

Geometric Principle of Spatial Remapping in the Hippocampal Formation

Daniel Zur

Geometric Principle of Spatial Remapping in the Hippocampal Formation

Research Thesis

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Medical Sciences

Daniel Zur

Submitted to the Senate
of the Technion — Israel Institute of Technology
Nisan 5786 Haifa April 2026

The Research Thesis Was Done in the Faculty of Medicine Under the Supervision of Dr. Hadas Benisty and Prof. Dori Derdikman.

The author of this thesis states that the research, including the collection, processing and presentation of data, addressing and comparing to previous research, etc., was done entirely in an honest way, as expected from scientific research that is conducted according to the ethical standards of the academic world. Also, reporting the research and its results in this thesis was done in an honest and complete manner, according to the same standards.

The generous financial help of the Technion is gratefully acknowledged.

Contents

List of Tables and Figures

Abstract	1
Abbreviations and Notations	3
1 Introduction	5
The Cognitive Map: Encoding Spatial Information	5
The Phenomenon of Remapping: Adaptation to Novel Environments	6
From Single-Cell Analysis to Population-Level Dynamics	6
Thesis Objectives: An Organizing Principle for Spatial Transformation	7
2 Experimental Framework	9
2.1 Experimental Setup	9
2.2 Diverse functional spatial cell types	10
3 Decoding Spatial Position from Population Dynamics	13
3.1 A recurrent neural network for position decoding	13
3.2 Low-dimensional, room-invariant latent representations	15
4 Geometric Principles of Population Remapping	17
4.1 Neural activity maps continuously between different spatial contexts	17
4.2 Coordinate transformation possesses affine properties	19
4.3 A unified geometric transformation governs remapping across diverse neural subpopulations	21
4.4 Single-cell tuning shifts exhibit heterogeneity compared to the global transformation	23
4.5 Spatial reflections are anchored to environmental geometry	24
5 Methods	25
5.1 Experimental Details and Preprocessing	25
5.1.1 Data Acquisition Pipeline	25
5.1.2 Behavioral Tracking and Neural Recordings	25
5.1.3 Single Cell Analysis	26

5.2	Position Decoding Model	28
5.3	Geometric Transformation Analysis	31
5.4	Comparison of NN model and PV analysis	34
5.5	Transformation consistency across neural subpopulations	34
5.6	Circular regression analysis	34
5.7	Single-cell conformity analysis	35
5.8	Reflection axis alignment with room geometry	35
5.9	General Methods	35
A	Supplementary Information	37
	Hebrew Abstract	i

List of Tables and Figures

2.1	Neuropixels recording setup.	9
2.2	Experimental A–B–A' paradigm.	10
2.3	Spatial cell types in MEC and SUBC.	11
2.4	Allocentric border cell examples.	11
3.1	Neural network position decoder architecture.	13
3.2	Decoder accuracy and robustness.	14
3.3	Latent space room invariance.	15
4.1	Continuous population remapping.	18
4.2	Affine structure of spatial remapping.	20
4.3	NN and PV affine comparison.	20
4.4	Remapping across neural subpopulations.	22
4.5	Single-cell conformity to global transformations.	23
4.6	Environmental anchoring of spatial reflections.	24
A.1	Recording regions by animal	37
A.2	Gallery of cross-room spatial transformations.	38
A.3	Single-session overview of tracking quality and behavior.	39
A.4	Unit inclusion for population analyses.	40
A.5	Cross-validation scheme.	40

Abstract

The remapping of spatial firing fields of individual cells in the hippocampal formation across environments is a well-studied phenomenon. Yet, the underlying organizing principle for this remapping remains unknown. While CA1 exhibits random remapping, regions such as the subicular complex (SUBC) and medial entorhinal cortex (MEC) demonstrate more structured transformations per cell.

Traditional analysis methods for remapping address the phenomenon by estimating spatial rate maps of cells across different environments, thereby treating neurons as independent units. Although effective for highly spatially tuned cells, this approach is inherently ineffective for untuned cells, and more importantly, ignores the coordinated network-level organization of the neuronal population.

Here, we propose a novel computational framework for studying representational remapping at the population level. Our architecture does not rely on prior models or knowledge and is based on nonlinear, time-dependent data-driven models trained to decode position from temporal sequences of ensemble activity. The system consists of two primary components: a recurrent encoder that maps population activity across all environments into a shared latent space, and a feedforward decoder head, one per environment, that estimates the mouse’s position from latent trajectories.

With this unique architecture, population-level remapping is derived by inferring how ensemble dynamics from one environment are interpreted by the decoder of another. We apply this analysis to arenas with diverse geometries to address the following questions: What are the characteristics of the learned cross-environment transformations? Are these transformations smooth or structured?

Notably, despite the decoder’s complexity, we find that a single low-dimensional affine map accurately captures the transformations across rooms. This principle remains consistent across all room geometries, recorded brain regions (SUBC and MEC), and among both spatially tuned and untuned sub-populations.

Altogether, these results demonstrate that the spatial population code is structured and governed by a simple transformation. This finding offers an organizing principle for spatial memory and establishes a foundation for decoding position in novel environments using learned transformation rules.

Abbreviations and Notations

Abbreviations

ABC	: Allocentric border cell
CA1	: Cornu Ammonis 1
EBC	: Egocentric border cell
HD	: Head direction (cell)
LFP	: Local field potential
LSTM	: Long short-term memory
MEC	: Medial entorhinal cortex
MLP	: Multi-layer perceptron
MSE	: Mean squared error
NN	: Neural network
PCA	: Principal component analysis
PRE	: Presubiculum
PV	: Population vector
ReLU	: Rectified linear unit
RMSE	: Root mean squared error
RV	: Rayleigh vector
SI	: Spatial information (metric)
SI cell	: Spatially modulated cell
SUB	: Subiculum
SUBC	: Subicular complex
SVM	: Support vector machine

Mathematical Symbols

$\mathbf{A}$	: Affine transformation matrix
$\mathbf{A}_{\text{lin}}$	: Linear component of affine transform
c	: Context (room) index
d^*, θ^*	: Preferred distance and angle (or direction)
$f(\cdot)$	: Shared temporal encoder function
$g_c(\cdot)$	: Context-specific decoder head
L	: Latent dimensionality
N	: Number of sessions (sample size) or total recorded neurons
$\mathbf{P}$	: Position trajectory matrix
$\hat{\mathbf{P}}$	: Predicted position trajectory matrix
$\mathbf{p}_t, \hat{\mathbf{p}}_t$	: True and predicted position vectors at time t
$\mathbf{R}$	: Time-binned population activity matrix
$\mathbf{r}_t$	: Population activity vector at time t
r	: Firing rate of a single cell (e.g., spatial or angular bin)
r_{peak}	: Peak firing rate
R_{decoder}^2	: Decoder accuracy (coefficient of determination)
R_{affinity}^2	: Affinity score (coefficient of determination)
$\mathbf{S}$	: Symmetric positive-definite scaling matrix
s	: Spike count (e.g., in a specific bin)
T	: Sequence length
$\mathbf{t}$	: Translation vector
$\mathbf{U}$	: Orthonormal matrix (rotation or reflection)
V	: Number of spatial bins
x, y	: Spatial coordinates (cm)
$\mathbf{Z}$	: Latent state matrix
$\Delta\theta_i$	: Per-cell angular shift between contexts
θ_{rot}	: Rotation angle
θ_{ref}	: Reflection axis angle
θ_{mode}	: Circular peak mode of single-cell shifts
τ	: Occupancy time in a specific bin (s)

Chapter 1

Introduction

The Cognitive Map: Encoding Spatial Information

The ability to navigate through space and remember the layout of environments is fundamental to survival. Over the past half-century, research on the hippocampal formation has established that the brain constructs an internal representation of space, known as a *cognitive map*, that enables flexible spatial behavior (O’Keefe and Nadel, 1978; Redish, 1999; Derdikman et al., 2014; Morris et al., 2024). The cornerstone of this map is the place cell: a hippocampal neuron that fires selectively at specific locations, providing an explicit neural code for position (O’Keefe and Dostrovsky, 1971).

The medial entorhinal cortex (MEC) harbors a diverse set of spatially tuned cell types. Grid cells fire at the vertices of a hexagonal lattice tiling the environment, providing a metric coordinate system (Hafting et al., 2005), and are organized into discrete modules with distinct spatial scales and orientations (Stensola et al., 2012). Allocentric border cells (ABCs) encode proximity to environmental boundaries (Solstad et al., 2008), egocentric border cells (EBCs) encode boundary distance and angle in a head-centered frame (Gofman et al., 2019), head direction (HD) cells signal allocentric facing direction (Ranck, 1985; Taube, 2007), speed cells scale linearly with running velocity (Kropff et al., 2015), and high spatial information (SI) cells exhibit broad positional modulation. Conjunctive cells that combine multiple spatial variables complete this repertoire (Sargolini et al., 2006; Moser et al., 2008).

The subicular complex (SUBC), comprising the subiculum (SUB) and presubiculum (PRE), occupies a pivotal anatomical position between the hippocampus and the entorhinal cortex. SUBC neurons exhibit positional coding (Sharp and Green, 1994), boundary-vector tuning (Lever et al., 2009), HD selectivity (Boccara et al., 2010), and task-dependent mixed selectivity (Ledergerber et al., 2021). Together, MEC and SUBC provide the core building blocks of the hippocampal formation’s cognitive map.

