

Disentangling Fact from Grid Cell Fiction in Trained Deep Path Integrators

Rylan Schaeffer¹, Mikail Khona², Sanmi Koyejo¹, and Ila Rani Fiete^{3,4}

¹Computer Science, Stanford

²Physics, MIT

³Brain and Cognitive Sciences, MIT

⁴McGovern Institute for Brain Research at MIT

Abstract

Work on deep learning-based models of grid cells suggests that grid cells generically and robustly arise from optimizing networks to path integrate, i.e., track one’s spatial position by integrating self-velocity signals. In previous work [27], we challenged this path integration hypothesis by showing that deep neural networks trained to path integrate almost always do so, but almost never learn grid-like tuning unless separately inserted by researchers via mechanisms unrelated to path integration. In this work, we restate the key evidence substantiating these insights, then address a response to [27] by authors of one of the path integration hypothesis papers [32]. First, we show that the response misinterprets our work, indirectly confirming our points. Second, we evaluate the response’s preferred “unified theory for the origin of grid cells” in trained deep path integrators [31, 33, 34] and show that it is at best “occasionally suggestive,” not exact or comprehensive. We finish by considering why assessing model quality through prediction of biological neural activity by regression of activity in deep networks [23] can lead to the wrong conclusions.

1 Introduction: Grid Cells & the Path Integration Hypothesis

The discovery of grid cells [13, 36] in the mammalian brain sparked considerable work on their mechanisms and origins. One approach to understanding grid cells is through the lens of optimization in deep neural networks. A number of works in this direction published in high-profile venues assert that grid cells generically and robustly emerge when recurrent networks are trained to path integrate (Nature [1], Neuron [34]) and machine learning conferences (Neural Information Processing Systems Spotlights [31, 23], International Conference on Learning Representations [5]):

- Main text of [1]: “Notably, therefore, our results show that grid-like representations reminiscent of those found in the mammalian entorhinal cortex emerge in a generic network trained to path integrate.”
- Abstract of [33]: “Here we forge an intimate link between the computational problem of path-integration and the existence of hexagonal grids, by demonstrating that such grids arise generically in biologically plausible neural networks trained to path integrate. Moreover, we develop a unifying theory for why hexagonal grids are so ubiquitous in path-integrator circuits.”
- Highlights of [34]: “RNNs trained to path integrate with nonnegative firing develop hexagonal grid cells.”

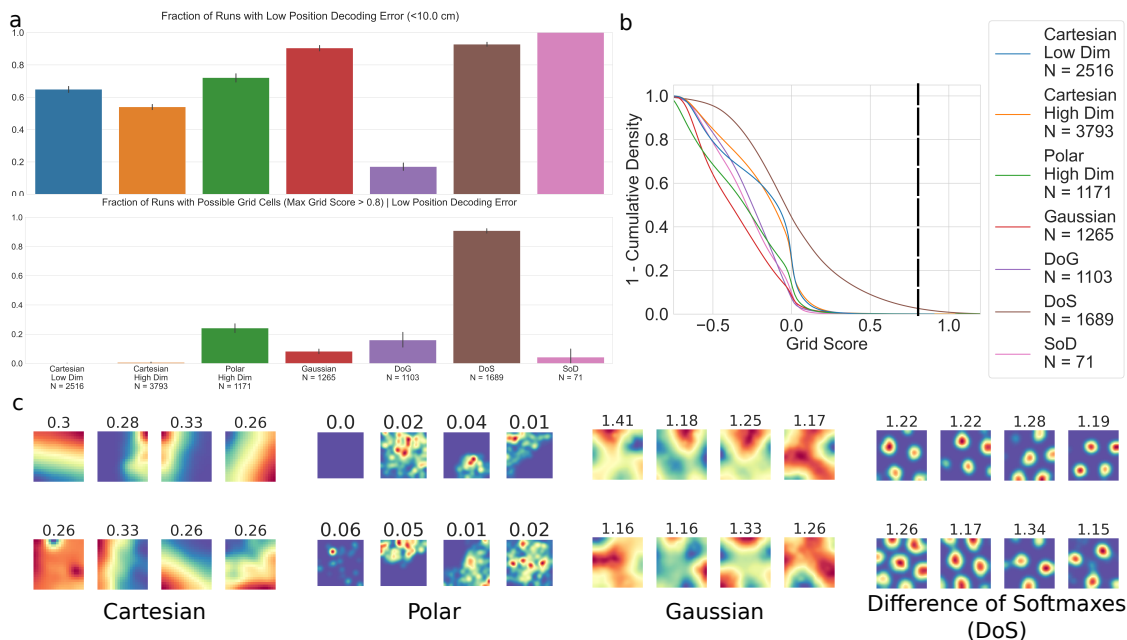


Figure 1: (a) Top: Most networks learn to path integrate, regardless of supervised target. Bottom: Few networks display possible grid-like tuning. (b) Survival functions of grid scores per supervised target. (c) Rate maps of top grid-scoring units in trained deep path integrators with i) Cartesian, ii) Polar, iii) Gaussian, iv) specifically selected (tuned) Difference-of-Softmaxes (DoS) targets. Only DoS supervised targets induce grid cells. Numbers above rate maps are 60° grid scores. Figure from [27].

- Abstract of [5]: “We trained recurrent neural networks (RNNs) to perform navigation tasks in 2D arenas based on velocity inputs. Surprisingly, we find that grid-like spatial response patterns emerge in trained networks.”
- Key figure in [34]: “Neural networks trained on normative tasks develop grid-like firing fields.”

These publications collectively paint the picture that simply learning to path integrate (possibly with a non-negativity constraint) surprisingly, generically, and robustly causes the networks to learn grid-like representations matching those found in the mammalian entorhinal cortex. We call this claim the *path integration hypothesis*. However, there are reasons to question this path integration hypothesis: the hypothesis is inconsistent with prior numerical results training deep recurrent neural networks to perform spatial tasks [15, 16] and contradicts theoretical work suggesting the grid code is special not just for path integration but also for its coding-theoretic properties [10, 35, 22, 38].

Motivated by these reasons, we rigorously and systematically investigated the path integration hypothesis in our paper “No Free Lunch from Deep Learning in Neuroscience: A Case Study through Models of the Entorhinal-Hippocampal Circuit” (NFL, for short) [27]. NFL combined large-scale numerical experiments and mathematical guidance to show that deep recurrent neural networks trained to path integrate frequently learn to do so, but almost never learn grid-like tuning, unless grid-like tuning is inserted via handcrafted supervised targets designed specifically to produce grid-like tuning. In this paper, we:

1. Restate the key evidence from NFL substantiating these conclusions.
2. Address a response to NFL by leading authors of the path integration hypothesis [32] (henceforth, “the Response”), explaining how the Response misinterprets NFL’s findings and in doing so validates NFL’s core findings.

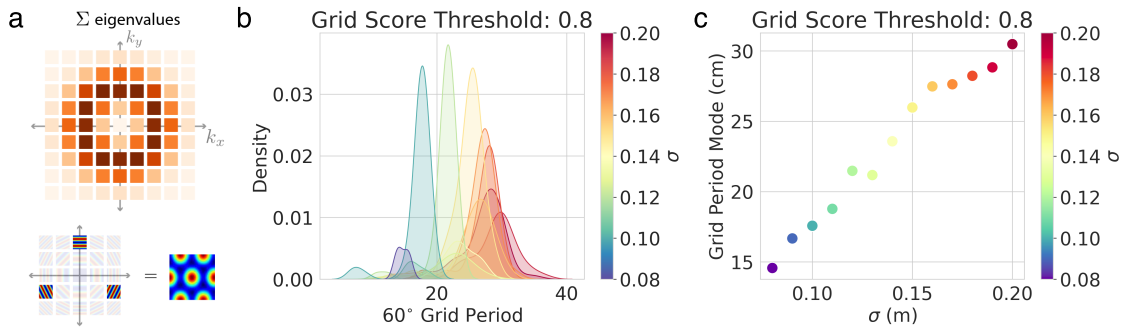


Figure 2: (a) The “Difference of Softmaxes” (DoS) supervised target is specifically chosen to place Fourier power on an annulus of a desired radius; doing so embeds grid-like tuning in the supervised targets that deep path integrators are then trained to predict. (b) The researchers’ choice of DoS targets inserts a single grid periodicity into the networks. (c) The choice of DoS hyperparameters sets the periodicity of the grid-like tuning. Subfigure (a) from [33], (b-c) from [27].

