Place Cells as Multi-Scale Position Embeddings: Random Walk Transition Kernels for Path Planning

Minglu Zhao^{1*}, Dehong Xu^{1*}, Deqian Kong^{1*}, Wen-Hao Zhang², Ying Nian Wu¹

¹UCLA ²UT Southwestern Medical Center

Abstract

The hippocampus supports spatial navigation by encoding cognitive maps through collective place cell activity. We model the place cell population as non-negative spatial embeddings derived from the spectral decomposition of multi-step random walk transition kernels. In this framework, inner product or equivalently Euclidean distance between embeddings encode similarity between locations in terms of their transition probability across multiple scales, forming a cognitive map of adjacency. The combination of non-negativity and inner-product structure naturally induces sparsity, providing a principled explanation for the localized firing fields of place cells without imposing explicit constraints. The temporal parameter that defines the diffusion scale also determines field size, aligning with the hippocampal dorsoventral hierarchy. Our approach constructs global representations efficiently through recursive composition of local transitions, enabling smooth, trap-free navigation and preplay-like trajectory generation. Moreover, theta phase arises intrinsically as the angular relation between embeddings, linking spatial and temporal coding within a single representational geometry.

1 Introduction

Place cells in the hippocampus are central to spatial cognition, firing selectively at specific locations as animals navigate their environment [1, 2]. Rather than modeling individual place fields, we propose viewing place cell populations as non-negative position embeddings derived from spectral decomposition of multi-step transition kernels that collectively encode spatial relationships. This population-level approach captures the hippocampus's role in forming cognitive maps, integrating metric and topological properties for flexible navigation [3].

Hippocampal place cells exhibit dynamic behaviors, adapting firing fields to environmental changes [4], scaling along the dorsoventral axis [5], and engaging in preplay before exploration [6]. These properties suggest a multi-scale, adaptive representation that supports navigation across diverse contexts—from narrow passages to open terrains. The hippocampus's ability to predict novel paths, as in preplay, and maintain robust navigation under environmental deformations further underscores its computational sophistication [7–9].

We formalize these insights by modeling place cell population as vector embeddings derived from a multi-step symmetric random walk. The transition matrix admits spectral decomposition, enabling construction of non-negative embeddings $h(x,\tau)$ such that the inner product between embeddings approximates the normalized transition probability: $\langle h(x,\tau), h(y,\tau) \rangle = q(y|x,\tau)$, where $h(x,\tau)$ is the embedding at location x, and $q(y|x,\tau)$ is the symmetric transition probability over time step τ . Remarkably, combining this inner product structure with non-negativity and orthogonality constraints

Project page: https://sites.google.com/view/place-cells. Code: https://github.com/mingluzhao/Place Cell.

^{*}Equal Contribution.

induces emergent sparsity—the embeddings naturally develop disjoint support sets, explaining why place cells exhibit localized firing fields without explicit regularization. The time parameter $\sqrt{\tau}$ defines a spatial scale hierarchy, mirroring hippocampal dorsoventral scaling [5]. $q(y|x,\tau)$ defines spatial adjacency between x and y at scale or resolution $\sqrt{\tau}$, and the pairwise adjacency relationships $(q(y|x,\tau), \forall x,y)$ are reduced into individual embeddings $(h(x,\tau), \forall x)$ that collectively form a multiscale, Euclideanized, and sparsified cognitive map of the environment. The Euclideanization emerges from inner product expansion of the transition kernel, the sparsification arises naturally from nonnegativity of the place cell responses and the orthogonality between embeddings with zero inner product, and the multi-scale hierarchy enables adaptive navigation across spatial resolutions. Efficient matrix squaring $(P_{2\tau}=P_{\tau}^2)$ computes global transitions from local ones (P_1) without past trajectory memorization, enabling preplay-like shortcut detection [10, 7].

Our framework uses gradient ascent on $q(y|x,\tau) = \langle h(x,\tau), h(y,\tau) \rangle$ with adaptive scale selection, choosing the time scale maximizing the gradient for trap-free, smooth trajectories. This produces robust navigation with properties like boundary avoidance, diffraction-like passage guidance, aligning with hippocampal navigation [11, 9, 3]. Continuous interpolation ensures smooth gradient fields, supporting natural paths. Due to the Euclideanization, path planning amounts to gradient descent on Euclidean distance at the selected scale, enabling what can be called "straight forward" path planning.

Beyond spatial coding, we propose a novel theoretical framework where theta phase naturally emerges from the population embedding structure, unifying spatial representation with temporal dynamics and offering a principled account of theta phase precession.

Appendix A provides an extensive review of related work.

Our contributions include: (1) reconceptualizing place cells as non-negative embeddings from spectral decomposition encoding transition probabilities at multiple scales; (2) demonstrating emergent sparsity explaining localized place fields; (3) modeling multi-scale, Euclideanized, and sparsified cognitive maps via time scale parameter $\sqrt{\tau}$ for transition probabilities; (4) employing efficient matrix squaring to build up multi-scale spatial relationships; (5) introducing adaptive gradient ascent for Euclideanized path planning at selected scale; (6) proposing theta phase formulation from population embeddings; and (7) demonstrating biological plausibility through properties like preplay. Bridging connectionist models [12, 13] and cognitive map theories [2, 9], our framework offers a scalable, biologically inspired model of hippocampal spatial navigation.

2 Method

2.1 Multi-Step Random Walk Transition Kernel

The foundation of our approach is a symmetric random walk on a discrete lattice over the environment (e.g., a 40×40 lattice), with a subset of lattice points belonging to the obstacles. This random walk serves as a mapping policy rather than a goal-reaching policy. Remarkably, this purely random exploration policy leads to optimal path planning without explicit policy optimization. We use τ to denote the time step of this random walk mapping policy, to avoid confusion with the time t of the planned trajectory. τ plays the role of scale, which is adaptively selected during path planning.

For a location x = (i, j) on a 2D lattice, let N(x) be the set of its unobstructed neighbors, i.e., neighbors that do not belong to obstacles (e.g., 4 nearest neighbors), we define the one-step transition probabilities as:

- For each unobstructed neighbor y of x: $p(y|x, \tau=1) = p_{\text{move}}$ (e.g., 1/4 in the case of 4 nearest neighbors)
- For self-transition: $p(x|x, \tau = 1) = 1 |N(x)| \cdot p_{\text{move}}$

where |N(x)| is the number of unobstructed neighbors of x. A critical property of this formulation is that the transition probabilities are symmetric, meaning $p(y|x,\tau=1)=p(x|y,\tau=1)$ for all locations x and y. This symmetry ensures that the random walk process preserves the bidirectional nature of spatial relationships, which is essential for creating a well-defined proximity metric.

The above random walk defines τ -step symmetric transition kernel $p(y|x,\tau)$ over time step τ , which measures the spatial adjacency between x and y at a spatial scale or resolution captured by τ .

2.2 Heat Diffusion, Geodesic Distance and Topological Connectivity

Our discrete random walk model establishes a connection to the continuous heat equation. For small spatial discretization dx and temporal discretization dt with $dx = \sqrt{dt}$, our discrete random walk converges to the heat diffusion equation with reflecting boundary conditions as $dx \to 0$ [14, 15]:

$$\frac{\partial p(y|x,\tau)}{\partial \tau} = \alpha \nabla^2 p(y|x,\tau) \quad \text{on} \quad \Omega \setminus \Omega_{\text{obstacles}}$$

$$\frac{\partial p(y|x,\tau)}{\partial n} = 0 \quad \text{on} \quad \partial \Omega_{\text{obstacles}}$$
(2)

$$\frac{\partial p(y|x,\tau)}{\partial n} = 0 \quad \text{on} \quad \partial \Omega_{\text{obstacles}}$$
 (2)

where Ω is the whole region, $\Omega_{\text{obstacles}}$ is the region of obstacles, and $\partial\Omega_{\text{obstacles}}$ is its boundary. α is the diffusion coefficient (different one-step transition $p(y|x, \tau=1)$ in discrete case leads to different α , e.g., $\alpha = 1/4$ for 4-nearest neighbor random walk).

A fundamental result from heat diffusion theory relates the short-time behavior of the heat kernel to geodesic distance. For small values of τ , the heat kernel has the asymptotic form [16, 17]:

$$p(y|x,\tau) \approx \frac{1}{4\pi\alpha\tau} \exp\left(-\frac{d_g^2(x,y)}{4\alpha\tau}\right)$$
 (3)

where $d_q(x,y)$ is the geodesic distance. For open domains, $d_q(x,y)$ becomes Euclidean distance, and $p(y|x,\tau) \sim \mathcal{N}(x,2\alpha\tau)$, a Gaussian distribution with variance $2\alpha\tau$ or standard deviation $\sqrt{2\alpha\tau}$, demonstrating that $\sqrt{\tau}$ correspond to the spatial scale. We take square root of τ to emphasize this scaling relationship between time and space.

This result demonstrates that $d_t(x,y) = -\tau \log q(y|x,\tau)$ approximates the squared geodesic distance for small values of τ , providing a fundamental connection between random walks and geodesic distances in complex environments.

As τ increases, the transition probability incorporates additional information about path multiplicity and global connectivity, creating a multi-scale representation of spatial relationships. In particular, for large τ , the transition probability is dominated by eigenvalues that are close to 1, whose eigenvectors depend on topological connectedness rather than local geometry.

Appendix D provides detailed background on heat equation. Appendix G explains topological properties for large τ .

2.3 Place Cells as Non-negative Spectral Embeddings

We formalize place cell population as vector embedding $h(x,\tau) \in \mathbb{R}^n$, where n is the number of place cells (e.g., n = 500). The inner product between embeddings approximates the normalized transition probability kernel:

$$\langle h(x,\tau), h(y,\tau) \rangle \approx q(y|x,\tau)$$
 (4)

where $q(y|x,\tau) = p(y|x,\tau)/\sqrt{p(x|x,\tau)\cdot p(y|y,\tau)}$ is the normalized transition probability, so that $||h(x,\tau)|| = 1$ due to normalization. For each cell $i, h_i(x,\tau) \ge 0$ for biological plausibility, and $h_i(x,\tau)$ is the response map or profile of cell i at spatial scale $\sqrt{\tau}$.

The above formulation can be derived through spectral decomposition of the transition matrix. Since the one-step transition matrix P_1 is symmetric by construction, its powers $P_{\tau}=P_1^{\tau}$ admit an eigendecomposition [18]: $P_{\tau}=Q\Lambda^{\tau}Q^T$, where Q is orthogonal $(Q^TQ=I)$ and $\Lambda=\mathrm{diag}(\lambda_1,\ldots,\lambda_n)$ contains eigenvalues $0\leq\lambda_i\leq 1$ (assuming the random walk is irreducible and aperiodic).

From this spectral decomposition, we can construct a spectral embedding by defining: $H_i(x,\tau)$ $\lambda_i^{\tau/2}Q_i(x)$ where Q_i is the *i*-th column of Q. This yields an exact representation of the transition probability through inner products:

$$p(y|x,\tau) = \sum_{i} H_i(x,\tau)H_i(y,\tau) = \langle H(x,\tau), H(y,\tau) \rangle$$
 (5)

where $H_i(x,\tau)$ is the *i*-th element of the vector $H(x,\tau)$. For the normalized transition probability $q(y|x,\tau)$, we define normalized embeddings:

$$h_{\text{spec}}(x,\tau) = \frac{H(x,\tau)}{\|H(x,\tau)\|} = \frac{H(x,\tau)}{\sqrt{p(x|x,\tau)}}$$
 (6)

which satisfy:

$$\langle h_{\text{spec}}(x,\tau), h_{\text{spec}}(y,\tau) \rangle = \frac{p(y|x,\tau)}{\sqrt{p(x|x,\tau) \cdot p(y|y,\tau)}} = q(y|x,\tau) \tag{7}$$

 $p(x|x,\tau)$ is constant in the open field and is smooth in general, so q is essentially a scaled version of p.

However, these spectral embeddings may contain negative components, conflicting with the biological constraint that neural firing rates must be non-negative. Horn's theorem [19, 20], which is built upon the above spectral decomposition, guarantees the existence of a non-negative $h(x,\tau)$, such that $\langle h(x,\tau), h(y,\tau) \rangle = q(y|x,\tau)$ for non-negative matrix factorization.

This formulation represents a fundamental shift from modeling individual place cells to modeling the place cell population as distributed position embedding. The embedding vector $h(x,\tau)$ represents the firing rates of all place cells at location x for spatial scale $\sqrt{\tau}$, capturing the idea that it is the pattern across the population—not the activity of any single cell—that encodes location.

The time parameter $\sqrt{\tau}$ serves as a fundamental unit of spatial resolution or scale, with larger values producing broader, more diffuse representations, and smaller values producing more localized representations. This naturally mirrors the variation in place field sizes observed along the dorsoventral axis of the hippocampus [5, 21].

See Appendix C for more details on spectral analysis. Appendix E provides analytical results for open field, where $q(y|x,\tau)$ is Gaussian, and elements of $h(x,\tau)$ exhibit Gaussian profiles over x.

2.4 Matrix Squaring, Learning, and Continuation

Let P_1 be the one-step transition matrix for the random walk on a discrete lattice. P_1 depends on obstacles in the environment and amounts to local one-step exploration, as detailed in subsection 2.1.