The Phenomenon of Remapping: Adaptation to Novel Environments

How does the cognitive map adapt to a new environment? When an animal moves between distinct domains, the neural representation reorganizes, a phenomenon termed *remapping* (Muller and Kubie, 1987; Colgin et al., 2008; Fenton, 2024). Two principal forms of remapping have been described: In *global remapping*, the spatial code is entirely reorganized: ensembles of active cells change and firing fields relocate, yielding apparently independent maps (Leutgeb et al., 2005). In *rate remapping*, firing fields persist but peak rates change, modulating the map without restructuring it (Leutgeb et al., 2005, 2004).

Remapping extends beyond the hippocampus. Grid cells in the MEC undergo coherent realignment by shifting phase, orientation, or scale during environmental transitions (Fyhn et al., 2007), and independent grid modules remap autonomously, providing a combinatorial basis for distinct hippocampal maps (Stensola et al., 2012; Kang et al., 2025). Environmental geometry strongly shapes this process: enclosure shape can anchor grid orientation (Krupic et al., 2015), rescale grid spacing (Barry et al., 2007), fragment grid patterns (Derdikman et al., 2009), and determine place field locations (O’Keefe and Burgess, 1996). Notably, the MEC establishes spatial maps upon first exposure to a novel environment, suggesting that pre-existing circuit architecture constrains new representations (Grieves et al., 2024).

The SUBC occupies a distinctive position in this landscape. Unlike the hippocampus, where maps differ markedly across environments, SUBC neurons produce more schematic, stable representations. This similar spatial framework across enclosures (Sharp, 2006; Lee et al., 2018) is consistent with a downstream integrator that translates hippocampal codes into a generalized format.

Despite decades of study, the rules governing how the population code transforms between environments remain poorly understood. Whether cross-environment transitions reflect unconstrained reorganization or obey systematic geometric rules remains a central open question (Fenton, 2024).

From Single-Cell Analysis to Population-Level Dynamics

The study of spatial navigation has historically centered on single-neuron analysis. Place, grid, and HD cells were identified by characterizing tuning curves and their responses to environmental manipulations. While foundational, these approaches cannot capture coordinated, network-wide reorganization during environmental transitions (Rule et al., 2020): single-neuron noise and variability obscure the systematic changes underlying map adaptation.

MEC and SUBC are organized as highly recurrent networks whose dynamics are best described by continuous attractors and low-dimensional neural manifolds (Wills et al., 2005; Burak and Fiete, 2009; Peyrache et al., 2017; Agmon and Burak, 2020; Gardner et al., 2022). Joint activity is constrained to a low-dimensional subspace

governed by network architecture; individual neurons cannot be understood in isolation.

A population-geometric perspective reveals that apparently random single-cell remapping often reflects coherent transformations of an underlying representational space (Aery Jones et al., 2024; Martín-Sánchez et al., 2025; Navas-Olive et al., 2025). Analyzing the full population demonstrates how the neural manifold shifts, rotates, or scales across environments, while a rigid internal structure is preserved across contexts.

Decoder-based population modeling goes beyond position readout: it reveals stable representational alignment across environments by overcoming single-cell noise (Rubin et al., 2019; Lee et al., 2025). This shift to population-level analysis is essential for uncovering the principles that govern cross-environmental generalization in the cognitive map.

Thesis Objectives: An Organizing Principle for Spatial Transformation

This thesis asks whether population-level remapping between environments obeys a principled geometric organization. We hypothesize that spatial-code transformations between rooms follow a low-dimensional family of geometric operations.

We recorded hundreds of neurons simultaneously in the MEC and SUBC using Neuropixels probes (Jun et al., 2017; Steinmetz et al., 2021) while mice navigated multiple rooms of varying shapes. We decoded position from population activity using a shared temporal encoder with room-specific decoder heads. Applying one room’s decoder to another room’s latent representation yielded continuous spatial correspondences.

We found that these correspondences were well described by low-dimensional affine transformations, which consist of rotation, reflection, scaling, shear, and translation. Three additional findings refine this perspective. First, this affine organization was consistent across both brain regions (MEC, SUBC) and functional subpopulations (ABCs, EBCs, HD cells, speed cells, and SI cells). This reveals a unified geometric framework, in contrast to the independent, subpopulation-specific transformations observed in hippocampal Cornu Ammonis 1 (CA1) (Navas-Olive et al., 2025). Second, reflection axes, when present, are systematically aligned with the orientation of the shared wall between rooms, thereby linking internal coordinate transformations to environmental geometry. Third, while the population as a whole obeyed this rigid geometric rule, the tuning shifts of individual neurons remained highly variable. This confirms that the organizing principle emerges strictly at the population level and is invisible to single-cell analysis.

These findings reveal a global organizing principle of representational alignment (Sucholutsky et al., 2023; Lee et al., 2025) in the hippocampal-formation cognitive map, providing empirical evidence for the theoretically predicted encoder-decoder remapping mechanism (Martín-Sánchez et al., 2025).

Chapter 2

Experimental Framework

2.1 Experimental Setup

To probe population-level remapping across distinct spatial contexts, we analyzed large-scale neuronal activity from freely moving mice. This dataset was acquired by the Derdikman Lab, led by Shai Abramson, using Neuropixels probes (Fig. 2.1) implanted in the medial entorhinal cortex (MEC), subiculum (SUB), and presubiculum (PRE).




Figure 2.1: Neuropixels recording setup. (a) Schematic of a mouse implanted with a Neuropixels probe for chronic extracellular recording during free exploration. (b) Coronal brain section showing reconstructed probe trajectories. Neuropixels 2.0 multi-shank probes are shown in blue and Neuropixels 1.0 single-shank probes in black; animal numbers are indicated next to each probe. (c) Representative extracellular voltage traces from adjacent channels showing simultaneous spiking activity (scale bars, 10 ms and 100 μ V). (d) Distribution of recording sites along the probe shanks spanning cortex, subiculum (SUB, cyan), presubiculum (PRE, cyan), and medial entorhinal cortex (MEC, red); Regions not included in the analysis due to small numbers of cells are marked in gray. From left to right, the schematics correspond to mice 1140 and 1150 (same coordinates), 1120, 7394, 1130, and 1170 (four shanks).

Mice explored a custom-built two-room arena (Fig. 2.2a) during daily sessions consisting of three consecutive 20-min epochs: exploration of room A, transition to room B, and a return to room A (A–B–A'; Fig. 2.2b). Across recording sessions, we systematically varied room geometries, generating varied pairs of environments with distinct combinations of straight and curved boundaries (Fig. 2.2c). Overall, we recorded 71 sessions from 6 animals, yielding a total of 9,009 single units in all sessions, out of which 3,855 units were in MEC, 3,070 units were in SUB, and 965 units were in PRE. To augment the number of cells for population analyses, we grouped the SUB and PRE together into the SUBC region (Table A.1 and Methods, Section 5.1.3).




Figure 2.2: **Experimental A–B–A' paradigm.** (a) Custom 120×120 cm two-room arena with a motorized tether system for unrestricted exploration. (b) Behavioral protocol: mice explored room A (0–20 min), then room B (20–40 min), and then returned to room A' (40–60 min). Traces show representative trajectories in room A (blue), room B (orange), and room A' (green). (c) Room geometries comprising diverse combinations of straight and curved boundaries in symmetric and asymmetric layouts.

2.2 Diverse functional spatial cell types

We first classified neurons by their tuning properties in each room to ask whether functional cell identity was preserved across environments. Units were classified according to quantitative tuning metrics, with significance assessed against shuffle-based null distributions (see Methods, Section 5.1.3). In both MEC and SUBC, we identified different functional classes, including allocentric border cells (ABCs; Fig. 2.3a), egocentric border cells (EBCs; Fig. 2.3b), head-direction (HD) cells (Fig. 2.3c), spatially selective cells with high spatial information (Fig. 2.3d), and speed-modulated cells (Fig. 2.3e). Many neurons exhibited conjunctive tuning, belonging to more than one functional class. Overlap among cells classified as one of the directionally-oriented classes (ABCs, EBCs, HD cells) revealed substantial conjunctive coding in both MEC




Figure 2.3: Spatial cell types in MEC and SUBC. (a)–(e) Representative examples of functional cell types. Each panel shows a spatial rate map (left) and a relevant tuning curve (right). (a) Allocentric border cell (ABC) with boundary-aligned firing fields (Rayleigh vector (RV) ≥ 0.3). (b) Egocentric border cell (EBC) tuned to boundary distance and angle in a head-centered frame ($RV \geq 0.3$). (c) Head-direction (HD) cell with allocentric preferred direction ($RV \geq 0.3$). (d) Spatially modulated cell with high spatial information ($SI > 1$). (e) Speed cell with linear relationship between firing rate and running speed ($R^2 > 0.6$). (f)–(h) Population summaries and overlap among functional cell types in the SUBC and MEC. Classification required metric significance ($P < 0.05$) relative to a shuffled null distribution. (f) Venn diagrams showing population overlap of ABC, EBC, and HD properties in SUBC and MEC. (g) Proportions of cells classified into each functional category by region. (h) An UpSet plot quantifying conjunctive coding (multi-modal tuning) across the population.

and SUBC regions (Fig. 2.3f–h). Across the broader population, speed-modulated cells were most prevalent in both MEC and SUBC, followed by ABCs, EBCs, and HD cells. Specifically, we identified $n = 158$ exclusively significant ABCs in MEC and 151 in SUBC; 85 EBCs in MEC and 38 in SUBC; and 253 high-spatial-information cells in MEC and 197 in SUBC. These cell classes enabled comparison of firing-rate maps across room geometries, as illustrated by representative ABCs (Fig. 2.4).




Figure 2.4: Allocentric border cell examples. Rate maps for rooms A and B across sessions with differing room geometries. Peak firing rates (Hz) are indicated.