3. Evaluate the validity of results from the “unified theory for the origin of grid cells” [31] and demonstrate that the Unified Theory is *at best* an occasionally suggestive description of trained deep path integrators.

2 Evidence Challenging the Path Integration Hypothesis

NFL challenged the path integration hypothesis by demonstrating that grid-like tuning does not emerge in deep recurrent networks trained to path integrate unless researchers insert grid-like tuning in supervised targets designed to produce grid-like tuning [27].

To substantiate these claims, NFL used large scale hyperparameter sweeps to train over 11,000 deep recurrent neural networks to path integrate. Most learned to successfully path integrate, but few learned *possible* grid-like tuning (Fig. 1a). Grid-like tuning appeared only when researchers chose a specific and otherwise hard-to-justify supervised target called “Difference-of-Softmaxes” (DoS) (Fig. 1ab) that was designed to produce grid-like tuning [31, 33] (Fig. 2a). This DoS target does not improve path integration performance: thus it is not a functionally motivated choice. The DoS is also not a biologically motivated choice [26]. NFL demonstrated that the choice of DoS inserts into the networks a single grid periodicity (Fig. 2b), and the grid periodicity is directly set by the researchers’ chosen DoS hyperparameters (Fig. 2c). Further, [31] notes that multiple grid modules, a key property of biological grid cells [36], do not emerge from their framework, which NFL confirmed, meaning one cannot evaluate whether the artificial grid modules follow biological grid module ratios.

Based on this and additional evidence, NFL concluded that grid-like tuning in trained deep path integrators is often a post-hoc result in which losses, target functions, and regularizers are explored until the result resembles grid cells.

3 The Response to NFL [32] Misinterprets Then Corroborates NFL [27]

The message of multiple prominent papers is that grid cells naturally result from a simple path integration task [31, 33, 23, 34]; these papers underemphasize or omit that *the* key ingredient for grid-like tuning in deep networks is researchers designing supervised targets to produce grid-like tuning. Even after the publication of NFL, a Neuron paper [34] neglected mentioning the critical importance of the DoG or DoS supervised targets in the Highlights, Summary and Introduction while emphasizing

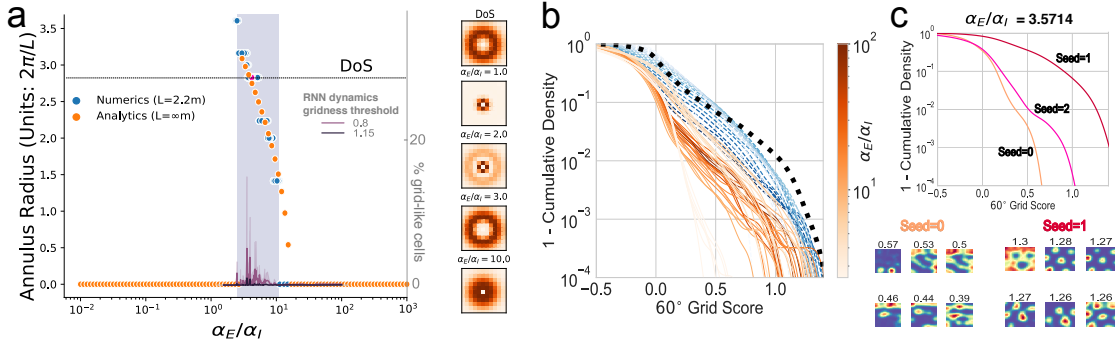


Figure 3: (a) In the Unified Theory, only a narrow range of hyperparameters σ_E, σ_I should (analytically and numerically) contain Fourier power on an annulus of an appropriate radius to create grid-like tuning (right). Networks trained on Difference-of-Gaussians (DoG) supervised targets do not develop grid-like tuning outside a narrow hyperparameter interval and inconsistently develop grid-like tuning inside the hyperparameter interval. (b) Across α_E/α_I DoG ratios, DoG networks (orange) generally score worse than filtered-and-thresholded noise (blue) and worse than DoS (black). (c) More densely sweeping within the theoretically feasible α_E/α_I DoG region, and choosing α_E and α_I magnitudes closely matching DoS, shown in (c) still showed only one ratio that did at least as well as a DoS. This result was true for only one out of three seeds: two “sibling” runs with otherwise identical settings produced poor gridness. All networks learn to path integrate. Subfigures from [27].

the less-central claim about path integration: “Here, we forge a link between the problem of path integration and the existence of hexagonal grids, by demonstrating that such grids arise in neural networks trained to path integrate under simple biologically plausible constraints. Moreover, we develop a unifying theory for why hexagonal grids are ubiquitous in path-integrator circuits.”

The Response to NFL [32] misinterprets NFL from its outset: “[NFL] presented a sequence of simulation results and some theoretical analysis suggesting prior work involved non-transparent fine-tuning.” NFL is not about fine-tuning. NFL is about prior work selecting improbable readout functions to drive grid-like tuning, but presenting the results as a tale of grid cells resulting naturally and directly from optimization for path integration.

The Response then corroborates NFL’s central findings at length (Sec. 2 in [32]). We summarize the “unified theory” for completeness (App. B) and highlight our key takeaway here: *The “unified theory” is not a theory of path integration, nor of deep recurrent neural networks, nor of learning dynamics of deep recurrent neural networks trained to path integrate. Rather, it is a theory of Fourier structure in supervised targets.* The “unified theory” starts with the static function approximation model of prior work [6], and studies learning representations \mathbf{g} to predict/reconstruct a given supervised target P :

$$\min_{\mathbf{g}} \|\mathbf{P} - \mathbf{g}\mathbf{g}^T \mathbf{P}\|^2 + \sigma(\mathbf{g}) \quad \text{subject to } \mathbf{g}^T \mathbf{g} = 1. \quad (1)$$

This optimization problem’s minima are set by the choice of supervised targets P , specifically through the autocorrelation matrix $\Sigma_{xx'} = (PP^T)_{xx'}$. The Unified Theory designs P using DoS tuning curves to embed grid-like tuning into Σ ’s Fourier structure (Fig. 2). These supervised targets drive networks (recurrent or feedforward) to develop grid-like tuning without regard to path integration. Indeed, the authors showed shallow nonlinear autoencoders can develop grid-like tuning when trained on these targets [31], despite the complete absence of a velocity signal.

