

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 EXponential Map Models as an Interpretable Framework for Generating Neural Spatial Representations

Anonymous authors
Paper under double-blind review

ABSTRACT

A fundamental challenge in neuroscience and AI is understanding how physical space is mapped into neural representations. While artificial neural networks can generate brain-like spatial representations, such as place and grid cells, their “black-box” nature makes it difficult to determine if these representations arise as general solutions or as artifacts of a chosen architecture, objective function, or training protocol. Critically, these models offer no guarantee that learned solutions for core navigational tasks, like path integration (updating position from self-motion), will generalize beyond their training data. To address these challenges, we introduce a first-principles framework based on an exponential map model. Instead of using deep networks or gradient-based optimization, the presented model uses generator matrices to map physical locations into neural representations through the matrix exponential, creating a transparent framework that allows us to identify several exact algebraic conditions underlying key properties of neural maps. We show that path invariance (ensuring location representations are independent of traversal route) and *exact* path integration is achieved if the generators commute, while translational invariance (maintaining consistent spatial relationships across locations) demand that generators produce orthogonal transformations. We also show that preserving the metric of flat space requires the eigenvalues of the generator matrices to form sets of roots of unity. Finally, we demonstrate that the proposed framework constructs diverse biologically relevant spatial tuning, including place cells, grid cells, and context-dependent remapping, and that the exponential map model corresponds to the on-manifold dynamics of a continuous attractor network. The framework we propose thus offers a transparent, theoretically-grounded alternative to “black-box” models, revealing the exact conditions required for a coherent neural map of space.

1 INTRODUCTION

A fundamental challenge in neuroscience and artificial intelligence is to understand the mapping from physical space to the representational space of neural population activity. In the mammalian brain, such representations are strongly associated with the hippocampal formation, which contains specialized neurons that encode spatial information. Most famously, place cells (O’Keefe & Dostrovsky, 1971) fire within specific, localized areas of an environment known as place fields, while grid cells (Hafting et al., 2005) fire in a periodic hexagonal pattern that tessellates the environment and is hypothesized to provide a neural metric for space (McNaughton et al., 2006; Moser & Moser, 2008; Ginosar et al., 2023). Together, these and other spatially-tuned cells form a high-dimensional representation of an animal’s location. This neuronal spatial map abruptly reorganizes in response to environmental changes, a phenomenon known as remapping, indicating that neurons also encode the environment’s identity (Leutgeb et al., 2004; Fyhn et al., 2007). While the firing patterns of these spatial neurons are well-characterized, the principles governing their emergence remain unclear.

In recent years, deep learning models, particularly recurrent neural networks (RNNs), trained to solve navigation tasks have been shown to learn representations that resemble biological place and grid cells (Banino et al., 2018; Cueva & Wei, 2018; Sorscher et al., 2023). Complementing these end-

054 to-end approaches, more structured architectures have demonstrated that such representations can
 055 arise from explicit theoretical constraints, such as the factorization of abstract structural knowledge
 056 from specific sensory experiences (Whittington et al., 2020). In parallel, theoretical frameworks
 057 rooted in reinforcement learning, most notably the Successor Representation (SR), have proposed
 058 that grid cells function as a low-dimensional eigenbasis for the environment’s transition matrix,
 059 effectively encoding a predictive map of space (Stachenfeld et al., 2017). Although these findings
 060 strongly suggest that spatial tuning is a normative solution to the demands of navigation, critical
 061 theoretical limitations remain. The “black-box” nature of deep neural networks makes it difficult
 062 to disentangle whether their learned representations reflect fundamental principles of navigation or
 063 are artifacts of a chosen architecture, objective function, or training protocol (see Fig. 1a) for an
 064 illustration). Furthermore, while SR and structured models improve interpretability, they share a
 065 fundamental limitation with deep learning approaches: they all rely on iterative optimization or
 066 statistical accumulation. Consequently, these methodologies do not offer an algebraic guarantee that
 067 navigational solutions will strictly generalize beyond the training data. In contrast, animals are able
 068 to seamlessly navigate vast, novel environments. To understand how biological brains solve these
 069 navigational challenges efficiently and robustly, there is a need for models that allow for exact and
 070 interpretable solutions to navigation problems.

071 In this work, we construct spatial representations using a first-principles framework based on the
 072 matrix exponential. Instead of relying on neural networks or gradient-based optimization, the presented
 073 approach builds representations from a transparent mathematical foundation. The core component
 074 of the model is a set of generator matrices that directly map a spatial location to a neural popula-
 075 tion firing rate vector. This construction enables us to derive the exact algebraic conditions required
 076 for a coherent neural spatial map. For a neural spatial map to be useful, it must support core nav-
 077 igational computations. One of the most fundamental is *path integration*, the process by which a
 078 navigator estimates its position by integrating self-motion cues. This process introduces a critical
 079 self-consistency problem: For the map to be coherent, the representation of a location must be in-
 080 dependent of the path taken to reach it. We show that path-independent representations required for
 081 reliable path integration are guaranteed if the model’s generator matrices commute. Furthermore,
 082 we find that equinorm representations, previously used as a learning constraint in neural networks
 083 (Schaeffer et al., 2023; Xu et al., 2022), arise naturally from generators that produce translationally
 084 invariant similarity structures—a desirable property for navigation in open-field environments. We
 085 also show that preserving the metric of flat space (Xu et al., 2022; Schøyen et al., 2025; Xu et al.,
 086 2025) requires the eigenvalues of the generator matrices to form sets of roots of unity on discrete
 087 rings in frequency space. Crucially, when these conditions are met, the framework recovers diverse,
 088 biologically observed tuning curves, such as place cells and grid cells, dictated strictly by the sym-
 089 metry of the generators. We also demonstrate that this framework can be seamlessly generalized
 090 from preserving the metric of space to preserving the similarity of more general inputs, which we
 091 use to model remapping, and even goal-oriented, multi-map navigation in an interpretable manner.
 092 A conceptual overview of the proposed framework and the key spatial map properties we address
 093 are presented in Fig. 1. Finally, we demonstrate that the exponential map can be interpreted as
 094 describing the on-manifold dynamics of a continuous attractor recurrent neural network.

095 Despite its simplicity, the proposed framework is powerful enough to construct a variety of biolog-
 096 ically plausible tuning curves, including place cells, grid cells, and context-dependent remapping,
 097 from the same underlying mechanism. By grounding spatial representations in a clear algebraic
 098 structure, the presented work provides a theoretically-grounded alternative to black-box models,
 099 revealing exact and interpretable principles that underpin a coherent neural map of space.

2 RESULTS & DISCUSSION

2.1 CONSTRUCTING SPATIAL REPRESENTATIONS WITH AN EXPONENTIAL MAP

100 Formally, a spatial representation is a map that assigns a neural population vector to every spatial
 101 location, as exemplified in Fig. 1b). For a 2D space with Cartesian coordinates (x, y) , the repre-
 102 sentation at a given point is a population vector $\mathbf{p}(x, y) \in \mathbb{R}^N$. Each of the N components of this
 103 vector corresponds to the firing rate of a neuron, making the vector a point in an N -dimensional
 104 state space that captures the activity of the entire neural ensemble. Building upon previous model-
 105 ing approaches Gao et al. (2021); McNamee et al. (2021); Xu et al. (2022), we define this map using
 106

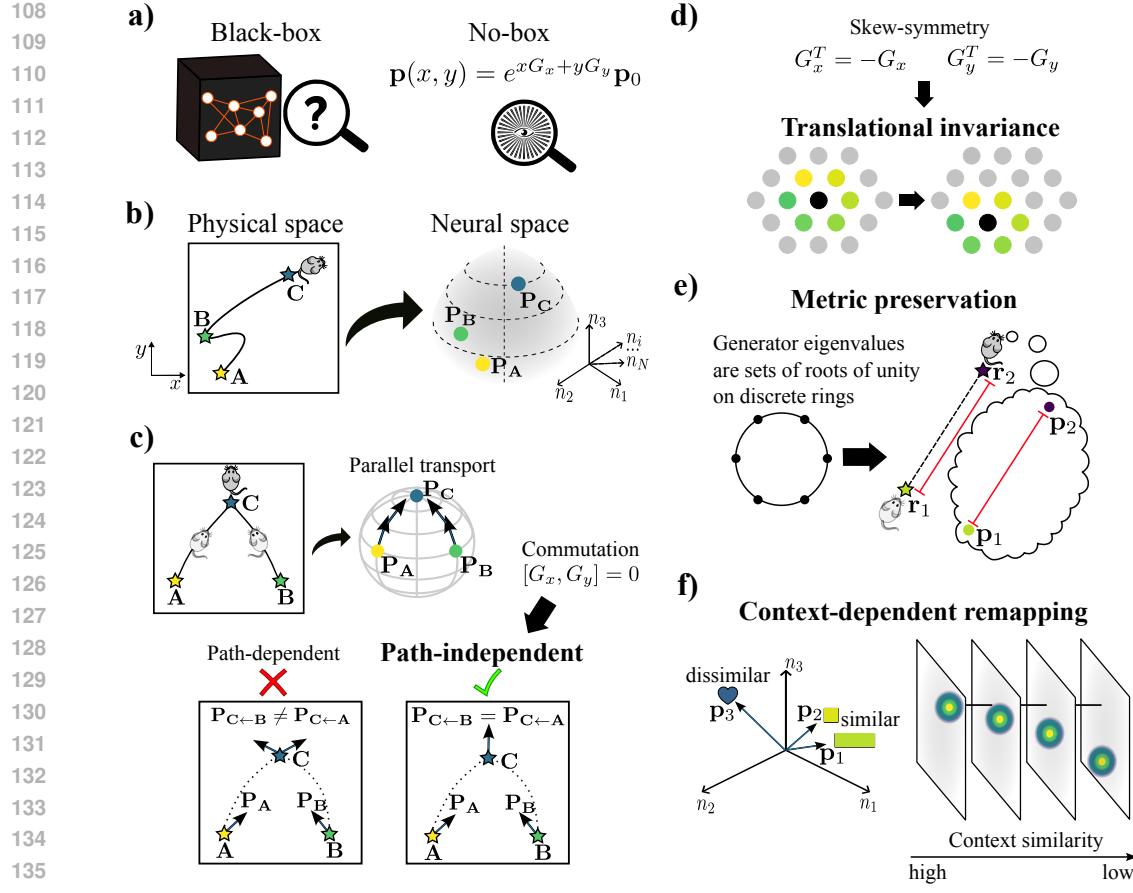


Figure 1: **The exponential map framework for interpretable spatial representations.** a) Deep learning models are “black-boxes” that learn spatial representations, but the underlying principles are obscured by the complexities of architecture, training, and objective functions. The exponential map model is a transparent “no-box” alternative, using generator matrices (G_x, G_y) to construct a population vector $\mathbf{p}(x, y)$. b) A neural population vector, which captures the activity of the entire neural ensemble, is assigned to every location, mapping physical space to a neural representational space. c) **Path Invariance:** Path integration can be viewed as a form of parallel transport, where a vector representing the neural representation is moved along a trajectory in a high-dimensional state space. Traversing a curved manifold can induce a net transformation in the vector at a point that is dependent on the traversal route. By imposing simple, interpretable algebraic constraints on the model’s generators, we can directly enforce fundamental properties. Path invariance is guaranteed if the generators commute. d) **Translational Invariance:** Making the generator matrices skew-symmetric ($G^T = -G$) imposes several biologically-relevant properties on the representation. First, it ensures that spatial relationships are consistently maintained across locations (translational invariance). Second, skew-symmetric generators produce orthogonal representations, meaning the population vector $\mathbf{p}(x, y)$ maintains a constant norm across the entire space. e) **Metric Preservation:** Preserving the geometry of flat space requires the generator eigenvalues to form sets of roots of unity on discrete rings in frequency space, which, for certain symmetry orders gives rise to grid-like patterns. f) **Remapping:** Generalizing the framework to non-spatial inputs allows for modeling remapping. The stacked sheets represent spatial maps for distinct values of a non-spatial context signal, s . As s changes, the peak of activity shifts location even if the physical position (x, y) is constant. This remapping preserves similarity: similar context values (neighboring sheets) result in spatially proximal firing fields.

the matrix exponential:

$$\mathbf{p}(x, y) = e^{xG_x + yG_y} \mathbf{p}_0, \quad (1)$$

162 where $G_x, G_y \in \mathbb{R}^{N \times N}$ are *generator matrices* for the cardinal directions and $\mathbf{p}_0 = \mathbf{p}(0, 0)$ is the
 163 representation at some origin point. The generator matrices define how locations in physical space
 164 translate into transformations in the high-dimensional neural state space. The exponential map then
 165 composes these transformations to “transport” the origin vector, \mathbf{p}_0 to a population vector at any
 166 target location (x, y) .

167 While Eq. (1) provides a constructive method for generating a spatial map, without further con-
 168 straints, an arbitrary choice of generators could produce a map ill-suited for navigation. For instance,
 169 the representation could prove trivial (all locations map to the same vector) or ambiguous (multiple
 170 locations map to the same vector). As we will show, the power of this framework lies in its trans-
 171 parency, allowing us to derive precise algebraic conditions on G_x and G_y that guarantee properties
 172 essential to navigation.

173

174 2.2 FROM SPATIAL REPRESENTATION TO PATH INTEGRATION AND PATH INDEPENDENCE

175

176 Path integration is a crucial skill possessed by most animals, wherein one’s location is inferred by
 177 integrating past location and self-motion information. In the context of the representation defined in
 178 Eq. (1), path integration is realized if the representation at a new location can be derived from the
 179 current representation via a displacement operator Q :

180

$$\mathbf{p}(x + \Delta x, y + \Delta y) = Q(\Delta x, \Delta y)\mathbf{p}(x, y). \quad (2)$$

181

182 Intuitively, we can say that we can perform path integration, if, for any past location (x, y) and
 183 the corresponding population vector $\mathbf{p}(x, y)$ we can arrive at the correct population vector $\mathbf{p}(x +$
 184 $\Delta x, y + \Delta y)$ at the new location $(x + \Delta x, y + \Delta y)$ through some operation Q that only depends on
 185 the displacement. Inserting our spatial representation from Eq. (1), we find that we want

186

$$e^{(x+\Delta x)G_x+(y+\Delta y)G_y}\mathbf{p}_0 = Q(\Delta x, \Delta y)e^{xG_x+yG_y}\mathbf{p}_0.$$

187

188 This equality suggests that we want

189

$$Q(\Delta x, \Delta y) = e^{\Delta x G_x + \Delta y G_y}.$$

190

191 However, the exponential function in Eq. (1) is a matrix exponential, which behaves differently from
 192 the regular exponential function. In particular, the Baker-Campbell-Hausdorff formula dictates that

193

$$e^A e^B = e^{A+B+\frac{1}{2}[A,B]+\frac{1}{12}[A,[A,B]]+\frac{1}{12}[B,[B,A]]+\dots},$$

194

195 where $[A, B] = AB - BA$ is the commutator between matrices A and B . However, this immediately
 196 reveals that if the generator matrices G_x, G_y commute, $[G_x, G_y] = 0$, then Eq. (2) is automatically
 197 satisfied for any displacement, as

198

$$[aG_x + bG_y, a'G_x + b'G_y] = (ab' - a'b)[G_x, G_y],$$

199

200 for all (a, b) and (a', b') . Thus, if the generators commute, the model can path integrate exactly
 201 and indefinitely. An important effect of this choice is that the representation is path-invariant (as
 202 illustrated in Fig. 1c), meaning that the population vector at a point does not depend on the path taken
 203 to it. This is also demonstrated explicitly in Appendix B. Going forward, we therefore demand that
 204 G_x, G_y commute, which ensures that the representation \mathbf{p} is path-integration compatible, as enacted
 205 by Eq. (2). Next, we demonstrate that commuting generator matrices enable an explicit construction
 206 that allows us to specify the similarity structure of the spatial representation.