Chapter 3

Decoding Spatial Position from Population Dynamics

3.1 A recurrent neural network for position decoding

To investigate population remapping without imposing linear or time-invariant assumptions, we developed an unbiased analysis framework that captures nonlinear and temporally structured relationships across contexts. The pipeline comprised two components: a recurrent temporal encoder based on Long Short-Term Memory (LSTM) networks (Hochreiter and Schmidhuber, 1997) that maps population activity across all rooms to a shared latent space, and a room-specific multilayer perceptron (MLP) decoder that transforms the latent coordinates to physical positions within each room (see Methods, Section 5.2, Fig. 3.1). Note that this is reminiscent of previous models of the hippocampus as a recurrent encoder of sensory and other inputs with emergent place and time cells (Benna and Fusi, 2021; Wang et al., 2024, 2025; Yu et al., 2026).




Figure 3.1: **Neural network position decoder architecture.** A shared temporal encoder maps population activity sequences from all rooms (A, B, A') to a common low-dimensional latent space. Distinct decoder heads map this latent state to position in each room's coordinate frame. The encoder was trained on activity traces from rooms A and B, while each decoder head was trained solely using data from its respective room (blue: A; orange: B; green: A').

We designed this architecture so it would: 1) provide a data-driven mapping from population dynamics to position; 2) disentangle the shared temporal code from room-specific spatial information; 3) allow us to explore how population activity related to a particular room would be decoded in any of the three rooms (A, B, and A'). First, we evaluated the performance of our pipeline in decoding the animal's position from population dynamics, on a per-room basis (models were trained using 10-fold cross-validation with 10 independent repetitions; Fig. 3.2a). Example trajectories demonstrate reliable tracking of position over time and accurate reconstruction of spatial paths (Fig. 3.2b). We evaluated the robustness of our approach by examining the effect of temporal integration. Instantaneous baseline models struggled to reliably map population activity to spatial position: linear regression performed poorly (median $R^2_{\text{decoder}} = 0.13$) and a MLP network provided only a moderate improvement (0.40) compared to the full temporally integrated architecture (0.61) (Fig. 3.2c). The accuracy of our full architecture improved rapidly as the temporal context of the sequence increased up to a few seconds before plateauing (Fig. 3.2d), leading us to select a 5-second temporal window processed by a recurrent LSTM block for the final model. To ensure our spatial mappings were




Figure 3.2: Decoder accuracy and robustness. Boxplots show median and interquartile range (IQR); dots represent individual sessions. **(a)** 10-fold cross-validation scheme. A continuous recording session is partitioned into temporally separated training (80%), validation (10%), and test (10%) segments. **(b)** Example: ground truth coordinates (black) and the decoder predictions (blue; mean $\pm$ standard deviation across repetitions) during 5 minutes. Right: Spatial 2D view of the trajectory color-coded by time. **(c)** Comparison to baseline models ($N = 71$ sessions). **(d)** Decoder accuracy as a function of sequence length (mean $\pm$ s.e.m., $N = 28$ sessions). **(e)** Ensemble decoding stability. Comparison of individual model repetitions (gray) with the ensemble average (blue) ($N = 71$ sessions). **(f)** Decoding accuracy (R^2_{decoder}) and position decoding error (RMSE, cm) for Rooms A, B, and A' ($N = 71$ sessions).

robust to noise, we adopted an ensemble decoding strategy by averaging predictions across multiple independent model repetitions. This approach consistently improved decoding accuracy (by ~ 0.10) and reduced position error (by ~ 2.6 cm) relative to single model runs, providing stable position estimates (Fig. 3.2e). Finally, we evaluated overall performance across all sessions ($N = 71$). The ensemble model achieved high decoding accuracy (median $R^2_{\text{decoder}} = 0.71$, inter-quartile range [IQR] [0.46–0.82] for room A; 0.72 [0.55–0.83] for room B; 0.75 [0.48–0.82] for room A'; Fig. 3.2f). Correspondingly, position decoding errors were low (16.4 cm [12.2–21.4] for room A; 14.4 cm [11.8–19.9] for room B; 15.7 cm [11.8–19.5] for room A'). These results confirm that the full architecture reliably maps population activity to spatial position.

3.2 Low-dimensional, room-invariant latent representations

To determine whether the shared temporal encoder produces a compressed representation for neural dynamics without introducing contextual separation, we quantified the variability of raw neural activity traces relative to the latent representation using Principal Component Analysis (PCA). For the raw data, 71 ± 37 principal components were required to capture 95% of its variance. In contrast, the latent representation required a much smaller number of 9 ± 2 components to capture the variability of the 64-dimensional latent representation (mean $\pm$ s.d. across sessions; Fig. 3.3a). Crucially, a linear Support Vector Machine (SVM) classifier based on the first three principal components of this compressed latent representation yielded near-chance misclassification rates (latent: 0.39 ± 0.01 ; mean $\pm$ s.d.; pairwise chance = 0.5), unable to discriminate environmental context better than when based on the raw activity (0.40 ± 0.01 ; Fig. 3.3b). Overall, these results suggest that the shared latent space effectively reduced the dimension of the representation while maintaining room-invariance.




Figure 3.3: Latent space room invariance. (a) Comparison of cumulative variance explained by principal components for the latent representation (**left**) versus raw population activity (**right**). The latent representation requires substantially fewer components to reach 95% variance ($N = 71$ sessions). (b) First three principal components colored by room identity (blue: A; orange: B; green: A') for latent (**left**) and raw activity (**right**) from a representative session, demonstrating that neither space allows linear separation of room identity.

Chapter 4

Geometric Principles of Population Remapping

4.1 Neural activity maps continuously between different spatial contexts

We next used our architecture to determine how the encoding of neural dynamics related to one environment is transformed to represent another. We passed the neural activity recorded in a source room through the temporal encoder, thereby generating latent trajectories in the shared space. These latent coordinates were then decoded using both the source-room and target-room decoder heads. In this way, the same underlying neural activity produced two temporally matched spatial sequences: the estimated position of the mouse in the source room and the corresponding position decoded in the target room’s frame of reference. This approach allowed us to define, at each time point, a pair of corresponding spatial coordinates that link the two environments (see Methods, Section 5.3; Fig. 4.1a).

To confirm that the inferred geometric structure is robust and independent of the decoding method, we compared our neural network (NN)-derived mappings against conventional population-vector (PV) correlation analysis (see Methods, Section 5.3). For a systematic comparison of these approaches, we quantified the preservation of local spatial structure using a graph-based smoothness score (Eq. 5.2; see Methods, Section 5.3). This metric evaluates whether neighboring bins in the source room map to contiguous locations in the target room. Comparing across identical room geometries ($A \rightarrow A'$) and distinct rooms ($A \rightarrow B$), we found that NN mappings were the smoothest (mean smoothness score $\pm$ s.d.: 0.98 ± 0.03 for $A \rightarrow A'$, 0.88 ± 0.08 for $A \rightarrow B$). PV mappings were significantly less smooth (0.68 ± 0.17 for $A \rightarrow A'$, 0.49 ± 0.13 for $A \rightarrow B$; $P < 0.001$), though both methods significantly outperformed a random null baseline (Fig. 4.1b). Importantly, the coordinate transformations obtained by the NN and PV methods were well correlated, and the correlation increased with decoding accuracy

(Fig. 4.1c). This high agreement confirms that the inferred geometric structure is not an artifact of the decoding method.

We visualized the spatial structure of this remapping by coloring the source-room decoded positions by their horizontal (x) or vertical (y) coordinates, and projecting these colors onto the corresponding decoded positions in the target room (Fig. 4.1d). The obtained mappings showed that spatial representations undergo structured geometric transformations. Although the analysis was highly nonlinear and data-driven, the resulting geometric transformations were usually simple, as quantified in the following sections. For comparison, we performed an equivalent visual analysis using the PV-based correspondences (Fig. 4.1e). While PV-based maps exhibited lower resolution and greater scatter, they showed similar patterns (see more examples in Supplementary Fig. A.2).




Figure 4.1: Continuous population remapping. (a) Schematic of the cross-decoding framework. Neural population activity maps through a shared temporal encoder to a common latent state, which is then read out by room-specific decoders to infer cross-room correspondences. (b) Smoothness quantification. Comparison of spatial mapping smoothness (Laplacian-based score, see Methods, Section 5.3) across three conditions: Null (random bin mapping), PV (population-vector mapping), and NN (decoder-based mapping), for $A \rightarrow B$ (blue) and $A \rightarrow A'$ (green). (c) Geometric agreement between NN and PV mappings quantified by Pearson correlation between target coordinates inferred by the NN decoder and PV analysis, plotted versus mean decoder accuracy (average R^2_{decoder} across rooms) for mappings $A \rightarrow B$ (blue) and $A \rightarrow A'$ (green). Triangular markers indicate the example sessions displayed in (d) and (e). (d) Visualization of decoder-based remapping. Source-room positions colored by horizontal (x) or vertical (y) coordinate (top), mapped to decoded target-room locations (bottom), revealing smooth, geometry-preserving gradients. (e) Comparison with PV remapping. Equivalent gradient analysis using PV correlations. Mappings are noisier but exhibit similar geometric structure.