We calculate P_{τ} for a discrete set of τ , $\mathcal{T} = \{\tau = 2^k, k = 1, ..., K\}$, via $P_{2\tau} = P_{\tau}^2$. The matrix squaring is very efficient for calculating P_{τ} for $\tau \in \mathcal{T}$, and these P_{τ} correspond to explorations of different spatial scales $\sqrt{\tau}$, where the adjacent spatial scales in \mathcal{T} have a ratio $\sqrt{2}$, mirroring grid cell module scaling (\sim 1.4–1.7) [22]. In addition to matrix squaring, one can design any discrete sequence $\mathcal{T} = \{t_k, k = 1, ..., K\}$, and calculate P_{τ} for $\tau \in \mathcal{T}$ using matrix multiplication.

We learn a separate population of place cells $h(x,\tau)$ for each $\tau\in\mathcal{T}$, by minimizing the least squares error:

$$\mathcal{L} = \sum_{x,y} \left[q(y|x,\tau) - \langle h(x,\tau), h(y,\tau) \rangle \right]^2 \tag{8}$$

We learn $h(x,\tau)$ over the discrete lattice, where $\sum_{x,y}$ in (8) is over all pairs of points on the lattice. We optimize this objective using the AdamW optimizer [23], where after each iteration, for each x, we set the negative elements of $h(x,\tau)$ to 0, and then normalize $h(x,\tau)$ so that $||h(x,\tau)|| = 1$.

The learning reduces pairwise adjacency relationships $(q(y|x,\tau), \forall x, y)$ into individual embeddings $(h(x,\tau), \forall x)$, which collectively form a map of the environment.

After learning, we can use bi-linear interpolation to make $h(x,\tau)$ a continuous map over x. As a result, $q(y|x,\tau) = \langle h(x,\tau), h(y,\tau) \rangle$ also becomes continuous, approximating the transition kernel of the continuous heat equation. The continuous $q(y|x,\tau)$ can then elegantly guide path planning in continuous space instead of discrete lattice.

2.5 Euclideanized Cognitive Map

Geometrically, $\langle h(x,\tau), h(y,\tau) \rangle$ is the cosine of the angle between the normalized vectors $h(x,\tau)$ and $h(y,\tau)$. Moreover, we have

$$\frac{1}{2}||h(x,\tau) - h(y,\tau)||^2 = 1 - \langle h(x,\tau), h(y,\tau) \rangle = 1 - q(y|x,\tau)$$
(9)

That is, the angle or the Euclidean distance between $h(x,\tau)$ and $h(y,\tau)$ encodes proximity or adjacency between x and y. Thus, our method geometrizes transition probability.

 $(h(x,\tau),\forall x)$ is a 2D manifold in the high-dimensional embedding space. This 2D manifold is an embedding of the 2D physical space. Because of the flexibility endowed by the high-dimensional embedding space, the path between x and y on the manifold is essentially "Euclideanized" even though the physical path is far from being Euclidean.

2.6 Emergent Sparsity from Non-negativity and Orthogonality

The conjunction of non-negativity and orthogonality induces emergent sparsity in the representational code. Let $h(x,\tau) \in \mathbb{R}^n_+$ denote the representation at position x and scale $\sqrt{\tau}$, such that $\langle h(x,\tau), h(y,\tau) \rangle = q(y|x,\tau)$, where $q(y|x,\tau)$ measures spatial adjacency at scale $\sqrt{\tau}$. If two positions x and y are non-adjacent $(q(y|x,\tau)=0)$, the corresponding vectors must be orthogonal. Because all components of $h(x,\tau)$ are non-negative, orthogonality implies disjoint support—no component can be active in both representations. Consequently, the joint requirement of non-negativity and orthogonality forces most components to be zero, yielding sparse, localized activity patterns. As a result, each place cell fires at a localized place, and the population of place cells collectively tile the environment. The symbolic representation thus emerges automatically from population-based representation.

The scale parameter τ modulates this sparsity: for small τ , most position pairs are orthogonal, producing highly localized "place fields"; for large τ , overlaps increase and place fields broaden. Thus, the locality of place-cell activation follows directly from the geometry of non-negative, inner-product spectral embeddings. Appendix F provides technical details.

2.7 Straight Forward Path Planning with Adaptive Scale Selection

We use adaptive gradient following for goal-directed navigation. When navigating from a current position x to a target location y, we select the next position x_{next} from the neighborhood $\partial(x)$ where $\partial(x) = \{z : z = x + \Delta r(\cos\theta, \sin\theta)\}$, Δr is the step size, and θ is discretized into equally spaced direction in $[0, 2\pi)$. We use $\partial(x)$ for continuous x for path planning to differentiate from N(x) in the discrete lattice for random walk.

The path planning algorithm is as follows:

• Compute the gradient of the normalized transition probability for each neighbor $z \in \partial(x)$ and each scale τ :

$$\Delta(z,\tau) = q(y|z,\tau) - q(y|x,\tau) = \langle h(y,\tau), h(z,\tau) \rangle - \langle h(y,\tau), h(x,\tau) \rangle \tag{10}$$

• Select the scale τ^* that provides the strongest directional signal:

$$\tau^* = \operatorname*{max} \max_{\tau \in \mathcal{T}} \max_{z \in \partial(x)} \Delta(z, \tau) \tag{11}$$

• Choose the next position that maximizes the gradient at the selected scale:

$$x_{\text{next}} = \underset{z \in \partial(x)}{\arg \max} \Delta(z, \tau^*)$$
 (12)

Note that in the above algorithm, x and z are continuous, because $h(x, \tau)$ is made continuous in x with bi-linear interpolation after learning on discrete lattice.

 $\Delta(z,\tau)$ measures the reduction in the angle or squared Euclidean distance. The path planning seeks maximal reduction in angle or Euclidean distance. Therefore we call it the straight forward path planning in the Euclideanized embedding space, where the vector $h(x,\tau)$ rotates straightly to $h(y,\tau)$ on the path.

This approach selects the scale τ^* that provides the clearest guidance for the current navigation step. Intuitively, larger scales provide better guidance for distant goals, while smaller scales offer more precise navigation for nearby goals. The adaptive scale selection mechanism automatically finds this optimal scale at each step, similar to choosing the most appropriate "ruler" for measuring at the current distance.

Now consider the idealized continuous limit where $\Delta r \to 0$, θ is continuous, and τ is continuous. The path planning algorithm follows the gradient $\nabla_x q(y|x,\tau) = \nabla_x q(x|y,\tau)$ where

 $q(y|x,\tau)=q(x|y,\tau)$ due to symmetry. The planed trajectory is the gradient ascent flow: $dx(t)/dt=\nabla_x q(x(t)|y,\tau(t))$, where t is the time on the planed trajectory, $\tau(t)=\arg\max_{\tau}\|\nabla_x q(x(t)|y,\tau)\|$ is the optimal scale at time t with the maximal gradient.

The gradient-based navigation framework has the following key properties that ensure computational efficiency and biological plausibility:

- 1. Adaptive scale selection dynamically adjusts $\sqrt{\tau^*} \propto d(x,y)$ for precise navigation in the open environment, mirroring hippocampal spatial tuning [11].
- 2. The unimodal smooth gradient field $\nabla_x q(x|y,\tau)$, with a unique maximum at the goal, ensures trap-free paths, aligning with hippocampal navigation [11, 9].
- 3. Planned paths match the shortest path for small τ , but prioritize topological connectivity for large τ , reflecting cognitive map robustness [9, 3].
- 4. Near obstacles, $\nabla_x q(x|y,\tau)$ flows parallel to boundaries, preventing collisions, akin to hippocampal obstacle avoidance [11].
- 5. Diffraction-like patterns guide trajectories through passages, resembling hippocampal maze navigation [11, 9].
- 6. Topological invariance maintains navigation under environmental deformations, mirroring hippocampal place cell encoding [3, 9].
- 7. Matrix squaring from local transitions predicts shortcuts, embodying hippocampal preplay [7, 8].

These properties, detailed in Appendix G, underpin the framework's alignment with neural mechanisms.

2.8 Theta-Phase Procession Based on Angle-Phase Duality

The hippocampus not only encodes spatial relationships but also organizes temporal dynamics through theta phase precession, where place cells fire at specific phases of the theta rhythm as an animal traverses their fields [24, 25]. We extend our population embedding framework to model this phenomenon.

As explained above, the non-negativity and inner product structure of the position embeddings induces localized sparse patterns, i.e., each place cell only fires at a localized field around a place, and together the place fields of all the place cells tile the environment. Let $\mu_i = \arg\max_x h_i(x,\tau)$, i.e., the center of the place field for cell i. The theta phase of cell i can be defined in terms of the angle between $h(x,\tau)$ and $h(\mu_i,\tau)$. As x approaches μ_i , the angle changes from $\pi/2$ to 0, and as x moves away from μ_i , the angle changes from 0 to $\pi/2$. Appendix H provides details.

2.9 Integrating Grid Cells

For scientific reductionism, we focus on place cells without incorporating grid cells. A natural extension would explore the relationship $h(x,\tau)=W(\tau)g(x)$, where g(x) represents the vector of grid cell activations and $W(\tau)$ is a learned transformation matrix, potentially unifying our framework with the complementary roles of place and grid cells in spatial representation [9]. Appendix I provides the formulation.

3 Experiments

We design experiments to evaluate both the biological plausibility of our position embedding framework and its functional capabilities in path planning. All experiments are conducted in a simulated environment. We first examine the positional representation and population density profiles in an open field. Then we create more complex environments by adding obstacles. For implementation details, see Appendix J; For additional experiment visualizations, see Appendix K.

3.1 Place Cell Representations in Open Field

We begin our numerical experiments in a simple open field environment to demonstrate fundamental properties. The environment consists of a 40×40 lattice grid. For the transition kernel, we

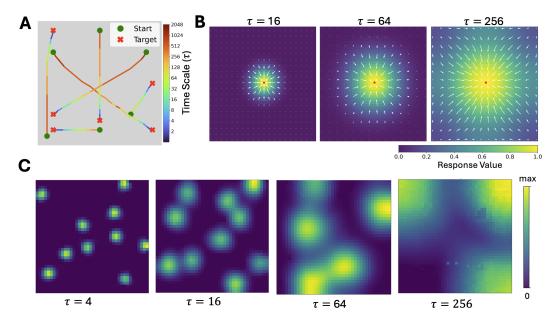


Figure 1: Place Cell Representations and Navigation in Open Field Environment. (A) Goal-directed path planning trajectories with adaptive scale selection (selected scale is color coded, red for big scale and blue for small scale). (B) Normalized transition probability kernels $q(y|x,\tau)$ at multiple scales with gradient vector fields. y is fixed at the center of the environment. (C) Learned activation patterns of $h(x,\tau)$ at different scales across randomly chosen cells, exhibiting Gaussian-like firing fields centered at specific locations within the open environment.

employ a 3×3 neighborhood, meaning each position has at most 8 neighboring points, with transition probability $p_{\text{move}}=1/9$ distributed uniformly among neighbors as well as self. Within this environment, we learn $h(x,\tau)$ embeddings across all lattice points. For hyperparameters, the total number of place cells is set to 500. The temporal dynamics are captured across multiple scales using the time parameter $\tau=2^k$, where k=1,2,...,11, allowing us to analyze system behavior across spatial scales.

3.1.1 Multi-Scale Transitions

To evaluate the normalized transition probability kernel $q(y|x,\tau)$, we visualize the learned Gaussian-like patterns across multiple temporal scales in Figure 1 (B). Each panel depicts the transition probability distribution for time scales $\tau=16,64,256$ with position y fixed at the environment center.

The gradient fields reveal a critical temporal dependency: as τ increases, gradient magnitudes intensify for locations distant from the center, while maintaining directional integrity toward the target. This multi-scale property is particularly significant for path planning applications, as it enables a multi-resolution navigational framework.

3.1.2 Spatial Activation Patterns

Our model is trained by minimizing Equation 8. The optimization results demonstrate remarkable fidelity in approximating normalized transition probabilities through inner products of population embedding vectors with correlation coefficients above 0.9 for all scales. Figure 1 (C) illustrates the learned activation patterns of $h(x,\tau)$ at scales $\tau=4,16,64,256$ for randomly chosen cells. The spatial receptive fields of individual units emerge naturally from our training process, exhibiting the characteristic Gaussian-like firing fields centered at specific locations within the open environment. This spatial tuning closely resembles the well-documented properties of hippocampal place cells observed in rodent navigation studies [1,2].

3.1.3 Path Planning and Adaptive Scale Selection

To evaluate the navigational capabilities of our model, we implemented the gradient-based path planning. We randomly selected start and target locations within the open field environment. At each step, the agent evaluates potential next positions by sampling orientation θ from 36 equally spaced directions within $[0,2\pi)$ and a fixed radius Δr of one grid unit. Even though our model is learned in a 40×40 discrete lattice, in path planning, we can reach continuous positions x, where the representation of its location $h(x,\tau)$ is calculated by linear interpolation of the 4 closest neighbors. Movement direction is determined by maximizing the gradient $q(y|z,\tau)-q(y|x,\tau)$ where y is the target location, x is the current position, and z represents each candidate next position. And for each step, we choose the z with the largest gradient. Our results in Figure 1 (A) demonstrate that the learned model consistently generates near-optimal trajectories.

A key feature of our position embedding model is the adaptive scale selection mechanism, where τ^* is selected with maximal gradient. The mechanism is illustrated in Figure 1 (A), where different colors represent the selected time scale τ^* at each navigational step. This reveals a systematic progression from coarse to fine scales as the agent approaches its goal. This pattern emerges naturally from the gradient fields associated with different scales.

When navigating to distant goals, larger scales ($\tau \geq 128$) are selected, which provides clearer guidance for long-range planning. As the agent reaches medium distances from the goal, the selected scales transition to intermediate values ($\tau = 16$ to $\tau = 64$), balancing directional guidance with increasing spatial resolution. In the final approach phase, the system converges on the smallest available scales ($\tau = 2$ to $\tau = 8$), which offered the most precise local guidance.