207

208

209 2.3 ORTHOGONAL TRANSFORMATIONS FOR EGOCENTRIC NAVIGATION

210

211 With a path integration-compatible model established, we turn to the properties of the representation
 212 itself. A critical metric for representational structure is the similarity between population vectors
 213 at distinct locations. Considering the path-integrating model defined in Eq. (2), we examine the
 214 normalized inner product (cosine similarity) between a population vector at (x, y) and a target vector
 215 at $(x + \Delta x, y + \Delta y)$:

216

$$C(x, y, \Delta x, \Delta y) = \frac{\mathbf{p}(x, y)^T Q(\Delta x, \Delta y) \mathbf{p}(x, y)}{|\mathbf{p}(x, y)| |Q(\Delta x, \Delta y) \mathbf{p}(x, y)|},$$

216 Substituting the exponential map form, using that $(e^A)^T = e^{A^T}$, and assuming commuting genera-
 217 tors, the similarity expression transforms to:

$$219 \quad C(x, y, \Delta x, \Delta y) = \mathbf{p}_0^T e^{x(G_x^T + G_x) + y(G_y^T + G_y) + \Delta x G_x + \Delta y G_y} \mathbf{p}_0 / Z, \quad (3)$$

220 where Z is the normalization factor from before. Inspection of Eq. (3) reveals a fundamental geo-
 221 metric condition. If the generator matrices G_x and G_y are skew-symmetric ($G^T = -G$), two key prop-
 222 erties emerge. First, the matrix exponential of a skew-symmetric matrix is orthogonal. This ensures
 223 that the transformation preserves the norm of the population vector everywhere ($|\mathbf{p}(x, y)| = |\mathbf{p}_0|$),
 224 fixing the normalization factor to $Z = |\mathbf{p}_0|^2$. Second, the position-dependent terms in the expo-
 225 nent vanish because $G^T + G = 0$, as illustrated in Fig. 1d). Consequently, the similarity becomes
 226 strictly translation invariant, depending solely on the displacement $(\Delta x, \Delta y)$. We therefore enforce
 227 the condition:

$$228 \quad G_x^T = -G_x \quad G_y^T = -G_y,$$

229 which yields the displacement-dependent similarity:

$$231 \quad C(\Delta x, \Delta y) = \frac{\mathbf{p}_0^T e^{\Delta x G_x + \Delta y G_y} \mathbf{p}_0}{|\mathbf{p}_0|^2}, \quad (4)$$

233 This property is functionally critical for navigation. In an open-field regime where no locations are
 234 inherently privileged, the similarity structure should remain consistent across the environment. Fur-
 235 thermore, it enables spatial inferences (such as distance estimation; see Section 2.4) without requir-
 236 ing absolute positional information, making orthogonal transformations an ideal basis for egocentric
 237 navigation.

238 Similarity translational invariance is a recurring feature in theoretical models of spatial coding. For
 239 instance, in Continuous Attractor Neural Networks (CANNs), the activity profile shifts across the
 240 neural sheet without changing shape, naturally preserving the similarity structure between displaced
 241 states (Burak & Fiete, 2009; McNaughton et al., 2006). Similarly, frameworks based on transition
 242 coding and the Successor Representation (SR) demonstrate that eigenvectors of the environment’s
 243 transition matrix, which essentially encode transition probabilities, naturally capture spatial period-
 244 icities and translational symmetries in open environments (Stachenfeld et al., 2017; Waniek, 2018).
 245 More recently, sequence coding models have shown that encoding trajectories via path integration
 246 leads to conformal isometries, where displacements in neural space differ from physical displace-
 247 ments only by a scale factor (Waniek, 2020; RG et al., 2025). Our framework unifies these obser-
 248 vations by deriving the property explicitly from the algebraic structure of the generator matrices.
 249 While previous approaches often obtain these representations through statistical learning of transi-
 250 tions or optimization of spatiotemporal sequences (Waniek, 2018), we identify the exact algebraic
 251 condition—skew-symmetry ($G^T = -G$)—that guarantees this property for high-dimensional vector
 252 representations. This condition ensures that the induced transformations are orthogonal, thereby
 253 maintaining a constant norm (an equinorm constraint) while rendering the similarity function purely
 254 dependent on displacement. This provides a rigorous algebraic description of the “shift-invariant”
 255 connectivity required by biological circuits, identifying skew-symmetry as the structural necessity
 256 for any path-integrating system that preserves representational similarity.

257 Given skew-symmetric generators G_x and G_y , we can decompose them into a canonical block di-
 258 agonal form via the spectral theorem:

$$258 \quad G_x = R^T \Sigma_x R \quad \text{and} \quad G_y = R^T \Sigma_y R, \quad (5)$$

260 where $R \in \mathbb{R}^{N \times N}$ is an orthogonal matrix, and Σ_x, Σ_y are real block diagonal matrices with 2×2
 261 skew-symmetric blocks. Assuming N is even, the non-zero entries of these blocks correspond to the
 262 imaginary parts of the eigenvalues, which appear in conjugate pairs $\pm(i\lambda_{i,x}, i\lambda_{i,y})$. This structure
 263 ensures that G_x and G_y commute, as the constituent 2×2 skew-symmetric blocks commute and
 264 $R^T R = I$. Under these conditions, the similarity function simplifies to:

$$266 \quad C(\Delta x, \Delta y) = \sum_i^N \alpha_{0,i}^2 \cos(\lambda_{i,x} \Delta x + \lambda_{i,y} \Delta y), \quad (6)$$

268 where $\alpha_0 = R \frac{\mathbf{p}_0}{|\mathbf{p}_0|}$ and $\lambda_{i,x}, \lambda_{i,y}$ are the imaginary parts of the i -th eigenvalues of G_x and G_y ,
 269 respectively (see Appendix E for the derivation).

270 The translational invariance derived here offers a significant functional advantage: the similarity
 271 depends only on the relative difference between inputs, not their absolute values. Consequently,
 272 the framework generalizes naturally to non-spatial variables. By introducing a context signal s
 273 corresponding to its own generator, the system can model “contextual displacements” independent
 274 of spatial displacements. As we explore in Section 2.5 (and illustrate in Fig. 1f), this property allows
 275 a single model to generate distinct, orthogonal spatial maps for different environmental contexts,
 276 mimicking hippocampal remapping (Leutgeb et al., 2004; Fyhn et al., 2007), simply by shifting the
 277 effective origin of the representation.

278 2.4 PRESERVING THE METRIC OF FLAT SPACE

280 Given a spatial representation and a notion of representational similarity, we determine what prop-
 281 erties the representation must possess to be geometrically consistent. Following Gao et al. (2021)
 282 and Xu et al. (2022), we posit that a foundational property of any spatial representation is the faith-
 283 ful translation of physical distances into distances on a neural manifold. In the open field, where
 284 all directions and locations are effectively equivalent, distances should not appear warped in any
 285 particular location or direction. Consequently, the metric induced by Eq. (1) must match the flat Eu-
 286 clidean metric, at least up to a constant factor (a condition known as conformal isometry (Xu et al.,
 287 2022)). Demanding metric preservation of a path-integrating, orthogonal representation, yields a
 288 simple condition on the spectra of the generators G_x and G_y : the representation preserves the flat
 289 metric if the eigenvalues form sets of roots of unity (see Appendix D for details and Fig. 1e for an
 290 illustration). Concretely, expressed in polar coordinates as $\lambda_{i,x} = k_i \cos(\phi_i)$ and $\lambda_{i,y} = k_i \sin(\phi_i)$
 291 (where λ are the imaginary parts of the eigenvalues), a flat metric-preserving representation satisfies

$$292 \sum_j^{N/2} \rho_j^2 e^{2i\phi_j} = 0,$$

295 where $\rho_i = \alpha_i k_i$ is shared by conjugate eigenvalues (see Appendix D). Geometrically, this implies
 296 that the eigenvalue angles ϕ_i must be evenly spaced on discrete rings in the frequency domain. For
 297 a given ring of radius ρ_m with symmetry order M , the eigenvalues are distributed as

$$299 \phi_i \rightarrow \varphi_i + \pi \frac{i}{M},$$

301 with $i = 0, 1, \dots, M - 1$.

302 To visualize the spatial representation produced by a particular symmetry M , we note that for or-
 303 thogonal matrices, the entries of the population vector \mathbf{p} constitute superpositions of 2D plane waves,
 304 denoted α (see Appendix C). The specific interference pattern is determined by the choice of the
 305 orthogonal matrix R , as $\mathbf{p}_0 = R^T \alpha$. Figure 2 shows example representations \mathbf{p} and plane waves
 306 α for different symmetries M , using a randomly sampled R (see Appendix A for details). For a
 307 single ring, lower-order symmetries such as $M = 2$ and 3 produce plane wave mixtures oriented at
 308 90° and 60° , respectively. This results in grid-like representations: square-type grids for $M = 2$
 309 and hexagonal grids for $M = 3$. Higher-order symmetries, for example, $M = 4$, yield increas-
 310 ingly complex interference patterns, resembling the honeycomb-like structures recently observed in
 311 RNNs trained for dual-agent path integration (Redman et al., 2024). Notably, some representations
 312 may not appear purely grid-like due to the random mixing induced by R . As M increases, the repre-
 313 sentation becomes heterogeneous and lacks obvious periodicity. However, while the spatial tuning
 314 curves are strongly influenced by R , the similarity function depends only on α . With increasing M ,
 315 the similarity becomes approximately radial; for $M = 20$, the similarity function approximates a
 316 Bessel function, as predicted in Appendix F.

317 Beyond uncovering a general condition for metric preservation, we find that the admissible solu-
 318 tions support the modular organization of grid cells observed in the brain (Stensola et al., 2012).
 319 Analyzing the similarity function (see Appendix F) reveals that if modules with the same spacing
 320 form roots of unity in their orientation, and the symmetry of the module orientation is co-prime
 321 to that of the pattern, the representation becomes head-direction independent over a large spatial
 322 range. Furthermore, if the relative spacing of different modules is proportional to the zeros of the
 323 Bessel function J_0 , the similarity function approximates a Fourier-Bessel series. This allows for
 324 the construction of radially symmetric similarity functions capable of tuning navigation to specific
 325 length scales. Intriguingly, the average ratio of successive low-order Bessel function zeros used to

construct such a series falls within the variability range of experimentally observed grid cell module spacings (approximately $\sqrt{2}$) (Stensola et al., 2012), suggesting a link between the roots-of-unity algebraic structure and the organization of the entorhinal cortex.

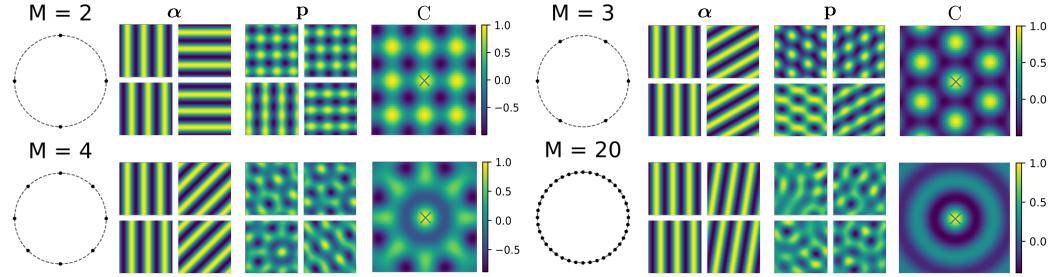


Figure 2: **Metric-preserving representations and roots of unity.** Example plane waves (α) and corresponding representations (p) alongside the similarity function (C) relative to the origin (black cross) for representations whose eigenvalues form sets of roots of unity, with symmetries M on a single ring. For each eigenvalue (imaginary part indicated by black dot), the corresponding conjugate eigenvalue is also shown. Representations were formed using generators with a single set root-of-unity solution with varying M , a random orthogonal matrix R , and $\mathbf{p}_0 = R^T \mathbf{1}$.

2.5 SIMILARITY PRESERVATION AND REMAPPING

With the generators defined in Eq. (5) and a choice of similarity function, we can generate spatial representations for a single environment up to a choice of orthogonal transformation R . However, animals are capable of distinctly encoding a variety of spatial and non-spatial information, such as room identity or olfactory cues, through remapping. In this section, we extend our model to this larger class of representations.

This generalization follows from the observation that between-representation similarities depend solely on the *spatial* displacement between them. If we encode non-spatial information while fixing the spatial location, representational similarities depend only on the change in the non-spatial input. We consider a global scalar signal s , such as a context variable, encoded identically to spatial coordinates:

$$\mathbf{p}(x, y, s) = e^{xG_x + yG_y + sG_s} \mathbf{p}_0. \quad (7)$$

The representation is coupled to the non-spatial signal via a generator matrix G_s . As with the spatial generators, we let $G_s = R^T \Sigma_s R$. The similarity between representations for two distinct context signals s and s' at a fixed location is then:

$$\mathbf{p}(x, y, s)^T \mathbf{p}(x, y, s') = \mathbf{p}^T e^{(s-s')G_s} \mathbf{p} = \mathbf{p}^T e^{\Delta s G_s} \mathbf{p}.$$

This inherits the form of the spatial similarity function. Consequently, comparing across contexts reveals that representations change even as spatial location remains fixed, mimicking the remapping behavior of spatial cells (Leutgeb et al., 2004; Fyhn et al., 2007).