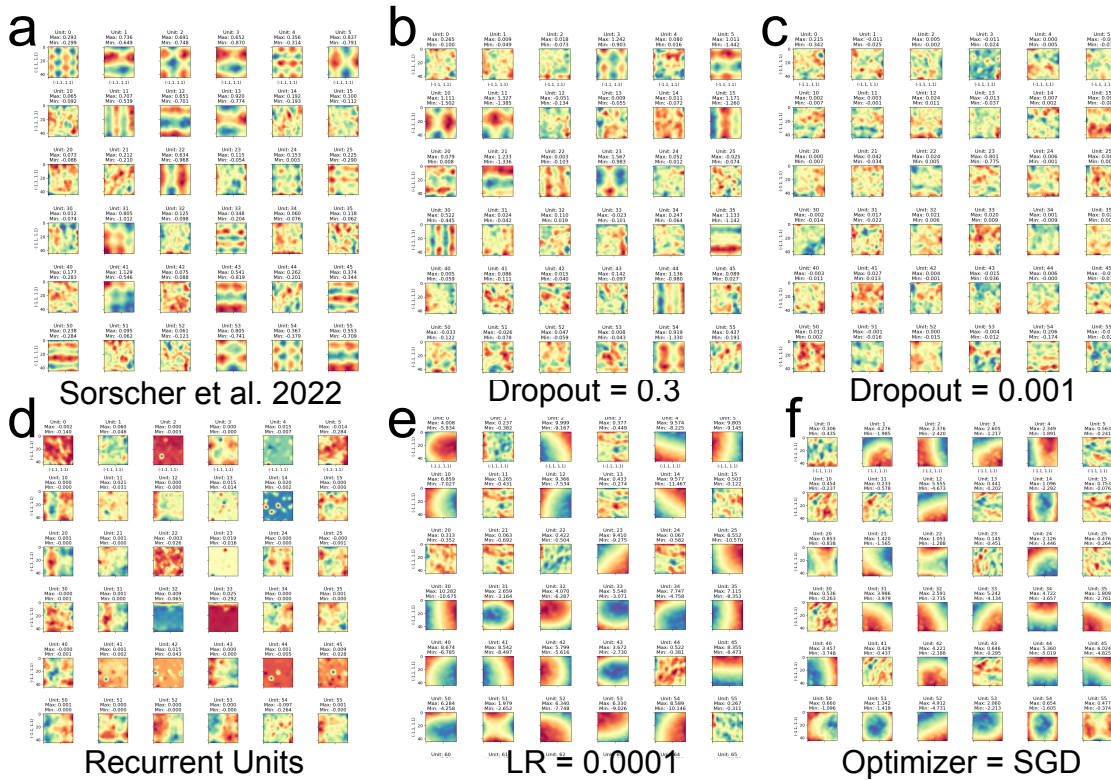


Figure 4: (a) Code by Ganguli et al. [34] to show that Gaussian targets can give (square) lattices uses a more-elaborate architecture with several additional components adapted from previous work [1], rather the architecture of their earlier publications[31, 33, 23]. (b-c) Reducing or removing dropout from this model causes lattice-like responses to disappear without affecting path integration performance. (d) The recurrent units of the architecture also display no lattices. (e-f) Changing other implementation details, such as the learning rate or optimizer, also causes the lattices to disappear without affecting path integration performance. See also Figs. 11, 12, 13

4 Evidence Challenging the Unified Theory [31]

The Response to NFL claimed that the Unified Theory [31], as a complete theory for the emergence of grid cells in deep path integrators, already contained the results shown in NFL – this is a point we disagree with. To see why, we turn to the Unified Theory. First, as noted in the previous section, the Unified Theory of grid cell emergence in deep path integrators is not that – it is a theory of the Fourier structure of supervised targets. Second, two critical assumptions made by the Unified Theory (i.e. place cells are translationally invariant as a population, and place cells have center-surround tuning individually) were recently shown to be biologically unlikely [26]. Here, we set aside these two issues and explore the results of the Unified Theory and their relevance to deep path integrators, which we find to be at most “occasionally suggestive,” not exact or comprehensive.

Fourier structure in supervised targets is necessary but insufficient. The Unified Theory predicts that supervised targets with Fourier power on an annulus of an appropriate radius should result in grid-like tuning in path integrating deep networks. We derive the spectrum for Difference-of-Gaussians (DoG) targets and show that only a narrow slice of hyperparameters possess the appropriate radius (Fig. 3a). But even when networks are trained inside this narrow interval, they only occasionally develop grid-like tuning (Fig. 3ab). The very specific Difference-of-Softmaxes (DoS) leads to grid-like




Figure 5: Deep path integrators with ReLU nonlinearities and Gaussian readouts, sweeping Gaussian scale, do not learn lattices: RNNs (Red) Gaussian readout scale was swept $\sigma_E \in [5, 50]$ cm in increments of 2.5 cm. For negative control, we consider Uniform $(-1, 1)$ noise filtered-then-ReLU-thresholded (Blue), and for positive control, we consider RNNs with ReLU nonlinearities trained on ideal-width (12 cm) DoS readouts (Black). RNN+ReLU+Gaussian 60° score distributions are dominated by almost all noise filter widths, and the ideal-width DoS 60° score distribution dominates most noise filter widths (especially at the tail). All other hyperparameters match [34]’s Page e2.

tuning more often than DoG target (Fig. 3b), but grid-like tuning still does not reliably emerge and is seed dependent (Fig. 3c). These findings show that the Unified Theory’s conditions are not sufficient. To reiterate, the criticism is not that deep learning requires hyperparameter tuning, but rather that a sequence of highly-specific and biologically-implausible choices orthogonal to path integration make clear that grid cell tuning was the goal of the entire setup, and that grid cells did not simply pop out as a natural consequence of path integration.

Gaussian-like supervised targets do not produce grid-like tuning without additional elements. A key point of contention is whether Gaussian supervised targets produce grid-like tuning. Early work claimed yes [1, 31, 33], but NFL showed that Gaussian supervised targets almost never produce lattices (either hexagonal or square) [27]. This finding of NFL was contested by [32], which claimed that with specific hyperparameters Gaussian targets should lead to grid cells in deep path integrators, and released new code in support. Here we present three new results showing that Gaussian supervised targets in deep path integrators do not on their own produce lattices. This debate matters for two reasons. Firstly, Gaussian supervised targets might be easier to justify biologically than DoG and DoS targets, and secondly, whether Gaussian targets produce grids sheds light on the size of the space of grid-producing supervised targets.

Firstly, we deconstruct the new code released by the authors of [34], which produces (square) lattices from Gaussian targets. This code uses a different architecture than the deep recurrent networks used in their prior work [31, 33, 23] and is similar to the model of [1]. It is described by the authors as: “a significantly more complex LSTM architecture [...] The “grid cells” in this architecture are not recurrently connected, but reside in a dense layer immediately following the LSTM. Importantly, the grid cells are now subject to dropout at a rate of 0.5.” We found that the high dropout rate, together with specific choices of high learning rates and optimizers, are pivotal to achieving lattice-like responses when the target functions are Gaussian (Fig. 4, App. A, Figs. 11, 12, 13), even though good path integration can be achieved without these features. This also bolsters NFL’s critique of the narrative of grid cell emergence in [1].

Secondly, returning to the authors’ original code from [23], we densely sweep the recommended Gaussian receptive field, testing (a) recurrent networks with ReLU nonlinearities that the Unified Theory predicts should learn hexagonal lattices and (b) Tanh nonlinearities that the Unified Theory




Figure 6: Deep path integrators with Gaussian readouts and Tanh nonlinearity achieve lower grid score distributions than low-pass-filtered-then-ReLU-thresholded noise: RNNs’ (Red) Gaussian readout scale was swept $\sigma_E \in [5, 50]$ cm in increments of 2.5 cm. For negative control, we consider Uniform(0, 1) noise filtered(Blue), and for positive control, we consider RNNs with ReLU nonlinearities trained on ideal-width (12 cm) DoS readouts. RNN+ReLU+Gaussian 90° score distributions are dominated by almost all noise filter widths, and the ideal-width DoS 90° score distribution dominates most noise filter widths (especially at the tail). All other hyperparameters match [34]’s Table on Page e2.

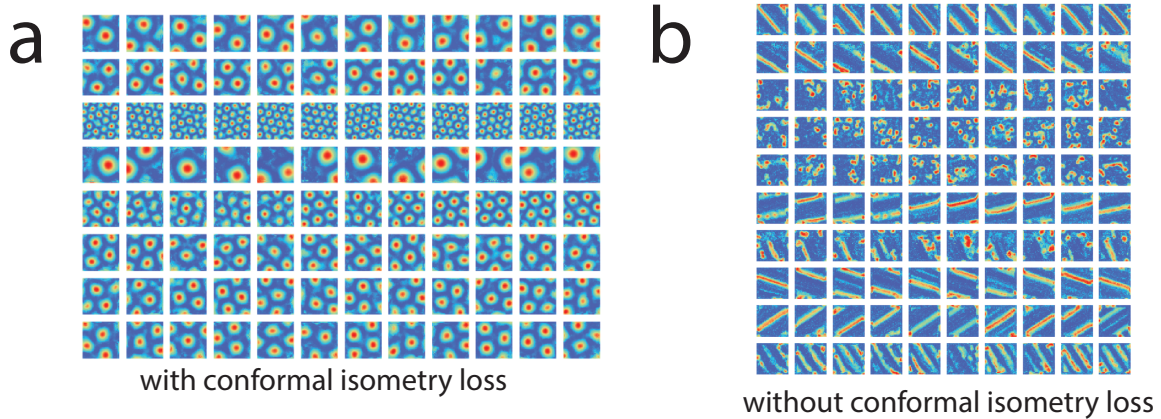


Figure 7: Independent research by Xu et al. 2022 [39] shows Gaussian supervised targets do not create grid-like tuning [39]. In order to obtain lattices with Gaussians, Xu et al. 2022 proposed a novel conformal isometry loss (left). Networks trained without the loss, i.e., solely with Gaussian supervised targets do not learn lattice-like units (right). Figure included with permission from authors.




Figure 8: Using DoS supervised targets with ideal hyperparameters, switching from Tanh to ReLU increases 60° scores (left), but also increases 90° scores (center). Example square rate maps (right). 3 seeds $\in \{0, 1, 2\}$ per nonlinearity; all other hyperparameters matched [31, 34]’s code defaults.