This adaptive scale selection mirrors the progressive engagement of different regions along the dorsoventral axis of the hippocampus during navigation, as observed in rodent studies [5, 21].

3.2 Place Cells in Complex Environments

To investigate the robustness and adaptability of our method, we extend our analysis to complex environmental geometries that more closely resemble naturalistic navigation scenarios. These environments incorporate obstacles and boundaries that fundamentally alter the adjacency relationships between locations, requiring the model to learn representations that respect environmental constraints rather than simple Euclidean distances.

3.2.1 Environment and Experiment Setup

To evaluate performance across different environments, we implement three distinct mazes. Detailed visualization is shown in Figure 2.

- U-shaped Maze: A corridor structure with a single 180-degree turn, creating a simple non-convex navigation challenge.
- S-shaped Maze: A serpentine corridor with multiple turns, requiring longer detour paths around obstacles.
- Four-room Maze: A compartmentalized environment with four chambers connected by narrow doorways, representing a hierarchically structured space.

To accommodate environmental complexity, we modify our random walk dynamics to incorporate spatial constraints imposed by obstacles and boundaries. For any location x, transitions to neighboring locations y that are prohibited by the obstacles setting $p(y|x,\tau=1)=0$ for these forbidden transitions. Consequently, the self-transition probability is adjusted to maintain normalization $p(x|x,\tau=1)=1-|N(x)|\cdot p_{\text{move}}$. The key difference between open field and obstacle-containing environments lies in the value of N(x). In open fields, N(x) consistently includes all 8 neighboring locations (except at boundaries). With obstacles, N(x) represents only the subset of accessible neighboring locations from position x, which could be smaller than 8. We maintain $p_{\text{move}}=1/9$ across all environments for consistency. Despite these modifications to the underlying transition dynamics, our learning approach remains consistent. We compute the normalized transition probabilities $q(y|x,\tau)$ based on these constraint-respecting dynamics and train our model to learn the embedding function $h(x,\tau)$ using the same objective function as in the open field.

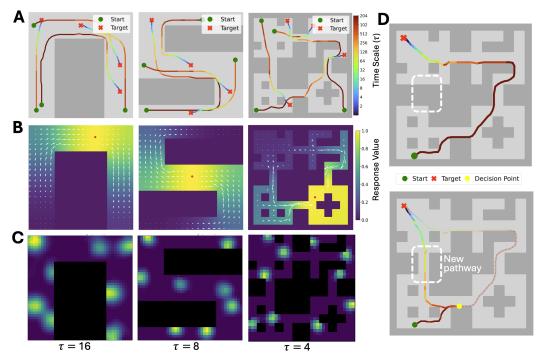


Figure 2: Place Cells in Complex Maze Environments. (A) Path planning through obstacle-containing environments. (B) Topologically-informed transition kernels $q(y|x,\tau)$ with gradient fields. Target y is marked as a red point. (C) Randomly sampled place cell profiles at multiple spatial scales. (D) Remapping with environmental modification.

3.2.2 Navigation Efficiency

We evaluated our model's navigation capabilities across multiple environments, conducting 50 trials per environment with randomly sampled start and goal positions from non-obstacle areas. Our model achieved 100% success rate across all tested environments.

We compare against the Bug algorithm [26] as a baseline since it prioritizes finding feasible paths rather than strictly optimizing for the shortest path. The Bug algorithm moves directly toward the goal until encountering an obstacle, then follows the obstacle's boundary until it can resume its direct path. We implement the oracle-enhanced version [27] that knows which boundary-following direction yields the shortest path for each obstacle, establishing an upper bound on performance.

To quantify efficiency, we use Success weighted by inverse Path Length (SPL) [28], where values near 1.0 indicate near-optimal paths and values exceeding 1.0 suggest our method found shortcuts the Bug algorithm missed. Details of SPL calculation appear in Appendix J. In simpler environments (U-shape), our method performs comparably to the oracle-enhanced Bug algorithm. However, in complex environments (S-shape and Four-room), our approach signif-

Table 1: Path planning results.

Environment	Success	SPL (†)
Open field	100%	0.991 ± 0.04
Ū-shape	100%	0.919 ± 0.25
S-shape	100%	1.392 ± 0.98
Four-room	100%	1.519 ± 1.17

icantly outperforms the baseline, discovering more efficient paths that the Bug algorithm fails to identify even with oracle assistance.

Table 2: Path planning in Four-Room environment.

Metrics	Bug (w/ Oracle)	Bug (w/o Oracle)	A* Search	Random Walk	Our Method
Success (%) SPL (↑)	$100\% \ 0.66 \pm 0.51$	$14\% \\ 0.15 \pm 0.00$	100% 1.08 ± 0.50	$2\% \\ 0.01 \pm 0.00$	100% 1.00

We also compared our method with several different baselines in the Four-Room environment, with each trial limited to 50,000 steps. The results in Table 2 show that baseline approaches like the Bug algorithm (without oracle guidance) and Random Walk fail to solve most of the challenging scenarios.

3.2.3 Transition and Learned Profile in Complex Environments

To evaluate the adaptive capabilities of our approach, we examine the transition kernel $q(y|x,\tau)$ and the resulting place cell activation profiles $h(x,\tau)$ in environments containing complex obstacles. Figure 2 (B) presents the normalized transition probability distributions $q(y|x,\tau)$ for strategically selected locations, with a large scale to reveal long-range spatial relationships. The diffraction-like patterns, where probability flow encounters obstacles, reveal how spatial information propagates through available pathways while respecting environmental constraints.

The gradient fields illustrate navigation potential, showing how an agent could efficiently pass obstacles. Importantly, these transition probabilities capture the true topological structure of complex environments rather than simple Euclidean distances, a critical property for realistic navigation where direct paths are often blocked. The emergent probability gradients provide a natural mechanism for guiding optimal path planning that automatically adapts to environmental geometry.

We visualized place cell profiles at multiple spatial scales across complex maze environments in Figure 2 (C). Each panel displays a small subset of non-overlapping place cells at different scales τ , emphasizing the emergent sparsity and localized nature of individual receptive fields. Despite increased environmental complexity, the learned place cell population maintains comprehensive spatial coverage throughout accessible regions. At smaller scales, place fields are tightly localized; at larger scales, fields broaden while maintaining their localized tiling structure. This multi-scale coverage property demonstrates the model's robust ability to develop effective spatial representations regardless of environmental geometry, a critical feature for reliable navigation in diverse settings.

3.2.4 Remapping Properties

We evaluated our model's ability to adapt to environmental modifications, analogous to the remapping phenomenon observed in rodent place cells. Previous neurobiological studies have shown that hippocampal place fields reorganize in geometry-dependent ways when familiar environments are altered through elongation or restructuring.

We conducted a two-phase experiment with the environment shown in the upper panel of Figure 2 (D). After training our model to convergence (2000 iterations) on this initial configuration, we modified the arena by removing several obstacles to introduce a novel shortcut path, as shown in the lower panel of Figure 2 (D). The intervention created a minimal physical alteration but significantly changed the environment's topological structure.

We then fine-tuned our pre-trained model on this modified environment with just 50 iterations at a reduced learning rate (5e-4), allowing the position embeddings to adapt efficiently to the new spatial configuration. When testing path planning to identical target locations, the fine-tuned model successfully identified and utilized the newly available shortcuts at critical decision points where pathways diverged.

4 Conclusion

This paper reconceptualizes hippocampal place cells as population embeddings approximating multistep random walk transition kernels, where $\langle h(x,\tau),h(y,\tau)\rangle=q(y|x,\tau)$, offering a biologically plausible model of spatial navigation [29, 30]. The time parameter $\sqrt{\tau}$ defines a multi-scale representation, mirroring dorsoventral place field scaling [5], adaptability [4], and remapping [31]. Gradient ascent on $q(y|x,\tau)$ with adaptive scale selection produces trap-free trajectories, guided by boundary avoidance, diffraction-like passage navigation, aligning with hippocampal cognitive maps [11, 9, 3]. Efficient matrix squaring $(P_{2\tau}=P_{\tau}^2)$ computes global transitions from local ones (P_1) without past trajectory memorization, enabling preplay-like shortcut detection [10, 7, 8]. Bridging connectionist models [12, 13] and cognitive map theories [2, 9], our framework captures hippocampal navigation's dynamic properties, providing a scalable, computationally efficient model for complex environments.

Acknowledgments

Y. W. is partially supported by NSF DMS-2415226, DARPA W912CG25CA007 and research gift funds from Amazon and Qualcomm.

References

- [1] John O'Keefe and Jonathan Dostrovsky. The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain research*, 1971.
- [2] John O'Keefe and Lynn Nadel. *The Hippocampus as a Cognitive Map*. Oxford University Press, Oxford, UK, 1978.
- [3] Edward C. Tolman. Cognitive maps in rats and men. *Psychological Review*, 55(4):189–208, 1948.
- [4] R. U. Muller and J. L. Kubie. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, 7(7):1951–1968, 1987.
- [5] Kirsten B. Kjelstrup, Trygve Solstad, Vegard H. Brun, Torkel Hafting, Stefan Leutgeb, Menno P. Witter, Edvard I. Moser, and May-Britt Moser. Finite scale of spatial representation in the hippocampus. *Science*, 321(5885):140–143, 2008.
- [6] George Dragoi and György Buzsáki. Temporal encoding of place sequences by hippocampal cell assemblies. *Neuron*, 50(1):145–157, 2006.
- [7] George Dragoi and Susumu Tonegawa. Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, 469(7330):397–401, 2011.
- [8] H. Freyja Ólafssdóttir, Caswell Barry, Aman B. Saleem, Demis Hassabis, and Hugo J. Spiers. Hippocampal place cells construct reward related sequences through unexplored space. *eLife*, 4:e06063, 2015.
- [9] E. I. Moser, E. Kropff, and M.-B. Moser. Place cells, grid cells, and the brain's spatial representation system. *Annual Review of Neuroscience*, 31:69–89, 2008.
- [10] James R. Norris. *Markov Chains*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1998.
- [11] Bruce L McNaughton, Francesco P Battaglia, Ole Jensen, Edvard I Moser, and May-Britt Moser. Path integration and the neural basis of the cognitive map'. *Nature Reviews Neuroscience*, 7(8):663–678, 2006.
- [12] Andrea Banino, Caswell Barry, Benigno Uria, Charles Blundell, Timothy Lillicrap, Piotr Mirowski, Alexander Pritzel, Martin J Chadwick, Thomas Degris, Joseph Modayil, et al. Vector-based navigation using grid-like representations in artificial agents. *Nature*, 557(7705):429–433, 2018.
- [13] James CR Whittington, Timothy H Muller, Shirley Mark, Guifen Chen, Caswell Barry, Neil Burgess, and Timothy EJ Behrens. The tolman-eichenbaum machine: Unifying space and relational memory through generalization in the hippocampal formation. *Cell*, 183(5):1249–1263, 2020.
- [14] Bernt Øksendal. Stochastic Differential Equations: An Introduction with Applications. Springer, 6 edition, 2003.
- [15] Alexander Grigoryan. Heat Kernel and Analysis on Manifolds. AMS/IP Studies in Advanced Mathematics. American Mathematical Society/International Press, 2009.
- [16] S. R. S. Varadhan. On the behavior of the fundamental solution of the heat equation with variable coefficients. *Communications on Pure and Applied Mathematics*, 20(2):431–455, 1967.
- [17] Elton P Hsu. On the principle of not feeling the boundary for diffusion processes. *Journal of the London Mathematical Society*, 51(2):373–382, 1995.

- [18] Carl D. Meyer. Matrix Analysis and Applied Linear Algebra. SIAM, 2000.
- [19] Roger A. Horn. Doubly stochastic matrices and the diagonal of a rotation matrix. *American Journal of Mathematics*, 76(3):620–630, 1954.
- [20] Abraham Berman and Robert J Plemmons. *Nonnegative matrices in the mathematical sciences*. SIAM, 1994.
- [21] Bryan A. Strange, Menno P. Witter, Ed S. Lein, and Edvard I. Moser. Functional organization of the hippocampal longitudinal axis. *Nature Reviews Neuroscience*, 15(10):655–669, 2014.
- [22] H. Stensola, T. Stensola, T. Solstad, K. Frøland, M.-B. Moser, and E. I. Moser. The entorhinal grid map is discretized. *Nature*, 492(7427):72–78, 2012.
- [23] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint* arXiv:1711.05101, 2017.
- [24] John O'Keefe and Michael L. Recce. Phase relationship between hippocampal place units and the eeg theta rhythm. *Hippocampus*, 3(3):317–330, 1993.
- [25] William E. Skaggs, Bruce L. McNaughton, Matthew A. Wilson, and Carol A. Barnes. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, 6(2):149–172, 1996.
- [26] Vladimir J Lumelsky and Alexander A Stepanov. Path-planning strategies for a point mobile automaton moving amidst unknown obstacles of arbitrary shape. *Algorithmica*, 2(1):403–430, 1987.
- [27] Erik Wijmans, Manolis Savva, Irfan Essa, Stefan Lee, Ari S Morcos, and Dhruv Batra. Emergence of maps in the memories of blind navigation agents. *AI Matters*, 9(2):8–14, 2023.
- [28] Peter Anderson, Angel Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, et al. On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*, 2018.
- [29] Matthew A Wilson and Bruce L McNaughton. Dynamics of the hippocampal ensemble code for space. *Science*, 261(5124):1055–1058, 1993.
- [30] Thomas B Christophel, P Christiaan Klink, Bernhard Spitzer, Pieter R Roelfsema, and John-Dylan Haynes. The distributed nature of working memory. *Trends in cognitive sciences*, 21(2):111–124, 2017.
- [31] Stefan Leutgeb, Jill K Leutgeb, Carol A Barnes, Edvard I Moser, Bruce L McNaughton, and May-Britt Moser. Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science*, 309(5734):619–623, 2005.
- [32] Eliott Robert Joseph Levy, Simón Carrillo-Segura, Eun Hye Park, William Thomas Redman, José Rafael Hurtado, Sue Yeon Chung, and André Antonio Fenton. A manifold neural population code for space in hippocampal coactivity dynamics independent of place fields. *Cell reports*, 42(10), 2023.
- [33] Richard J Gardner, Eivind Hermansen, Marius Pachitariu, Yoram Burak, Nils A Baas, Benjamin A Dunn, May-Britt Moser, and Edvard I Moser. Toroidal topology of population activity in grid cells. *Nature*, 602(7897):449–454, 2022.
- [34] John P Cunningham and Byron M Yu. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, 17(11):1500–1509, 2014.
- [35] Arno Onken, Jian K Liu, PP Chamanthi R Karunasekara, Ioannis Delis, Tim Gollisch, and Stefano Panzeri. Using matrix and tensor factorizations for the single-trial analysis of population spike trains. *PLoS computational biology*, 12(11):e1005189, 2016.
- [36] Dileep George, Rajeev V Rikhye, Nishad Gothoskar, J Swaroop Guntupalli, Antoine Dedieu, and Miguel Lázaro-Gredilla. Clone-structured graph representations enable flexible learning and vicarious evaluation of cognitive maps. *Nature Communications*, 12(1):2392, 2021.