Encoding non-metric information, such as context, raises a distinct challenge: unlike physical space, there is no intrinsic metric to preserve. Instead, we require that similar context signals produce similar representations, while dissimilar contexts result in orthogonal ones. Prior studies have demonstrated that such similarity preservation yields localized receptive fields, resembling biological place fields, when applied to spatial inputs (Sengupta et al., 2018; Pettersen et al., 2024). To implement this algebraically, we observe that the similarity function may be written as:

$$C(\Delta s) = \sum_i \alpha_{0,i}^2 \cos(\Delta s \lambda_{i,s}), \quad (8)$$

where $\lambda_{i,s}$ denotes the imaginary part of the i th eigenvalue of G_s (assuming fixed spatial location). Since this expression is derived from the cosine similarity, it is strictly bounded to the interval $[-1, 1]$. To achieve similarity preservation, we seek a profile $C(\Delta s)$ that decays with increasing

Δs towards a baseline level where inputs are deemed dissimilar. We can approximate a wide class of such functions by recognizing that Eq. (8) is effectively a cosine series with non-negative coefficients. Specifically, it may be viewed as a discrete approximation of the inverse Fourier transform of a symmetric function with a non-negative Fourier spectrum:

$$f(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} F(k) e^{ikx} dk = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} F(k) \cos(kx) dk \\ \approx \frac{1}{\sqrt{2\pi}N} \sum_i^N \frac{F(k_i)}{p(k_i)} \cos(k_i x),$$

where the summation represents a Monte Carlo estimate of the integral using importance sampling with density $p(k_i)$.

For example, to approximate a Gaussian similarity function $f(x) = e^{-\sigma^2 x^2}$ (which ensures that dissimilar contexts become decorrelated), we sample the eigenvalues $\lambda_{i,s}$ from the Fourier transform of the Gaussian, that is, a normal distribution $\mathcal{N}(0, 2\sigma^2)$. Setting the coefficients $\alpha_{0,i}^2 = 1/N$ (achieved if $\mathbf{p}_0 = \frac{1}{\sqrt{N}} R^T \mathbf{1}$) yields:

$$C(\Delta s) \approx e^{-\sigma^2 \Delta s^2}.$$

Going forward, we set $\sigma = 1$.

To demonstrate this remapping behavior, we instantiated a metric-preserving spatial representation consisting of 10 sets of identical root-of-unity solutions and extended it to encode a non-spatial signal s via Eq. (7). Multiple sets were used to ensure the Monte Carlo estimate approximates the Gaussian sufficiently well. The resulting between-context similarity is shown in Fig. 3. Similarities decay with increasing context dissimilarity, and the rate maps of unit activity shift between contexts, reproducing remapping dynamics (Fyhn et al., 2007). Crucially, spatial similarities are preserved for any fixed s , maintaining the grid-like structure derived in the previous section.

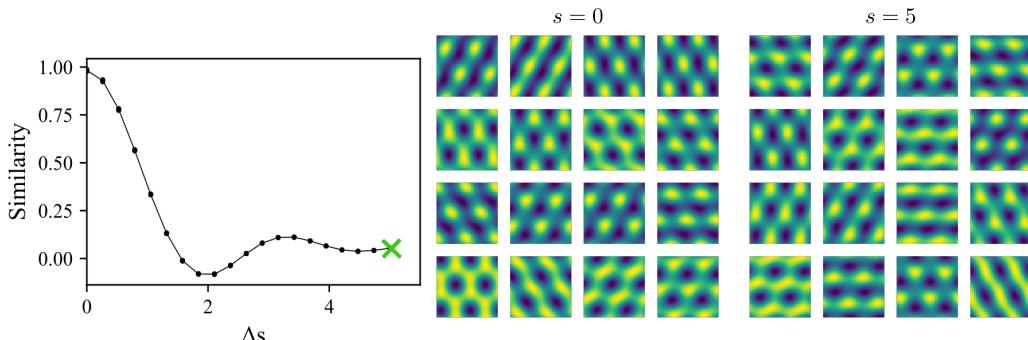


Figure 3: Context-dependent remapping and similarity preservation. Between-context similarity as a function of context separation. Also inset are example rate maps for two distant context values ($s = 0, s = 5$) corresponding to $\Delta s = 5$. Representations were formed using generators with 10 identical root-of-unity solutions with $M = 3$, a random orthogonal matrix R , and $\mathbf{p}_0 = R^T \mathbf{1}$.

Finally, we find that relaxing the metric preservation requirement in the spatial domain, demanding only similarity preservation via Gaussian-sampled eigenvalues for G_x and G_y , results in an approximate Gaussian spatial similarity (see Appendix G). In this regime, spatial representations become heterogeneous and localized, resembling hippocampal place fields (O’Keefe & Dostrovsky, 1971). Thus, by modulating the similarity function, the exponential map model can generate the diverse range of spatial tuning curves observed in the brain. Notably, these spatial tuning curves are compatible with the Probabilistic Population Codes (PPC) framework (Ma et al., 2006; Beck et al., 2008). Specifically, if neural variability follows a distribution in the exponential family, for example, Poisson-like, with linear sufficient statistics, the Gaussian-like tuning curves derived here can support optimal Bayesian inference via linear integration. In this context, the algebraically derived representations effectively define the tuning kernel within the PPC formalism.

432 3 BIOLOGICAL REALIZATION AND FUNCTIONAL NAVIGATION

434 While derived from algebraic principles, the exponential map framework maps directly onto bi-
 435 ological mechanisms. We demonstrate that our model emerges from the dynamics of Continuous
 436 Attractor Neural Networks (CANNs) (Appendix H), can learn and even generalize experimental data
 437 (Appendix I), and supports robust, interpretable, multimap goal-oriented navigation (Appendix J).

439 3.1 EMERGENCE FROM ATTRACTOR DYNAMICS AND LOCAL LEARNING

441 The exponential map model emerges naturally as the effective on-manifold dynamics of a gain-
 442 modulated CANN flowing towards a hypersphere attractor (see Appendix H.1). By analyzing the
 443 Lyapunov energy of the network, we find that the dynamics decompose into a stabilizing non-linear
 444 term, which constrains activity to the manifold, and a linear transport term. On the manifold, the
 445 state evolves according to $\dot{\mathbf{z}} = U(\mathbf{v})\mathbf{z}$, where U is the velocity-dependent skew-symmetric compo-
 446 nent of the recurrent weights. This formulation physically identifies our algebraic generator matrices
 447 with the synaptic connectivity of the circuit. Consequently, the algebraic constraints derived in this
 448 work map directly to biological connectivity patterns. The skew-symmetry condition, required for
 449 translational invariance, corresponds to the asymmetric component of the recurrent weights. The
 450 commutativity requirement ensures that the time-ordered integration of synaptic inputs simplifies
 451 to a state-independent update (a vanishing Magnus expansion), providing a dynamical definition of
 452 path integration (see Appendix H.2). Thus, the exponential map offers a rigorous description of how
 453 recurrent networks perform exact temporal integration without trajectory-dependent errors. Further-
 454 more, we show that such connectivity matrices need not be hard-coded; rather, they can emerge
 455 naturally via a biologically plausible local learning rule that exploits time-lagged anti-symmetric
 456 correlations. Finally, we establish that for hexagonal grid cells, the grid spacing scales linearly
 457 with the attractor network’s time constant. This finding offers a normative explanation for the grid-
 458 scale hierarchy observed along the dorsal-ventral axis of the medial entorhinal cortex (MEC) (see
 459 Appendix H.3).

460 3.2 LEARNING FROM DATA AND FUNCTIONAL NAVIGATION

461 To validate the generative ability of the framework, we trained the model to reproduce experimental
 462 grid cell rate maps by minimizing reconstruction error subject to a commutation penalty (see Ap-
 463 pendix I). The learned model successfully extrapolates grid patterns beyond the training boundaries,
 464 demonstrating that it captures the intrinsic algebraic structure underlying the biological data.

465 Functionally, the framework enables robust navigation by leveraging principles from Hyperdimen-
 466 sional Computing (HDC) (see Appendix J). Salient locations can be aggregated via “bundling” into
 467 memory vectors $\mathbf{p}_R = \sum \mathbf{p}(x_i, y_i)$, creating a similarity landscape that supports navigation via
 468 gradient ascent. Furthermore, context-dependent remapping serves as a “binding” operation, ef-
 469 fectively orthogonalizing representations across contexts. This allows multiple reward maps to be
 470 superimposed within a single neural population, facilitating context-specific retrieval and flexible
 471 goal-oriented navigation.

473 4 CONCLUSION

475 In this work, we introduced a first-principles framework for generating neural spatial representations
 476 using an exponential map model. By leveraging generator matrices and the matrix exponential, we
 477 bypassed the “black-box” nature of deep learning models, enabling a transparent and theoretically-
 478 grounded investigation into the principles of neural navigation. We derived the exact algebraic con-
 479 ditions required for a coherent map of space. Specifically, we demonstrated that commuting genera-
 480 tors are necessary to guarantee path-independent representations, a critical requirement for accurate
 481 path integration. Furthermore, we showed that constraining generators to be skew-symmetric pro-
 482 duces orthogonal transformations, yielding representations with translationally invariant similarity
 483 structures, an ideal property for egocentric navigation in open-field environments. We also estab-
 484 lished that preserving the flat metric of Euclidean space requires the generator eigenvalues to form
 485 sets of roots of unity on discrete rings in the frequency domain. Despite its mathematical simplicity,
 the proposed framework constructs a diverse range of biologically plausible spatial tuning curves,

486 including grid cells and place cells, and models context-dependent remapping by extending these
 487 principles to non-spatial inputs. This work offers an interpretable alternative to conventional deep
 488 learning approaches, revealing the fundamental mathematical structures that may underpin how the
 489 brain represents and navigates through space.

492 5 LIMITATIONS AND FUTURE WORK

494 While our framework provides a transparent account of how coherent spatial maps can be formed,
 495 it has several limitations that open avenues for future research.

497 First, the current model is primarily developed for navigation in flat, open-field environments. Animals,
 498 however, must navigate complex, curved, and obstacle-laden spaces. Future work should
 499 explore how the generator framework can be extended to represent non-Euclidean geometries. This
 500 may involve introducing position-dependent or non-commuting generators that reflect the local
 501 topology and geometry of the environment.

502 Second, our remapping model currently uses a scalar context signal. Generalizing this to handle
 503 high-dimensional, structured inputs, such as visual scenes or complex sensory cues, is a critical next
 504 step. The algebraic structure naturally supports vector-valued generators ($s \cdot G_s$); in this regime,
 505 similarity preservation implies preserving the semantic distances between high-dimensional inputs,
 506 allowing for the modeling of how environmental identity and spatial location are integrated into
 507 a unified representation. This would bridge the gap between our algebraic approach and the rich,
 508 multi-modal inputs processed by biological and artificial systems.

509 Third, while geometric considerations fix most model parameters, the choice of the orthogonal matrix
 510 R remains a degree of freedom. In this work, we restricted our analysis to randomly sampled
 511 matrices, which strongly influence the resulting tuning curves by mixing the underlying plane waves.
 512 Notably, this choice can be dissociated from the representational similarity structure: as shown in
 513 the remapping analysis, selecting an appropriate initial vector p_0 renders the similarity function C
 514 independent of R . This suggests that while individual tuning curves depend on R , the overall geo-
 515 metry of the neural map does not. Future work should investigate whether biological constraints—
 516 such as metabolic energy efficiency (Cueva & Wei, 2018), non-negativity (Sorscher et al., 2023),
 517 or extrinsic distance preservation (Xu et al., 2025)—mandate specific matrices R . In particular, Xu
 518 et al. (2025) demonstrated that hexagonal symmetry ($M = 3$) is optimal for preserving extrinsic
 519 distances (global Euclidean distance) in periodic representations, potentially explaining the prevalence
 520 of hexagonal grids over square lattices. Additionally, improved quadrature rules for similarity
 521 function approximation warrant exploration. For instance, the eigenfunctions of the Laplacian can
 522 construct optimal truncated Fourier series (Bronstein et al., 2021), suggesting that coefficient se-
 523 lection strategies beyond our Monte Carlo approach may yield superior approximations, while also
 524 selecting for periodic solutions. Together, these constraints could drive generated representations
 525 toward the specific hexagonal or sparse place-bound tunings observed in the brain.

526 Finally, while we propose exact algebraic conditions for properties like path integration and metric
 527 preservation, our framework is primarily descriptive rather than prescriptive regarding their acquisi-
 528 tion. Although we have outlined a local learning mechanism for the emergence of skew-symmetric
 529 weights in Appendix H, a critical direction is to extend the framework to incorporate biological noise
 530 and imperfect commutation. Future research should focus on developing learning rules that yield
 531 approximately commuting weight matrices from realistic synaptic plasticity, while quantifying how
 532 deviations from perfect commutation accumulate into path integration errors. This would bridge the
 533 gap between the mathematical idealization and noisy neural circuits, establishing tolerance bounds
 534 for biological navigation systems.

535 6 CODE AVAILABILITY AND DISCLOSURES

536 All code used to generate the results and figures in this work will be made available upon publication.

537 Large language models were used in writing this paper, with usage limited to improving writing and
 538 readability.

540 REFERENCES

542 M. I. Anderson and K. J. Jeffery. Heterogeneous modulation of place cell firing by changes in
543 context. *The Journal of Neuroscience*, 23(26):8827–8835, 2003. doi: 10.1523/JNEUROSCI.
544 23-26-08827.2003.

545 A. Banino, C. Barry, B. Uria, C. Blundell, T. Lillicrap, P. Mirowski, A. Pritzel, M. J. Chadwick,
546 T. Degris, J. Modayil, G. Wayne, H. Soyer, F. Viola, B. Zhang, R. Goroshin, N. Rabinowitz,
547 R. Pascanu, C. Beattie, S. Petersen, A. Sadik, S. Gaffney, H. King, K. Kavukcuoglu, D. Hassabis,
548 R. Hadsell, and D. Kumaran. Vector-based navigation using grid-like representations in artificial
549 agents. *Nature*, 557(7705):429–433, 2018. doi: 10.1038/s41586-018-0102-6.

550 J. M. Beck, W. J. Ma, R. Kiani, T. Hanks, A. K. Churchland, J. Roitman, M. N. Shadlen, P. E.
551 Latham, and A. Pouget. Probabilistic population codes for bayesian decision making. *Neuron*, 60
552 (6):1142–1152, 2008. doi: 10.1016/j.neuron.2008.09.021.

553 S. Blanes, F. Casas, J.-A. Oteo, and J. Ros. The magnus expansion and some of its applications.
554 *Physics reports*, 470(5-6):151–238, 2009. doi: 10.1016/j.physrep.2008.11.001.