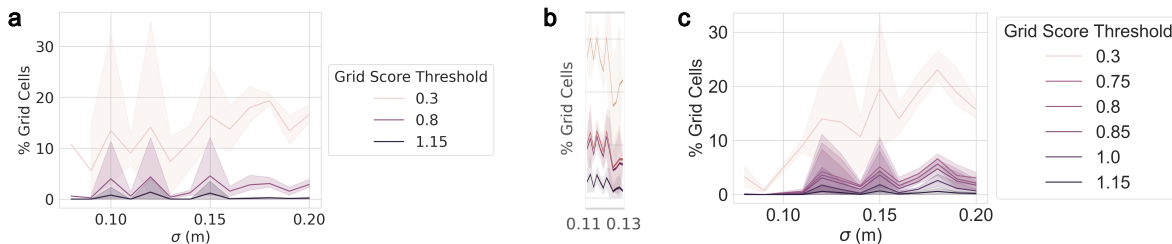


Figure 9: (a) NFL discovered small alterations to the DoS target width σ_E can result in the disappearance of grid units [27]. (b) The Response disputes this result [32], albeit over a much narrower range of σ_E values. (c) Rerunning our previous sweep but with larger networks (4096 units, previously 1024 units) to exactly match newly published hyperparameters [34] again shows that grid-like tuning is highly sensitive to hyperparameters. All networks achieve low position decoding error.

predicts should learn square lattices. Following [31], we used filtered and thresholded noise as the null distribution to determine when observed patterns were more significant than chance in all cases. Across nonlinearities and receptive field widths, the sweeps show networks do not learn grid-like tuning from Gaussian targets (Figs 5, 6).

Thirdly, we highlight recent independent research by Xu et al. 2022 [39] confirming that Gaussian supervised targets do not produce grid-like tuning (Fig. 7).

Contrary to Unified Theory, non-negativity produces hexagonal *and* square lattices. A key result of the Unified Theory [34, 31] is that non-negativity favors hexagonal lattices over square lattices: “Once a nonlinear constraint such as non-negativity is added, the optimization favors a single type of map corresponding to hexagonal grid cells.” We find that the default hexagonal grid-forming hyperparameters consistently produce networks that simultaneously learn both hexagonal *and* square lattices (Fig. 8): Replacing Tanh with ReLU for non-negativity causes an upward shift in hexagonal 60° scores, but also causes a larger upward shift in square 90° scores *in the same networks*. Visual examination of units with high 90° scores in ReLU networks confirms square tuning curves. Training here and in NFL [27] is $20\times$ longer than in [34], so insufficient training cannot explain the result.

Sensitivity to supervised target hyperparameters. NFL reported that even with DoS supervised targets, the emergence of grid-like tuning is highly sensitive to the receptive field width σ_E (Fig. 9A) when sweeping over the range 8cm to 20cm [27]. The Response stated that DoS readouts are robust to σ_E [32] but showed only a narrow range of 11cm to 13cm (Fig. 9B). The only noted difference in published hyperparameters between the NFL result, which was based on code from [23], and the




Figure 10: Three independent research labs studying three different brain circuits in three different species across three different modalities found a consistent relationship: test R^2 in neural regressions is approximately an affine function of the log of the effective dimensionality (Eqn. 2) of the artificial neural activations. Figures from (a) Spatial Navigation in Mouse Medial Entorhinal Cortex [27], (b) Vision in Macaque IT Cortex [8] (note the log-X scaling), and (c) Audition in Human fMRI [37].

Response was the number of hidden units (ours: 1024 units; their Response: 4096 units). We find that the NFL result on sensitivity to supervised target field width holds in 4096-unit networks (Fig. 9C).

5 Why Can Neural Regressions Be So Confidently Wrong?

A widely used method to test models in modern computational neuroscience is neural regression [40, 18]: predicting biological neural recordings from activity in trained deep networks, typically using linear regression. Nayebi et al. 2021 [23] used this technique to accurately predict tetrode recordings from mouse medial entorhinal cortex (MEC) from the path integrating networks and thus to conclude that their networks were excellent models of MEC.

NFL hypothesized that trained deep networks might achieve higher neural predictivity scores than competitor models simply because trained deep networks provide higher dimensional, richer basis functions that could be able to generate closer fits to the data (and by extension better within-sample predictions of the data) than lower dimensional, poorer basis functions supplied by alternative models [27]. NFL found initial supporting evidence: the reported test Pearson correlation (termed “neural predictivity”) approximately followed a linear-log relationship with the effective dimensionality (sometimes termed intrinsic dimensionality or participation ratio) [21, 25, 24] of the networks’ representations (Fig. 10a), defined as:

$$\text{Participation Ratio a.k.a. Effective Dimensionality} \stackrel{\text{def}}{=} \frac{(\sum_r \lambda_r)^2}{\sum_r \lambda_r^2}. \quad (2)$$

NFL then conjectured similar results would be found in other modalities. Two independent research groups studying two different modalities (vision and audition) in two different brain circuits (visual ventral stream and auditory cortex) across two different species (macaque and human) using two different recording technologies (ECoG, fMRI) found a quantitatively similar relationship: in vision, Elmoznino et al. 2022 [8] found the same approximate linear-log relationship in regressions between convolutional networks and macaque IT Cortex (Fig 10b), and in audition, Tuckute et al. 2022 [37] found the same linear-log relationship regressions between deep networks and human fMRI responses (Fig 10c).

NFL and these results call the neural regression methodology into question. Han, Poggio & Cheung 2023 [14] sought to validate the neural regression methodology by asking how well the methodology can identify the correct network if the correct network is included within the set of candidate models, and concluded that the results were poor. Canatar et al. 2023 [3] provided a theory showing candidate networks can achieve high predictivity scores so long as candidate networks have certain spectral properties relative to the target networks, without any assumption of a mechanistic, biological or

otherwise meaningful relationship between the candidate and target networks. One possibility is that neural regressions are spiritually similar to Ramsey Theory [17]: networks trained with more parameters, data and compute are more likely to offer richer basis functions such that there will exist a subset of basis functions able to predict whatever neural recordings, and we know that both underparameterized and overparameterized linear regression can generalize [29, 30]. Mathematically formalizing and proving such a possibility for neural regressions is non-trivial, but collectively, these results call for further investigation into what exactly the neural regression methodology tells us about neural representations.

6 Discussion and Future Directions

NFL, the Response, and the current document together show that path integration is an insufficient condition for the emergence of grid cells in deep neural networks. Contrary to the claims made in the various papers on grid cell emergence in deep networks, successfully training networks to path integrate does not automatically lead to grid cells – unrelated and special ingredients must be added to obtain grid cells. Thus, to obtain grid cells in such networks requires effectively baking them in, with the researcher implicitly using grid cells as a major part of their own loss function in their investigations. *While there is no harm in explicitly taking such an approach, and indeed such approaches can be enlightening when the conditions of emergence are carefully explored and delineated, our evidence shows it is incorrect to claim that grid cells drop out of deep networks when simply trained to path integrate.*

If path integration is an insufficient task for grid cells to emerge, what might be sufficient? Prior work showing that grid cells are a powerful neural code with exponential capacity and error correction abilities [10, 35, 22, 38] suggests a potential direction: our hypothesis is that the grid code emerges by directly optimizing the learned representation for space, in a coding-theoretic sense, so that coding states are well-distributed within the available coding volume. Work in this direction exists [12, 11, 7, 39], and a recent work of ours, “Self-Supervised Learning of Representations for Space Generates Multi-Modular Grid Cells” [28], demonstrates that training deep neural networks with self-supervised learning to both path integrate and optimize the packing of the code in the coding space can generate multi-modular grid cells - without any supervised targets.

7 Acknowledgements

IRF is supported by the Simons Foundation through the Simons Collaboration on the Global Brain, the ONR, the NSF, and the K. Lisa Yang ICoN Center. SK acknowledges support from NSF grants No. 2046795, 1934986, 2205329, and NIH 1R01MH116226-01A, NIFA 2020-67021-32799, the Alfred P. Sloan Foundation, and Google Inc. RS is supported by a Stanford Data Science Scholarship. MK is supported in part by a MathWorks Science Fellowship and the MIT Department of Physics.