4.2 Coordinate transformation possesses affine properties

Visual observation of the obtained coordinate mappings (Fig. 4.1d) revealed a large proportion of transformations that could be described by a rotation, stretch, and reflection of the coordinates. We therefore asked to what extent the transformation in both methods could be described as affine, namely comprising translation, rotation, possible reflection, scaling, and shear (Fig. 4.2a). To quantify this, we fitted a global affine transformation to the cross-room position correspondences. We defined an affinity score as the coefficient of determination (R_{affinity}^2 , ranging between 0 and 100%) of this fit (see Methods, Section 5.3). This score evaluates how well the complex spatial remapping is captured by simple geometric operations, with higher values indicating that the transformation between environments is predominantly affine. We applied a polar decomposition to the linear component of these fitted affine matrices to independently isolate anisotropic scaling from orthogonal rotations or reflections, allowing us to extract explicit rotation (θ_{rot}) and reflection-axis (θ_{ref}) angles. Representative examples illustrating how the remapping can be broken down into these sequential operations are shown in Fig. 4.2b. For instance, the mapping between environments could be decomposed into an orthogonal rotation coupled with scaling and translation, or alternatively, into a reflection across a specific spatial axis combined with scaling and translation. This demonstrates how complex population-level remapping can be reduced to a compact set of basic affine operations.

The joint distribution of decoding accuracy and affinity scores (Fig. 4.2c) effectively partitioned the data into distinct regimes. First, in the low-accuracy regime, characterized by low decoding accuracy ($R_{\text{decoder}}^2 < 0.3$), the affinity scores were predominantly low, reflecting the inherent impossibility of recovering precise geometric transformations from noisy position estimates. Second, once decoding became reliable ($R_{\text{decoder}}^2 \geq 0.3$), the distribution bifurcated: 58 sessions (83%) reached high accuracy, of which a minority remained in the low-affinity quadrant, while the majority (81%; $N = 47$) converged into a distinct high-affinity cluster ($R_{\text{affinity}}^2 \geq 0.5$). This separation indicates that whenever the spatial code is robust, the remapping between rooms is predominantly affine. In subsequent analyses in this section, we quantified the affinity of these transformations for all reliably decoded sessions ($R_{\text{decoder}}^2 \geq 0.3$; $N = 58$ sessions for $A \rightarrow B$, $N = 54$ sessions for $A \rightarrow A'$ and $B \rightarrow A'$). To characterize the specific parameters of the transformations (such as rotation and reflection angles), we restricted our subsequent analyses to the high-confidence regime ($R_{\text{decoder}}^2 \geq 0.3$ and $R_{\text{affinity}}^2 \geq 0.5$) where the geometric structure is well-defined and can be reliably estimated.




Figure 4.2: **Affine structure of spatial remapping.** (a) Schematic of 2D affine transformation components. (b) Representative examples of affine decompositions for spatial mappings. The source room representation (dashed blue) is sequentially transformed (magenta) to align with the target room (dashed orange). Examples demonstrate both a rotation-dominated transformation (top) and a reflection-dominated transformation (bottom), each incorporating scaling and translation. (c) Joint distribution of decoder accuracy and affinity score for A→B mapping. Marginal histograms show distributions. Dashed lines indicate empirical inclusion thresholds ($R^2_{\text{decoder}} \geq 0.3$; $R^2_{\text{affinity}} \geq 0.5$) derived from the natural bifurcation of the data.

Affinity scores for mapping between rooms (A→B mappings) were significantly higher for NN than those derived from PV-based correspondences (0.62 [0.54–0.77] for NN vs. 0.27 [0.17–0.38] for PV; Fig. 4.3a), reflecting the NN decoder’s superior noise robustness. Similarly, for returning to the same room (A→A’), the affinity scores were higher for the NN (0.93 [0.87–0.95] for NN vs. 0.64 [0.40–0.78] for PV). Despite these differences in fit quality, both methods showed the same qualitative pattern: A→A’ mappings yielded higher affinity than A→B mappings. We then compared the affine parameters recovered by the two approaches.




Figure 4.3: **NN and PV affine comparison.** (a) Comparison of affinity scores between NN and PV methods for A→B and control A→A’ mappings. Boxplots show the median and IQR, with individual session scores overlaid. Dashed line indicates inclusion threshold ($R^2_{\text{affinity}} \geq 0.5$ for NN). (b) Agreement of transformation rotation angles between NN and PV methods. Dashed line: identity; solid line indicates circular-linear regression ($P < 0.001$, slope 1.02, RMSE 26.8°, $N = 23$). (c) Agreement of reflection-axis angles between NN and PV methods. Dashed line: identity; solid line indicates circular-linear regression ($P < 0.001$, slope 0.98, RMSE 10.0°, $N = 10$).

For both NN and PV-based correspondences, we decomposed the affine transformations into rotation, reflection, scaling, shear, and translation components (as depicted in Fig. 4.2b). For this comparison, we included only sessions in which both methods agreed on whether the affine transformation was a rotation or a reflection (70% agreement: $N = 23$ rotation, $N = 10$ reflection; see Methods, Section 5.4). We found that rotation angles (θ_{rot}) were strongly correlated between methods (circular-linear regression: $R^2 = 0.61$, $P < 0.001$; Fig. 4.3b), indicating that both approaches recover similar remapping rotations. For sessions with reflection components, reflection-axis angles (θ_{ref}) showed good alignment between methods ($R^2 = 0.90$, $P < 0.001$; Fig. 4.3c). This cross-method agreement confirms that the geometric transformations resolved by the NN decoder reflect fundamental population-level remapping that remains detectable, despite increased noise, using more traditional PV-based analysis.

4.3 A unified geometric transformation governs remapping across diverse neural subpopulations

We next asked whether the geometric structure of remapping discussed above reflects a global network property or is specific to particular functional modules. We retrained our room decoding architecture (Fig. 3.1) separately for each subpopulation defined by brain region (SUBC, MEC) and cell type (ABC, EBC, HD, SI, Speed, see Methods, Section 5.5). SUBC cells enabled decoding accuracy similar to that of the full population, albeit with fewer cells (0.68 [0.49–0.76] for SUBC vs. 0.72 [0.48–0.82] for the full population, Fig. 4.4a). The MEC cells, however, yielded somewhat lower values (0.50 [0.23–0.70]). Across functional cell types, ABC and SI cells yielded the best decoding (0.57 [0.35–0.69] for ABC, 0.57 [0.34–0.67] for SI), followed by Speed (0.50 [0.23–0.71]), EBC (0.47 [0.22–0.62]), and HD cells (0.38 [0.16–0.55]).

Moreover, affinity scores for $A \rightarrow B$ remained stable across subpopulations (median R^2_{affinity} was 0.62 [0.54–0.77] for the full population, 0.66 [0.55–0.76] for SUBC and 0.63 [0.43–0.74] for MEC, and 0.49–0.58 across ABC, EBC, SI, and Speed, with HD lower at 0.27 [0.24–0.56], Fig. 4.4b). Control $A \rightarrow A'$ mappings yielded high affinity across all groups (medians 0.81–0.93; Fig. 4.4c). Most sessions exceeded the affinity threshold ($R^2_{\text{affinity}} \geq 0.5$) for the full population (81%), for SUBC (81%), and for MEC (68%), with moderate rates for specific cell types (50–59%), indicating that the geometric structure of cross-room transformations is a robust property shared throughout diverse cell types and brain regions. The affine parameters learned per subpopulation were consistent with those of the full population, both for rotation ($R^2 = 0.68$, $P < 0.001$; Fig. 4.4d) and reflection ($R^2 = 0.97$, $P < 0.001$; Fig. 4.4e).




Figure 4.4: **Remapping across neural subpopulations.** (a) Decoder accuracy among neural subpopulations: full population (‘All’), regions (MEC, SUBC), and functional types (ABC, EBC, HD, SI, Speed). Dashed lines mark the median for the full population as baseline (magenta) and the decoder accuracy threshold ($R^2_{\text{decoder}} \geq 0.3$; black). (b) Affinity scores across neural subpopulations for A→B mapping. (c) Affinity scores across neural subpopulations for control A→A' mapping. Dashed lines mark the full-population median (magenta) and the affinity score threshold ($R^2_{\text{affinity}} \geq 0.5$; black). (d) Consistency of transformation rotation angles across subpopulations. Dashed line: identity; solid line indicates significant circular-linear regression ($P < 0.001$, slope 0.98, RMSE 23.7°, $N = 131$). (e) Consistency of transformation reflection angles across subpopulations. Dashed line: identity; solid line indicates significant circular-linear regression ($P < 0.001$, slope 1.02, RMSE 5.7°, $N = 54$).

4.4 Single-cell tuning shifts exhibit heterogeneity compared to the global transformation

Finally, we tested whether single-cell tuning shifts align with the population-level transformation. To this end, we compared each cell's ABC directional shift with the session-level rotation angle (θ_{rot}), restricting the analysis to sessions best described by a rotation ($N = 26$ for $A \rightarrow B$ and $N = 49$ for $A \rightarrow A'$; see examples in Fig. 4.5a). Defining alignment as an angular difference of $\leq 30^\circ$ between the predominant single-cell shift angle (the circular mode of individual shifts, θ_{mode}) and the global affine transform (θ_{rot}), we found that 9 of 26 $A \rightarrow B$ sessions (34.6%) and 32 of 49 $A \rightarrow A'$ sessions (65.3%) showed alignment. Generally, we conclude that the transformation of single cells was heterogeneous in comparison to the global transformation in the $A \rightarrow B$ case, and less so when returning to the same room (Fig. 4.5b).




Figure 4.5: Single-cell conformity to global transformations. (a) Per-session distributions of ABC angular shifts ($\Delta\theta_i$) in rotation sessions (left: $A \rightarrow B$; right: $A \rightarrow A'$). Vertical lines indicate the global rotation angle (θ_{rot} , magenta) and the circular peak mode of single-cell shifts (θ_{mode} , red). (b) Session-level conformity. Scatter plots relating the single-cell shift mode (θ_{mode}) to the global rotation angle (θ_{rot}) for $A \rightarrow B$ (top; $N = 26$ sessions) and $A \rightarrow A'$ (bottom; $N = 49$ sessions). Dashed lines: identity.