- [37] Rajkumar Vasudeva Raju, J Swaroop Guntupalli, Guangyao Zhou, Carter Wendelken, Miguel Lázaro-Gredilla, and Dileep George. Space is a latent sequence: A theory of the hippocampus. *Science Advances*, 10(31):eadm8470, 2024.
- [38] Kimberly L Stachenfeld, Matthew M Botvinick, and Samuel J Gershman. The hippocampus as a predictive map. *Nature Neuroscience*, 20(11):1643–1653, 2017.
- [39] Bruno A Olshausen and David J Field. Wavelet-like receptive fields emerge from a network that learns sparse codes for natural images. *Nature*, 381:607–609, 1996.
- [40] Tony Lindeberg. A computational theory of visual receptive fields. *Biological Cybernetics*, 107(6):589–635, 2013.
- [41] Uğur M Erdem, Michael J Milford, and Michael E Hasselmo. A hierarchical model of goal directed navigation selects trajectories in a visual environment. *Neurobiology of learning and memory*, 117:109–121, 2015.
- [42] Daniel Bush, Caswell Barry, Daniel Manson, and Neil Burgess. Using grid cells for navigation. *Neuron*, 87(3):507–520, 2015.
- [43] Ida Momennejad, Evan M Russek, Jin H Cheong, Matthew M Botvinick, Nathaniel D Daw, and Samuel J Gershman. The successor representation in human reinforcement learning. *Nature Human Behaviour*, 1(9):680–692, 2017.
- [44] Yu Fan Chen, Shih-Yuan Liu, Miao Liu, Justin Miller, and Jonathan P How. Motion planning with diffusion maps. In 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 1423–1430. IEEE, 2016.
- [45] Siyuan Huang, Zan Wang, Puhao Li, Baoxiong Jia, Tengyu Liu, Yixin Zhu, Wei Liang, and Song-Chun Zhu. Diffusion-based generation, optimization, and planning in 3d scenes. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 16750–16761, 2023.
- [46] Zilu Liang, Simeng Wu, Jie Wu, Wen-Xu Wang, Shaozheng Qin, and Chao Liu. Distance and grid-like codes support the navigation of abstract social space in the human brain. *Elife*, 12:RP89025, 2024.
- [47] Albert J Wakhloo, Will Slatton, and SueYeon Chung. Neural population geometry and optimal coding of tasks with shared latent structure. *arXiv preprint arXiv:2402.16770*, 2024.
- [48] Peter Dayan. Improving generalization for temporal difference learning: The successor representation. *Neural computation*, 5(4):613–624, 1993.
- [49] Samuel J Gershman. The successor representation: its computational logic and neural substrates. *Journal of Neuroscience*, 38(33):7193–7200, 2018.
- [50] Surya Ganguli and Haim Sompolinsky. Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis. *Annual Review of Neuroscience*, 35:485–508, 2012.
- [51] Chris Eliasmith. *How to build a brain: A neural architecture for biological cognition*. Oxford University Press, 2012.
- [52] Rishidev Chaudhuri, Albert Geróin, and Bruno Averbeck. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature Neuroscience*, 22(9):1512–1520, 2019.
- [53] Bernhard Schölkopf and Alexander J Smola. *Learning with kernels: support vector machines, regularization, optimization, and beyond.* MIT press, 2002.
- [54] Shun-ichi Amari. Information geometry and its applications. Springer, 2016.
- [55] Vincent Dutordoir, Alan Saul, Zoubin Ghahramani, and Fergus Simpson. Neural diffusion processes. In *International Conference on Machine Learning*, pages 8990–9012. PMLR, 2023.

- [56] Ronald R Coifman and Stéphane Lafon. Diffusion maps. Applied and computational harmonic analysis, 21(1):5–30, 2006.
- [57] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.
- [58] Michael M. Yartsev and Nachum Ulanovsky. Representation of three-dimensional space in the hippocampus of flying bats. *Science*, 340(6130):367–372, 2013.
- [59] Randall J. LeVeque. Finite Difference Methods for Ordinary and Partial Differential Equations: Steady-State and Time-Dependent Problems. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2007.
- [60] J. R. Norris. Path integrals and heat kernel asymptotics. *Probability Theory and Related Fields*, 109(1):1–33, 1997.
- [61] Dori Derdikman, Jonathan R Whitlock, Albert Tsao, Marianne Fyhn, Torkel Hafting, May-Britt Moser, and Edvard I Moser. Fragmentation of grid cell maps in a multicompartment environment. *Nature Neuroscience*, 12(10):1325–1332, 2009.
- [62] Patrick D Rich, Hsin-Pei Liaw, and Albert K Lee. Large environments reveal the statistical structure governing hippocampal representations. *Science*, 345(6198):814–817, 2014.
- [63] Dmitri Burago, Yuri Burago, and Sergei Ivanov. *A Course in Metric Geometry*. Graduate Studies in Mathematics. American Mathematical Society, 2001.
- [64] Anita Kamondi, László Acsády, Xiao-Jing Wang, and György Buzsáki. Theta oscillations in somata and dendrites of hippocampal pyramidal cells in vivo: Activity-dependent phaseprecession of action potentials. *Hippocampus*, 8(3):244–261, 1998.
- [65] Christopher D Harvey, Forrest Collman, Daniel A Dombeck, and David W Tank. Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature*, 461(7266):941–946, 2009.
- [66] Torkel Hafting, Marianne Fyhn, Sturla Molden, May-Britt Moser, and Edvard I. Moser. Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052):801–806, 2005.
- [67] Dehong Xu, Ruiqi Gao, Wenhao Zhang, Xue-Xin Wei, and Ying Nian Wu. On conformal isometry of grid cells: Learning distance-preserving position embedding. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [68] Marianne Fyhn, Torkel Hafting, Alessandro Treves, May-Britt Moser, and Edvard I Moser. Hippocampal remapping and grid realignment in entorhinal cortex. *Nature*, 446(7132):190–194, 2007.

Contents

1	Introduction	1
2	Method	2
3	Experiments	6
4	Conclusion	10
A	Related Work	15
В	Limitations	17
C	Spectral Analysis of Multi-Step Random Walk	17
D	Heat Diffusion with Reflecting Boundaries (Background)	19
E	Open-Field Environment	20
F	Emergent Sparsity from Non-negativity and Orthogonality	22
G	Path Planning Properties	23
Н	Theta Phase Modeling	28
Ι	Grid Cells Integration	29
J	Implementation Details	30
K	Additional Experiment Results	30

Sections and subsections that contain mostly background materials are so marked in the titles. The remaining sections and subsections contain novel developments.

A Related Work

A.1 Place Cell Models and Representations

The study of hippocampal place cells has a rich history since their discovery by [1]. Computational models of place cells have evolved from simple Gaussian tuning curves [29] to more sophisticated approaches. Several models have explored population-level representations of place cells, including manifold embeddings [32] and latent space models [33]. However, these approaches typically treat place cells as separate entities encoding specific locations rather than as collective embeddings encoding transition probabilities.

Matrix factorization approaches to neural population activity have been applied to various brain regions [34, 35], though rarely with the specific mathematical connection to random walk processes proposed in our work. Recent work by [13] proposed a successor representation framework for place cells, which shares some conceptual similarities with our transition probability approach but differs in the specific mathematical formulation and implementation.

Hidden Markov Models (HMMs) have also been applied to model hippocampal spatial coding. The Clone-Structured Causal Graph model (CSCG) [36] treats space as a latent sequence and uses a structured graph with cloned nodes to disambiguate aliased sensory observations in different contexts. Related work [37] further develops sequence-based models of hippocampal function.

A.2 Multi-scale Spatial Representations

The variation in place field sizes along the dorsoventral axis of the hippocampus has been extensively documented [5, 21, 22], but computational models that explicitly address this multi-scale organiza-

tion remain limited. Models incorporating scale in place cell representations include hierarchical approaches [38], wavelet-like representations [39], and scale-space theories [40] borrowed from computer vision.

Our approach differs by deriving the multi-scale representation directly from the time parameter of a random walk process, providing a principled connection between scale and exploration time that has not been previously exploited in place cell models.

A.3 Navigation and Path Planning

Biologically-inspired navigation algorithms have drawn on various hippocampal properties. Vector-based navigation models [41, 42] and successor representation approaches [38, 43] have demonstrated effective navigation capabilities. Diffusion-based path planning algorithms [44, 45] share mathematical similarities with our heat equation formulation but lack the direct connection to neural representations.

Recent work by [46] and [47] has emphasized the role of population coding in navigation but without the specific inner product relationship and transition probability framework we propose.

A.4 Successor Representation

Our work is closely related to the Successor Representation (SR) [48, 49], as both frameworks build upon powers of a transition matrix to construct a representation of space. The SR is typically defined as a discounted sum, $M = \sum_{k=0}^{\infty} (\gamma T)^k$, which yields the expected discounted future occupancy of all states from any given starting state, and the discount factor γ implicitly sets a single, predictive temporal horizon. In contrast, our approach utilizes the discrete powers of the transition matrix directly, $P_{\tau} = P^{\tau}$, to explicitly model the transition probabilities at multiple, distinct time horizons, where the parameter τ corresponds directly to the time scale of a random walk. This multi-scale representation aligns more closely with the observed functional gradient of place field sizes along the hippocampal dorsoventral axis.

The primary novelty of our framework lies in learning a biologically-constrained, non-negative matrix factorization of these transition kernels, where vector embeddings $h(x,\tau)$ are learned such that their inner product reconstructs the transition probability, $\langle h(x,\tau), h(y,\tau) \rangle = q(y|x,\tau)$. A crucial theoretical insight is that the combination of this inner product objective with non-negativity and orthogonality constraints for distant locations gives rise to **emergent sparsity**; the embeddings are forced to have disjoint support sets, providing a geometric explanation for the localized firing fields of place cells without explicit regularization. This multi-scale architecture confers significant functional advantages, enabling an adaptive gradient-ascent-based navigation policy that selects the optimal scale τ^* at each step to generate smooth, trap-free trajectories. This approach is not only more scalable, but also offers a more direct navigation mechanism that bypasses the need for explicit value function computation common in SR-based reinforcement learning agents.

A.5 Inner Product Spaces in Neural Representation

Inner product spaces as a basis for neural computation have been explored by several researchers [50, 51]. More recently, [52] proposed that neural populations represent probability distributions through their activation patterns, with some conceptual overlap with our approach.

Navigational planning in inner product spaces has connections to kernel methods in machine learning [53] and information geometry [54], though these connections have been underexplored in neuroscience. Our work bridges these fields by explicitly relating inner products between place cell populations to transition probabilities derived from random walk processes.

A.6 Heat Equation and Diffusion Models in Neuroscience

The connection between neural dynamics and diffusion processes has been explored in various contexts [55]. The specific relationship between heat kernels and geodesic distances, which forms a foundation for our approach, has strong connections to manifold learning [56] and dimensionality reduction techniques [57].

Our work builds upon these ideas by applying them specifically to place cell populations and spatial navigation, providing a novel bridge between diffusion processes, neural representations, and navigational behavior.

B Limitations

B.1 Lack of Sensory Integration

Our current model does not incorporate sensory inputs for spatial navigation[11, 2]. Real hippocampal place cells integrate multimodal sensory information including visual landmarks, self-motion cues from path integration, and boundary detection. This sensory integration is fundamental to how place cells form their spatial receptive fields and anchor representations to environmental features.

Without sensory processing, our model cannot explain how place fields emerge during initial exposure to novel environments. In our framework, position embeddings $h(x,\tau)$ are learned from pre-defined transition probabilities that already encode environmental structure, rather than being built up through sensory experience. We also cannot model perceptual anchoring, where place cells maintain stable firing relative to visual landmarks, or how sensory conflicts (such as navigation in darkness) affect place cell stability. Future extensions incorporating vision-based inputs, self-motion integration, and boundary vector cell signals would significantly enhance biological plausibility.