References

- [1] A. Banino, C. Barry, B. Uribe, C. Blundell, T. Lillicrap, P. Mirowski, A. Pritzel, M. J. Chadwick, T. Degris, J. Modayil, G. Wayne, H. Soyer, F. Viola, B. Zhang, R. Goroshin, N. Rabinowitz, R. Pascanu, C. Beattie, S. Petersen, A. Sadik, S. Gaffney, H. King, K. Kavukcuoglu, D. Hassabis, R. Hadsell, and D. Kumaran. Vector-based navigation using grid-like representations in artificial agents. *Nature*, 557(7705):429–433, May 2018.
- [2] Y. Burak and I. R. Fiete. Accurate path integration in continuous attractor network models of grid cells. *PLoS computational biology*, 5(2):e1000291, 2009.
- [3] A. Canatar, J. Feather, A. Wakhloo, and S. Chung. A spectral theory of neural prediction and alignment. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [4] M. C. Cross and P. C. Hohenberg. Pattern formation outside of equilibrium. *Reviews of modern physics*, 65(3):851, 1993.
- [5] C. J. Cueva and X.-X. Wei. Emergence of grid-like representations by training recurrent neural networks to perform spatial localization. *International Conference on Learning Representations*, page 19, 2018.
- [6] Y. Dordek, D. Soudry, R. Meir, and D. Derdikman. Extracting grid cell characteristics from place cell inputs using non-negative principal component analysis. *eLife*, 5:e10094, Mar. 2016. Publisher: eLife Sciences Publications, Ltd.
- [7] W. Dorrell, P. E. Latham, T. E. Behrens, and J. C. Whittington. Actionable neural representations: Grid cells from minimal constraints. *arXiv preprint arXiv:2209.15563*, 2022.
- [8] E. Elmoznino and M. F. Bonner. High-performing neural network models of visual cortex benefit from high latent dimensionality, July 2022. Pages: 2022.07.13.499969 Section: New Results.
- [9] B. Ermentrout. Neural networks as spatio-temporal pattern-forming systems. *Reports on progress in physics*, 61(4):353, 1998.
- [10] I. R. Fiete, Y. Burak, and T. Brookings. What Grid Cells Convey about Rat Location. *Journal of Neuroscience*, 28(27):6858–6871, July 2008. Publisher: Society for Neuroscience Section: Articles.
- [11] R. Gao, J. Xie, X.-X. Wei, S.-C. Zhu, and Y. N. Wu. On path integration of grid cells: group representation and isotropic scaling. *Advances in Neural Information Processing Systems*, 34:28623–28635, 2021.
- [12] R. Gao, J. Xie, S.-C. Zhu, and Y. N. Wu. Learning grid cells as vector representation of self-position coupled with matrix representation of self-motion. *International Conference on Learning Representations*, 2019.
- [13] T. Hafting, M. Fyhn, S. Molden, M.-B. Moser, and E. I. Moser. Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052):801–806, 2005.
- [14] Y. Han, T. A. Poggio, and B. Cheung. System identification of neural systems: If we got it right, would we know? In *International Conference on Machine Learning*, pages 12430–12444. PMLR, 2023.
- [15] I. Kanitscheider and I. Fiete. Training recurrent networks to generate hypotheses about how the brain solves hard navigation problems. *arXiv:1609.09059 [q-bio]*, Dec. 2016. arXiv: 1609.09059.
- [16] I. Kanitscheider and I. Fiete. Emergence of dynamically reconfigurable hippocampal responses by learning to perform probabilistic spatial reasoning. Dec. 2017.

- [17] M. Katz and J. Reimann. *An introduction to Ramsey theory*, volume 87. American Mathematical Soc., 2018.
- [18] S.-M. Khaligh-Razavi and N. Kriegeskorte. Deep supervised, but not unsupervised, models may explain it cortical representation. *PLoS computational biology*, 10(11):e1003915, 2014.
- [19] M. Khona, S. Chandra, and I. R. Fiete. From smooth cortical gradients to discrete modules: A biologically plausible mechanism for the self-organization of modularity in grid cells. *bioRxiv*, pages 2021–10, 2022.
- [20] M. Khona and I. R. Fiete. Attractor and integrator networks in the brain. *Nature Reviews Neuroscience*, 2022.
- [21] A. Litwin-Kumar, K. D. Harris, R. Axel, H. Sompolinsky, and L. F. Abbott. Optimal Degrees of Synaptic Connectivity. *Neuron*, 93(5):1153–1164.e7, Mar. 2017.
- [22] A. Mathis, A. V. M. Herz, and M. Stemmler. Optimal Population Codes for Space: Grid Cells Outperform Place Cells. *Neural Computation*, 24(9):2280–2317, Sept. 2012.
- [23] A. Nayebi, A. Attinger, M. Campbell, K. Hardcastle, I. Low, C. S. Mallory, G. Mel, B. Sorscher, A. H. Williams, S. Ganguli, L. Giocomo, and D. Yamins. Explaining heterogeneity in medial entorhinal cortex with task-driven neural networks. In *Advances in Neural Information Processing Systems*, volume 34, pages 12167–12179. Curran Associates, Inc., 2021.
- [24] S. Recanatesi, S. Bradde, V. Balasubramanian, N. A. Steinmetz, and E. Shea-Brown. A scale-dependent measure of system dimensionality. *Patterns*, 3(8), 2022.
- [25] S. Recanatesi, G. K. Ocker, M. A. Buice, and E. Shea-Brown. Dimensionality in recurrent spiking networks: Global trends in activity and local origins in connectivity. *PLoS computational biology*, 15(7):e1006446, 2019.
- [26] R. Schaeffer, M. Khona, A. Bertagnoli, S. Koyejo, and I. Fiete. Testing assumptions underlying a unified theory for the origin of grid cells. *Neural Information Processing Systems Workshops: UniReps, NeurReps, AI for Science*, 2023.
- [27] R. Schaeffer, M. Khona, and I. R. Fiete. No Free Lunch from Deep Learning in Neuroscience: A Case Study through Models of the Entorhinal-Hippocampal Circuit. page 2022.08.07.503109, Aug. 2022. Section: New Results.
- [28] R. Schaeffer, M. Khona, T. Ma, C. Eyzaguirre, S. Koyejo, and I. R. Fiete. Self-supervised learning of representations for space generates multi-modular grid cells. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [29] R. Schaeffer, M. Khona, Z. Robertson, A. Boopathy, K. Pistunova, J. W. Rocks, I. R. Fiete, and O. Koyejo. Double descent demystified: Identifying, interpreting & ablating the sources of a deep learning puzzle. *arXiv preprint arXiv:2303.14151*, 2023.
- [30] R. Schaeffer, Z. Robertson, A. Boopathy, M. Khona, I. Fiete, A. Gromov, and S. Koyejo. Divergence at the interpolation threshold: Identifying, interpreting & ablating the sources of a deep learning puzzle. *NeurIPS 2023 Workshop on Mathematics of Modern Machine Learning*, 2023.
- [31] B. Sorscher, G. C. Mel, S. Ganguli, and S. A. Ocko. A unified theory for the origin of grid cells through the lens of pattern formation. *Advances in Neural Information Processing Systems*, 2019.
- [32] B. Sorscher, G. C. Mel, A. Nayebi, L. Giocomo, D. Yamins, and S. Ganguli. When and why grid cells appear or not in trained path integrators, Nov. 2022. Pages: 2022.11.14.516537 Section: Contradictory Results.

- [33] B. Sorscher, G. C. Mel, S. A. Ocko, L. Giacomo, and S. Ganguli. A unified theory for the computational and mechanistic origins of grid cells. Technical report, bioRxiv, Dec. 2020. Section: New Results Type: article.
- [34] B. Sorscher, G. C. Mel, S. A. Ocko, L. M. Giacomo, and S. Ganguli. A unified theory for the computational and mechanistic origins of grid cells. *Neuron*, 0(0), Oct. 2022. Publisher: Elsevier.
- [35] S. Sreenivasan and I. Fiete. Grid cells generate an analog error-correcting code for singularly precise neural computation. *Nature Neuroscience*, 14(10):1330–1337, Oct. 2011. Number: 10 Publisher: Nature Publishing Group.
- [36] H. Stensola, T. Stensola, T. Solstad, K. Frøland, M.-B. Moser, and E. I. Moser. The entorhinal grid map is discretized. *Nature*, 492(7427):72–78, 2012.
- [37] G. Tuckute, J. Feather, D. Boebinger, and J. H. McDermott. Many but not All Deep Neural Network Audio Models Predict Auditory Cortex Responses and Exhibit Hierarchical Layer-Region Correspondence. In *2022 Conference on Cognitive Computational Neuroscience*, San Francisco, 2022. Cognitive Computational Neuroscience.
- [38] X.-X. Wei, J. Prentice, and V. Balasubramanian. A principle of economy predicts the functional architecture of grid cells. *Elife*, 4:e08362, 2015.
- [39] D. Xu, R. Gao, W.-H. Zhang, X.-X. Wei, and Y. N. Wu. Conformal isometry of lie group representation in recurrent network of grid cells. *arXiv preprint arXiv:2210.02684*, 2022.
- [40] D. L. Yamins, H. Hong, C. F. Cadieu, E. A. Solomon, D. Seibert, and J. J. DiCarlo. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the national academy of sciences*, 111(23):8619–8624, 2014.