4.5 Spatial reflections are anchored to environmental geometry

In a subset of sessions, the best-fitting affine transformation included a reflection ($N = 15$ sessions; see Methods, Section 5.8). To test whether reflections are geometrically constrained, we compared the orientations of reflection axes to the direction of the wall shared between the rooms. We defined alignment as an angular difference $\leq 30^\circ$ relative to the shared wall (see Methods, Section 5.8). Reflection axes aligned systematically with room geometry in 14 out of 15 sessions (93.3%; Fig. 4.6). Horizontal and diagonal geometry sessions deviated by at most 15.9° from the expected axis ($N = 9$), while vertical geometry sessions ($N = 6$) ranged up to 33.5° , with 5 of 6 falling within 21.1° . This systematic alignment indicates that reflections are anchored to environmental boundaries rather than constituting arbitrary transformations.




Figure 4.6: **Environmental anchoring of spatial reflections.** Alignment of reflection axes with room geometry. Reflection axes are grouped by the orientation of the shared wall (vertical, horizontal, or diagonal ($\pm 45^\circ$)). Vertical lines indicate the expected axis direction (red).

Chapter 5

Methods

5.1 Experimental Details and Preprocessing

5.1.1 Data Acquisition Pipeline

Electrophysiological recordings were performed in six adult male C57BL/6J mice using Neuropixels 1.0 and 2.0 probes (Jun et al., 2017; Steinmetz et al., 2021). Probes were secured in a custom 3D-printed encapsulation for stable chronic implantation (Fig. 2.1a). The SUBC and MEC were targeted based on preoperatively planned trajectories (Peters, 2022; Wang et al., 2020) (Fig. 2.1b,d). All procedures were approved by the Animal Care and Use Committee of the Technion – Israel Institute of Technology. To enable unrestricted movement in complex-shaped arenas, a motorized XY stage was suspended above a 120×120 cm experimental arena to keep the tethered recording cable directly above the animal and prevent entanglement (Fig. 2.2a). Animal rotations were accommodated by thin twisted tether wires, which were manually untwisted after each trial. A stage-mounted camera (15 Hz) and a closed-loop controller used DeepLabCut (Mathis et al., 2018) to track the animal’s position (tail–body junction) and head direction (ears and nose). Absolute coordinates were obtained by composing the stage’s physical position with the camera’s within-frame tracking. When needed, the XY stage could be manually controlled via a MATLAB interface to bias animals toward undersampled regions of the arena gently. A secondary fixed overhead camera (8.3 Hz) provided independent monitoring for experimental oversight and validation. We synchronized video streams and electrophysiology using external pulses delivered every 15 ± 2 s.

5.1.2 Behavioral Tracking and Neural Recordings

Behavioral variables (position and head direction) were extracted from DeepLabCut outputs (Mathis et al., 2018). Brief tracking gaps where the detection likelihood fell below 30% were linearly interpolated. We excluded time points where running speed exceeded 20 cm/s (implausible under tethering) or where an interpolated gap spanned

≥ 3.0 s or ≥ 15.0 cm. Sessions were required to have sufficient behavioral coverage of the arena for inclusion (Supplementary Fig. A.3).

Neural signals were recorded continuously at 30 kHz using a National Instruments system (NI PXIe-6374). For Neuropixels 1.0, local field potentials (LFPs) and spikes were processed separately; for Neuropixels 2.0, LFPs were computed from spike data via digital filtering. Spike sorting was performed with Kilosort 2.5 (Steinmetz et al., 2021), followed by automated curation with Bombcell (Fabre et al., 2023). We restricted downstream analyses to high-quality units: those labeled good by Kilosort and not flagged by Bombcell as noise. We further excluded units with < 500 total spikes or with mean firing rates outside the 0.1–50 Hz range, establishing the final units (Supplementary Fig. A.4). Behavioral traces and neural recordings were aligned to a common time base using the recorded synchronization events. Sessions were included only if they met quality-control criteria for minimum unit yield.

5.1.3 Single Cell Analysis

Rate maps

Spatial firing-rate maps (rate maps) were computed by discretizing the arena into 4×4 cm spatial bins and assigning each detected spike to the corresponding bin. Occupancy was estimated as the number of video frames in which the animal’s head position fell within a bin, multiplied by the frame period (40 ms), and firing rate was computed as spike count divided by occupancy time. Bins with < 0.2 s occupancy or < 2 spikes were excluded. Rate maps were smoothed with a 2D Gaussian kernel ($\sigma = 2$ bins).

Cell classification

Statistical significance for spatial metrics was assessed by comparing the observed scores against a null distribution generated via 1,000 random circular time shifts of the spike train relative to behavior.

Spatial information We calculated spatial information (SI; bits per spike) for each cell as

$$\text{SI} = \sum_i p_i r_i \log_2 \left(\frac{r_i}{\bar{r}} \right),$$

where r_i is the firing rate in spatial bin i , p_i is the occupancy probability of bin i , and $\bar{r}$ is the mean firing rate across bins. Neurons were classified as SI cells if $\text{SI} > 1$.

Head direction For each neuron, we computed a head-direction tuning curve by discretizing head direction into 6° bins. For bin j , firing rate was

$$r_j = \frac{s_j}{\tau_j},$$

where s_j is the number of spikes fired when head direction fell within bin j , and τ_j is the total time spent in bin j . The peak firing rate was $r_{\text{peak}} = \max_j(r_j)$.

Directionality was quantified using the mean resultant vector length (Rayleigh vector length) across all angular bins,

$$\text{RV} = \frac{\sqrt{\left(\sum_j r_j \cos \theta_j\right)^2 + \left(\sum_j r_j \sin \theta_j\right)^2}}{\sum_j r_j},$$

where θ_j is the center angle of bin j . The preferred head direction was

$$\theta^* = \text{atan2}\left(\sum_j r_j \sin \theta_j, \sum_j r_j \cos \theta_j\right).$$

Neurons were classified as HD cells if $\text{RV} > 0.3$ (Sanguinetti-Scheck and Brecht, 2020; Taube, 2007).

Allocentric border-vector cells For each cell and behavioral epoch, we computed at each time step (i) the distance from the animal’s head to the nearest arena border d , and (ii) the allocentric angle from the head to the nearest wall θ . Angles were discretized into 10° bins. Distances were discretized into non-uniform bins (0–2 cm, 2–5 cm, 5–15 cm, 15–25 cm, and 25–40 cm) to provide higher resolution near walls.

For each (distance bin k , angle bin m) pair, firing rate was

$$r_{k,m} = \frac{s_{k,m}}{\tau_{k,m}},$$

where $s_{k,m}$ is the number of spikes fired while the animal occupied bin (k, m) , and $\tau_{k,m}$ is the total time spent in that bin. Bins with $\tau_{k,m} < 0.2$ s were set to zero ($r_{k,m} = 0$), and the map was smoothed with a Gaussian kernel ($\sigma = 2$ bins). The preferred wall distance and angle were defined by

$$(d^*, \theta^*)_{\text{allo}} = \arg \max_{k,m} (r_{k,m}).$$

Neurons were classified as ABC cells if their Rayleigh score was > 0.3 .

Egocentric border-vector cells For each cell and each video frame, we computed the egocentric distance to the nearest wall as a function of wall bearing relative to the animal’s head direction. Egocentric angles θ_{ego} were discretized into 10° bins spanning -180° to 180° . For each time point, distances were obtained by ray-casting from the head position at angles θ_{ego} and taking the closest wall intersection, then discretized into non-uniform bins (0–2 cm, 2–5 cm, 5–15 cm, 15–25 cm, and 25–40 cm).

For each (distance bin k , angle bin m) pair, firing rate was

$$r_{k,m} = \frac{s_{k,m}}{\tau_{k,m}},$$

where $s_{k,m}$ is the number of spikes fired while the animal occupied bin (k, m) , and $\tau_{k,m}$ is the total time spent in that bin. Bins with $\tau_{k,m} < 0.2$ s were set to zero and the map was smoothed with a Gaussian kernel ($\sigma = 2$ bins). The preferred egocentric distance and angle were

$$(d^*, \theta^*)_{\text{ego}} = \arg \max_{k,m} (r_{k,m}).$$

Neurons were classified as EBCs if their Rayleigh score was > 0.3 .

Speed cells Speed modulation was assessed by relating firing rate to running speed estimated from tracking. Instantaneous speed (cm/s) was computed in 1-s non-overlapping bins from the position trace; missing values were linearly interpolated, and negative values were set to zero. Speed bins (1 cm/s width) were defined across the observed range. For each speed bin, the mean firing rate was computed as the mean spike rate (spikes/s) across all 1-s time bins whose speeds fell within the interval.

The relationship between mean firing rate and speed was quantified using ordinary least-squares linear regression, and R^2 was used as an index of speed modulation. Neurons were classified as speed cells if $R^2 > 0.6$.

5.2 Position Decoding Model

We developed a decoding model based on population dynamics to obtain a quantitative description of mapping across rooms. We chose a model class that is highly expressive and recurrent so that it could, in principle, capture arbitrary history-dependent nonlinear relationships between firing patterns and position within each room. This allowed us to test whether the between-rooms relationship remained simple even when the within-room decoding model was maximally general. Neural and behavioral time series were aligned to a common 20 Hz sampling grid, with population activity computed as the raw spike count within a 250 ms sliding window centered at each time step.