B.2 Limited to 2D Static Environments

Our experiments and analyses are restricted to two-dimensional static environments, whereas biological spatial navigation operates in three-dimensional space and continuously changing contexts[5, 58]. Studies in flying bats have revealed volumetric place cells encoding 3D positions, and terrestrial animals navigating multi-level structures require three-dimensional representations. Extension to 3D introduces complications our framework does not address: vertical movements and gravity-dependent asymmetries, different place field scaling across dimensions, and substantially increased computational complexity. Dynamic environments would require time-varying transition probabilities and continuous updating mechanisms our framework does not provide.

B.3 Online Sequential Learning

Another limitation is our lack of online learning over sequential experiences to model temporal dynamics and memory consolidation[7, 8]. Our batch optimization learns $h(x,\tau)$ across all spatial locations simultaneously, fundamentally differing from how biological place cells develop through experience. Biological place cell formation is incremental: initial weak or unstable spatial tuning gradually sharpens through repeated exploration via activity-dependent synaptic plasticity. Our framework does not capture this gradual emergence from unstructured initial conditions. Our matrix squaring $P_{2\tau}=P_{\tau}^2$ achieves global integration instantaneously based on environmental structure, not through iterative learning over sleep-wake cycles analogous to biological replay.

Despite these limitations, our framework offers valuable theoretical insights into computational principles underlying hippocampal spatial navigation, particularly the role of multi-scale transition probabilities, emergence of sparsity from geometric constraints, and the connection between random walk dynamics and cognitive maps.

C Spectral Analysis of Multi-Step Random Walk

This appendix provides a detailed spectral analysis of the multi-step random walk transition kernel, establishing connections between our random walk model and position embeddings.

C.1 Eigendecomposition of the Transition Matrix (Background)

Since the one-step transition matrix P_1 is symmetric by construction, it admits an eigendecomposition:

$$P_1 = Q\Lambda Q^T, (13)$$

with orthogonal Q ($Q^TQ = I$) and diagonal $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$, where $0 \le \lambda_i \le 1$, where we assume the random walk is irreducible and aperiodic. The multi-step transition matrix is:

$$P_1^{\tau} = Q\Lambda^{\tau}Q^T \tag{14}$$

The eigenvalues are typically ordered as $1=\lambda_1>\lambda_2\geq\lambda_3\geq\ldots$, with the first eigenvalue corresponding to the stationary distribution and subsequent eigenvalues capturing spatial patterns at increasing levels of detail.

C.2 Position Embeddings from Spectral Decomposition

The spectral decomposition provides a natural position embedding. If we define:

$$H_i(x,\tau) = \lambda_i^{\tau/2} Q_i(x) \tag{15}$$

in the discrete case where $Q_i(x)$ is the i-th column of Q, or:

$$H_i(x,\tau) = e^{\lambda_i \tau/2} \phi_i(x) \tag{16}$$

in the continuous case (where λ_i and ϕ_i are respectively eigenvalues and eigenfunctions of the Laplacian), then the transition probability can be expressed as:

$$p(y|x,\tau) = \sum_{i} H_i(x,\tau)H_i(y,\tau) = \langle H(x,\tau), H(y,\tau) \rangle$$
 (17)

This provides a closed-form expression for position embeddings that exactly reproduce the transition probabilities through inner products.

C.3 Normalization and Non-Negative Embeddings

To obtain normalized embeddings, we define:

$$h_{\text{spec}}(x,\tau) = \frac{H(x,\tau)}{\|H(x,\tau)\|} = \frac{H(x,\tau)}{\sqrt{p(x|x,\tau)}}$$
 (18)

This normalized embedding satisfies:

$$\langle h_{\text{spec}}(x,\tau), h_{\text{spec}}(y,\tau) \rangle = \frac{p(y|x,\tau)}{\sqrt{p(x|x,\tau) \cdot p(y|y,\tau)}} = q(y|x,\tau)$$
(19)

However, $h_{\rm spec}(x,\tau)$ may contain negative components, which conflicts with the biological constraint that neural firing rates must be non-negative. This is where Horn's theorem becomes relevant.

C.4 Non-Negative Matrix Factorization (Background)

Horn's theorem [19] provides the theoretical foundation for obtaining non-negative embeddings from our transition matrices: If A is a symmetric matrix with non-negative entries, then there exists a non-negative matrix B such that $A = BB^T$.

We provide a sketch of the proof:

Proof sketch of Horn's theorem. Given a symmetric non-negative matrix A, consider its spectral decomposition $A = U\Sigma U^T$. Let $A = \sum_i \sigma_i u_i u_i^T$, where σ_i are eigenvalues and u_i are eigenvectors.

For each outer product $u_i u_i^T$ (which may have negative entries), we can express it as a linear combination of non-negative rank-1 matrices: $u_i u_i^T = \sum_j c_j v_j v_j^T$, where v_j are non-negative vectors and c_j are coefficients.

By appropriate selection of v_j (e.g., using vertices of the hypercube defined by the signs of u_i), we can ensure that $c_j \geq 0$ when $\sigma_i > 0$. This allows us to express A as a sum of non-negative rank-1 matrices, which can be arranged as $A = BB^T$ where B has non-negative entries.

For our normalized transition matrix $Q_{\tau} = [q(y|x,\tau)]$, which is symmetric with non-negative entries, Horn's theorem guarantees the existence of a non-negative matrix H_{τ} such that $Q_{\tau} = H_{\tau}H_{\tau}^{T}$. The rows of H_{τ} provide our desired non-negative position embeddings $h(x,\tau)$.

D Heat Diffusion with Reflecting Boundaries (Background)

This appendix provides a detailed mathematical derivation of the connection between our discrete random walk model and the continuous heat equation, establishing the relationship between transition probabilities and geodesic distances.

D.1 From Discrete Random Walk to Reflecting Heat Equation

We begin with a discrete random walk on a two-dimensional integer grid. For simplicity, we assume the one-step transition $p(y|x,\tau=1)$ of the random walk is to move to one of 4 nearest neighbors with $p_{\text{move}}=1/4$. We assume this simplest $p(y|x,\tau=1)$ in our theoretical derivations in all the relevant sections in the Appendix. Similar results can be obtained for more general $p(y|x,\tau=1)$, with a different diffusion coefficient α .

Let $(i, j) \in \mathbb{Z}^2$ denote discrete spatial coordinates, and $k \in \mathbb{Z}_{\geq 0}$ denote discrete time steps. To connect with continuous diffusion, we introduce spatial and temporal units:

$$x = i \cdot dx, \quad y = j \cdot dx, \quad \tau = k \cdot dt$$
 (20)

where dx is the spacing between adjacent grid points and dt is the time step. Following standard diffusion scaling, we set $dx = \sqrt{dt}$, which ensures convergence to a well-defined limit as $dx \to 0$ [14].

The discrete random walk has the following transition probabilities:

$$p((i',j')|(i,j),k=1) = p_{\text{move}} = \frac{1}{4} \quad \text{for each unobstructed neighbor } (i',j') \text{ of } (i,j) \tag{21}$$

$$p((i,j)|(i,j),k=1) = 1 - N(i,j) \cdot p_{\text{move}} = 1 - \frac{N(i,j)}{4}$$
(22)

where N(i, j) is the number of unobstructed neighbors of location (i, j) (maximum 4 in a 2D grid with 4-connectivity).

D.2 Derivation of the Heat Equation

To derive the continuous limit, for interior points (those not adjacent to obstacles), the discrete update rule is:

$$p(i, j, k + 1) = p(i, j, k)(1 - 4p_{\text{move}}) + p_{\text{move}}[p(i + 1, j, k) + p(i - 1, j, k) + p(i, j + 1, k) + p(i, j - 1, k)]$$

$$= p(i, j, k) + p_{\text{move}}[p(i + 1, j, k) + p(i - 1, j, k) + p(i, j + 1, k) + p(i, j - 1, k) - 4p(i, j, k)]$$
(23)

Dividing both sides by dt and using $p_{\text{move}} = 1/4$:

$$\frac{p(i,j,k+1) - p(i,j,k)}{dt} = \frac{1}{4} \cdot \frac{1}{dt} [p(i+1,j,k) + p(i-1,j,k) + p(i,j+1,k) + p(i,j-1,k) - 4p(i,j,k)]$$
(24)

Using the standard finite difference approximation for the Laplacian [59]:

$$\nabla^2 p(i,j,k) \approx \frac{1}{dx^2} [p(i+1,j,k) + p(i-1,j,k) + p(i,j+1,k) + p(i,j-1,k) - 4p(i,j,k)]$$
(25)

Substituting $dx^2 = dt$ and taking the limit as $dt \to 0$:

$$\lim_{dt \to 0} \frac{p(i, j, k+1) - p(i, j, k)}{dt} = \frac{1}{4} \nabla^2 p(x, y, t)$$
 (26)

This yields the heat equation with diffusion coefficient $\alpha = 1/4$:

$$\frac{\partial p(x,y,\tau)}{\partial \tau} = \alpha \nabla^2 p(x,y,\tau) \tag{27}$$

D.3 Reflecting Boundary Conditions

Our discrete random walk enforces reflecting boundary conditions, preventing probability flow into obstacles, mirroring hippocampal obstacle avoidance [11]. Consider a boundary point (i,j) with an obstacle at (i+1,j), so the number of valid neighbors is N(i,j)=3 (points (i-1,j), (i,j+1), (i,j-1)). The self-transition probability is $p_{\text{stay}}=1-\frac{N(i,j)}{4}=\frac{1}{4}$, redistributing probability that would flow to the obstacle back to (i,j).

The probability update for (i, j) at time step k + 1 is:

$$p(i,j,k+1) = \frac{1}{4}p(i,j,k) + \frac{1}{4}\left[p(i-1,j,k) + p(i,j+1,k) + p(i,j-1,k)\right]$$
(28)

This can be rewritten as:

$$p(i,j,k+1) = p(i,j,k) + \frac{1}{4} \left[p(i-1,j,k) + p(i,j+1,k) + p(i,j-1,k) - 3p(i,j,k) \right]$$
 (29)

In finite difference methods, the reflecting condition $\frac{\partial p}{\partial n}\Big|_{(i+1,j)} = 0$ is enforced using a ghost point at (i+1,j), setting p(i+1,j,k) = p(i,j,k) to ensure zero normal flux [59]. Substituting into the interior Laplacian update (equation 30):

$$p(i,j,k+1) = p(i,j,k) + \frac{1}{4} \left[p(i-1,j,k) + p(i+1,j,k) + p(i,j+1,k) + p(i,j-1,k) - 4p(i,j,k) \right]$$
(30)

yields equation (29), as p(i+1,j,k) = p(i,j,k) reduces the neighbor terms to three. This substitution ensures the boundary update aligns with the random walk's mechanism, where $p_{\text{stay}} = \frac{1}{4}$ assigns zero probability to obstacle transitions, maintaining probability conservation and enabling smooth navigation around obstacles, as observed in hippocampal place cell activity [11].

D.4 Connection to Geodesic Distance

Varadhan's formula [16] establishes a deep relationship between the heat kernel and geodesic distance. For a complete Riemannian manifold M with heat kernel $p(x, y, \tau)$, Varadhan proved that:

$$\lim_{\tau \to 0} -4t \log p(x, y, \tau) = d_g^2(x, y) \tag{31}$$

where $d_q(x, y)$ is the geodesic distance between points x and y.

While Varadhan's original large deviation principle [16] applies to smooth manifolds without boundaries, extensions to domains with reflecting boundaries [60] ensure that the short-time behavior of the heat kernel $p(y,\tau|x,0)$ reflects the geodesic distance $d_g(x,y)$, the shortest path within Ω avoiding obstacles. For our normalized transition probability: a similar asymptotic relationship holds:

$$\lim_{\tau \to 0} -\tau \log q(y|x,\tau) = \frac{d_g^2(x,y)}{4\alpha},\tag{32}$$

where $\alpha=1/4$ is the diffusion coefficient. This follows because $p(x|x,\tau)$ and $p(y|y,\tau)$, influenced by boundary reflections, have asymptotic forms that cancel in the logarithm as $\tau\to 0$, leaving the geodesic term dominant.

As τ increases, our distance metric transitions from approximating geodesic distance to incorporating more global aspects of the domain's connectivity, creating a multi-scale representation that seamlessly integrates local metric information with global connectivity structure. Section G explains eigen analysis of connectivity.

E Open-Field Environment

In unbounded, obstacle-free environments (open fields), the symmetric random walk simplifies to an isotropic diffusion process, offering a theoretical justification for the Gaussian tuning of place cells

observed in such settings [1, 29]. As the grid discretization refines $(dx \to 0, dt \to 0, dx = \sqrt{dt})$, the transition probability $p(y|x,\tau)$ converges to the heat equation's fundamental solution in 2D free space:

$$p(y|x,\tau) = \frac{1}{4\pi\alpha\tau} \exp\left(-\frac{\|y-x\|^2}{4\alpha\tau}\right)$$
 (33)

where $\alpha=1/4$ is the diffusion coefficient, and $\|y-x\|^2$ is the squared Euclidean distance. Since $p(x|x,\tau)=\frac{1}{4\pi\alpha\tau}$ is position-independent, the normalized transition probability becomes:

$$q(y|x,\tau) = \frac{p(y|x,\tau)}{\sqrt{p(x|x,\tau)p(y|y,\tau)}} = \exp\left(-\frac{\|y-x\|^2}{4\alpha\tau}\right) = \exp\left(-\frac{\|y-x\|^2}{\tau}\right)$$
(34)

This Gaussian kernel, with variance $\sigma^2 = 2\alpha\tau = \tau/2$, reflects the diffusive spread of the random walk and mirrors the approximately Gaussian firing fields of hippocampal place cells in open environments. Given $q(y|x,\tau) = \langle h(x,\tau), h(y,\tau) \rangle$, the position embeddings $h(x,\tau)$ must reproduce this Gaussian decay. We construct $h(x,\tau)$ with Gaussian components, providing a mathematical basis for why place cell population activity exhibits Gaussian profiles in open fields.