555 M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković. Geometric deep learning: Grids,
556 groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021. URL <https://arxiv.org/abs/2104.13478>.

557 Y. Burak and I. R. Fiete. Accurate path integration in continuous attractor network models of grid
558 cells. *PLoS Computational Biology*, 5(2):e1000291, 2009. doi: 10.1371/journal.pcbi.1000291.

559 C. J. Cueva and X.-X. Wei. Emergence of grid-like representations by training recurrent neural net-
560 works to perform spatial localization. In *International Conference on Learning Representations*,
561 2018. URL <https://openreview.net/forum?id=B17JTOe0->.

562 W. Dorrell, P. E. Latham, T. E. J. Behrens, and J. C. R. Whittington. Actionable neural representa-
563 tions: Grid cells from minimal constraints. In *International Conference on Learning Representa-
564 tions*, 2023. URL <https://openreview.net/forum?id=xfqDe72zh41>.

565 M. Fyhn, T. Hafting, A. Treves, M.-B. Moser, and E. I. Moser. Hippocampal remapping and grid
566 realignment in entorhinal cortex. *Nature*, 446(7132):190–194, 2007. doi: 10.1038/nature05601.

567 R. Gao, J. Xie, X.-X. Wei, S.-C. Zhu, and Y. N. Wu. On path integration of grid cells:
568 Group representation and isotropic scaling. In *Advances in Neural Information Processing
569 Systems*, 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/f076073b2082f8741a9cd07b789c77a0-Paper.pdf.

570 R. J. Gardner, E. Hermansen, M. Pachitariu, Y. Burak, N. A. Baas, B. A. Dunn, M.-B. Moser, and
571 E. I. Moser. Toroidal topology of population activity in grid cells. *Cell*, 602(7895):123–128,
572 2022. doi: 10.1038/s41586-021-04268-7.

573 G. Ginosar, J. Aljadeff, L. Las, D. Derdikman, and N. Ulanovsky. Are grid cells used for navigation?
574 on local metrics, subjective spaces, and black holes. *Neuron*, 111(12):1858–1875, 2023. doi:
575 10.1016/j.neuron.2023.03.027.

576 L. M. Giocomo and M. E. Hasselmo. Time constants of h current in layer ii stellate cells differ
577 along the dorsal to ventral axis of medial entorhinal cortex. *Journal of Neuroscience*, 28(38):
578 9414–9425, 2008. doi: 10.1523/JNEUROSCI.3196-08.2008.

579 L. M. Giocomo, E. A. Zilli, E. Fransén, and M. E. Hasselmo. Temporal frequency of subthreshold
580 oscillations scales with entorhinal grid cell field spacing. *Science*, 315(5819):1719–1722, 2007.
581 doi: 10.1126/science.1139207.

582 L. M. Giocomo, S. A. Hussaini, F. Zheng, E. R. Kandel, M.-B. Moser, and E. I. Moser. Grid cells use
583 hcn1 channels for spatial scaling. *Cell*, 147(5):1159–1170, 2011. doi: 10.1016/j.cell.2011.08.051.

584 T. Hafting, M. Fyhn, S. Molden, M.-B. Moser, and E. I. Moser. Microstructure of a spatial map in
585 the entorhinal cortex. *Nature*, 436(7052):801–806, 2005. doi: 10.1038/nature03721.

594 P. Kanerva. Hyperdimensional computing: An introduction to computing in distributed representa-
 595 tion with high-dimensional random vectors. *Cognitive Computation*, 1(2):139–159, 2009. doi:
 596 10.1007/s12559-009-9009-8.

597

598 D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint*
 599 *arXiv:1412.6980*, 2017. URL <https://arxiv.org/abs/1412.6980>.

600 S. Leutgeb, J. K. Leutgeb, A. Treves, M.-B. Moser, and E. I. Moser. Distinct ensemble codes in
 601 hippocampal areas CA3 and CA1. *Science*, 305(5688):1295–1298, 2004. doi: 10.1126/science.
 602 1100265.

603

604 W. J. Ma, J. M. Beck, P. E. Latham, and A. Pouget. Bayesian inference with probabilistic population
 605 codes. *Nature neuroscience*, 9(11):1432–1438, 2006. doi: 10.1038/nn1790.

606 D. C. McNamee, K. L. Stachenfeld, M. M. Botvinick, and S. J. Gershman. Flexible modulation of
 607 sequence generation in the entorhinal–hippocampal system. *Nature Neuroscience*, 24(6):851—
 608 862, 2021. doi: 10.1038/s41593-021-00831-7.

609

610 B. L. McNaughton, F. P. Battaglia, O. Jensen, E. I. Moser, and M.-B. Moser. Path integration and
 611 the neural basis of the ‘cognitive map’. *Nature Reviews Neuroscience*, 7(8):663–678, 2006. doi:
 612 10.1038/nrn1932.

613 E. I. Moser and M.-B. Moser. A metric for space. *Hippocampus*, 18(12):1142–1156, 2008. doi:
 614 10.1002/hipo.20483.

615

616 S. A. Ocko, K. Hardcastle, L. M. Giocomo, and S. Ganguli. Emergent elasticity in the neural code
 617 for space. *Proceedings of the National Academy of Sciences*, 115(50):E11798–E11806, 2018.
 618 doi: 10.1073/pnas.1805959115.

619 J. O’Keefe and J. Dostrovsky. The hippocampus as a spatial map. preliminary evidence from
 620 unit activity in the freely-moving rat. *Brain Research*, 34(1):171–175, 1971. doi: 10.1016/
 621 0006-8993(71)90358-1.

622

623 A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein,
 624 L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chil-
 625 amkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-
 626 performance deep learning library. In *Advances in Neural Information Processing Systems*,
 627 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf.

628

629 M. Pettersen, F. Rogge, and M. E. Lepperød. Learning place cell representations and context-
 630 dependent remapping. In *Advances in Neural Information Processing Systems*, 2024. URL
 631 <https://openreview.net/forum?id=7ESHFpqjNO>.

632

633 W. T. Redman, F. Acosta, S. Acosta-Mendoza, and N. Miolane. Not so griddy: Internal repre-
 634 sentations of RNNs path integrating more than one agent. In *Advances in Neural Information
 635 Processing Systems*, 2024. URL <https://openreview.net/forum?id=dsMSWUBN8f>.

636 R. RG, G. A. Ascoli, N. M. Sutton, and H. Dinnenberg. Spatial periodicity in grid cell firing is
 637 explained by a neural sequence code of 2-d trajectories. *eLife*, 13:RP96627, 2025. doi: 10.7554/
 638 eLife.96627.

639

640 R. Schaeffer, M. Khona, T. Ma, C. Eyzaguirre, S. Koyejo, and I. Fiete. Self-
 641 supervised learning of representations for space generates multi-modular grid
 642 cells. In *Advances in Neural Information Processing Systems*, 2023. URL
 643 https://proceedings.neurips.cc/paper_files/paper/2023/file/4846257e355f6923fc2a1fbe35099e91-Paper-Conference.pdf.

644

645 V. S. Schøyen, K. Beshkov, M. B. Pettersen, E. Hermansen, K. Holzhausen, A. Malthe-Sørenssen,
 646 M. Fyhn, and M. E. Lepperød. Hexagons all the way down: grid cells as a conformal isometric
 647 map of space. *PLoS computational biology*, 21(2):e1012804, 2025. doi: 10.1371/journal.pcbi.
 1012804.

648 A. Sengupta, C. Pehlevan, M. Tepper, A. Genkin, and D. Chklovskii. Manifold-tiling localized
 649 receptive fields are optimal in similarity-preserving neural networks. In *Advances in Neural Infor-*
 650 *mation Processing Systems*, 2018. URL [https://proceedings.neurips.cc/paper_](https://proceedings.neurips.cc/paper_files/paper/2018/file/ee14c41e92ec5c97b54cf9b74e25bd99-Paper.pdf)
 651 [files/paper/2018/file/ee14c41e92ec5c97b54cf9b74e25bd99-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2018/file/ee14c41e92ec5c97b54cf9b74e25bd99-Paper.pdf).

652 B. Sorscher, G. C. Mel, S. A. Ocko, L. M. Giocomo, and S. Ganguli. A unified theory for the
 653 computational and mechanistic origins of grid cells. *Neuron*, 111(1):121–137.e13, 2023. doi:
 654 10.1016/j.neuron.2022.10.003.

655 K. L. Stachenfeld, M. M. Botvinick, and S. J. Gershman. The hippocampus as a predictive map.
 656 *Nature Neuroscience*, 20(11):1643–1653, 2017. doi: 10.1038/nn.4650.

657 H. Stensola, T. Stensola, T. Solstad, K. Frøland, M.-B. Moser, and E. I. Moser. The entorhinal grid
 658 map is discretized. *Nature*, 492(7427):72–78, 2012. doi: 10.1038/nature11649.

659 P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski,
 660 P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman,
 661 N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W.
 662 Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R.
 663 Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors.
 664 SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:
 665 261–272, 2020. doi: 10.1038/s41592-019-0686-2.

666 N. Waniek. Hexagonal grid fields optimally encode transitions in spatiotemporal sequences. *Neural*
 667 *computation*, 30(10):2691–2725, 2018. doi: 10.1162/neco_a_01122.

668 N. Waniek. Transition scale-spaces: A computational theory for the discretized entorhinal cortex.
 669 *Neural computation*, 32(2):330–394, 2020. doi: 10.1162/neco_a_01255.

670 J. C. R. Whittington, T. H. Muller, S. Mark, G. Chen, C. Barry, N. Burgess, and T. E. J. Behrens. The
 671 Tolman-Eichenbaum Machine: Unifying space and relational memory through generalization in
 672 the hippocampal formation. *Cell*, 183(5):1249–1263.e23, 2020. doi: 10.1016/j.cell.2020.10.024.

673 D. Xu, R. Gao, W. Zhang, X.-X. Wei, and N. Wu. Conformal isometry of lie group representation in
 674 recurrent network of grid cells. In *NeurIPS 2022 Workshop on Symmetry and Geometry in Neural*
 675 *Representations*, 2022. URL <https://openreview.net/forum?id=FszPdSkvGjz>.

676 D. Xu, R. Gao, W. Zhang, X.-X. Wei, and Y. Wu. On conformal isometry of grid cells:
 677 Learning distance-preserving position embedding. In *International Conference on Representa-*
 678 *tion Learning*, 2025. URL [https://proceedings.iclr.cc/paper_files/paper_](https://proceedings.iclr.cc/paper_files/paper_2025/file/4c751ed6cacd53971a1183dcd7821d8c-Paper-Conference.pdf)
 679 [2025/file/4c751ed6cacd53971a1183dcd7821d8c-Paper-Conference.pdf](https://proceedings.iclr.cc/paper_files/paper_2025/file/4c751ed6cacd53971a1183dcd7821d8c-Paper-Conference.pdf).

680 K. Zhang. Representation of spatial orientation by the intrinsic dynamics of the head-direction
 681 cell ensemble: a theory. *Journal of Neuroscience*, 16(6):2112–2126, 1996. doi: 10.1523/
 682 JNEUROSCI.16-06-02112.1996.

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

APPENDIX

A METHODS

All computational simulations were implemented using the matrix exponential operator in PyTorch (Paszke et al., 2019). Random orthogonal matrices were generated using the `ortho_group` module from the SciPy library (Virtanen et al., 2020), which samples uniformly from the orthogonal group $O(N)$. For experiments involving root-of-unity solutions, the spatial domain was set to 20×20 to capture the full periodicity of the representation. For the similarity-preserving experiments, the domains were defined as $s \in [0, 5]$ for the non-spatial context signal and $(x, y) \in [-2, 2]$ for the spatial coordinates, commensurate with the length scale of the target Gaussian similarity function.

B COMMUTING GENERATORS PRODUCE PATH-INDEPENDENT REPRESENTATIONS

Since the spatial representation is defined by the exponential map in Eq. (1), we can enforce specific geometric properties by constraining the generators G_x and G_y . For example, Schaeffer et al. (2023) proposed that representations should be path-independent, that is, the representation at a given location must not depend on the trajectory taken to reach it. In the exponential map formalism, this condition is satisfied if generators commute. To illustrate this, consider two distinct paths to a point D starting from A : a path $A \rightarrow B \rightarrow D$ and a path $A \rightarrow C \rightarrow D$. The resulting representations are generated by composing the transformations for each segment:

$$\mathbf{p}_{ABD} = e^{\Delta x_{BD}G_x + \Delta y_{BD}G_y} e^{\Delta x_{AB}G_x + \Delta y_{AB}G_y} \mathbf{p}_A,$$

$$\mathbf{p}_{ACD} = e^{\Delta x_{CD}G_x + \Delta y_{CD}G_y} e^{\Delta x_{AC}G_x + \Delta y_{AC}G_y} \mathbf{p}_A,$$

where \mathbf{p}_A denotes the representation at A . If G_x and G_y commute, their linear combinations also commute. By the Baker-Campbell-Hausdorff (BCH) formula, the composite transformations then combine into a single exponential summing the exponents:

$$\mathbf{p}_{ABD} = e^{(\Delta x_{AB} + \Delta x_{BD})G_x + (\Delta y_{AB} + \Delta y_{BD})G_y} \mathbf{p}_A = \mathbf{p}_{ACD}.$$

Here, the final state depends only on the net displacement from the initial location, which is identical for both paths.

Conversely, if the generators do not commute, path dependence arises from non-vanishing terms in the BCH expansion. For matrices U and V , the expansion is given by:

$$e^U e^V = e^{U + V + \frac{1}{2}[U, V] + \frac{1}{12}([U, [U, V]] - [V, [U, V]]) + \dots}.$$

Applying this to the path $A \rightarrow B \rightarrow D$, we define $U = \Delta x_{AB}G_x + \Delta y_{AB}G_y$ and $V = \Delta x_{BD}G_x + \Delta y_{BD}G_y$. The representation becomes:

$$\mathbf{p}_{ABD} = e^U e^V \mathbf{p}_A = e^{U + V + \frac{1}{2}[U, V] + \dots} \mathbf{p}_A.$$

The commutator $[U, V]$ expands to:

$$[U, V] = [\Delta x_{AB}G_x + \Delta y_{AB}G_y, \Delta x_{BD}G_x + \Delta y_{BD}G_y].$$

Using the linearity of the commutator and the property $[G_x, G_y] = -[G_y, G_x]$, this simplifies to:

$$[U, V] = (\Delta x_{AB}\Delta y_{BD} - \Delta y_{AB}\Delta x_{BD})[G_x, G_y].$$

Crucially, if $[G_x, G_y] \neq 0$, the exponent includes a term proportional to the cross product of the path segment displacements (geometrically related to the area enclosed by the path components).