Notations. Let $t \in \{1, \dots, T\}$ be the relative time step within a single sequence, where T denotes sequence length. Let $\mathbf{R} \in \mathbb{R}^{N \times T}$ denote a sequence of population activity, where N is the number of cells, and the t -th column is a population vector $\mathbf{r}_t \in \mathbb{R}^N$. We take $\mathbf{P} \in \mathbb{R}^{2 \times T}$ to be the corresponding position trajectory, where $\mathbf{p}_t = (x_t, y_t)^\top$, the t -th column, contains the x and y coordinates. The coordinate origin was defined at the top-left corner of the arena.

Model. We trained a neural-network decoder to predict position from population activity (Fig. 3.1). The model factorizes into: (i) a shared temporal encoder f that maps neural activity to a low-dimensional latent trajectory, and (ii) room-specific decoder heads g_c (for room $c \in \{A, B, A'\}$) that map the shared latent trajectory to positions in the coordinate frame of each room. For a given sequence $\mathbf{R}$ sampled from room c :

$$\begin{aligned}\mathbf{Z} &= f(\mathbf{R}) \in \mathbb{R}^{L \times T}, \\ \hat{\mathbf{P}} &= g_c(\mathbf{Z}) \in \mathbb{R}^{2 \times T},\end{aligned}$$

where L is the latent dimensionality. The temporal encoder, f , is a unidirectional recurrent long short-term memory (LSTM) layer (512 hidden units) followed by a multi-layer perceptron (MLP) network (two hidden layers of 256 and 128 units, ReLU (rectified linear unit) activation), yielding a latent trajectory $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_T] \in \mathbb{R}^{L \times T}$ where $L = 64$. For each room c , we trained a dedicated model to decode position from the shared latent coordinates per time step:

$$\hat{\mathbf{p}}_t = g_c(\mathbf{z}_t), \quad \hat{\mathbf{p}}_t \in \mathbb{R}^2.$$

Each decoder head consisted of four hidden layers of 32, 32, 16, and 8 units, with ReLU activations and a final linear layer to produce the estimated position, $(\hat{x}_t, \hat{y}_t)$.

Training. Sequence samples were constructed from $T = 100$ consecutive timestamps (spanning 5 s) with a stride of one timestamp ($\Delta t = 50$ ms). Position coordinates were normalized to the $[0, 1]^2$ interval via min-max scaling relative to room boundaries.

We trained the model to minimize the mean squared error (MSE) between predicted and observed positions. To enforce robust temporal integration while heavily penalizing errors at later timesteps that possess maximal sequence context, we computed an exponentially weighted loss over all time bins t , where $w_t = \exp(\lambda t) / \sum_{s=1}^T \exp(\lambda s)$ with $\lambda = 0.01$. The sequence-weighted prediction loss was

$$\text{Loss} = \sum_{t=1}^T w_t \|\hat{\mathbf{p}}_t - \mathbf{p}_t\|_2^2. \quad (5.1)$$

We computed the loss over the full sequence $\hat{\mathbf{P}}$, whereas we evaluated performance using only the final prediction $\hat{\mathbf{p}}_T$. Training was performed in two stages to prevent the shared representation for the two visits to room A from biasing toward the repeated environment. In the first phase, the temporal encoder and room-specific heads for rooms A and B were trained jointly, excluding room A' . In the second phase, the shared modules were frozen, and only the decoder head for room A' was trained. Model weights were optimized using Adam with a learning rate of 1×10^{-3} , $p = 0.1$ dropout, and weight decay of 1×10^{-4} . Early stopping was triggered if validation performance did not improve for 50 consecutive epochs (maximum: 500 epochs).

Cross-Validation. We evaluated model performance using 10-fold cross-validation with 10 repetitions. For each room, we partitioned the data into 10 contiguous temporal segments. In each fold, one segment served as the test set, while the final one-ninth of each remaining segment was reserved for validation (Fig. 3.2a). The remaining data constituted the training set. We prevented data leakage across overlapping temporal windows by using only sequences fully contained within a single contiguous segment. To make the results robust to the relative temporal alignment between rooms, we repeated the entire 10-fold procedure 10 times, each time circularly shifting the fold assignments of room B relative to room A (Supplementary Fig. A.5). Overall, this yielded $10 \text{ folds} \times 10 \text{ repetitions} = 100$ trained models per session.

Evaluation. We evaluated test performance using only the final timestamp prediction ($\hat{\mathbf{p}}_T$) of each sample sequence, concatenating these predictions to reconstruct a continuous, full-length trajectory for each of the 10 cross-validation repetitions. We quantified decoding accuracy using two metrics, both computed as a pooled R^2 across x and y coordinates: (i) **Mean R^2** : The R^2 was calculated for each repetition individually, then averaged across the repetitions. (ii) **Ensemble R^2** : We first computed an “ensemble trajectory” (mean $\hat{\mathbf{P}}$) by averaging the predicted coordinates across the repetitions at each time step, then calculated the R^2 of this mean trajectory against the true trajectory $\mathbf{P}$. Unless otherwise stated, reported performance metrics refer to the ensemble R^2 and the corresponding Root Mean Squared Error (RMSE, in cm). Furthermore, all downstream geometric transformation analyses used the ensemble predictions to maximize stability.

Latent space dimensionality and separability

To characterize the structure of the learned representations, we extracted latent trajectories ($\mathbf{Z}$) by feeding the full dataset of population activity sequences through the shared temporal encoder, f . We quantified the dimensionality of the population code using Principal Component Analysis (PCA) on the concatenated latent vectors ($\mathbf{z}$), determining the number of principal components required to explain 95% of the total variance (mean $\pm$ s.d. across models). For comparison, we applied identical PCA to the raw spike count vectors ($\mathbf{r}$, model input).

To quantify geometric overlap between room representations, we assessed the linear separability of room identity. For each room pair, we trained a linear Support Vector Machine (SVM) to predict room identity from the first three principal components. To test for separability rather than generalization, we evaluated classification accuracy on the full dataset (all samples). We reported the misclassification rate ($1 - \text{accuracy}$) as a measure of overlap. Chance level was defined as the majority-class baseline ($1 - \max(n_1, n_2)/(n_1 + n_2) \approx 0.5$). This analysis was repeated using the first three principal components of the raw neural activity.

Comparison with Baseline Models

We compared our full architecture against two instantaneous ($T = 1$; Fig. 3.2c) baselines: (i) a **linear regression baseline** mapping instantaneous population activity ($\mathbf{r}$) to spatial coordinates ($\mathbf{p}$); and (ii) a **multilayer perceptron baseline** with three hidden layers of 512, 256, and 32 units, and ReLU activation. These models test the extent to which position can be decoded from instantaneous activity using simple linear or complex nonlinear transformations, serving as controls for the full architecture’s temporal integration.

5.3 Geometric Transformation Analysis

We quantified the geometric structure of spatial remapping by computing position correspondences as described below between room coordinate frames. We evaluated these transformations for all ordered room pairs ($A \rightarrow A'$, $A \rightarrow B$, $B \rightarrow A'$).

Deriving spatial correspondences

To ensure robust estimates, we established cross-environment correspondences using two complementary approaches: a continuous coordinate mapping derived from the neural network (NN) decoder and a discrete spatial bin assignment based on population-vector (PV) correlation.

NN-based mapping. We utilized our trained model to infer the correspondence between spatial representations across rooms. For a given source room, we provided the decoder with neural activity recorded while the animal occupied that room, and first decoded the position in that room. We next presented the same activity pattern to the decoder while specifying the identity of a second, target room, thereby obtaining the position in the target room that was most consistent with that activity pattern under the learned model. This yielded a set of matched coordinate pairs ($\hat{\mathbf{p}}_{\text{source},t}$, $\hat{\mathbf{p}}_{\text{target},t}$) for each time step t (Fig. 4.1a).

PV-based mapping. For each spatial bin in each room (34×34 grid, occupancy ≥ 0.2 s), we computed the average firing rate of the population within that bin. We assigned each source-room spatial bin to the target-room bin that maximized the Pearson correlation between the population vectors of the two bins. Bins were considered matched if their maximum correlation exceeded 0.2. The strength of this mapping was evaluated by averaging PV correlations across all matched bin pairs.

Comparison of mapped coordinates

To directly compare the spatial transformations inferred by the NN decoder and PV analysis (Fig. 4.1d,e), we discretized the continuous NN decoder output to match the 34

$\times 34$ bin resolution of the PV analysis. For each source-room spatial bin, we identified all decoded positions falling within that bin, computed the centroid of their corresponding target-room coordinates, and assigned it to the nearest target spatial bin.