Theorem 1 (Gaussian Embeddings in Open Fields). In an unbounded 2D open field, where $q(y|x,\tau) = \exp\left(-\frac{\|y-x\|^2}{\tau}\right)$, there exists a position embedding $h(x,\tau) \in \mathbb{R}^n$ with non-negative components such that $\langle h(x,\tau), h(y,\tau) \rangle = q(y|x,\tau)$, and each component $h_i(x,\tau)$ is a Gaussian function of x with variance $\tau/2$ per dimension.

Proof. The transition kernel $q(y|x,\tau)=\exp\left(-\frac{\|y-x\|^2}{\tau}\right)$ is a positive definite Gaussian kernel with variance $\tau/2$ (since $4\alpha\tau=\tau$ for $\alpha=1/4$), consistent with the Gaussian tuning of place cells in open fields. We construct $h(x,\tau)\in\mathbb{R}^n$ directly as:

$$h_i(x,\tau) = c_i \exp\left(-\frac{\|x - \mu_i\|^2}{\tau}\right)$$
(35)

where $\mu_i \in \mathbb{R}^2$ are fixed anchor points (e.g., a uniform grid), $c_i > 0$ are constants, and the variance per dimension is $\tau/2$. The inner product is:

$$\langle h(x,\tau), h(y,\tau) \rangle = \sum_{i=1}^{n} h_i(x,\tau) h_i(y,\tau) = \sum_{i=1}^{n} c_i^2 \exp\left(-\frac{\|x - \mu_i\|^2 + \|y - \mu_i\|^2}{\tau}\right)$$
(36)

Rewrite the exponent:

$$||x - \mu_i||^2 + ||y - \mu_i||^2 = ||x - y||^2 + 2\left\|\frac{x + y}{2} - \mu_i\right\|^2$$
(37)

so:

$$\langle h(x,\tau), h(y,\tau) \rangle = \exp\left(-\frac{\|x-y\|^2}{\tau}\right) \sum_{i=1}^n c_i^2 \exp\left(-\frac{2\|\frac{x+y}{2} - \mu_i\|^2}{\tau}\right)$$
 (38)

For a dense set of μ_i (e.g., a grid with spacing $\ll \sqrt{\tau}$), the sum approximates a constant over a local region around (x+y)/2, as the terms $\exp\left(-\frac{2\left\|\frac{x+y}{2}-\mu_i\right\|^2}{\tau}\right)$ form a kernel density estimate. Set $c_i^2=\frac{1}{n}$; as $n\to\infty$, the sum converges to a constant C (proportional to the density of μ_i), yielding:

$$\langle h(x,\tau), h(y,\tau) \rangle \approx C \exp\left(-\frac{\|x-y\|^2}{\tau}\right)$$
 (39)

Adjust C=1 by scaling c_i (e.g., $c_i=\sqrt{\frac{1}{Cn}}$), ensuring $\langle h(x,\tau),h(y,\tau)\rangle=q(y|x,\tau)$. Each $h_i(x,\tau)$ is Gaussian with variance $\tau/2$ per dimension, and $h_i(x,\tau)\geq 0$, satisfying the theorem.

This open-field case not only validates our model against biological data but also serves as a baseline to explore how environmental structure (e.g., obstacles) perturbs these Gaussian properties, as observed in constrained settings [61, 62].

F Emergent Sparsity from Non-negativity and Orthogonality

Let $h(x,\tau) \in \mathbb{R}^n_+$ denote the non-negative representation of position $x \in \mathcal{X}$ at scale $\sqrt{\tau}$, and let

$$\langle h(x,\tau), h(y,\tau) \rangle = q(y \mid x,\tau)$$
 (40)

where $q(\cdot | \cdot, \tau) \ge 0$ is a symmetric multi-step transition kernel (adjacency at scale $\sqrt{\tau}$). Define the (undirected) transition graph $\mathcal{G}_{\tau} = (\mathcal{X}, E_{\tau})$ by

$$(x,y) \in E_{\tau} \iff q(y \mid x,\tau) > 0$$
 (41)

Note that this transition graph is for τ steps instead of one step. Write $supp(v) = \{i \in [n] : v_i > 0\}$ for the support of a non-negative vector $v \in \mathbb{R}^n_+$, and put

$$S_i(\tau) = \{x \in \mathcal{X} : h_i(x,\tau) > 0\}$$
 (the active set of coordinate *i* at scale τ) (42)

Let $\deg_{\tau}(x)$ denote the degree of x in \mathcal{G}_{τ} , $\Delta(\tau) = \max_{x} \deg_{\tau}(x)$ the maximum degree, and

$$\rho(\tau) = \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} \deg_{\tau}(x) \tag{43}$$

the mean neighborhood size.

Lemma 2 (Non-negativity + orthogonality \Rightarrow disjoint support). If $h(x,\tau), h(y,\tau) \in \mathbb{R}^n_+$ and $\langle h(x,\tau), h(y,\tau) \rangle = 0$, then

$$\operatorname{supp}(h(x,\tau)) \cap \operatorname{supp}(h(y,\tau)) = \varnothing \tag{44}$$

Proof. By non-negativity,

$$0 = \langle h(x,\tau), h(y,\tau) \rangle = \sum_{i=1}^{n} h_i(x,\tau) h_i(y,\tau)$$

$$(45)$$

is a sum of non-negative terms. Hence each term vanishes: $h_i(x,\tau) h_i(y,\tau) = 0$ for all $i \in [n]$, and thus no coordinate is simultaneously positive in both vectors.

Proposition 3 (Clique structure and sparsity bound). Assume (40) and that

$$q(y \mid x, \tau) = 0$$
 whenever $(x, y) \notin E_{\tau}$ (46)

Then for each coordinate $i \in [n]$:

- (i) $S_i(\tau)$ induces a clique in \mathcal{G}_{τ} ; i.e., every pair $x, y \in S_i(\tau)$ is adjacent in \mathcal{G}_{τ}
- (ii) Consequently,

$$|S_i(\tau)| \le 1 + \min_{x \in S_i(\tau)} \deg_{\tau}(x) \le 1 + \Delta(\tau) \tag{47}$$

(iii) The average number of active coordinates per position obeys

$$\frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} \|h(x, \tau)\|_{0} = \frac{1}{|\mathcal{X}|} \sum_{i=1}^{n} |S_{i}(\tau)| \le \frac{n (1 + \Delta(\tau))}{|\mathcal{X}|}$$
(48)

Moreover, since $\Delta(\tau) = \max_x \deg_{\tau}(x)$ and $\rho(\tau)$ is the mean degree, any a priori bound of the form $\Delta(\tau) \leq C \rho(\tau)$ (with graph-dependent constant $C \geq 1$) yields

$$\frac{1}{|\mathcal{X}|} \sum_{\tau} \|h(x,\tau)\|_{0} \leq \frac{n \left[1 + C \rho(\tau)\right]}{|\mathcal{X}|}$$
(49)

Proof. (i) Take any $x,y \in S_i(\tau)$ with $x \neq y$. Then $h_i(x,\tau) > 0$ and $h_i(y,\tau) > 0$, so by Lemma 2 they cannot be orthogonal. By the hypothesis $q(y \mid x,\tau) = \langle h(x,\tau), h(y,\tau) \rangle$, orthogonality occurs exactly when $(x,y) \notin E_\tau$. Therefore $(x,y) \in E_\tau$. Since the choice of x,y was arbitrary, $S_i(\tau)$ induces a clique.

(ii) Fix any $x^* \in S_i(\tau)$. By (i), every $y \in S_i(\tau) \setminus \{x^*\}$ must be adjacent to x^* , so $S_i(\tau) \subseteq \{x^*\} \cup \mathcal{N}_{\tau}(x^*)$, where $\mathcal{N}_{\tau}(x^*)$ is the neighborhood of x^* . Thus $|S_i(\tau)| \leq 1 + \deg_{\tau}(x^*)$. Minimizing over $x^* \in S_i(\tau)$ gives the first inequality; the second follows from $\deg_{\tau}(x^*) \leq \Delta(\tau)$.

(iii) The identity

$$\sum_{x \in \mathcal{X}} \|h(x, \tau)\|_{0} = \sum_{i=1}^{n} |S_{i}(\tau)|$$
 (50)

is a double-counting equality (both sides count the number of pairs (x,i) with $h_i(x,\tau)>0$). Using (ii) and summing over i yields $\sum_i |S_i(\tau)| \leq n \, [1+\Delta(\tau)]$, and dividing by $|\mathcal{X}|$ establishes (48). If an a priori comparison $\Delta(\tau) \leq C \, \rho(\tau)$ holds for the graph family under consideration, the final bound follows immediately. \square

Geometric specialization. If $\mathcal{X} \subset \mathbb{R}^m$ is η -separated (minimum inter-point distance $\geq \eta > 0$) and \mathcal{G}_{τ} is a geometric threshold graph

$$(x,y) \in E_{\tau} \iff ||x-y|| \le \sqrt{\tau}$$
 (51)

then each closed neighborhood fits inside a ball $B(x,\sqrt{\tau})$ of radius $\sqrt{\tau}$. A standard packing argument implies

$$\deg_{\tau}(x) + 1 \le C_m \left(\frac{\sqrt{\tau}}{\eta}\right)^m \tag{52}$$

for a constant C_m depending only on the ambient dimension m. Consequently,

$$|S_i(\tau)| \leq C_m \left(\frac{\sqrt{\tau}}{\eta}\right)^m \qquad \frac{1}{|\mathcal{X}|} \sum_x ||h(x,\tau)||_0, \leq \frac{dC_m}{|\mathcal{X}|} \left(\frac{\sqrt{\tau}}{\eta}\right)^m \tag{53}$$

Thus, as τ increases, the allowable clique size (hence the bound on average support) grows polynomially with the geometric volume of the $\sqrt{\tau}$ -ball, quantitatively linking scale to reduced sparsity (broader fields).

Corollary 4 (Scale-dependent localization). Assume $\tau \mapsto \mathcal{G}_{\tau}$ is monotone in the sense that $\tau_1 < \tau_2$ implies $E_{\tau_1} \subseteq E_{\tau_2}$. Then $\Delta(\tau)$ and $\rho(\tau)$ are non-decreasing in τ , and the upper bounds in Proposition 3 are non-decreasing in τ , i.e., expected sparsity decreases with scale, corresponding to larger place fields.

Proof. Monotonicity $E_{\tau_1} \subseteq E_{\tau_2}$ implies $\deg_{\tau_1}(x) \leq \deg_{\tau_2}(x)$ for all x, hence $\Delta(\tau_1) \leq \Delta(\tau_2)$ and $\rho(\tau_1) \leq \rho(\tau_2)$. Apply Proposition 3 (iii).

Interpretation. Lemma 2 enforces disjoint supports for non-adjacent locations; Proposition 3 shows that each coordinate i can only support a *clique* of mutually adjacent positions and bounds the size of that clique by local neighborhood size. In geometric environments, the bound scales with the volume of a $\sqrt{\tau}$ -ball, making the τ -field-size relationship explicit. Corollary 4 then formalizes the empirical trend: smaller $\tau \Rightarrow$ higher sparsity (smaller fields); larger $\tau \Rightarrow$ lower sparsity (larger fields).

G Path Planning Properties

This appendix elaborates on the properties of the gradient-based navigation framework introduced in Section 2.7.

G.1 Scale Transition Dynamics

The optimal scale $\tau^* = \arg\max_t \|\nabla_x q(x|y,\tau)\|$ adapts dynamically to the agent's distance from the goal, ensuring efficient navigation. In an open field, we formalize this selection with the following theorem.

Theorem 5 (Optimal Scale Selection in an Open Field). In an open field with a Gaussian transition kernel $p(x|y,\tau) = \frac{1}{4\pi\alpha\tau} \exp\left(-\frac{\|x-y\|^2}{4\alpha\tau}\right)$, where $\alpha = 1/4$, the optimal time scale τ^* that maximizes the gradient magnitude $\|\nabla_x q(x|y,\tau)\|$, with $q(x|y,\tau) = p(x|y,\tau)/\sqrt{p(x|x,\tau)\cdot p(y|y,\tau)}$, is:

$$\tau^* = \frac{d^2(x,y)}{4\alpha} = d^2(x,y) \tag{54}$$

where d(x, y) = ||x - y|| is the Euclidean distance.

