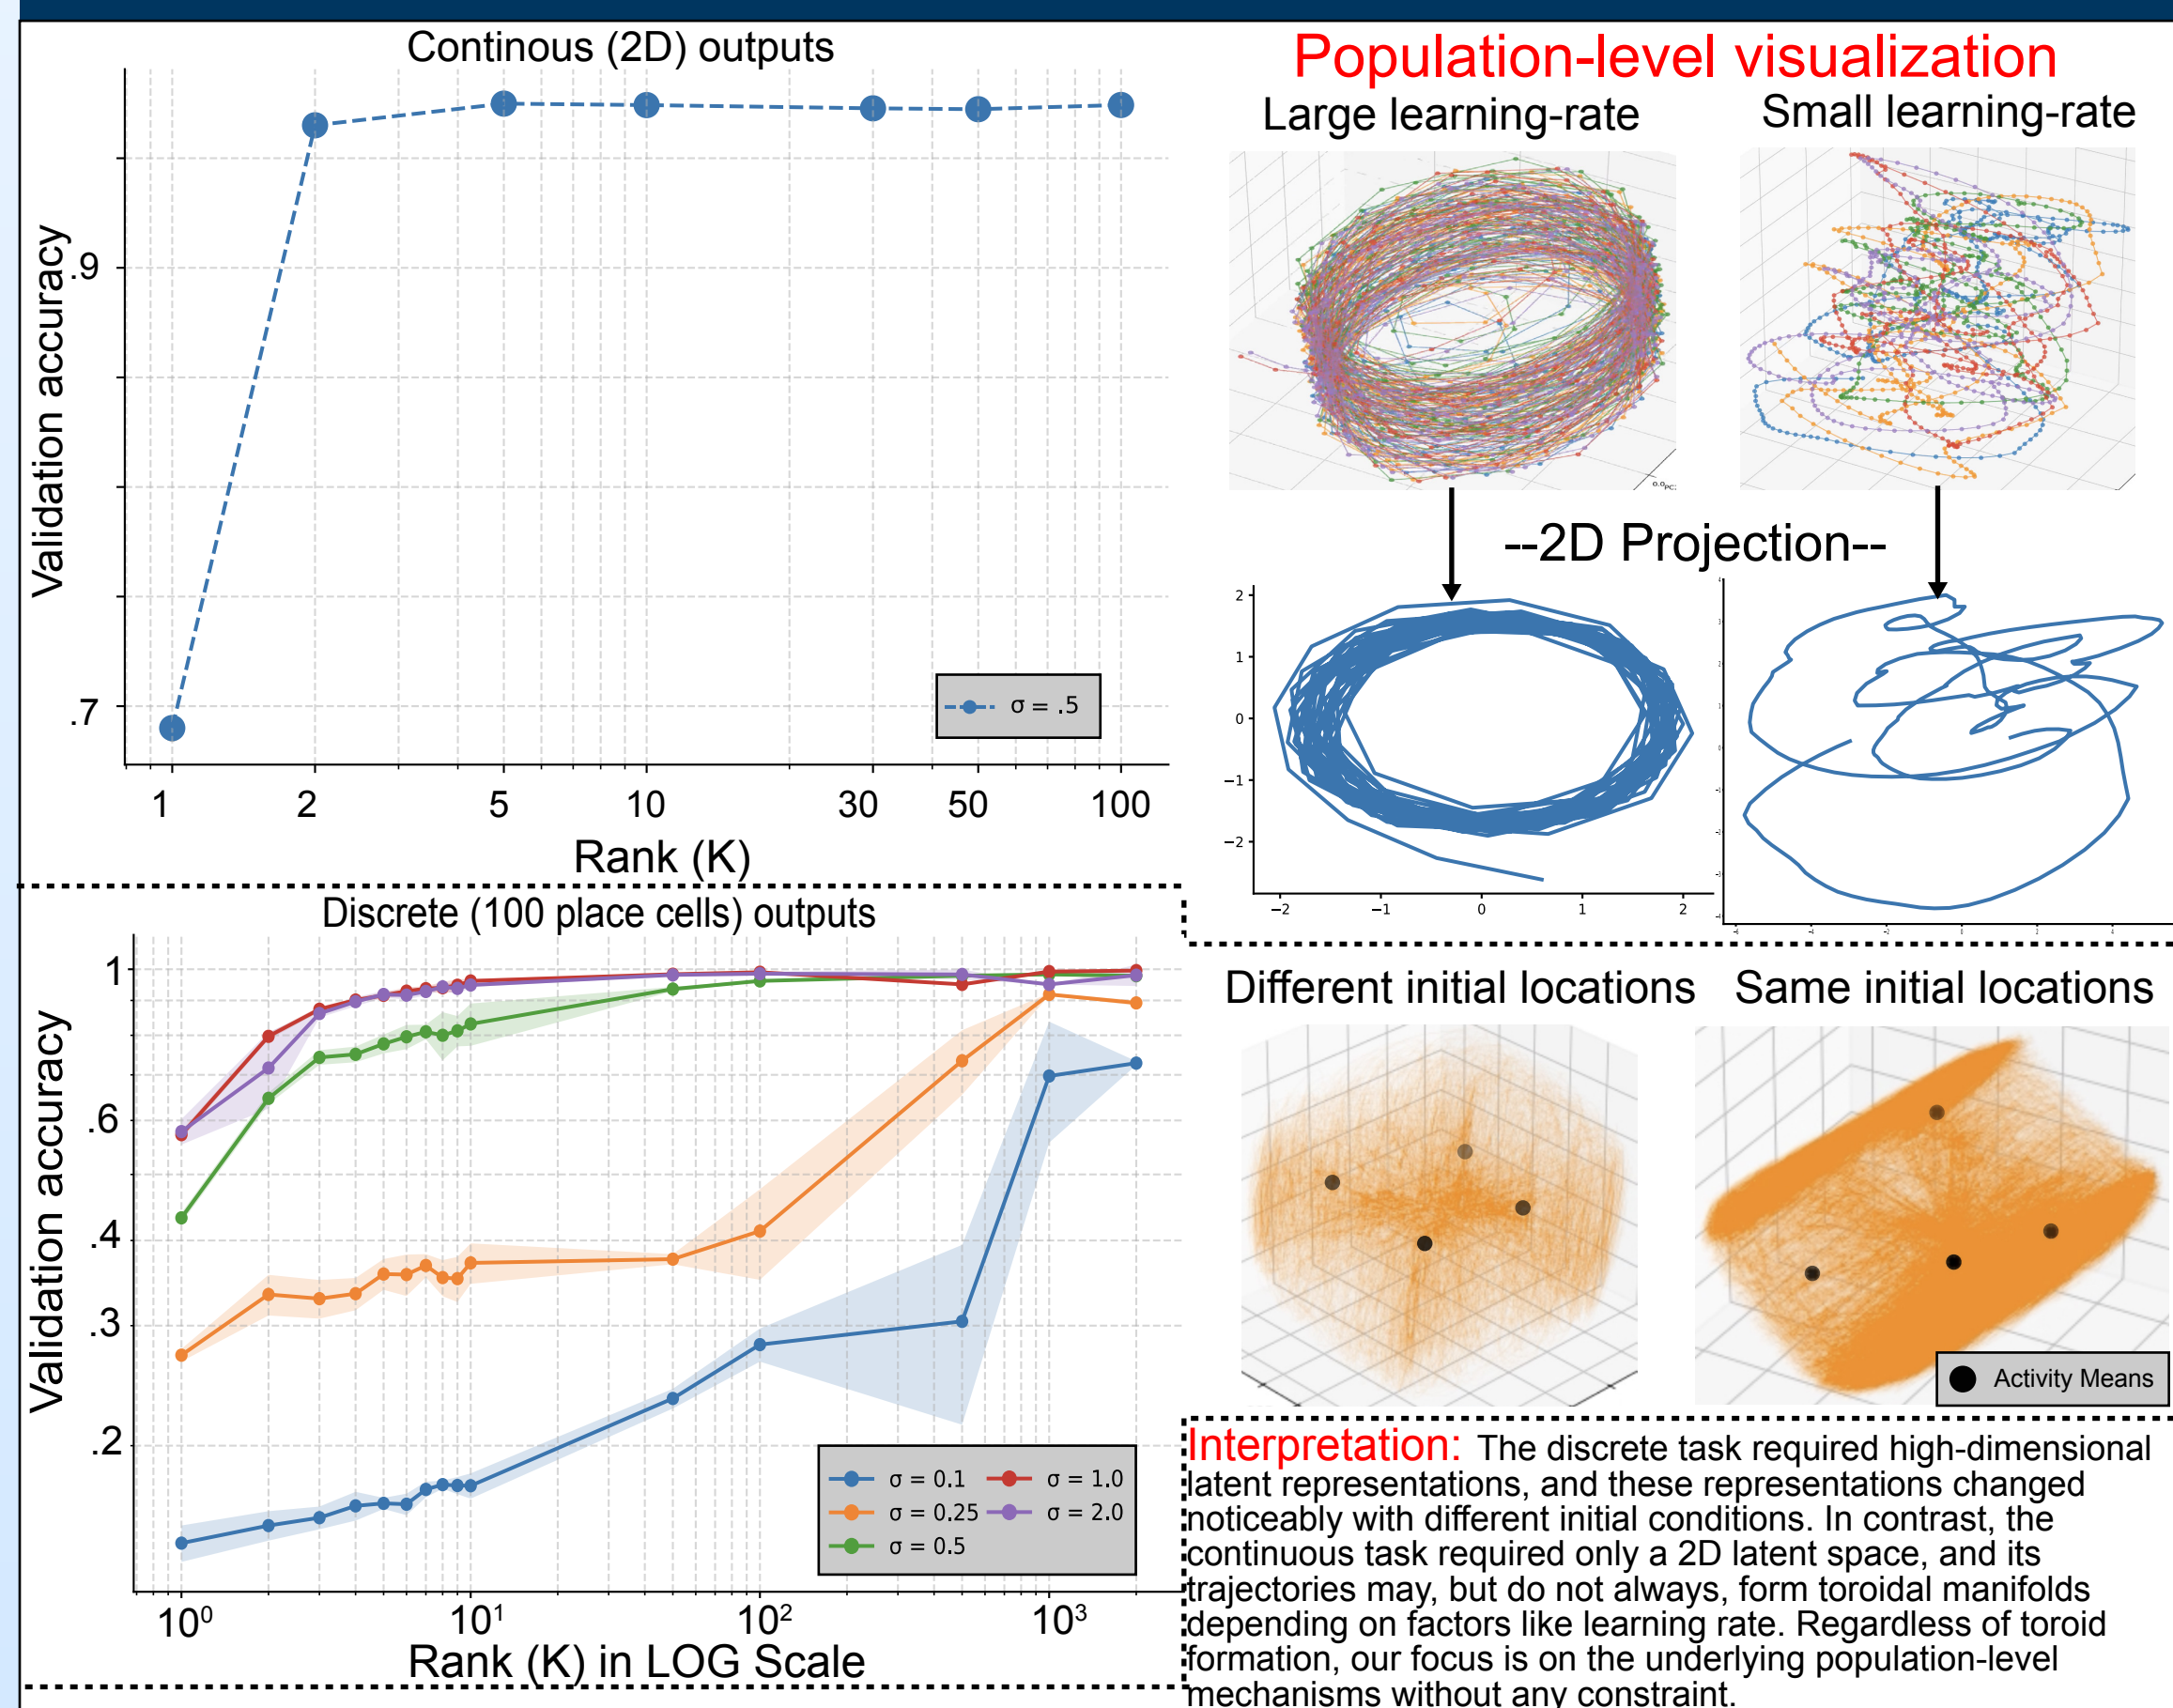
Such that: $f_\mu(x) \approx \sum_{k=1}^K w_k(\mu) \phi_k(x)$. where $w_k(\mu)$ defines output weight from the k th basis function to produce μ th place cell.