Smoothness. To quantify the local continuity of the mapping (Fig. 4.1b), we used the graph Laplacian quadratic form, also known as the graph Dirichlet energy. We built a graph based on the spatial bins of the source room and then measured the Dirichlet energy of the target coordinates on this graph, relative to the source coordinates. By construction, this measure is small if the mapping preserves the same local neighbors. Formally, we constructed a graph with V nodes representing the spatial bins in the source room, and each node is connected to its 8 adjacent spatial neighboring bins. Denote $\mathbf{A}$ as the adjacency matrix indicating which nodes are connected (edge value = 1) and which are not (=0), and $\mathbf{D}$ as the degree matrix, where $\mathbf{D}_{i,i} = \sum_j \mathbf{A}_{i,j}$. Following standard spectral graph theory (Chung, 1997), we computed the normalized graph Laplacian $\tilde{\mathbf{L}} = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$. The coordinate column vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^V$ are treated as graph signals on this graph (Shuman et al., 2013). We defined the energy changes for the horizontal and vertical coordinates as:

$$\Delta E_x = \mathbf{x}_{\text{target}}^\top \tilde{\mathbf{L}} \mathbf{x}_{\text{target}} - \mathbf{x}_{\text{source}}^\top \tilde{\mathbf{L}} \mathbf{x}_{\text{source}}, \quad \Delta E_y = \mathbf{y}_{\text{target}}^\top \tilde{\mathbf{L}} \mathbf{y}_{\text{target}} - \mathbf{y}_{\text{source}}^\top \tilde{\mathbf{L}} \mathbf{y}_{\text{source}}.$$

We combined these terms into a normalized smoothness score:

$$\text{Smoothness} = \exp \left(-\frac{\Delta E_x + \Delta E_y}{2V \|\tilde{\mathbf{L}}\|_F} \right), \quad (5.2)$$

where $\|\tilde{\mathbf{L}}\|_F$ is the Frobenius norm of the normalized Laplacian matrix. We used the denominator $2V \|\tilde{\mathbf{L}}\|_F$ as a per-session normalization term to place energy changes on a comparable scale across sessions with varying numbers of bins and graph structures. The exponential transform was used to obtain a bounded, interpretable score. Higher values indicate that the mapping introduces less additional local variation relative to the source coordinate field. This metric was calculated for both the NN and PV mappings. Only sessions with at least 100 valid bins were included. Statistical significance was assessed by comparing the observed smoothness score against a null distribution generated via 1,000 random bin-to-bin mappings. Scores were considered significant if they exceeded the 95th percentile of this null distribution (one-sided test, $P < 0.05$).

Affine Transformation estimation

We estimated the optimal global affine transformation mapping from the source room to the target room. We used least squares to fit an affine transformation to the position correspondences. We then decomposed the resulting mapping into interpretable components (rotation, reflection, stretching, and translation). The goodness of fit of the estimated transformation was quantified by the mean squared error (MSE) between

the affine-transformed source coordinates and the corresponding target coordinates, as well as by the coefficient of determination (R_{affinity}^2). We hereafter refer to this R_{affinity}^2 as the **affinity score**. High affinity scores indicate that the relationship between room representations is well-approximated by a global affine model.

Affine transformation decomposition. We decomposed the estimated affine transformations into explicit geometric components. In homogeneous coordinates, the transformation maps a point $\tilde{\mathbf{p}} = [x, y, 1]^\top$ to $\tilde{\mathbf{p}}' = \mathbf{A}\tilde{\mathbf{p}}$, where $\mathbf{A} \in \mathbb{R}^{3 \times 3}$ is defined as:

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{\text{lin}} & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{bmatrix} \quad (5.3)$$

$\mathbf{A}_{\text{lin}} \in \mathbb{R}^{2 \times 2}$ encodes linear transformations (rotation, reflection, scaling) and $\mathbf{t} \in \mathbb{R}^2$ represents translation.

Polar decomposition. We factored the linear component $\mathbf{A}_{\text{lin}}$ (Eq. 5.3) using the left polar decomposition:

$$\mathbf{A}_{\text{lin}} = \mathbf{S}\mathbf{U}$$

where $\mathbf{U}$ is an orthonormal matrix and $\mathbf{S}$ is a symmetric positive-definite scaling matrix. The left-polar form implies that scaling occurs along axes defined in the target frame. The determinant of $\mathbf{U}$ determined whether the transformation involved a rotation ($\det(\mathbf{U}) > 0$) or a reflection ($\det(\mathbf{U}) < 0$).

Rotation and reflection angles. For rotations ($\det(\mathbf{U}) > 0$), we computed the rotation angle:

$$\theta_{\text{rot}} = \text{atan2}(U_{21}, U_{11}) \quad (\text{wrapped to } [0^\circ, 360^\circ)).$$

For reflections ($\det(\mathbf{U}) < 0$), we identified the reflection axis $\mathbf{v} = (v_x, v_y)^\top$ as the eigenvector of $\mathbf{U}$ with eigenvalue $+1$. The axis orientation was computed as:

$$\theta_{\text{ref}} = \text{atan2}(v_y, v_x) \quad (\text{wrapped to } [0^\circ, 180^\circ)).$$

Anisotropic scaling. We decomposed the symmetric matrix $\mathbf{S}$ to obtain scaling factors $s_1 \geq s_2 > 0$ and their corresponding orthogonal principal axes. These factors quantify expansion or compression along specific directions in the target frame.

Significance of affinity scores

We compared the obtained affinity scores, R_{affinity}^2 , against a null distribution generated via 1,000 cyclic time-shifts of the target data relative to the source. This preserves temporal structure within each room while disrupting cross-room correspondence. Observed affinity scores were considered statistically significant if they exceeded the 95th

percentile of the null distribution ($P < 0.05$). Additionally, we applied a minimum threshold of $R_{\text{affinity}}^2 \geq 0.5$ to classify a transformation as structurally affine.

5.4 Comparison of NN model and PV analysis

We compared the transformation parameters obtained from PV analysis to those obtained from the NN model. We included sessions meeting decoder and affinity criteria ($R_{\text{decoder}}^2 \geq 0.3$, $R_{\text{affinity}}^2 \geq 0.5$) where both methods agreed on the transformation type (rotation or reflection) and yielded valid angle estimates (Fig. 4.3).

We extracted rotation angles (θ_{rot}), reflection angles (θ_{ref}), and scaling factors (s_1, s_2) from the affine models fitted by both methods. Orientation parameters were standardized to $[0^\circ, 360^\circ)$ for rotations and $[0^\circ, 180^\circ)$ for reflection. We assessed consistency using circular-linear regression (evaluating slope and circular R^2 ; see Circular regression analysis below) and quantified angular error using RMSE.

5.5 Transformation consistency across neural subpopulations

We retrained our model per subpopulation where neurons were partitioned by region (MEC, SUBC) or by functional class (ABC, EBC, HD, SI, Speed) as defined in single-cell analysis (Section 5.1.3). Cells with mixed selectivity contributed to multiple functional subgroups; subgroups with fewer than 10 cells were discarded. A majority of sessions achieved $R_{\text{decoder}}^2 \geq 0.3$. Angular consistency of the resulting transformations between subpopulations and the full ensemble was then evaluated (see Circular regression analysis below).

5.6 Circular regression analysis

We quantified the consistency of transformation parameters (rotation and reflection angles) across methods (PV vs. NN) and neural subpopulations, using circular regression framework as follows. We modeled the relationship as a directional fit $\theta_2 \approx a \cdot \theta_1 + b$ by minimizing the circular loss

$$\text{Loss}_{\text{circ}} = \sum [1 - \cos(\text{wrap}(\theta_2 - (a \cdot \theta_1 + b)))]$$

using the Nelder-Mead algorithm. For reflection axes (180° periodicity), angles were doubled prior to analysis, and the resulting intercept was halved. Given the nonlinear nature of this loss, statistical significance was assessed by comparing the circular R^2 against a null distribution generated via 1,000 random permutations of θ_2 across pairs.

5.7 Single-cell conformity analysis

To determine whether global geometric transformations emerge from coherent single-cell remapping, we analyzed directional tuning shifts in ABCs. We included rotation sessions that met the decoder and affinity criteria ($R_{\text{decoder}}^2 \geq 0.3$, $R_{\text{affinity}}^2 \geq 0.5$) and passed the tuning-prominence filter (see below). We computed the distribution of angular shifts for all ABCs ($n \geq 10$ cells per session). Let θ^* denote the preferred allocentric wall angle for an ABC (Section 5.1.3). The shift for cell i was

$$\Delta\theta_i = (\theta_{\text{target},i}^* - \theta_{\text{source},i}^*) \pmod{360^\circ}.$$

We assessed conformity and quantified the alignment between the session-level single-cell shift mode (θ_{mode}) and the global affine rotation (θ_{rot}). Define θ_{mode} as the circular peak of the smoothed $\Delta\theta_i$ histogram (15° bins, Gaussian $\sigma = 1$ bin). Sessions were classified as conforming if the angular separation between θ_{mode} and θ_{rot} was $\leq 30^\circ$. We evaluated conformity only in sessions exhibiting a clear population-level shift. Specifically, we required the peak prominence of the smoothed shift distribution to exceed its median baseline by a factor of ≥ 1 .

5.8 Reflection axis alignment with room geometry

For transformations characterized as reflections ($\det(\mathbf{U}) < 0$), we evaluated the alignment of the reflection axis θ_{ref} relative to the shared wall orientation. We included sessions meeting decoder and affinity criteria ($R_{\text{decoder}}^2 \geq 0.3$, $R_{\text{affinity}}^2 \geq 0.5$). We formally defined a reflection as geometrically aligned (or anchored) to the room geometry if the angular deviation between the reflection axis and the shared wall was $\leq 30^\circ$.

5.9 General Methods

Statistical analysis

We performed all analyses using Python 3.x with PyTorch (Paszke et al., 2019) for model training, NumPy and Pandas for data manipulation, and Weights & Biases (W&B) (Biewald, 2020) for experiment tracking. We performed statistical comparisons across rooms and conditions using cross-validated estimates with appropriate corrections for multiple comparisons. We computed standard errors of the mean (SEM) across cross-validation folds to quantify uncertainty in performance estimates.

Computational resources

Models were trained on NVIDIA GPUs. Typical training time for a single fold was approximately 5 ± 3 minutes depending on dataset size. All training runs were tracked using W&B, including hyperparameters, training curves, and performance metrics.

Appendix A

Supplementary Information

Table A.1: **Recording regions by animal.** Session counts per animal in SUBC, MEC, or both (SUBC+MEC).