Proof. The transition kernel is:

$$p(x|y,\tau) = \frac{1}{4\pi\alpha\tau} \exp\left(-\frac{d^2}{4\alpha\tau}\right), \quad d = d(x,y)$$
 (55)

Normalization gives:

$$p(x|x,\tau) = p(y|y,\tau) = \frac{1}{4\pi\alpha\tau}, \quad \sqrt{p(x|x,\tau) \cdot p(y|y,\tau)} = \frac{1}{4\pi\alpha\tau}$$
 (56)

Thus:

$$q(x|y,\tau) = \frac{p(x|y,\tau)}{\frac{1}{4\pi\alpha\tau}} = \exp\left(-\frac{d^2}{4\alpha\tau}\right)$$
 (57)

The gradient is:

$$\nabla_x q(x|y,\tau) = \exp\left(-\frac{d^2}{4\alpha\tau}\right) \cdot \left(-\frac{2(x-y)}{4\alpha\tau}\right) = -\frac{x-y}{2\alpha\tau} \exp\left(-\frac{d^2}{4\alpha\tau}\right)$$
 (58)

The magnitude is:

$$\|\nabla_x q(x|y,\tau)\| = \frac{d}{2\alpha\tau} \exp\left(-\frac{d^2}{4\alpha\tau}\right)$$
 (59)

Maximize $f(\tau) = \frac{d}{2\alpha\tau} \exp\left(-\frac{d^2}{4\alpha\tau}\right)$. Compute:

$$\ln f(\tau) = \ln \left(\frac{d}{2\alpha}\right) - \ln \tau - \frac{d^2}{4\alpha\tau} \tag{60}$$

Differentiate:

$$\frac{\partial \ln f}{\partial \tau} = -\frac{1}{\tau} + \frac{d^2}{4\alpha \tau^2} = 0 \implies \tau = \frac{d^2}{4\alpha} \tag{61}$$

For $\alpha = 1/4$, $\tau^* = d^2$. The second derivative at τ^* :

$$\frac{\partial^2 \ln f}{\partial \tau^2} = \frac{1}{\tau^2} - \frac{d^2}{2\alpha \tau^3}, \quad \text{at } \tau = \frac{d^2}{4\alpha}, \quad \frac{d^2}{2\alpha \tau} = 2, \quad \frac{\partial^2 \ln f}{\partial \tau^2} = -\frac{1}{\tau^2} < 0 \tag{62}$$

confirms a maximum.

In an open field, $\tau^* \propto d^2(x,y)$, so τ^* decreases as the agent approaches the goal, focusing on finer spatial scales for precision.

G.2 Properties of the Gradient Field

The gradient field $\nabla_x q(x|y,\tau)$ drives navigation. We highlight three properties ensuring a smooth, trap-free path to the goal y.

First, $p(x|y,\tau)$ satisfies the heat equation with reflecting boundary conditions:

$$\frac{\partial p(x|y,\tau)}{\partial \tau} = \alpha \nabla^2 p(x|y,\tau), \quad \alpha = 1/4, \quad \frac{\partial p(x|y,\tau)}{\partial n} = 0 \text{ on } \partial \Omega_{\text{obstacles}}.$$
 (63)

For fixed y and τ , $p(x|y,\tau)$ is smooth in x, as the heat kernel (e.g., $p(x|y,\tau) = \frac{1}{4\pi\alpha\tau} \exp\left(-\frac{\|x-y\|^2}{4\alpha\tau}\right)$ in an open field) is infinitely differentiable [15]. It has a unique maximum at x=y.

Second, $p(x|x,\tau)$ is smooth in x for fixed τ . In an open field, $p(x|x,\tau) = \frac{1}{4\pi\alpha\tau}$ is constant, while in general, $p(x|x,\tau)$ varies smoothly due to the heat kernel's differentiability, reflecting the domain's geometry near obstacles.

Third, since the random walk is symmetric, $\nabla_x q(x|y,\tau) = \nabla_x q(y|x,\tau)$. The gradient field of $q(y|x,\tau)$ is smooth, as $q(y|x,\tau) = q(x|y,\tau)$ inherits the smoothness of $p(x|y,\tau)$, and has a unique maximum at x=y.

These properties ensure that $\nabla_x q(y|x,\tau)$ forms a smooth field with a unique maximum at the goal, producing a continuous, trap-free path toward y. This mirrors the hippocampus's efficient navigation, where place cells encode smooth, goal-directed trajectories [11].

G.3 Planned Path vs. Shortest Path

The planned path follows the gradient $\nabla_x q(x|y,\tau)$, and $p(x|y,\tau)$ satisfies the heat equation with reflecting boundary conditions. We compare this path to the shortest (geodesic) path from x to the goal y.

In an open field, the planned path is a straight line. The transition kernel is:

$$p(x|y,\tau) = \frac{1}{4\pi\alpha\tau} \exp\left(-\frac{\|x-y\|^2}{4\alpha\tau}\right), \quad \alpha = 1/4$$
 (64)

Since $p(x|x,\tau) = p(y|y,\tau) = \frac{1}{4\pi\alpha\tau}$, the normalization gives:

$$q(x|y,\tau) = \exp\left(-\frac{\|x-y\|^2}{4\alpha\tau}\right) \tag{65}$$

The gradient is:

$$\nabla_x q(x|y,\tau) = -\frac{x-y}{2\alpha\tau} \exp\left(-\frac{\|x-y\|^2}{4\alpha\tau}\right) \propto -(x-y)$$
 (66)

Thus, $\frac{dx(t)}{dt} = \nabla_x q(x|y,\tau)$ follows a straight line from x to y, matching the Euclidean shortest path [63].

For small τ , $q(x|y,\tau)$ aligns with the geodesic distance:

$$p(x|y,\tau) \sim \frac{1}{(4\pi\alpha\tau)^{d/2}} \exp\left(-\frac{d_g^2(x,y)}{4\alpha\tau}\right)$$
 (67)

where $d_g(x,y)$ is the geodesic distance [16]. Since $q(x|y,\tau) \propto p(x|y,\tau)$ in an open field, $\nabla_x q(x|y,\tau) \propto -\nabla_x d_g^2(x,y)$, ensuring the planned path follows the shortest route, even around obstacles.

For large τ , $q(x|y,\tau)$ emphasizes global topological connectivity over local geometric details, often producing paths superior to the shortest. The transition matrix $P_{\tau} = P_1^{\tau}$ has spectral decomposition:

$$P_{\tau} = Q\Lambda^{\tau}Q^{T} \tag{68}$$

where Q is orthogonal, and $\Lambda = \operatorname{diag}(\lambda_1, \dots, \lambda_n)$ contains eigenvalues $0 \le \lambda_i \le 1$ [56]. For large τ , dominant eigenvalues $(\lambda_i \approx 1)$ amplify global connectivity:

$$p(x|y,\tau) = [P_{\tau}]_{xy} \approx \sum_{i:\lambda_i \approx 1} \lambda_i^{\tau} Q_i(x) Q_i(y)$$
(69)

The eigenvectors Q_i corresponding to $\lambda_i \approx 1$ encode topological features, such as major pathways and passage connectivity, invariant to deformations like stretching [56]. Thus, $\nabla_x q(x|y,\tau)$ prioritizes routes with high connectivity (e.g., wider corridors or multiple paths), potentially deviating from the shortest path to favor robust, flexible trajectories [10]. For example, in a maze, the planned path may choose a longer but more reliable route through a well-connected passage, avoiding narrow or risky shortcuts.

This topological focus aligns with neuroscience observations, where hippocampal place cells encode abstract connectivity (e.g., room layouts) over precise metrics, enabling robust navigation in complex environments [11, 9, 3]. Paths driven by topology are often "better" than the shortest, as they prioritize accessibility and adaptability, reflecting cognitive strategies in spatial tasks.

The planned path's alignment with the shortest path for small τ and its topological robustness for large τ mirror the hippocampus's ability to balance precision and global structure, ensuring efficient navigation across diverse environments.

G.4 Boundary Layer Effects

Near obstacles, the gradient field $\nabla_x q(x|y,\tau)$ aligns parallel to boundaries, preventing collisions. This behavior arises from the reflecting boundary condition of the heat equation governing $p(x|y,\tau)$, with $\partial p(x|y,\tau)/\partial n=0$ on $\partial \Omega_{\text{obstacles}}$. The condition ensures the normal component of $\nabla_x p(x|y,\tau)$ vanishes at boundaries, creating a flow parallel to obstacles, with strength modulated by the goal's position and time scale τ .

G.5 Diffraction-Like Patterns

Diffraction, in the context of our navigation framework, refers to the bending of the gradient field $\nabla_x q(x|y,\tau)$ around obstacles, such as corners or narrow passages, analogous to how waves bend around edges in optics or acoustics. This phenomenon arises because the transition probability $p(x|y,\tau)$, governed by the heat equation, diffuses probability mass through available paths, concentrating gradients toward openings like passages or around corners. These diffraction-like patterns guide trajectories through high-gradient regions, ensuring efficient navigation in complex environments.

The recursive nature of the random walk generates diffraction-like patterns:

$$p(x|y, \tau + 1) = \sum_{z} p(x|z, \tau) p(z|y, 1)$$
(70)

This convolution reflects the diffusion of probability from y to x through intermediate points z, allowing $p(x|y,\tau)$ to "bend" around obstacles by accumulating contributions from multiple paths [15]. Near a corner or passage, the heat kernel explores paths that wrap around the obstacle, creating a gradient field that converges toward navigable routes.

For a narrow passage of width w, the transition probability approximates the geodesic distance $d_g(x,y)$ through the passage. For small to moderate τ :

$$q(x|y,\tau) \approx \exp\left(-\frac{d_g^2(x,y)}{4\alpha\tau}\right)$$
 (71)

The gradient is:

$$\nabla_x q(x|y,\tau) \approx -\frac{1}{2\alpha\tau} \exp\left(-\frac{d_g^2(x,y)}{4\alpha\tau}\right) \nabla_x d_g^2(x,y)$$
 (72)

Since $\nabla_x d_g^2(x,y)$ points along the geodesic path through the passage, $\nabla_x q(x|y,\tau)$ converges toward the opening, creating a funneling effect. The width of this convergence region scales with $\sqrt{\tau}$, as the heat kernel's diffusion spreads over a distance proportional to $\sqrt{4\alpha\tau}$. For small τ , the gradient tightly focuses on the passage, ensuring precise navigation. For large τ , the broader diffusion enables early detection of passages from greater distances, smoothing trajectories by integrating multiple nearby paths.

At sharp corners, diffraction-like patterns emerge similarly. The heat kernel accumulates probability along paths bending around the corner, creating a curved gradient field that guides the agent past the obstacle. Unlike passages, corners lack a single geodesic path, so the gradient field spreads, resembling optical diffraction patterns where intensity peaks near edges. This spreading is more pronounced for large τ , as the heat kernel explores longer, circuitous routes.

These patterns ensure robust navigation in complex environments. The funneling effect through passages and curved trajectories around corners mirror hippocampal navigation, where place cells

encode paths that navigate mazes or avoid obstacles [11, 9]. The τ -dependent scaling reflects the hippocampus's ability to adjust focus, balancing precision for small τ with global awareness for large τ , enhancing adaptability in varied spatial contexts [16].

G.6 Topological Invariance

Topological invariance ensures that navigation behavior, driven by the gradient field $\nabla_x q(x|y,\tau)$, remains consistent under continuous deformations (e.g., stretching, bending) that preserve the environment's connectivity, such as the number of passages or loops. This robustness mirrors the hippocampus's cognitive maps, which encode abstract spatial relationships invariant to physical distortions [3].

The transition probability $p(x|y, \tau)$ arises from a random walk on a discrete grid, with one-step transition matrix P_1 . For time τ , the transition matrix is:

$$P_{\tau} = P_1^{\tau} \tag{73}$$

with spectral decomposition:

$$P_{\tau} = Q\Lambda^{\tau}Q^{T} \tag{74}$$

where Q is orthogonal $(Q^TQ = I)$, and $\Lambda = \operatorname{diag}(\lambda_1, \ldots, \lambda_n)$ contains eigenvalues $0 \le \lambda_i \le 1$ [56]. The eigenvectors Q_i and eigenvalues λ_i encode topological features, such as connectivity (number of connected components), bottlenecks (e.g., passages), and cycles. The second eigenvalue (λ_2) quantifies global connectivity, with higher values indicating tighter bottlenecks, while eigenvectors capture local structures like loops [15].

For large τ , dominant eigenvalues ($\lambda_i \approx 1$) amplify global topological features:

$$P_{\tau} \approx \sum_{i:\lambda_i \approx 1} \lambda_i^{\tau} Q_i Q_i^T \tag{75}$$

emphasizing coarse connectivity (e.g., major pathways between regions). The eigenvectors corresponding to $\lambda_i \approx 1$ represent low-frequency modes that partition the environment into connected regions, invariant to metric distortions like stretching a corridor. This is because Q_i depends on the graph's adjacency structure, not precise distances, ensuring stability under deformations that preserve connectivity [56]. For example, stretching a corridor adjusts transition probabilities proportionally, but the eigenvectors Q_i retain the same connectivity patterns, as they reflect the graph's topology.

The gradient field is $\nabla_x q(x|y,\tau)$, where $p(x|y,\tau) = [P_{\tau}]_{xy}$. Since P_{τ} 's eigenvectors are topologically invariant, $\nabla_x q(x|y,\tau)$ preserves navigation paths (e.g., through passages) across equivalent environments. For large τ , the dominance of $\lambda_i \approx 1$ ensures that $\nabla_x q(x|y,\tau)$ prioritizes global connectivity, guiding trajectories along major routes regardless of local geometry.

This invariance aligns with hippocampal navigation, where place cells encode connectivity (e.g., room layouts) over precise metrics, enabling robust path planning under environmental changes [9]. The focus on topology for large τ , driven by eigenvectors tied to $\lambda_i \approx 1$, reflects cognitive maps that emphasize structural relationships, as observed in rodent navigation [11, 3].

G.7 Hippocampal Preplay and Shortcut Detection

Our path planning framework discovers novel shortcuts using only local exploration, mirroring hippocampal preplay, where neural sequences predict unexplored paths [7, 8]. This process relies on efficient computation of transition probabilities and adaptive scale selection, without memorizing past successful paths.