Similarly, for the path $A \rightarrow C \rightarrow D$ with operators W and Z , the commutator $[W, Z]$ introduces different displacement cross-terms. Consequently, the higher-order corrections differ between the two paths, and $\mathbf{p}_{ABD} \neq \mathbf{p}_{ACD}$. Thus, commutativity of the generators is a necessary condition for the representation to depend solely on the net displacement.

756 **C FROM GENERATORS TO REPRESENTATIONS**
 757

758 The structure of the spatial representation in Eq. (1) allows for an explicit decomposition of in-
 759 dividual cell responses. Assuming unit-norm \mathbf{p}_0 and commuting, skew-symmetric generators in
 760 block-diagonal form (as in Eq. (5)), the dynamics in the rotated basis $\boldsymbol{\alpha} = R\mathbf{p}$ are given by:
 761

$$762 \boldsymbol{\alpha}(x, y) = e^{x\Sigma_x + y\Sigma_y} \boldsymbol{\alpha}_0,$$

763 where $\boldsymbol{\alpha}_0 = R\mathbf{p}_0$. In this basis, the matrix exponential reduces to a block-diagonal matrix populated
 764 by 2×2 rotation matrices. Following Dorrell et al. (2023), the representation decomposes into
 765 independent rotations within distinct 2D subspaces. For the i -th block ($i = 1, \dots, N/2$), the update
 766 rule is:
 767

$$768 \boldsymbol{\alpha}^i = \begin{pmatrix} \cos(\Omega_i) & -\sin(\Omega_i) \\ \sin(\Omega_i) & \cos(\Omega_i) \end{pmatrix} \boldsymbol{\alpha}_0^i$$

769 where $\boldsymbol{\alpha}^i \in \mathbb{R}^2$ represents the state vector within the i -th subspace. The rotation angle $\Omega_i =$
 770 $x\lambda_{i,x} + y\lambda_{i,y}$ couples the spatial displacement to the generator eigenvalues.
 771

772 Letting α_1^i and α_2^i denote the two components of the i -th block (corresponding to the conjugate
 773 eigenvalue pair), the matrix multiplication yields:
 774

$$775 \alpha_1^i = \alpha_{0,1}^i \cos(\Omega_i) - \alpha_{0,2}^i \sin(\Omega_i), \\ 776 \alpha_2^i = \alpha_{0,1}^i \sin(\Omega_i) + \alpha_{0,2}^i \cos(\Omega_i).$$

777 By applying the harmonic addition theorem, these components can be rewritten as phase-shifted
 778 sinusoids:
 779

$$780 \alpha_1^i = A_i \cos(x\lambda_{i,x} + y\lambda_{i,y} + \omega_i), \\ 781 \alpha_2^i = A_i \sin(x\lambda_{i,x} + y\lambda_{i,y} + \omega_i),$$

782 where the amplitude is $A_i = \sqrt{(\alpha_{0,1}^i)^2 + (\alpha_{0,2}^i)^2}$ and the phase is $\omega_i = \arctan(\alpha_{0,2}^i/\alpha_{0,1}^i)$.
 783

784 These equations demonstrate that, in the canonical basis, each component acts as a 2D plane wave
 785 with orientation and frequency determined by $\lambda_{i,x}$ and $\lambda_{i,y}$, and a phase shift ω_i along the wave
 786 direction. Since the observed neural representation is given by $\mathbf{p} = R^T \boldsymbol{\alpha}$, the firing rate of each
 787 neuron consists of a linear superposition (mixture) of these plane waves.
 788

789 **D METRIC PRESERVATION**
 790

791 Consider the representation along a parametrized trajectory $\mathbf{r}(t) = (x(t), y(t))$:

$$793 \mathbf{p}(x(t), y(t)) = e^{x(t)G_x + y(t)G_y} \mathbf{p}_0.$$

794 The length of this trajectory in the representational space is given by the path integral of the line
 795 element $ds = |\mathbf{p}|$. By the chain rule, the differential change in the representation is:
 796

$$797 d\mathbf{p} = \left(\frac{\partial \mathbf{p}}{\partial x} \frac{dx}{dt} + \frac{\partial \mathbf{p}}{\partial y} \frac{dy}{dt} \right) dt$$

800 The trajectory length L can be expressed in terms of the induced metric g_{ij} as:
 801

$$802 L = \int_0^S \sqrt{|\mathbf{p}|^2} = \int_0^T \sqrt{\sum_{ij} g_{ij} \frac{dr_i}{dt} \frac{dr_j}{dt}} dt.$$

803 Comparing with the squared line element, we can then simply read off the induced metric g induced
 804 metric, as
 805

$$806 g = - \begin{pmatrix} \mathbf{p}_0^T G_x^2 \mathbf{p}_0 & \mathbf{p}_0^T G_x G_y \mathbf{p}_0 \\ \mathbf{p}_0^T G_x G_y \mathbf{p}_0 & \mathbf{p}_0^T G_y^2 \mathbf{p}_0 \end{pmatrix},$$

807 where we have used the fact that $\mathbf{p}^T G_{r_i}^T G_{r_j} \mathbf{p} = -\mathbf{p}_0^T G_{r_i} G_{r_j} \mathbf{p}_0$ due to the skew-symmetry of the
 808 generator matrices.
 809

We further simplify this using the block-diagonal decomposition $G = R^T \Sigma R$. Since Σ contains 2×2 skew-symmetric blocks, $G_x^2 = R^T D_x R$, where D_x is a diagonal matrix containing the squared imaginary part of a given eigenvalue $-\lambda_{i,x}^2$. The mixed term $G_x G_y$ similarly diagonalizes to entries $-\lambda_{i,x} \lambda_{i,y}$. Letting $\alpha_0 = R p_0$, the metric becomes:

$$g = - \begin{pmatrix} \alpha_0^T D_x \alpha_0 & \alpha_0^T D_{xy} \alpha_0 \\ \alpha_0^T D_{xy} \alpha_0 & \alpha_0^T D_y \alpha_0 \end{pmatrix}.$$

Substituting this back into the path length integral yields:

$$L = \int_0^T \sqrt{\sum_{i=1}^N \alpha_{0,i}^2 (\lambda_{i,x}^2 \dot{x}^2 + 2\lambda_{i,x} \lambda_{i,y} \dot{x} \dot{y} + \lambda_{i,y}^2 \dot{y}^2)} dt$$

To preserve the flat Euclidean metric (that is, $g = \sigma^2 I$), we require the off-diagonal terms to vanish and the diagonal terms to be equal:

$$\sum_{i=1}^N \alpha_{0,i}^2 \lambda_{ix}^2 = \sum_{i=1}^N \alpha_{0,i}^2 \lambda_{iy}^2, \quad \text{and} \quad \sum_{i=1}^N \alpha_{0,i}^2 \lambda_{ix} \lambda_{iy} = 0.$$

Introducing polar coordinates for the eigenvalues $\lambda_{i,x} = k_i \cos \phi_i$ and $\lambda_{i,y} = k_i \sin \phi_i$, and defining $\rho_i = \alpha_{0,i} k_i$, these conditions become:

$$\begin{aligned} \sum_{i=1}^N \rho_i^2 \cos^2(\phi_i) &= \sum_{i=1}^N \rho_i^2 \sin^2(\phi_i) \\ \sum_{i=1}^N \rho_i^2 \cos \phi_i \sin \phi_i &= 0. \end{aligned}$$

Using trigonometric identities, this system simplifies to requiring that the weighted sum of phasors vanishes at double the angle:

$$\sum_{i=1}^N \rho_i^2 \cos(2\phi_i) = 0 \quad \text{and} \quad \sum_{i=1}^N \rho_i^2 \sin(2\phi_i) = 0,$$

which is equivalent to the complex condition

$$\sum_{i=1}^N \rho_i^2 e^{2i\phi_i} = 0.$$

Since eigenvalues appear in conjugate pairs ($\phi_j^* = \phi_j + \pi$), and $e^{2i(\phi+\pi)} = e^{2i\phi}$, the sum effectively runs over $N/2$ independent pairs. Assuming the simplest case, with equal weighting $\rho_i = \rho$, the condition reduces to finding a set of angles such that:

$$Z = \sum_{j=1}^{N/2} e^{2i\phi_j} = 0.$$

This condition is satisfied if the $N/2$ eigenvalue pairs are partitioned into subsets (modules), where each subset forms a collection of *roots of unity* that sums to zero. Specifically, for a single module of symmetry order M , we require M eigenvalue pairs (consuming $2M$ dimensions of the total N) with angles distributed uniformly on the circle:

$$\phi_j = \pi \frac{j}{M}, \quad j = 0, 1, \dots, M-1.$$

Consequently, the full high-dimensional representation can be constructed as a linear combination of such sets. For a system with multiple modules $k = 1 \dots K$, each with radius ρ_k , symmetry M_k , and orientation φ_k , the total sum vanishes if each module vanishes individually:

$$Z = \sum_{j=1}^J \rho_j^2 e^{2i\varphi_j} \sum_{m=0}^{M_j-1} e^{2\pi i \frac{m}{M_j}}.$$

864 In other words, for each radius ρ , there can be multiple rotated sets of roots of unity, each with its
 865 own rotational symmetry.

866 A comparison with the explicit form of the representation in Appendix C reveals a striking parallel to
 867 the modular organization of grid cells in the medial entorhinal cortex (Hafting et al., 2005; Stensola
 868 et al., 2012), which are similarly organized into modules defined by grid spacing (ρ), orientation
 869 (φ), and symmetry (M).

871 E SIMILARITY FUNCTION DERIVATION

872 We derive the explicit form of the representational similarity starting from the Eq. (4), We assume
 873 the generators are skew-symmetric and commute, admitting the decomposition

$$874 \quad G_x = R^T \Sigma_x R \quad \text{and} \quad G_y = R^T \Sigma_y R,$$

875 where R is a shared orthogonal matrix. Using the identity $e^{P^{-1}AP} = P^{-1}e^A P$, we rewrite the
 876 similarity expression as

$$877 \quad C(\Delta x, \Delta y) = \left(R \frac{\mathbf{p}_0}{|\mathbf{p}_0|} \right)^T e^{\Delta x \Sigma_x + \Delta y \Sigma_y} \left(R \frac{\mathbf{p}_0}{|\mathbf{p}_0|} \right) \\ 878 = \boldsymbol{\alpha}_0^T e^{\Delta x \Sigma_x + \Delta y \Sigma_y} \boldsymbol{\alpha}_0,$$

880 where we define the rotated unit vector $\boldsymbol{\alpha}_0 \equiv R \frac{\mathbf{p}_0}{|\mathbf{p}_0|}$. Note that $\sum_i \alpha_{0,i}^2 = 1$ due to the orthogonality
 881 of R .

882 The exponent matrix $\Omega = \Delta x \Sigma_x + \Delta y \Sigma_y$ retains the block-diagonal, skew-symmetric structure of
 883 the generators. In the power series expansion of the matrix exponential, even powers Ω^{2n} result in
 884 diagonal matrices (as the square of a 2×2 skew-symmetric block is diagonal), while odd powers
 885 Ω^{2n+1} remain skew-symmetric. Since a quadratic form $\mathbf{x}^T A \mathbf{x}$ vanishes for any skew-symmetric
 886 matrix A , odd terms do not contribute to the similarity. The expansion therefore reduces to a sum
 887 over even powers:

$$888 \quad C(\Delta x, \Delta y) = \boldsymbol{\alpha}_0^T \left(\sum_n \frac{(-1)^n}{(2n)!} D^{2n} \right) \boldsymbol{\alpha}_0,$$

889 where D is a diagonal matrix with entries $\theta_i = \lambda_{i,x} \Delta x + \lambda_{i,y} \Delta y$, and $\lambda_{i,\cdot}$ denotes the imaginary part
 890 of the corresponding eigenvalue. Recognizing the Taylor series for the cosine function, the matrix
 891 sum converges to a diagonal matrix with entries $\cos(\theta_i)$. Consequently, the similarity simplifies to:

$$892 \quad C(\Delta x, \Delta y) = \sum_i^N \alpha_{0,i}^2 \cos(\lambda_{i,x} \Delta x + \lambda_{i,y} \Delta y).$$

903 F DESIGNING SPATIAL SIMILARITY FUNCTIONS

904 We established that the similarity function Eq. (6) takes the general form of a weighted sum of
 905 cosines. To understand the structure of the resulting similarity function, we rewrite the expression
 906 in polar coordinates using $x = r \cos \theta$, $y = r \sin \theta$, and $\lambda_{i,x} = k_i \cos \phi_i$, $\lambda_{i,y} = k_i \sin \phi_i$:

$$907 \quad C = \sum_i \alpha_{0,i}^2 \cos(k_i r \cos(\theta - \phi_i)),$$

908 Applying the Jacobi-Anger expansion

$$909 \quad \cos(z \cos(\omega)) = J_0(z) + 2 \sum_{n=1}^{\infty} (-1)^n J_{2n}(z) \cos(2n\omega) \\ 910 \\ 911 = J_0(z) + 2 \sum_{n=1}^{\infty} (-1)^n \Re \{ J_{2n}(z) e^{2in\omega} \},$$

918 where $J_n(z)$ is the n -th Bessel function of the first kind, we obtain:
 919

$$920 \quad C(r, \theta) = \sum_j \alpha_{0,j}^2 J_0(k_j r) + 2 \sum_{n=1}^{\infty} (-1)^n \Re \left\{ e^{2in\theta} \sum_j \alpha_{0,j}^2 J_{2n}(k_j r) e^{2in\phi_j} \right\}.$$

$$921$$

$$922$$

923 This decomposes the similarity into a purely radial component (the J_0 term) and a mixed term
 924 dependent on head direction θ .
 925

926 Further simplification relies on the structure of the eigenvalues. If the representation preserves the
 927 flat metric (see Appendix D), the eigenvalues form discrete roots-of-unity constellations. Assuming
 928 constant weighting $\alpha_{0,j}$ on each ring j , the inner phasor sum becomes:
 929

$$930 \quad \sum_j \alpha_{0,j}^2 J_{2n}(k_j r) e^{2in\phi_j} = \sum_j \alpha_{0,j}^2 J_{2n}(k_j r) e^{2in\varphi_j} \sum_{m=0}^{M_j-1} e^{2\pi imn/M_j},$$

$$931$$

932 where the geometric series $\sum_{m=0}^{M_j-1} e^{2\pi imn/M_j}$ vanishes unless n is a multiple of the symmetry order
 933 M_j . Consequently, angular dependence only arises at harmonic orders $n = \ell M_j$. For large M , the
 934 similarity function becomes approximately isotropic (head-direction independent).
 935