Then, following Sorscher et al (2023, 2019), the optimal set of basis functions are defined to minimize the following: $E_\mu \int_{\Omega} \left[f_\mu(x) - \sum_{k=1}^K w_k(\mu) \phi_k(x) \right]^2 dx$. (mean squared error for all positions for each place cell)

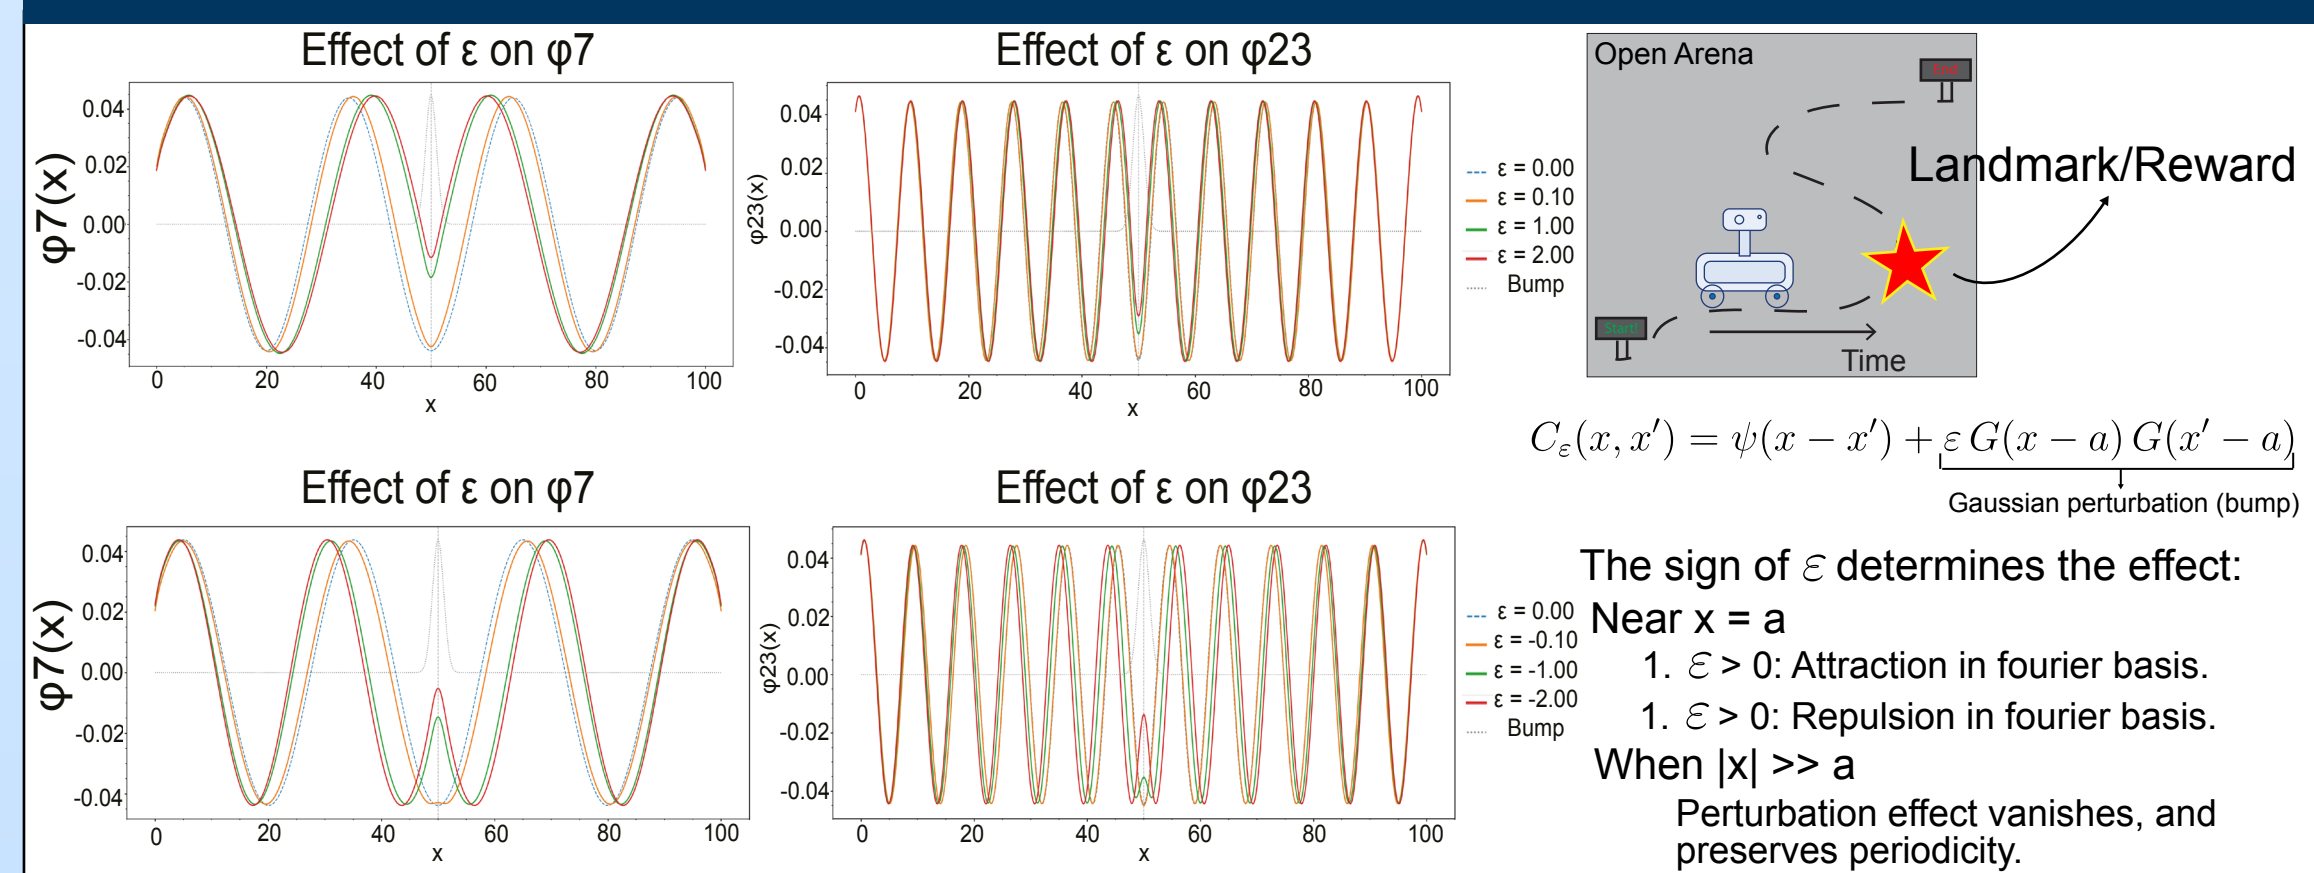
This problem reduces to the classical **Karhunen-Loève (KL) expansion**, or functional principal component analysis (Ramsay and Silverman, 2005). The solution is determined by the correlation-kernel: $C(x, x') = \mathbb{E}_\mu [f_\mu(x) f_\mu(x')]$, whose eigenfunctions provide the optimal basis: $\int_{\Omega} C(x, x') \phi_k(x') dx' = \lambda_k \phi_k(x)$. Compact spatial representations are obtained by using the leading eigenfunctions of $C(x, x')$ as the basis functions, whereas the error in the reconstruction is obtained by summing over the eigenvalues λ_{κ} that belong to the eigenfunctions not picked as the basis function. **(More basis functions --> Less reconstruction error)**