A Further Investigation of Ganguli et al.’s Gaussian Result

We include additional results from investigating the newly released code of Ganguli et al. 2022 [32] that claims to produce square lattices from Gaussian supervised targets. We find that a combination of additional implementation details unrelated to path integration performance are critical for lattice-like tuning. Specifically, (1) high dropout probability e.g., 50% (Fig. 11), (2) high learning rate, e.g., 0.1 (Fig. 12) and (3) choice of optimizer, e.g., Adam (Fig. 13) increases the 90° grid score distribution without improving path integration performance.

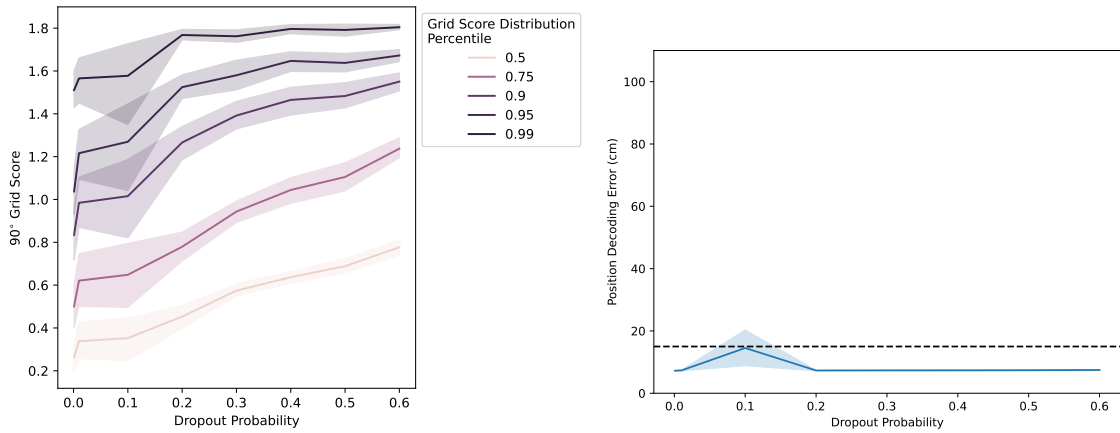


Figure 11: Left: High dropout probability increases the 90° grid score distribution. Right: High dropout does not increase path integration performance.




Figure 12: Left: High learning rate increases the 90° grid score distribution. Right: High learning rate does increase path integration performance.

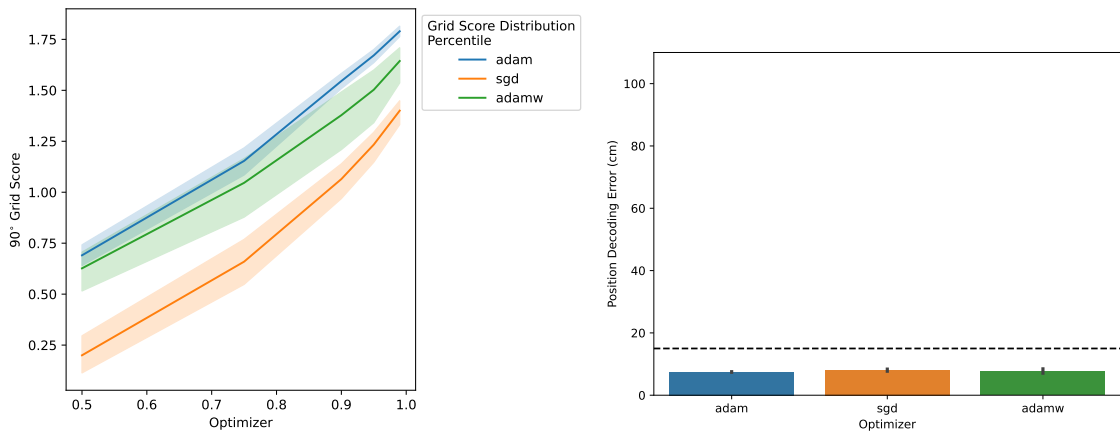


Figure 13: Left: Specific choice of optimizer increases the 90° grid score distribution. Right: Specific choice of optimizer does not increase path integration performance.

B Reproduction & commentary on the Unified Theory of [31]

Here, we reproduce the theory of [31, 33], highlighting the simplifying assumptions made by the authors, to provide an explanation for the mismatch between theory and simulations of trained path integrating networks. Further, we reproduce, very briefly, the theory of pattern formation in first-principles continuous attractor models, highlighting similarities and differences between these two theories.

This theory does not deal with dynamics of path integration or learning dynamics of a deep recurrent network from initialization, but rather concerns the problem of readout reconstruction. This leads us to the first assumption:

Assumption 1 (A1): The hypothetical hidden representation $G \in \mathbb{R}^{n_x \times n_g}$ is some function of space. Here n_x is the number of spatial locations and n_g is the number of hidden units. This is a subtle but significant assumption because, for recurrent networks given velocity inputs, the networks' representations are not a function of space, but rather develop into a function of space (i.e., builds a continuous attractor) over the course of training. For a better understanding of why the assumption of building a continuous manifold of fixed points is significant, see literature of the theory of continuous attractors, which is briefly reviewed in [20].

Under A1, consider a feedforward mapping $\hat{P} \stackrel{\text{def}}{=} GW$ where $W \in \mathbb{R}^{n_g \times n_p}$. Here n_p is the number of readout units. One can define the readout reconstruction error as the mean square loss between the readout target $P \in \mathbb{R}^{n_x \times n_p}$ and prediction $\hat{P} \stackrel{\text{def}}{=} GW$:

$$\mathcal{E}(G, W) \stackrel{\text{def}}{=} \|P - \hat{P}\|_F^2 = \|P - GW\|_F^2 \quad (3)$$

Assumption 2 (A2): Linear readout W relaxes, reaching its optimum much faster than G changes, so that we can replace W with its optimal ordinary least squares value for fixed G, P :

$$W^*(G, P) = (G^T G)^{-1} G^T P \quad (4)$$

Substituting $W^*(G, P)$ into the loss for W yields the error as a function of P and G :

$$\mathcal{E}(G, P) = \|P - G(G^T G)^{-1} G^T P\|_F^2 \quad (5)$$

Assumption 3 (A3): G 's columns can be made orthonormal i.e. $G^T G = I_{n_g}$. This means that each grid unit has the same average firing rate over all space, and that any 2 grid cells do not overlap. There is no provided reason why the dynamics of (stochastic) gradient descent in RNNs, with or without regularization, can or should find this solution.

Then, we can write down a Lagrangian for this optimization problem, using a Lagrange multiplier λ , for this constraint:

$$\mathcal{L} = -\mathcal{E}(G, P) - \lambda(G^T G - I_{n_g}) \quad (6)$$

We can then set $G^T G$ to I in the error term and the Lagrangian is now written as:

$$\mathcal{L} = -\|P - GG^T P\|_F^2 - \lambda(G^T G - I_{n_g}) \quad (7)$$

$$\mathcal{L} = -\text{Tr}[(P - GG^T P)^T (P - GG^T P)] - \lambda(G^T G - I_{n_g}) \quad (8)$$

Here, the identity $\|M\|_F^2 = \text{Tr}(M^T M)$ has been employed.