936 We can further suppress low-order angular terms by requiring the orientation offsets φ_j of different
 937 modules to also form a root-of-unity constellation. Specifically, if we sum over a set of orientations
 938 indexed by l such that:
 939

$$940 \quad \sum_j \alpha_{0,j}^2 J_{2n}(k_j r) e^{2in\phi_j} = \sum_j \alpha_{0,j}^2 J_{2n}(k_j r) \sum_l e^{2in\varphi_l} \sum_{m=0}^{M-1} e^{2\pi imn/M},$$

$$941$$

942 and the orientations satisfy the condition $\sum_l e^{2in\varphi_l} = \sum_{l=0}^{N-1} e^{2\pi iln/N}$, then non-zero terms persist
 943 only when n is a multiple of N and simultaneously a multiple of M . If M and N are coprime, the
 944 lowest order angular dependence is pushed to $n = MN$. In this case, the full similarity function is:
 945

$$946 \quad C(r, \theta) = \sum_j \alpha_{0,j}^2 J_0(k_j r) + 2MN \sum_{\ell=1}^{\infty} (-1)^{\ell NM} \alpha_{0,j}^2 J_{2\ell MN}(k_j r) \cos(2\ell MN\theta).$$

$$947$$

948 The magnitude of the angular terms is governed by the high-order Bessel functions. For small
 949 arguments z , the Bessel function of order γ behaves as:
 950

$$951 \quad J_{\gamma}(z) \approx \frac{1}{\Gamma(\gamma+1)} \left(\frac{z}{2}\right)^{\gamma}$$

$$952$$

953 Here, the order is $\gamma = 2\ell MN$. For large MN , this term decays rapidly near the origin ($z\sqrt{\gamma+1}$).
 954 Consequently, for a large range of displacements r , the angular terms vanish, and the similarity
 955 becomes effectively radial:
 956

$$957 \quad C(r, \theta) \approx \sum_j \alpha_{0,j}^2 J_0(k_j r).$$

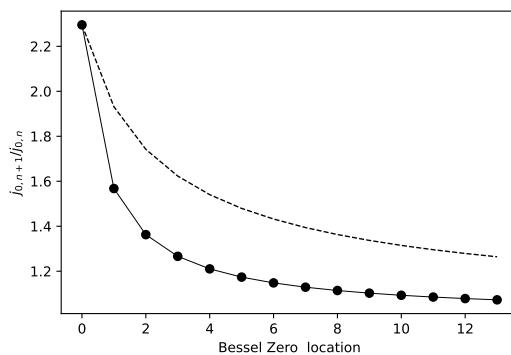
$$958$$

959 Intriguingly, this connection offers a theoretical prediction for grid spacing ratios. The ratio of
 960 successive zeros of the Bessel function J_0 converges toward an average value close to $\sqrt{2}$ when
 961 including low-order zeros (Fig. 4). This falls precisely within the variability range of grid module
 962 spacing ratios observed experimentally (Stensola et al., 2012). While grid scale ratios are often
 963 assumed to be $\sqrt{2}$ to maximize spatial range, our framework suggests they may arise from the
 964 optimal approximation of a radial similarity kernel via a Fourier-Bessel expansion. Finally, we
 965 note that while this ratio determines relative scaling, the Fourier-Bessel series inherently defines an
 966 absolute length scale related to the domain of approximation. This suggests that the absolute grid
 967 scale may be set by the size of the region the animal needs to encode reliably.
 968

969 G SIMILARITY-PRESERVING SPATIAL REPRESENTATIONS

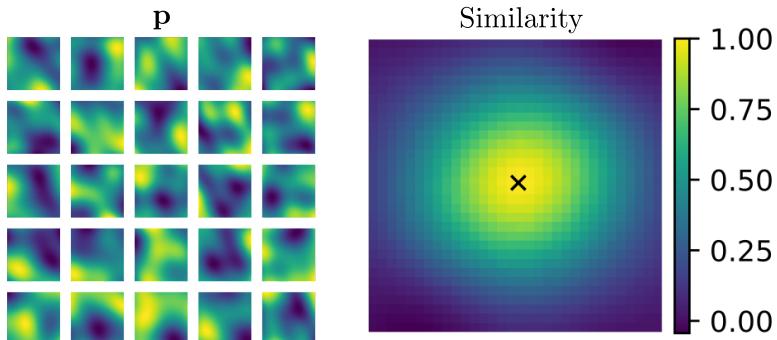
$$970$$

971 By relaxing the strict requirement for metric preservation, we can explore representations designed
 972 to preserve similarity structure. Following the approach outlined in Section 2.5, we consider the case



989 where generator eigenvalues are sampled from a normal distribution. This choice yields an approx-
990 imate Gaussian similarity function. To demonstrate that this generalizes to spatial representations,
991 we simulated a population of $N = 256$ units with generator eigenvalues for G_x and G_y sampled
992 from $\mathcal{N}(0, 2)$. The results are shown in Fig. 5. Unlike the periodic, grid-like patterns character-
993 istic of low-order roots-of-unity solutions (metric preservation), these units exhibit heterogeneous,
994 spatially localized tuning curves reminiscent of hippocampal place fields (O’Keefe & Dostrovsky,
995 1971). Given that place cells are known to encode both spatial and non-spatial cues, such as ol-
996 factory context, (Anderson & Jeffery, 2003), this result suggests that the context-dependent model
997 in Eq. (7) could be naturally extended to model conjunctive representations of space and context,
998 consistent with recent theoretical proposals (Pettersen et al., 2024).

999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012



H BIOLOGICAL INTERPRETATION

1022
1023

1024 The exponential map framework provides a rigorous algebraic description of spatial representa-
1025 tions. Here, we demonstrate that this framework is not merely an abstraction but emerges as the
1026 on-manifold dynamics of a gain-modulated Continuous Attractor Neural Network (CANN).

1026
1027

H.1 CONTINUOUS ATTRACTOR FORMULATION

1028
1029

We consider the general dynamics of a recurrent neural network Zhang (1996); Ocko et al. (2018) governed by:

1030
1031

$$\tau \frac{d\mathbf{z}}{dt} = W[\mathbf{v}(t)]\sigma(\mathbf{z}(t)), \quad (9)$$

1032
1033
1034
1035
1036
1037
1038
1039

where $\mathbf{z} \in \mathbb{R}^N$ is the neural state vector and τ is the effective time constant. Here, $W[\mathbf{v}]$ represents the effective state-transition matrix modulated by an external input \mathbf{v} , such as 2D velocity, and σ is a generalized non-linear function. Note that this formulation subsumes standard leaky-integrator models (for example, $\tau \dot{\mathbf{z}} = -\mathbf{z} + J\phi(\mathbf{z})$) if the decay term is absorbed into the effective interaction $W\sigma(\mathbf{z})$. We investigate the case where network activity flows toward a low-dimensional manifold \mathcal{M} , specifically a unit hypersphere, consistent with the orthogonal transformations derived in our algebraic framework. To determine necessary conditions for a hypersphere attractor, we consider the Lyapunov energy function:

1040
1041

$$E(t) = \frac{1}{2}(\|\mathbf{z}(t)\|^2 - 1)^2, \quad (10)$$

1042
1043
1044

assuming a unit radius ($R = 1$) for simplicity. This energy quantifies the deviation of the state from the manifold surface. The time evolution of the energy along a trajectory in the neural state space is determined via the chain rule. Letting $u = \|\mathbf{z}\|^2 - 1$ denote the deviation from the manifold, we have:

1045
1046

$$\dot{E} \equiv \frac{dE}{dt} = \frac{dE}{du} \frac{du}{dt} = u\dot{u}.$$

1047
1048

The rate of change of the squared norm is derived using the product rule on the inner product $\mathbf{z}^T \mathbf{z}$:

1049
1050

$$\frac{d}{dt}(\|\mathbf{z}\|^2) = \frac{d}{dt}(\mathbf{z}^T \mathbf{z}) = \dot{\mathbf{z}}^T \mathbf{z} + \mathbf{z}^T \dot{\mathbf{z}} = 2\mathbf{z}^T \dot{\mathbf{z}},$$

1051
1052

where the last step follows from the symmetry of the Euclidean inner product ($\mathbf{a}^T \mathbf{b} = \mathbf{b}^T \mathbf{a}$). Combining these terms yields:

1053
1054

$$\dot{E} = (\|\mathbf{z}\|^2 - 1)2\mathbf{z}^T \dot{\mathbf{z}}.$$

1055

Substituting the effective network dynamics from Eq. (9), the energy derivative becomes:

1056
1057

$$\tau \dot{E} = (\|\mathbf{z}\|^2 - 1)2\mathbf{z}^T W[\mathbf{v}]\sigma(\mathbf{z}).$$

1058
1059

This form of the energy derivative implies a specific condition for convergence. If the effective interaction satisfies

1060

$$W[\mathbf{v}]\sigma(\mathbf{z}) = -(\|\mathbf{z}\|^2 - 1)M(\mathbf{v})\mathbf{z},$$

1061
1062

where M is an input-dependent matrix chosen such that the symmetric form is positive definite, then the energy derivative becomes:

1063
1064

$$\tau \dot{E} = -2(\|\mathbf{z}\|^2 - 1)^2 \mathbf{z}^T M(\mathbf{v}) \mathbf{z} \leq 0.$$

1065
1066
1067

Since the quadratic form $\mathbf{z}^T M \mathbf{z}$ is positive, the energy is strictly non-increasing, driving any non-zero neural state toward the hypersphere ($\dot{E} = 0$ only when $\|\mathbf{z}\| = 1$). A simple sufficient condition is for the symmetric part of M to have positive eigenvalues.

1068
1069
1070

We observe that this attractor network admits a fundamental symmetry: the evolution of the energy is invariant under the transformation

1071
1072

$$M(\mathbf{v}) \rightarrow M(\mathbf{v}) + \hat{U}(\mathbf{v}),$$

1073
1074
1075
1076

provided \hat{U} is skew-symmetric ($\hat{U}^T = -\hat{U}$). This invariance holds because the quadratic form of any skew-symmetric matrix vanishes identically ($\mathbf{z}^T \hat{U} \mathbf{z} = 0$), meaning \hat{U} does not contribute to the energy derivative. By choosing $\hat{U}(\mathbf{v}) = (\|\mathbf{z}\|^2 - 1)^{-1}U(\mathbf{v})$ and setting $\sigma(\mathbf{z}) = (\|\mathbf{z}\|^2 - 1)\mathbf{z}$, the general hypersphere attractor admits the effective dynamics:

1077
1078

$$\tau \dot{\mathbf{z}} = U(\mathbf{v})\mathbf{z} + M\sigma(\mathbf{z}). \quad (11)$$

1079

This decomposition reveals two distinct functional components: the non-linear term (scaled by M) enforces the attractor dynamics normal to the manifold, while the linear term (scaled by U) drives

1080 transport along the manifold. Whenever the state is on the manifold, $\mathbf{z}^* \in \mathcal{M}$, the attractor contribution vanishes ($\sigma(\mathbf{z}^*) = 0$). The on-manifold dynamics then reduce to pure transport:
 1081
 1082

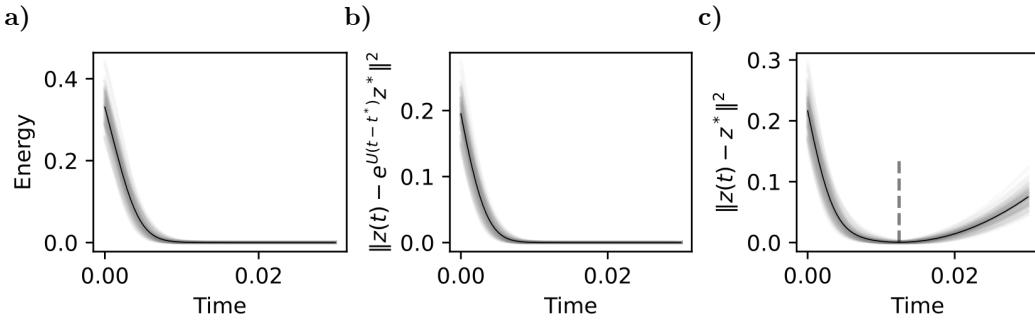
$$1083 \tau \dot{\mathbf{z}}^*(t) = U(\mathbf{v}(t))\mathbf{z}^*(t).$$

1084 For constant velocity inputs, this linear system has the exact solution:
 1085
 1086

$$1087 \mathbf{z}^*(t) = e^{\frac{t}{\tau}U(\mathbf{v})}\mathbf{z}^*(0).$$

1088 Thus, the exponential map exactly describes the trajectory of the neural state on the attractor manifold.
 1089 This framework can be realized by a simple recurrent neural network where U is a gain-modulated skew-symmetric matrix (encoding path integration) and M is a positive definite matrix (enforcing stability).
 1090

1091 To validate this derivation, we simulated the dynamics of Eq. (11) using Euler integration. We
 1092 constructed the connectivity matrices as $U = A - A^T$ and $M = B^T B$, where $A, B \in \mathbb{R}^{N \times N}$
 1093 were sampled from a normal distribution. These choices ensure that U is skew-symmetric and M
 1094 is positive definite. We simulated 100 trajectories initialized randomly within a hypercube of side
 1095 length 0.1, using a time step $dt = 5 \cdot 10^{-5}$ and total duration $T = 0.03$ (for example, seconds,
 1096 but units are arbitrary). For each trajectory, we computed the Lyapunov energy defined in Eq. (10).
 1097 Points satisfying the condition $E \leq \varepsilon = 10^{-4}$ were classified as “close points,” representing
 1098 states effectively on the manifold. Subsequently, we computed the deviation between the evolving
 1099 network state and the theoretical exponential map trajectory initiated from the first identified close
 1100 point. The results are shown in Fig. 6 and demonstrate three key properties. First, the Lyapunov
 1101 decays to zero for all trajectories, indicating that the manifold is attractive and that our construction
 1102 effectively realizes an attractor network (Fig. 6a). Second, the network state on the manifold is
 1103 accurately described by the exponential map; the Euclidean distance between the simulated state
 1104 and the theoretical exponential evolution approaches zero after convergence (Fig. 6b). Finally, to
 1105 demonstrate that the network state evolves on the manifold, we computed the deviation between the
 1106 network state and the first close point. As shown in Fig. 6c, this deviation increases after the close
 1107 point is reached, indicating that the network moves away from \mathbf{z}^* over time, in a manner consistent
 1108 with the dynamics captured by the exponential map.
 1109



1121 **Figure 6: Exponential maps describe the on-manifold dynamics of a recurrent attractor net-
 1122 work.** **a)** Time evolution of the Lyapunov energy of the hypersphere attractor network for 100
 1123 trajectories randomly initialized in a hypercube. The gray lines represent individual trajectories and
 1124 the dark line represents the average. **b)** Squared distance between the network state \mathbf{z} and the theo-
 1125 retical exponential map evolution of the first close point on the attractor (\mathbf{z}^*), occurring at $t = t^*$. **c)**
 1126 Squared distance between network state \mathbf{z} and the static first close point on the attractor (\mathbf{z}^*), show-
 1127 ing divergence due to transport. The dashed line indicates the average time to convergence $\langle t^* \rangle$.
 1128

1129 H.2 CONNECTION TO THE PATH-INTEGRATING EXPONENTIAL MAP

1130 As we saw in the previous section, we can decompose the weight matrix into a baseline component
 1131 that sustains the attractor bump and a skew symmetric component that determines the on-manifold
 1132 dynamics of the attractor network. To bridge this general result with the specific exponential map
 1133 structure proposed in the main text, we adopt the mechanism of velocity-driven updating standard

1134 in continuous attractor models (Burak & Fiete, 2009; McNaughton et al., 2006). Specifically, we
 1135 define the transport operator U to be linearly modulated by the self-motion signal $\mathbf{v}(t)$.
 1136

1137 By defining $U(\mathbf{v}(t)) = v_x(t)U_x + v_y(t)U_y$, the linearized dynamics become:

$$1138 \tau \dot{\mathbf{z}} = (v_x U_x + v_y U_y) \mathbf{z}$$

1140 The solution to this time-varying matrix differential equation involves the time-ordered exponential
 1141 (\mathcal{T}):

$$1142 \mathbf{z}(t) = \mathcal{T} \left\{ \exp \left(\int_0^t \frac{1}{\tau} (v_x(\xi)U_x + v_y(\xi)U_y) d\xi \right) \right\} \mathbf{z}(0).$$

1144 Mathematically, the time-ordered exponential expands into the infinite series known as the Magnus
 1145 expansion (Blanes et al., 2009). This series includes integral terms involving nested commutators of
 1146 the operators at different time points, such as, $[U(\mathbf{v}(\xi_1)), U(\mathbf{v}(\xi_2))]$. Expanding this term reveals
 1147 that it is proportional to the commutator of the basis matrices, $[U_x, U_y]$.