The one-step transition matrix P_1 defines probabilities of moving between adjacent grid points. Multi-step transition matrices $P_{\tau} = P_1^{\tau}$ for scales $\tau = 2^k$ (k = 1, 2, ..., K) are computed via matrix squaring:

$$P_2 = P_1^2, \quad P_4 = P_2^2, \quad P_8 = P_4^2, \quad \dots, \quad P_{2^k} = P_{2^{k-1}}^2$$
 (76)

requiring $\log_2 \tau$ operations.

The path planning algorithm adaptively selects $\tau^* = \arg\max_t \|\nabla_x q(x|y,\tau)\|$, guiding trajectories via $\nabla_x q(x|y,\tau^*)$. For regions A and B connected by an unexplored passage, local exploration

defines P_1 . Matrix squaring yields P_{τ} , and adaptive τ -selection ensures $q(a|b,\tau^*)>0$, discovering the shortcut without prior traversal.

This process—matrix squaring to compute P_{τ} and adaptive τ -selection—embodies hippocampal preplay, predicting connectivity from local data, akin to place cell sequence pre-activation. The framework's efficiency and reliance on P_1 enhance its biological plausibility, supporting robust navigation.

H Theta Phase Modeling

We propose a local definition of theta phase that embeds spatial adjacency within the phase angle, capturing the observed precession range of 270° to 90° while enabling navigation across a tessellated environment. This formulation leverages an external 8 Hz theta rhythm for temporal pacing, where t denotes real time during navigation, distinct from the random walk time scale τ , and aligns with distributed hippocampal computation without requiring individual cells to explicitly encode spatial centers.

H.1 From Inner Product Angle to Theta Phase

The emergent sparsity from non-negativity and orthogonality constraints endows each cell in $h(x,\tau)$ a localized place field, so that each cell i has a place field centered at μ_i . The place fields of all the cells tile the whole environment. At the current position x(t) at real time t, the adjacency between x(t) and each μ_i is measured by the angle between $h(x(t),\tau)$ and $h(\mu_i,\tau)$, and this angle can then be embedded as the theta phase.

Specifically, define the adjacency of the current position x(t) to the field center of the *i*-th place cell at scale τ :

$$a_i(t) = \langle h(x(t), \tau), h(\mu_i, \tau) \rangle \tag{77}$$

where $h(x(t), \tau) \in \mathbb{R}^n$ is the population embedding (firing rates of n place cells) at position x(t), and $h(\mu_i, \tau)$ represents the embedding at the center μ_i , where

$$h_i(\mu_i, \tau) = \max_{\tau \in \mathcal{X}} h_i(x, \tau) \tag{78}$$

We interpret $h(\mu_i, \tau) = [w_{1i}, w_{2i}, \dots, w_{ni}]$ as the recurrent connection weights from the population to neuron i, such that:

$$a_i(t) = \sum_{j=1}^{n} h_j(x(t), \tau) w_{ji}$$
(79)

representing the net recurrent input to neuron i. This corresponds to the normalized transition probability $q(x(t)|\mu_i, \tau)$, and in open fields (Appendix E), it follows a Gaussian profile:

$$a_i(t) = \exp\left(-\frac{\|x(t) - \mu_i\|^2}{\tau}\right) \tag{80}$$

The theta phase is defined as:

$$\phi_i(x(t), \tau) = \pi + \operatorname{sgn}\left(\frac{da_i(t)}{dt}\right) \cdot \left(\frac{\pi}{2} - \arccos(a_i(t))\right)$$
(81)

where: $\frac{da_i(t)}{dt} = \sum_{j=1}^n \frac{dh_j(x(t),\tau)}{dt} w_{ji}$ is the temporal rate of change of this recurrent input, reflecting movement toward (> 0) or away (< 0) from μ_i as encoded by the weights. $\arccos(a_i(t)) \in [0,\pi/2]$ is the angle between $h(x(t),\tau)$ and $h(\mu_i,\tau)$, and this angle is embedded into the phase, possibly by sinusoidal inhibitory mechanism [64, 65].

The phase progresses as:

- 1. At field entry $(a_i \to 0, \frac{da_i}{dt} > 0)$: $\phi_i = \frac{3\pi}{2}$ (270°), late phase.
- 2. At field center ($a_i=1, \frac{da_i}{dt}=0$): $\phi_i=\pi$ (180°), mid-phase.
- 3. At field exit ($a_i \to 0$, $\frac{da_i}{dt} < 0$): $\phi_i = \frac{\pi}{2}$ (90°), early phase.

This does not require neuron i to encode μ_i explicitly; rather, $h(\mu_i, \tau)$ emerges as a distributed representation within recurrent connectivity, learned through experience and computed collectively, with ϕ_i storing $a_i(t)$ as an angular signal.

H.2 Implications for Navigation

This phase also enables navigation by tessellating the environment with $\{\mu_i\}$. Each μ_i anchors a tile, and ϕ_i signals adjacency: 180° indicates proximity to μ_i , while 270° and 90° mark entry and exit. For a goal at μ_g , navigating toward $\phi_g = \pi$ (max a_g) guides movement, with $\{\phi_i\}$ identifying intermediate μ_i to approach μ_g . This discretizes continuous navigation into a phase-driven sequence over $\{\mu_i\}$, complementing our gradient-based approach with a recurrently-driven strategy. During rest, iterating ϕ_i at 8 Hz drives replay, with $\frac{da_i}{dt}$ reflecting virtual transitions, unifying coding and navigation.

By tessellating the environment with $\{\mu_i\}$ and guiding navigation via $\{\phi_i\}$, this approach unifies spatial representation, temporal coding, and path planning within the embedding framework, offering an elegant, biologically plausible mechanism rooted in recurrent connectivity and the random walk model.

I Grid Cells Integration

I.1 Grid Cells as a Conformal, Multi-Scale Basis

While place cells form the primary focus of our model, we establish a connection to grid cells in the medial entorhinal cortex. Grid cells exhibit hexagonal firing patterns across multiple scales [66, 9] and are thought to provide a universal metric that feeds into the more adaptive place cell system.

Let g(x) be the vector formed by the firing rates of the population of grid cells at position x. The high dimensional g(x) forms an embedding of 2D position x. The grid cell representation g(x) provides an effective basis for learning the transformation matrix $W(\tau)$ to produce place cell embeddings $h(x,\tau)=W(\tau)g(x)$ that approximate g(x). Its effectiveness stems from conformal, isotropic, and multi-scale properties of grid cells, aligning with the hippocampus's multi-resolution encoding [22].

While place cell representation $h(x,\tau)$ preserves the proximity within the environment, g(x) preserves the local distance and has the property of conformal isometry.

Conformal isometry [67] is defined as:

$$||g(x+dx) - g(x)|| = s||dx|| + o(||dx||)$$
(82)

where s is a scaling factor, and $o(\|dx\|)$ denotes higher-order terms vanishing as $\|dx\| \to 0$, ensuring local distance and angle preservation and isotropic scaling.

s plays the role of metric. Grid cell population consists of multiple modules, with each module corresponding to a metric s, mirroring the multi-scale property of place cells.

I.2 Implementation of Grid-to-Place Transformation

The place cell embedding is modeled as $h(x,\tau) = W(\tau)g(x)$, or more generally $h(x,\tau) = \text{Norm}(\text{ReLU}(W(\tau)g(x) + b(\tau)))$ where the transformation matrix $W(\tau)$ is learned to optimize the proximity-preserving property:

$$\langle W(\tau)g(x), W(\tau)g(y)\rangle \approx q(y|x,\tau)$$
 (83)

The matrix $W(\tau)$ serves a dual purpose: (1) it captures the appropriate spatial scale by weighting the contribution of each grid cell module to match the place cell scale τ , and (2) it adjusts for deviations from isotropy and isometry in grid cells. g(x) facilitates fast adaptation to the environment in that learning $W(\tau)$ may be more efficient than learning $h(x,\tau)$ directly.

To further enhance adaptability, we can backpropagate gradients to g(x). This allows the grid cell representation to deform in response to environmental constraints, simulating the dynamic remapping observed in biological grid cells [68].

This formulation provides a computational account of how place field embeddings adapt to environmental constraints based on a flexible metric provided by grid cells. The transformation matrix $W(\tau)$, combined with the learnable parameters of g(x), effectively functions as a computational cognitive map, translating the multi-scale grid metric into place representations that reflect both the appropriate spatial scale and deviations from isometry and isotropy.

J Implementation Details

All the models were trained on a single NVIDIA A6000 GPU for 2,000 iterations with a learning rate of 0.001. For fine-tuning, we use 50 iterations with a 0.0005 learning rate. All models contain 500 cells. For batch size, we learn all combinations of different locations in the field for each iteration. The training time is less than 5 minutes.

We evaluate planning using two metrics: (1) Success, whether the agent reaches within 1 unit of the goal, and (2) Success weighted by inverse Path Length (SPL) [28]. For an agent's path $[x_1, \ldots, x_T]$ with initial geodesic distance d_i for episode i (as computed by baseline algorithms), we compute:

$$l_i = \sum_{\tau=2}^{T} \|x_{\tau} - x_{\tau-1}\|_2 \tag{84}$$

Then SPL for episode i is defined as:

$$SPL_i = Success_i \cdot \frac{d_i}{l_i} \tag{85}$$

We report SPL as the average of SPL_i across all episodes. In Table 1, we treat Bug algorithm with oracle as the baseline and in Table 2, our method serves as the baseline.

K Additional Experiment Results

K.1 Additional Path Planning Results in Four-Room Environment

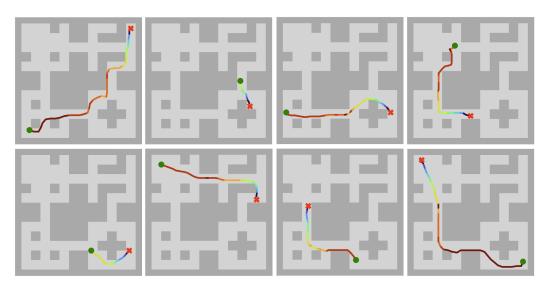


Figure 3: Successful navigation across eight randomly sampled start-goal configurations in the four-room environment.

To further demonstrate the robustness of our path planning framework, here we provide additional navigation trials in the four-room environment with randomly sampled start and goal locations.

Figure 3 shows eight scenarios spanning diverse spatial configurations. In all cases, our gradient-based path planning with adaptive scale selection successfully generated smooth, obstacle-avoiding trajectories that reached the goal. These results, combined with the quantitative evaluation in Table 1 showing 100% success rate across 50 random trials, demonstrate that our model achieves reliable navigation across diverse start-goal configurations even in complex multi-room environments.

K.2 Remapping

We evaluated our model's ability to adapt to bidirectional environmental modifications in the four-room environment (our most complex tested scenario). In addition to the obstacle removal experiment shown in Figure 2 (D), we conducted the complementary scenario: initially, rooms were connected allowing direct northward navigation; we then added an obstacle blocking this direct pathway and fine-tuned the model with 100 iterations at learning rate 0.01. The agent successfully adapted, discovering alternative routes through adjacent rooms to reach the target (see Figure 4). Together, these experiments demonstrate topology-dependent remapping – a well-documented phenomenon where place fields reorganize in response to changes in spatial connectivity. By both removing obstacles (creating novel shortcuts) and adding obstacles (blocking direct paths), we show that minimal fine-tuning enables the position embeddings to adapt to environmental changes that alter topological connectivity, whether exploiting new connections or navigating around new barriers. This bidirectional adaptability aligns with hippocampal responses to environmental modifications.

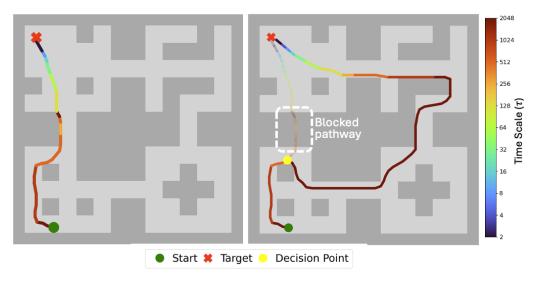


Figure 4: Remapping and path adaptation in response to added obstacles in the four-room environment.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The claims in the abstract and introduction (Section 1) reflect the contributions and scope of the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes].

Justification: We explicitly discuss limitations of our approach in Appendix B.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We provide the full set of assumptions and complete proofs for our theoretical results in the Appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Section 3 specifies all training and evaluation details. We also include link to our experiment code to ensure reproduceability.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived
 well by the reviewers: Making the paper reproducible is important, regardless of
 whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
- (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Our code is publicly available at https://github.com/mingluzhao/Place_Cell.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Section 3 specifies all training and evaluation details.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We provide experiment results with statistical significance in Section 3.

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We mention the compute usage and training details in the Section 3 and the Appendix J.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: Our research fully conforms to the NeurIPS Code of Ethics. We conduct purely algorithmic simulations with no human data or private information.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: There is no societal impact of the work performed.

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our work presents a computational algorithm with no high-risk data or models; safeguards are not applicable.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with
 necessary safeguards to allow for controlled use of the model, for example by requiring
 that users adhere to usage guidelines or restrictions to access the model or implementing
 safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: Our paper does not use existing assets.

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the
 package should be provided. For popular datasets, paperswithcode.com/datasets
 has curated licenses for some datasets. Their licensing guide can help determine the
 license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We do not propose new datasets or pretrained models; our work focuses on an algorithmic method.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No human subjects or crowdsourced data were involved in this research.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: No human subjects research was conducted.

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [No]

Justification: No large language models were used as a component of the core methodology of this research.

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.