The trace term can be simplified further using the cyclic permutation property of Trace: $\text{Tr}(ABC) = \text{Tr}(CAB)$,

$$\begin{aligned} \text{Tr}[(P - GG^T P)^T (P - GG^T P)] &= \text{Tr}(P^T P) + \text{Tr}[(GG^T P)^T (GG^T P)] - \text{Tr}(P^T GG^T P) - \text{Tr}(P^T GG^T P) \\ &= \text{Tr}(P^T P) + \text{Tr}[G^T (PP^T) G (G^T G)] - 2\text{Tr}(G^T PP^T G) \end{aligned}$$

Here $\Sigma \stackrel{\text{def}}{=} \frac{1}{n_p} PP^T \in \mathbb{R}^{n_x \times n_x}$ is the readout spatial correlation matrix. We can also drop the G independent term above. Using the trace identity again, this term simplifies to $\text{Tr}(G^T PP^T G)$. Hence the total simplified Lagrangian is then:

$$\mathcal{L} = \text{Tr}[G^T \Sigma G - \lambda(G^T G - I_{n_g})] \quad (9)$$

Considering gradient learning dynamics, one gets the following evolution equation for G :

$$\frac{d}{dt}G = \nabla_G \mathcal{L} \Rightarrow \frac{d}{dt}G = \Sigma G - \lambda G \quad (10)$$

[31, 33] then simplify further analysis by considering a single grid unit. This corresponds to replacing the $n_x \times n_g$ matrix G by the $n_x \times 1$ column vector g :

$$\frac{d}{dt}g = \Sigma g - \lambda g \quad (11)$$

Thus, this linear dynamical system captures how the pattern g of a unit evolves with gradient learning. Further, [31, 33] conclude the eigenvectors corresponding to the subspace of the *top eigenvalue* form the optimal pattern, since these eigenvectors will grow exponentially with the fastest rate.

Assumption 4 (A4): The readout spatial correlation Σ is translation-invariant over space i.e. $\Sigma_{x,x'} = \frac{1}{n_p} \sum_{i=1}^{n_p} p_i(x)p_i(x') = \frac{1}{n_p} \sum_{i=1}^{n_p} p_i(x+\Delta)p_i(x'+\Delta) = \Sigma_{x+\Delta,x'+\Delta} \forall \Delta$. This is a strong assumption on the readouts and not necessarily true in biology [26].

Assumption 5 (A5): The environment has periodic boundaries (or no boundaries, which corresponds to a continuum limit). [An alternative assumption in other parts of the derivations and numerics is (A5'): the assumption involves periodic boundary conditions with a small box size L ; the smallness of the box is important for the prediction of grids from Gaussian units. There are other biological plausibility concerns with making a specific choice for box size to explain grid cell responses, which exhibit high invariance to box size in familiar environments (see paragraph below, on ‘‘The nature of discretization effects’’).]

Under A4 and A5, the eigenmodes of Σ are exactly Fourier modes across space and thus form a periodic basis. The normalized eigenvectors are indexed by their wavelength, \mathbf{k} and are denoted $f_{\mathbf{k}}$ with corresponding eigenvalue $\lambda_{\mathbf{k}}$.

To calculate this eigenvalue, [31, 33] use Fourier theory:

$$\begin{aligned} \Sigma f_{\mathbf{k}} &= \lambda_{\mathbf{k}} f_{\mathbf{k}} \\ \implies \lambda_{\mathbf{k}} &= f_{\mathbf{k}}^\dagger \Sigma f_{\mathbf{k}} \end{aligned}$$

Here $f_{\mathbf{k}}^\dagger$ denotes the conjugate of the eigenvector $f_{\mathbf{k}}$.

Next, they rewrite in component form: $\Sigma_{x,x'} = 1/n_p (PP^T)_{x,x'} = 1/n_p \sum_{i=1}^{n_p} p_i(x)p_i(x')$

$$\begin{aligned} \lambda_{\mathbf{k}} &= f_{\mathbf{k}}^\dagger \Sigma f_{\mathbf{k}} = \sum_{i=1}^{n_p} \sum_{x,x'} \frac{1}{n_p} f_{\mathbf{k}}^*(x') p_i(x) p_i(x') f_{\mathbf{k}}(x) \\ &= \frac{1}{n_p} \sum_{i=1}^{n_p} \left(\sum_x p_i(x) f_{\mathbf{k}}^*(x) \right) \left(\sum_{x'} p_i(x') f_{\mathbf{k}}(x') \right) \\ &= \frac{1}{n_p} \sum_{i=1}^{n_p} \tilde{p}^*(k) \tilde{p}(k) = |\tilde{p}(k)|^2 \\ \implies \lambda_{\mathbf{k}} &= |\tilde{p}(k)|^2 \end{aligned}$$

So, [31, 33] conclude that the eigenvalue corresponding to eigenvectors with wavelength k is given by the corresponding power of the Fourier spectrum of the readout correlation matrix Σ . And thus the optimal pattern is the one which has the highest Fourier power in Σ .

Further, [31, 33] consider the effect of non-negativity perturbatively in the readout regression framework by phenomenologically adding a term to the Lagrangian, $= \int_x \sigma(g)dx$.

$$\mathcal{L} = \text{Tr}[G^T \Sigma G - \lambda(G^T G - I_{n_g})] + \int_x \sigma(g)dx \quad (12)$$

In Fourier space, [31, 33] show perturbatively that this amounts to a cubic interaction term, which is the leading order term that non-trivially distinguishes between nonlinearities such as ReLU and Sigmoid which break the $g \mapsto -g$ symmetry and nonlinearities such as Tanh which do not. Again, specializing to the single neuron Lagrangian,

$$\mathcal{L}_{int} = \int_{k, k', k''} g(k)g(k')g(k'')\delta(k + k' + k'')dkdk'dk''$$

This term effectively acts as a penalty for non-negativity. Here, it is important to point out that this cubic term appears not only for non-negativity, but rather *any* function that is not anti-symmetric. Negative activation functions such as slightly shifted Tanh can also have a cubic term. Thus, non-negativity is a special case, as has been noted in [31] Appendix B. We refer to this as **Assumption 6 (A6)**.

Under this assumption, [31, 33] conclude that the optimal pattern consists of a triplet of Fourier waves with equal amplitude and k -vectors that lie on an equilateral triangle, at 60° to each other.

Next, we examine the Fourier spectrum of a translationally invariant Gaussian readout $f(\Delta x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp(-(\Delta x^2)/2\sigma^2)$ under the assumptions of this theory. For simplicity and to provide intuition we write its Fourier transform in 1d, which is given by another Gaussian. The peak of the Fourier spectrum is at $k = 0$, or the DC, non-periodic mode:

$$\begin{aligned} \tilde{f}(k) &= \int_{\mathcal{R}} \frac{1}{\sqrt{2\pi\sigma^2}} \exp(-(\Delta x^2)/2\sigma^2) e^{ik\Delta x} d\Delta x \\ &= \exp(-k^2\sigma^2/2) \end{aligned}$$

In simulations in a finite environment of length L , the allowed frequency modes are discretized with bin-size $2\pi/L$ (**(A5')**). This is shown in Fig.14 as lattice points with the Fourier spectrum overlaid as in [31, 33]. Gaussian readouts produce a Fourier spectrum peaked at the central DC mode. This mode has no periodicity and thus the theory for a single hidden unit predicts no lattices in this hidden unit, in the continuum or small-box discrete limit.

Until now, all analysis was performed for a single grid unit. What happens in full multi-cell setting? For this case, [31] shows that the global optimum to the constrained optimization problem:

$$\max_G \text{Tr}(G^T \Sigma G) \text{ such that } G^T G = I_{n_g},$$

i.e. under **(A3)** involves the columns of G spanning the top n_g eigenmodes of Σ (Theorem B.2, Appendix B of [31]). Under **(A6)**, [31] also shows that with a Fourier spectrum consisting of a wide annulus (in the discrete setting this corresponds to a Fourier spectrum of rings of different radii), the optimum consists of a hierarchy of hexagonal maps, but only if the Fourier powers of these rings are exactly equal (Lemma B.3, Appendix B). For purely Gaussian spectra (corresponding to Gaussian readouts), the full theoretical solution of this problem depends on specific details of the power spectrum curve (i.e. the width of the Gaussian) since each discrete eigenmode has a different Fourier power which must be taken into account while constructing linear combinations of modes, and thus is not solvable analytically. Instead, [31] numerically simulates the Lagrangian dynamics with assumption **(A3)** to find a hierarchy of lattices. However, the number of different period lattices that result is related directly to and nearly as numerous as the number of cells. For a large number of neurons, such as the 4096 hidden units used in simulations[34, 32], and sufficient wide power spectra, this would mean the

optimum solutions would consist of likely hundreds of discrete frequencies (each corresponding to a grid module). This is to be contrasted with the 4-8 modules estimated to exist in rodents [36].