1148 For the system to perform exact path integration, the resulting population vector $\mathbf{p}(x, y)$ must de-
 1149 pend solely on the net accumulated displacement, independent of the specific velocity history $\mathbf{v}(t)$
 1150 or trajectory taken. This path-independence requires the time-ordered exponential to reduce to a
 1151 standard matrix exponential of the integrated inputs:

$$1152 \mathbf{z}(t) = \exp \left(\int_0^t \frac{1}{\tau} (v_x(\xi)U_x + v_y(\xi)U_y) d\xi \right) \mathbf{z}(0).$$

1155 This simplification occurs if the commutator terms in the Magnus expansion vanish. Consequently,
 1156 exact path integration imposes the algebraic constraint that the effective skew-symmetric matrices
 1157 must commute: $[U_x, U_y] = 0$.

1158 Under this condition, we can integrate the velocity inputs directly. Identifying the spatial coordinates
 1159 as $x(t) = \int_0^t v_x(\xi) d\xi$ and $y(t) = \int_0^t v_y(\xi) d\xi$, we recover the exponential map model:

$$1161 \mathbf{z}(t) = e^{xG_x + yG_y} \mathbf{z}(0),$$

1162 where the generator matrices are identified as $G_x = \frac{1}{\tau}U_x$ and $G_y = \frac{1}{\tau}U_y$.

1164 This derivation provides a rigorous physical interpretation of the abstract generators: they corre-
 1165 spond to the velocity-modulated synaptic connectivity U scaled by the inverse effective time con-
 1166 stant $1/\tau$. Furthermore, it demonstrates that the algebraic condition for path independence (com-
 1167muting generators) derived in the main text is dynamically equivalent to the requirement for a neural
 1168 circuit to perform exact temporal integration without trajectory-dependent errors.

1169 H.3 MEMBRANE DYNAMICS AND GRID SCALE HIERARCHIES

1171 The identification of the generator matrices as $G = \frac{1}{\tau}U$ in the previous section implies an intrinsic
 1172 coupling between the network’s time constant τ and the spatial scale of the resulting representation.
 1173 To derive this relationship, consider the path-integrating solution on the manifold:

$$1175 \mathbf{z}(x, y) = e^{\frac{1}{\tau}(xU_x + yU_y)} \mathbf{z}_0.$$

1176 As established in Appendix C, the activity of any unit in this representation decomposes into a
 1177 superposition of plane waves. For a metric-preserving representation, for example, $M = 3$, these
 1178 waves are generated by rotations in 2D subspaces. The phase angle Φ_i for the i -th subspace is given
 1179 by:

$$1180 \Phi_i(x, y) = \frac{1}{\tau}(x\lambda_{i,x} + y\lambda_{i,y}),$$

1182 where $\lambda_{i,\cdot}$ are the imaginary parts of the eigenvalues of the connectivity matrices U . Converting to
 1183 polar coordinates with spatial displacement r and head direction θ , the argument becomes:

$$1184 \Phi_i(r, \theta) = \frac{1}{\tau} r k_i \cos(\theta - \phi_i),$$

1185 where k_i represents the intrinsic frequency (magnitude of the eigenvalue) and ϕ_i the orientation
 1186 of the wave. The pattern repeats at integer multiples of 2π . Considering motion along the wave’s

1188 propagation direction ($\theta = \phi_i$), the fundamental spatial period Λ_i is determined by the condition
 1189 $\Phi_i(\Lambda_i, \phi_i) = 2\pi$. Solving for Λ_i yields:
 1190

$$1191 \quad \Lambda_i = \frac{2\pi\tau}{k_i}.$$

1193 This result demonstrates that the period of the constituent plane waves scales linearly with the net-
 1194 work's time constant τ .
 1195

1196 For a hexagonal grid cell representation ($M = 3$), the global grid pattern is constructed from the
 1197 interference of three such plane waves. Since the grid spacing corresponds to the distance between
 1198 pattern repetitions, it is geometrically constrained to be a fixed multiple of the constituent wave
 1199 periods. Consequently, the grid spacing λ_{grid} must also scale linearly with the time constant:
 1200

$$1201 \quad \lambda_{\text{grid}} \propto \tau.$$

1202 This algebraic derivation mirrors the known topographic organization of the medial entorhinal cor-
 1203 tex (MEC). Experimental evidence confirms that the membrane time constants of MEC stellate cells
 1204 increase along the dorsal-ventral axis (Giocomo et al., 2007; Giocomo & Hasselmo, 2008), corre-
 1205 latting precisely with the expansion of the grid scale. Furthermore, genetic perturbations that increase
 1206 the integrative time constant (such as HCN1 channel knockouts) cause a corresponding expansion
 1207 in grid spacing (Giocomo et al., 2011). Our framework thus provides a normative mathematical
 1208 explanation for this phenomenon: the spatial scale of the neural map is physically grounded in the
 1209 integration speed of the underlying neural substrate.
 1210

1210 H.4 EMERGENCE OF SKEW-SYMMETRIC CONNECTIVITY VIA LOCAL LEARNING

1212 We investigate a biologically plausible local learning rule capable of driving the weight matrix W
 1213 toward skew-symmetry. Motivated by the requirement that the symmetric component of the connec-
 1214 tivity must vanish to preserve the manifold energy, we propose the update rule:
 1215

$$1216 \quad \Delta W_{ij} = \eta(r_i(t)r_j(t+\tau) - r_j(t)r_i(t+\tau)) - \alpha W_{ij},$$

1217 where r_i and r_j denote pre- and postsynaptic activities, τ represents a small temporal lag, η is the
 1218 learning rate, and α is a local decay term.
 1219

1220 Decomposing the weight matrix into symmetric ($S = W + W^T$) and antisymmetric ($A = W - W^T$)
 1221 components reveals distinct evolutionary dynamics. Since the Hebbian term $H_{ij} = r_i(t)r_j(t+\tau) -$
 1222 $r_j(t)r_i(t+\tau)$ is inherently antisymmetric ($H_{ij} = -H_{ji}$), it contributes zero to the update of the
 1223 symmetric component. Consequently, the dynamics of S are governed solely by the decay term:
 1224

$$\Delta S = -\alpha S.$$

1225 This ensures that any initial symmetric connectivity decays asymptotically to zero. Conversely, the
 1226 antisymmetric component is reinforced by the Hebbian term:
 1227

$$1228 \quad \Delta A = 2\eta(r_i(t)r_j(t+\tau) - r_j(t)r_i(t+\tau)) - \alpha A.$$

1229 Thus, A is driven by the time-lagged anti-correlated activity patterns while being stabilized by the
 1230 decay α . In the steady state, the antisymmetric part converges to:
 1231

$$1232 \quad A^* = \frac{2}{\alpha} \langle r_i(t)r_j(t+\tau) - r_j(t)r_i(t+\tau) \rangle_t,$$

1233 where the brackets denote a time average. This implies that the learned skew-symmetric connectivity
 1234 is determined by the statistics of the time-lagged antisymmetric correlations in the network activity.
 1235 By controlling these correlations, the network can, in principle, learn specific generator structures.
 1236

1237 To validate this mechanism, we simulated the training of a 32×32 weight matrix. We initialized the
 1238 weights and the neural activity rates from a normal distribution, with parameters set to $\alpha = 0.001$
 1239 and $\eta = 0.01$. The evolution of the weights was tracked over 5000 iterations. The results, shown in
 1240 Fig. 7, demonstrate that the magnitude of the symmetric component decays to zero, while the skew-
 1241 symmetric component persists. This confirms that the proposed rule effectively filters out symmetric
 1242 connectivity, allowing skew-symmetric generators to emerge naturally from local plasticity rules.
 1243

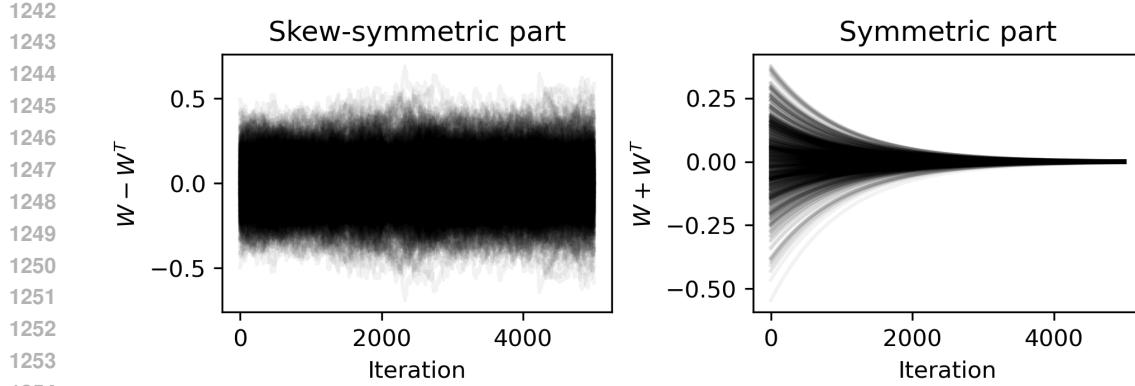


Figure 7: **Emergence of skew-symmetric connectivity via local learning.** Evolution of weight matrix components under the anti-symmetric Hebbian learning rule. The symmetric component (decaying curve) vanishes over time, while the skew-symmetric component (fluctuating curve) persists, driven by the correlation structure of the inputs.

I LEARNING EXPONENTIAL MAPS FROM DATA

To validate the practical applicability of our theoretical framework, we demonstrate that the exponential map model can learn to reproduce and generalize experimental neural data. Using publicly available recordings of grid cells from rats navigating an open field (Gardner et al., 2022), we trained the model to capture the generative structure of the spatial code.

We formulated a one-step predictive model where the predicted firing rate \hat{g} at a target location is generated from the rate at a source location via the matrix exponential. For a spatial displacement $(\Delta x, \Delta y)$ corresponding to small integer index steps, the update rule is:

$$\hat{g}[x + \Delta x, y + \Delta y] = e^{\Delta x G_x + \Delta y G_y} \hat{g}[x, y],$$

where G_x and G_y are learnable generator matrices. We optimized the model using stochastic gradient descent to minimize a composite loss function:

$$\mathcal{L} = \frac{1}{K} \sum_k^K \|\mathbf{g}_{target} - \hat{\mathbf{g}}_{pred}\|^2 + \frac{1}{N^2} \sum_{i,j}^{N^2} ((G_x G_y)_{ij} - (G_y G_x)_{ij})^2.$$

The first term represents the reconstruction error between the predicted activity and the experimental rate map g (averaged over a batch size K). The second term is a commutation penalty, which enforces the algebraic constraint $[G_x, G_y] \approx 0$. This constraint is critical for ensuring path independence and allows the model to learn consistent spatial maps without requiring computationally expensive multi-step training sequences.

For the training protocol, generators G_x and G_y were initialized as random uniform matrices scaled by $1/\sqrt{N}$, where N is the number of simulated neurons. Optimization was performed using the Adam optimizer (Kingma & Ba, 2017) with a learning rate of 0.001 for 20,000 iterations. Training samples were generated by selecting random starting locations within the experimental rate maps and taking single steps of up to two pixels; boundary effects were mitigated by reflecting steps that exceeded the rate map limits.

To evaluate the model’s generative capacity, we reconstructed global rate maps by path integrating from the center of the arena (origin) to a dense grid of position coordinates (x, y) :

$$\mathbf{p}(x, y) = e^{x G_x + y G_y} \mathbf{p}(0, 0).$$

To test extrapolation, we extended this coordinate grid to cover an area with twice the side length of the original experimental enclosure.

1296 The results, shown in Fig. 8, demonstrate that the exponential map not only reproduces the training
 1297 data but successfully generalizes the grid pattern well beyond the boundaries of the original recording
 1298 enclosure (indicated by the black square). The loss history (Fig. 9) confirms that the model
 1299 simultaneously minimizes reconstruction error and enforces commutativity. This indicates that the
 1300 commutation constraint enables the model to extract the robust, intrinsic algebraic structure of the
 1301 spatial representation directly from noisy biological data. We note, however, that the learned rate
 1302 maps exhibit more uniform peak firing rates than their biological counterparts. We hypothesize that
 1303 this uniformity may be an emergent consequence of the commutation penalty and the strict path
 1304 invariance it enforces. While this idealization abstracts away biological heterogeneity, which may
 1305 stem from noise or conjunctive inputs not modeled here, it allows the framework to robustly capture
 1306 the intrinsic spatial phase and periodicity of the grid pattern.