Essence: Under uniform coverage, $C(x, x')$ becomes *translation-invariant*, and the eigenfunctions turn into periodic **Fourier modes**, giving a direct link between **spatial resolution** (σ) and the number of **required latent dimensions** (K).

Results



Landmark Effects



Which constraints does prior work impose?

From a computational neuroscience perspective, though these prior work pave the way for future studies, they don't form a theoretical background on how grid cells emergence or how their different orientations (both in scale and phase) help to carry out the computations, Sorscher et al. (2019, 2023). More importantly, revealing grid cell-like modules within the constrained and manipulated settings (i.e., positive firing-rates and differences of Gaussians place-cells) leverages the obscure nature of spatial navigation in recurrent neural networks further, and rises more questions. From this viewpoint, path integration can be studied as an important problem in its own right and arguably a more central concern for the brain than producing fancy grid cells.

References

• Ben Sorscher, Gabriel Mel, Surya Ganguli, and Samuel Ocko. A unified theory for the origin of grid cells through the lens of pattern formation. *Advances in neural information processing systems*, 32, 2019.

• Ben Sorscher, Gabriel C Mel, Samuel A Ocko, Lisa M Giacomo, and Surya Ganguli. A unified theory for the computational and mechanistic origins of grid cells. *Neuron*, 111(1):121–137, 2023.

• James O. Ramsay and Bernard W. Silverman. *Functional data analysis*. Springer, 2005.

• McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M.-B. (2006). *Path integration and the neural basis of the cognitive map*. *Nature Reviews Neuroscience*, 7(8), 663–678.

• Fatih Dinc, Marta Blanco-Pozzo, David Klimdt, Francisco Acosta, Yiqi Jiang, Sadegh Ebrahimi, Adam Shai, Hidenori Tanaka, Peng Yuan, Mark J. Schnitzer, et al. *Latent computing by biological neural networks: A dynamical systems framework*. *arXiv preprint arXiv:2502.14337*, 2025.

Next Steps: Testing with actual rodents!

We train mice and record their behavior and neural activity on a circular track (both on landmark and uniform cases). Next, we will look at the Fourier basis' emergence both in RNNs and brains. We then compare the RNN's predictions to the actual neural data.