Until NFL, whether the sequence of assumptions in the theory of [31] holds over the course of gradient learning in path-integrating RNNs, and whether their derived global optimal solution is generically gradient-learnable, remained unanswered. Our simulations suggest a negative answer. We also note that in simulations, **(A3)** (the orthonormality of all hidden unit ratemaps), on which this analysis and proof rests, does not seem to hold [34], since there are multiple cells with the same phase and periodicity. Indeed, in the biological system, grid units are far from orthonormal – in each grid module, neurons densely (and nearly continuously) tile all phases, meaning that there are always pairs of cells with non-identical but very similar phases and highly overlapping fields.

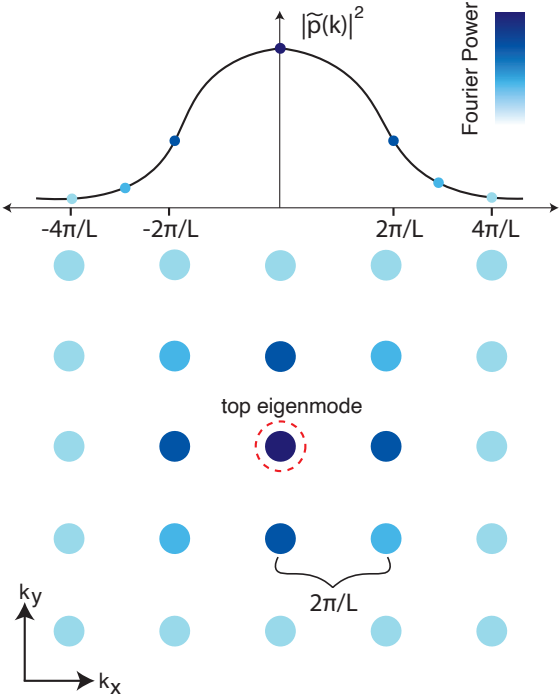


Figure 14: Fourier structure of Gaussian readouts

Theory of pattern formation in continuous attractor networks Next, we briefly summarize the theory of pattern formation in continuous attractor networks[19, 2]. For a full theoretical treatment see [4, 9]. Neurons are arranged on a sheet with sheet coordinates denoted by n . These models assume translation-invariant connectivity between neurons on a neural sheet: $W(n, n') = W(|n - n'|)$, which is specified by a kernel of interaction. For a network of N neurons, the equations describing their dynamics are given by the following system of equations:

$$\frac{ds(n, t)}{dt} + \frac{s(n, t)}{\tau} = \phi \left[\sum_{n'=1}^N W(n, n')s(n', t)dn' + B(n) \right] \quad (13)$$

We can take the number of neurons, $N \mapsto \infty$, making n a continuous coordinate on the attractor sheet and use neural field equations to describe the dynamics:

$$\frac{\partial s(n, t)}{\partial t} + \frac{s(n, t)}{\tau} = \phi \left[\int_{-\infty}^{+\infty} W(n, n')s(n', t)dn' + B(n) \right] \quad (14)$$

This connectivity structure sets up an approximate continuum of fixed points. With the correct kernel, the stable states are periodic. To find the periodicity of the formed pattern, we linearize around an unstable uniform state and use Fourier theory.

First, to identify the unstable uniform state:

$$\frac{s_0}{\tau} = \phi \left[\int_{-\infty}^{+\infty} W(n - n')s_0(n')dn' + B \right]. \quad (15)$$

$$s_0 = \frac{\tau B}{1 - \tau \bar{W}}, \text{ for a ReLU nonlinearity,} \quad (16)$$

where $\bar{W} = \int W(n - n')dn'$.

We then consider a perturbative analysis, by examining the evolution of $s(n, t) = s_0 + \epsilon(n, t)$.

We apply our analysis to the early time evolution of this initial condition, such that $\epsilon(n, t) \ll s_0$ is a small growing perturbation. Inserting our form of $s(n, t)$ in Eq. (14), we obtain

$$\frac{\partial \epsilon(n, t)}{\partial t} + \frac{\epsilon(n, t)}{\tau} = \phi'(\bar{W}s_0 + B) \int_{-\infty}^{\infty} W(n - n')\epsilon(n', t)dn' \quad (17)$$

Next, we posit the ansatz $\epsilon(n', t) = \epsilon e^{ik \cdot n' + \alpha(k)t}$, where $\alpha(k)$ denotes the wavelength-dependent growth rate of this ϵ perturbation.

$$\begin{aligned} \alpha(k) + 1/\tau &= \phi'(\bar{W}s_0 + B) \int_{-\infty}^{\infty} W(n - n')e^{-ik \cdot (n - n')}dn', \\ &= \phi'[\bar{W}s_0 + B]\mathcal{F}[W(n - n')], \\ &= \tilde{W}(k), \text{ for a ReLU non-linearity, assuming } \bar{W}s_0 + B > 0, \end{aligned}$$

where $\mathcal{F}[W(n - n')] = \tilde{W}(k)$ is the Fourier transform of the interaction kernel. Hence, the wavelength that will dominate the formed pattern on the attractor sheet is the one with highest growth rate, which is given by:

$$k^* = \operatorname{argmax}_k \tilde{W}(k), \quad (18)$$

Typically, in simulations of mechanistic continuous attractor models, a mexican-hat connectivity profile is used[2]. This corresponds to local excitation and broader inhibition in the network:

$$W(n, n') = W(|n - n'|) = W(\Delta n) = \alpha_E \exp\left(-\frac{(\Delta n)^2}{2\sigma_E}\right) - \alpha_I \exp\left(-\frac{(\Delta n)^2}{2\sigma_I}\right), \quad (19)$$

Here, σ_E, σ_I and α_E, α_I refer to the widths and amplitudes of the excitation and inhibition respectively.

We can use Fourier analysis to calculate the periodicity of the formed pattern:

$$\tilde{W}(k) = \mathcal{F}[W(n - n')] = \alpha_E \sigma_E \exp\left(-\frac{\sigma_E^2 k^2}{2}\right) - \alpha_I \sigma_I \exp\left(-\frac{\sigma_I^2 k^2}{2}\right)$$

Then, using Eq. 18, we find the Fourier mode dominating the pattern is given by:

$$[k^*]^2 = \frac{2}{\sigma_E^2 - \sigma_I^2} \log\left(\frac{\alpha_E \sigma_E^3}{\alpha_I \sigma_I^3}\right).$$

B.1 Connection between the two theories

To connect these two theories, [31] alternate between a Lagrangian gradient step (Eq.11), and a rectification step to impose non-negativity which is equivalent to applying a ReLU non-linearity. Further they assume a total sum constraint on each grid cell’s firing to reduce their dynamics to look similar to mechanistic attractor dynamics with a ReLU non-linearity. This is an assumption that we refer to as **Assumption 7 (A7)**:

$$\frac{d}{dt}G = \sigma(\Sigma G) - \lambda G \tag{20}$$

Thus, here the readout correlation matrix plays the role that the interaction matrix plays in mechanistic models.

The nature of discretization effects Now, we can see similarities between the readout reconstruction loss and continuous attractor networks. In the former case, the Fourier mode with highest power in the readout correlation function tends to dominate the pattern and in the latter, the Fourier mode with highest power in the interaction kernel determines the pattern.

An important difference is that in the readout reconstruction setting, the pattern is in *real space*, while in the continuous attractor network the pattern is in the *neural sheet* space. This difference is conceptually critical: It is reasonable to assume that the neural sheet is of a given and fixed size, but it is unreasonable to think of the physical space in which the animal moves as being of a given and fixed size.

The discretization and finite size lattice effects in the readout reconstruction setting are due to the size (and shape) of the physical room used in readout reconstruction and in trained path integrators. This can lead to artefactual lattice-like patterns in some cells as has been observed in [31, 33]. It also predicts that the size and shape of the formed lattice will be affected by the room; this is not true in biology for familiar environments of different sizes [13].

On the other hand, the continuous attractor network is subject to discretization effects at the level of number of neurons N – this is an invariant constraint on the circuit regardless of the current environment the animal is exploring. The number of neurons in MEC is also large, and thus the continuum limit approximation is a valid assumption.