1307 These results suggest several promising avenues for future inquiry. First, future work should compare
 1308 the remapping dynamics of metric-preserving model units (Fig. 3) with biological recordings
 1309 to determine if they exhibit the coherent remapping characteristic of entorhinal and hippocampal
 1310 ensembles (Fyhn et al., 2007). Second, fitting exponential map models to other cell types, such as
 1311 place cells, offers a pathway to deriving interpretable models of their underlying dynamics. Specif-
 1312 ically, it remains to be determined whether the exponential map can account for the heterogeneous,
 1313 apparently stochastic spatial arrangement of biological place fields. A compelling validation would
 1314 involve fitting the model to rate maps recorded in restricted enclosures and generating predictive
 1315 extrapolations for larger environments. Comparing these predictions against experimental data from
 1316 expanded arenas would rigorously test the model’s ability to capture the intrinsic generative structure
 1317 of the spatial code.

1318 J BINDING AND BUNDLING AS A BASIS FOR MULTIPLE-MAP, 1319 REWARD-ORIENTED NAVIGATION

1320 While our primary focus has been on the geometric and computational properties of spatial rep-
 1321 resentations, navigation involves more than localization. In this section, we demonstrate that the
 1322 exponential map framework supports reward-based navigation in an interpretable manner by draw-
 1323 ing on principles from Hyperdimensional Computing (HDC) (Kanerva, 2009).

1324 HDC encodes information using distributed representations based on high-dimensional vectors,
 1325 commonly termed hypervectors. A core property of such high-dimensional spaces is the “con-
 1326 centration of measure,” which ensures that randomly sampled vectors become nearly orthogonal
 1327 with high probability with increasing vector dimension. HDC systems exploit this phenomenon via
 1328 two primary operations. The first is *bundling*, or superposition, which aggregates vectors to form
 1329 a composite representation that remains similar to its inputs. The second is *binding*, an operation
 1330 that combines vectors to produce a result that is dissimilar, and effectively orthogonal, to its con-
 1331 stituents. We show that the exponential map naturally implements these operations, thereby tying
 1332 together spatial representation, memory, and goal-oriented behavior.

1334 J.1 BUNDLING: CONSTRUCTING REWARD MAPS

1335 Assuming a similarity-preserving representation, such as the Gaussian similarity derived in Sec-
 1336 tion 2.5, we can implement a simple memory mechanism by “bundling” the population vectors of
 1337 salient locations. Let $\{\mathbf{r}_i\}_{i=1}^{N_R}$ be a set of locations associated with a reward. A composite memory
 1338 vector \mathbf{p}_R is formed by summing the representations:

$$1339 \mathbf{p}_R = \frac{1}{N_R} \sum_i^{N_R} \mathbf{p}(\mathbf{r}_i).$$

1340 Querying the current location $\mathbf{p}(\mathbf{r})$ against this memory yields a similarity score C_R :

$$1341 C_R(\mathbf{r}) = \mathbf{p}(\mathbf{r})^T \mathbf{p}_R = \frac{1}{N_R} \sum_i^{N_R} \mathbf{p}(\mathbf{r})^T \mathbf{p}(\mathbf{r}_i).$$

1342 Since the pairwise similarity approximates a Gaussian, the resulting score $C_R(\mathbf{r}) \approx$
 1343 $\frac{1}{N_R} \sum_i e^{-\|\mathbf{r} - \mathbf{r}_i\|^2/\sigma^2}$ acts as a smooth reward density map (a mixture of Gaussians). This allows an

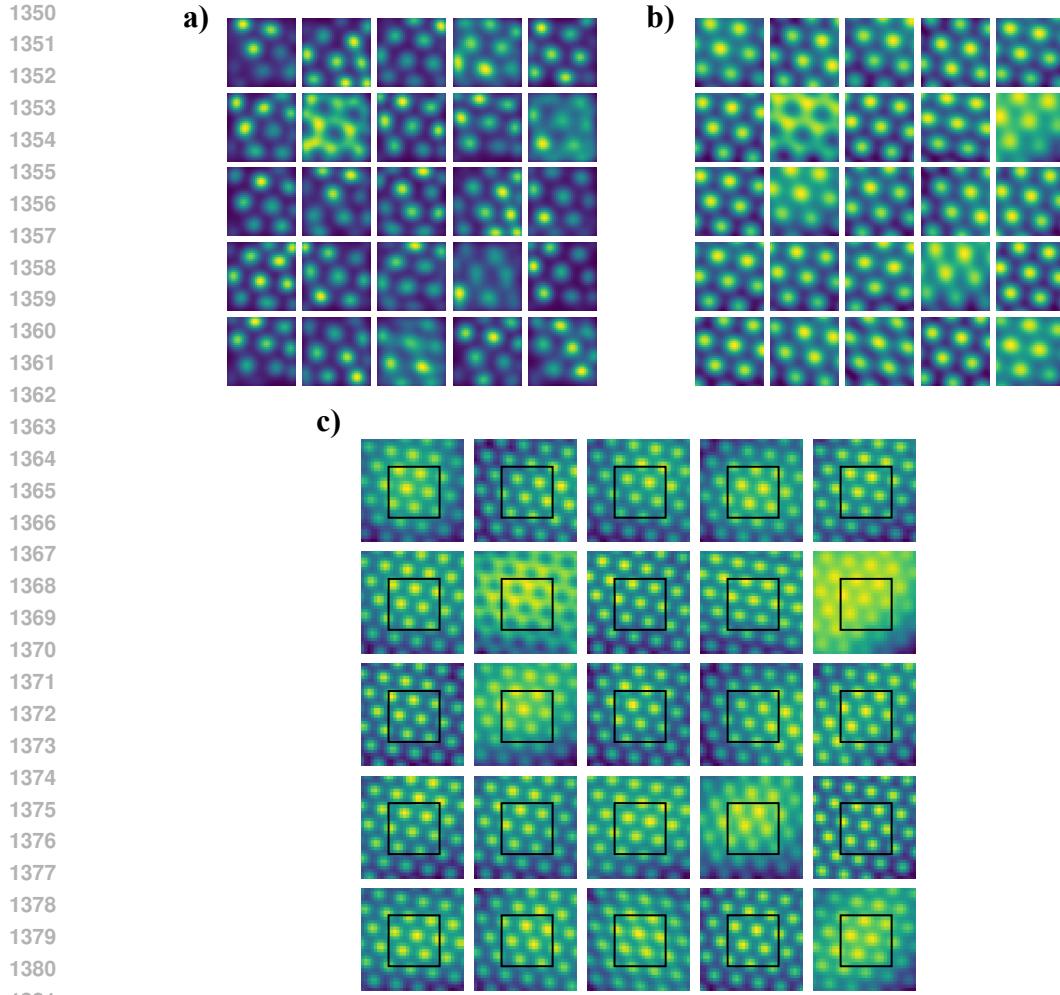


Figure 8: **Generative modeling and extrapolation of experimental grid cells.** **a)** Experimental grid cell rate maps used for training. **b)** Synthetic rate maps generated by the exponential map model within the training domain. **c)** Extrapolated rate maps in the extended environment. The model captures the grid structure within the training domain and successfully extrapolates the periodic pattern beyond the boundaries (indicated by the black outline). Rate maps are matched cell-to-cell across rows.

agent to navigate to rewards by simply following the gradient of the similarity surface:

$$\mathbf{r}_{t+1} \leftarrow \mathbf{r}_t + \eta \nabla_{\mathbf{r}} C_R,$$

where η is a step size. Biologically, this gradient ascent could be approximated by sampling local steps and moving in the direction of increasing similarity.

J.2 BINDING: CONTEXT-DEPENDENT MAP RETRIEVAL

While bundling creates a single reward map, complex navigation requires storing distinct maps for different contexts, for instance distinguishing between a “Food” context and a “Home” context. The standard HDC “binding” operation, often implemented via element-wise multiplication, orthogonalizes vectors. In our framework, context-dependent remapping functions as an intrinsic binding operation.

Recall that the context generator G_s produces orthogonal transformations. If the remapping is sufficiently strong, implying dissimilar contexts, the representations become nearly orthogonal. We can thus form a general, context-dependent memory vector $\mathbf{p}_{R,S}$ that bundles spatial locations across

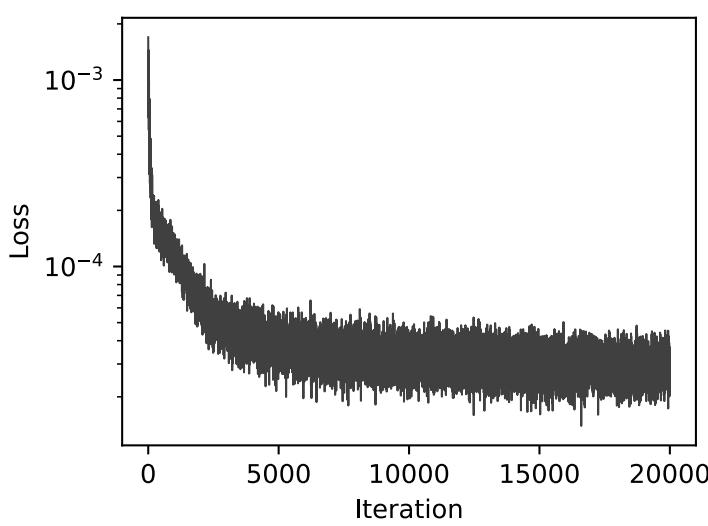


Figure 9: **Optimization and commutation dynamics.** Evolution of the total loss function during the optimization of the exponential map model on experimental grid cell data.

different contexts, such as rooms s_1 and s_2 :

$$\mathbf{p}_{R,S} = \sum_{i=1}^{N_1} \mathbf{p}(\mathbf{r}_i, s_1) + \sum_{j=1}^{N_2} \mathbf{p}(\mathbf{r}_j, s_2) \quad (12)$$

$$= e^{s_1 G_s} \sum_i \mathbf{p}(\mathbf{r}_i) + e^{s_2 G_s} \sum_j \mathbf{p}(\mathbf{r}_j). \quad (13)$$

When we query this composite memory with a current state in context s_1 , the orthogonality of the remapping filters out the interference from context s_2 :

$$\begin{aligned} C_{R,S}(\mathbf{r}, s_1) &= \mathbf{p}(\mathbf{r}, s_1)^T \mathbf{p}_{R,S} \\ &= \mathbf{p}(\mathbf{r}, s_1)^T \sum_i \mathbf{p}(\mathbf{r}_i, s_1) + \mathbf{p}(\mathbf{r}, s_1)^T \sum_j \mathbf{p}(\mathbf{r}_j, s_2) \\ &\approx C_{R1}(\mathbf{r}) + 0. \end{aligned}$$

Here, the cross-context term vanishes because the relative context shift $\Delta s = s_2 - s_1$ decorrelates the vectors (as shown in Fig. 3). This mechanism allows for the superposition of multiple, distinct cognitive maps within a single neural population, enabling context-specific retrieval without crosstalk.

Figure 10 illustrates this capability. We constructed a memory vector summing three target locations across two contexts (“Home” and “Food”). Querying with the appropriate context signal retrieves the correct spatial map, guiding gradient-based navigation to the relevant targets. This demonstrates that the exponential map framework naturally supports flexible, goal-oriented navigation through the algebraic composition of bundling and binding operations.

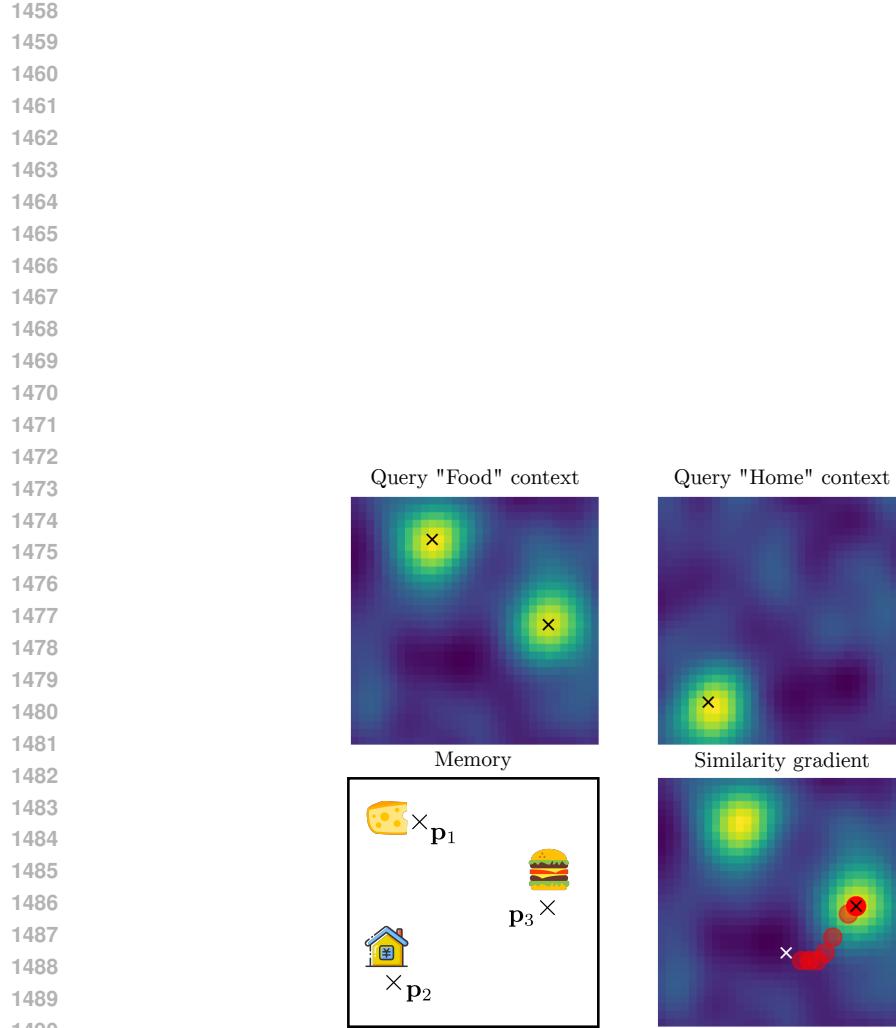


Figure 10: **Multi-map, reward-oriented navigation using exponential maps.** **Bottom left:** A composite memory vector stores the locations of three objects (Cheese, Hamburger, Home) belonging to two distinct contexts (Food and Home). **Top left:** Querying the memory with the “Food” context retrieves a similarity map highlighting the two food items. **Top right:** Querying with the “Home” context retrieves the home location. **Bottom right:** Gradient ascent on the retrieved “Food” similarity map generates a trajectory towards the rewards. The similarity function was constructed to be approximately Gaussian ($N = 300$).

1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511