Animal#	Region		
	SUBC	MEC	SUBC+MEC
1120		26	
1130	1		
1140	3		
1150			13
1170	24		
7394			4




Figure A.2: **Gallery of cross-room spatial transformations.** Comparison of global mapping structure inferred by neural network decoders (NN Mapping, center) and population vectors (PV Mapping, right) for representative sessions. Leftmost columns report session metrics: decoder accuracy (R^2_{decoder}) and affinity scores (R^2_{affinity}) for $A \rightarrow A'$ and $A \rightarrow B$. Within each method block, horizontal (x , left pair) and vertical (y , right pair) spatial gradients in the source room are projected onto the target room. NN mappings yield smooth, continuous gradient transformations (e.g., rotated, curved, or compressed bands), reflecting robust geometric correspondence across environments. PV mappings exhibit lower spatial precision, producing noisy coordinate assignments, yet preserve a similar global structure. Thus, while both methods capture related spatial transformations, NN decoding resolves the mapping with greater continuity and precision.




Figure A.3: Single-session overview of tracking quality and behavior. (a) Timeline of tracking states. The color strip encodes per-frame status: valid raw positions, short interpolated segments, long temporal gaps (≥ 3.0 s), long spatial jumps (≥ 15.0 cm), and missing coordinates. Traces show head-position X and Y . (b) Binary validity trace (1 = valid or short-interpolated, 0 = otherwise). Background shading indicates the experimental room segments (A: blue, B: orange, A': green). (c) Mean running speed (cm/s) in 10-s bins. (d) Distribution of binned speeds from (c). (e) Sine of head direction as a function of time (minutes). (f) Polar histogram of head-direction angles from (e). (g) Spatial trajectories across rooms. (h) Full trajectory (grey). Highlighted intervals mark extended tracking gaps (long-time or long-distance). (i) Spatial occupancy. Color scale is capped at the 95th percentile.




Figure A.4: **Unit inclusion for population analyses.**

Columns. All: initial sorted units ($N = 100\%$). Depth: units within targeted depth. Firing rate: mean rate within 0.1–50 Hz. Kilosort (included): labeled **good**. Bombcell (excluded): not labeled **noise**. Final: Kept (passed all criteria) vs. Dropped. **Metrics.** Intermediate Pass/Fail nodes reflect individual criteria, while the Final column reflects joint inclusion. All percentages are relative to N . **Colors.** Green and red ribbons trace units to the final Kept and Dropped sets.




Figure A.5: **Cross-validation scheme.** Schematic of the 10-fold cross-validation procedure with 10 repetitions. A continuous 60-minute recording session is partitioned into temporally separated training (blue, 80%), validation (gray, 10%), and test (red, 10%) segments. **(a)** The full schedule displays all fold-offset combinations. **(b)** A fold-wise view at a fixed repetition (offset) shows all 10 folds, with the held-out test window shifting sequentially over time. **(c)** A repetition-wise view at a fixed fold (fold 1) illustrates the same fold repeated across different offsets. The held-out window shifts through time across repetitions, minimizing bias from temporal autocorrelation.

Bibliography

- Emily A. Aery Jones, Lisa M. Giocomo, et al. Conserved hippocampal population geometry supports task generalization. *bioRxiv*, 2024. doi: 10.1101/2024.10.24.620127. Preprint.
- Harel Agmon and Yoram Burak. A theory of joint attractor dynamics in the hippocampus and the entorhinal cortex accounts for artificial remapping and grid cell over-dispersion. *eLife*, 9:e56894, 2020. doi: 10.7554/eLife.56894.
- Caswell Barry, Richard Hayman, Neil Burgess, and John O’Keefe. Experience-dependent rescaling of entorhinal grids. *Nature Neuroscience*, 10(6):682–684, 2007. doi: 10.1038/nn1905.
- Marcus K Benna and Stefano Fusi. Place cells may simply be memory cells: Memory compression leads to spatial tuning and history dependence. *Proceedings of the National Academy of Sciences*, 118(51):e2018422118, 2021.
- Lukas Biewald. Experiment tracking with weights and biases, 2020. URL <https://www.wandb.com/>. Software available from wandb.com.
- Charlotte N. Boccara, Francesca Sargolini, Vilde H. Thoresen, Trygve Solstad, Menno P. Witter, Edvard I. Moser, and May-Britt Moser. Grid cells in pre-and parasubiculum. *Nature Neuroscience*, 13(8):987–994, 2010. doi: 10.1038/nn.2596.
- Yoram Burak and Ila R. Fiete. Accurate path integration in continuous attractor network models of grid cells. *PLoS Computational Biology*, 5(2):e1000291, 2009. doi: 10.1371/journal.pcbi.1000291.
- Fan R. K. Chung. *Spectral Graph Theory*, volume 92 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, 1997.
- Laura Lee Colgin, Edvard I Moser, and May-Britt Moser. Understanding memory through hippocampal remapping. *Trends in neurosciences*, 31(9):469–477, 2008.
- 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. doi: 10.1038/nn.2396.

- Dori Derdikman, James J Knierim, et al. *Space, time and memory in the hippocampal formation*, volume 10. Springer, 2014.
- Julie M. J. Fabre, E. H. van Beest, A. J. Peters, M. Carandini, and K. D. Harris. Bomb-cell: automated curation and cell classification of spike-sorted electrophysiology data (version 1.0.0). Zenodo, 2023. URL <https://doi.org/10.5281/zenodo.8172822>.
- André A Fenton. Remapping revisited: how the hippocampus represents different spaces. *Nature Reviews Neuroscience*, 25(6):428–448, 2024.
- 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. doi: 10.1038/nature05601.
- Richard J. Gardner, Erik 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(7895):123–128, 2022. doi: 10.1038/s41586-021-04268-7.
- Xenia Gofman, Shanna Stewart, Nadav Nitzan, Charlotte L O’Brien, Elie Hwaun, and Dori Derdikman. Dissociation between postrhinal cortex and downstream parahippocampal regions in the representation of egocentric boundaries. *Current Biology*, 29(16):2751–2757, 2019. doi: 10.1016/j.cub.2019.07.007.
- Roddy M. Grieves, Liyu Lu, Salma Jedidi-Ayoub, V. Villette, Neil Burgess, and Caswell Barry. One-shot entorhinal maps enable flexible navigation in novel environments. *Nature*, 635(8038):399–406, 2024. doi: 10.1038/s41586-024-08034-3.
- 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. doi: 10.1038/nature03721.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. doi: 10.1162/neco.1997.9.8.1735. URL <https://doi.org/10.1162/neco.1997.9.8.1735>.
- James J. Jun, Nicholas A. Steinmetz, Joshua H. Siegle, Daniel J. Denman, Marius Bauza, Brian Barbarits, Albert K. Lee, Costas A. Anastassiou, Alexandra Andrei, Cengiz Aydin, Mladen Barbic, Timothy J. Blanche, Vincent Bonin, João Couto, Barundeb Dutta, Sergey L. Gratiy, Diego A. Gutnisky, Michael Häusser, Bill Karsh, Paul Ledochowitsch, Claudia Mora Lopez, Catalin Mitelut, Silke Musa, Michael Okun, Marius Pachitariu, Jan Putzeys, Philip D. Rich, Cyrille Rossant, Wei-lung Sun, Karel Svoboda, Matteo Carandini, and Kenneth D. Harris. Fully integrated silicon probes for high-density recording of neural activity. *Nature*, 551(7679):232–236, 2017. doi: 10.1038/nature24636. URL <https://doi.org/10.1038/nature24636>.

- L. Kang, B. Bhatt, and Loren M. Frank. Functional independence of entorhinal grid cell modules enables remapping in hippocampal place cells. *bioRxiv*, 2025. doi: 10.1101/2025.09.24.677985. Preprint.
- Emilio Kropff, James E Carmichael, May-Britt Moser, and Edvard I Moser. Speed cells in the medial entorhinal cortex. *Nature*, 523(7561):419–424, 2015.
- Julija Krupic, Marius Bauza, Stephen Burton, Caswell Barry, and John O’Keefe. Grid cell symmetry is shaped by environmental geometry. *Nature*, 518(7538):232–235, 2015. doi: 10.1038/nature14153.
- Debora Ledergerber, Claudia Battistin, Jan Sigurd Blackstad, Richard J Gardner, Menno P Witter, May-Britt Moser, Yasser Roudi, and Edvard I Moser. Task-dependent mixed selectivity in the subiculum. *Cell reports*, 35(8), 2021.
- Hyun-Woo Lee, Su-Min Lee, and Inah Lee. Neural firing patterns are more schematic and less sensitive to changes in background visual scenes in the subiculum than in the hippocampus. *Journal of Neuroscience*, 38(34):7392–7408, 2018. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.0156-18.2018. URL <https://www.jneurosci.org/content/38/34/7392>.
- J Quinn Lee, Alexandra T Keinath, Erica Cianfarano, and Mark P Brandon. Identifying representational structure in cal to benchmark theoretical models of cognitive mapping. *Neuron*, 113(2):307–320, 2025.
- Stefan Leutgeb, Jill K Leutgeb, Alessandro Treves, May-Britt Moser, and Edvard I Moser. Distinct ensemble codes in hippocampal areas ca3 and ca1. *Science*, 305(5688):1295–1298, 2004. doi: 10.1126/science.1100265.
- 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. doi: 10.1126/science.1114037.
- Colin Lever, Stephen Burton, Ali Jeewajee, John O’Keefe, and Neil Burgess. Boundary vector cells in the subiculum of the hippocampal formation. *Journal of Neuroscience*, 29(31):9771–9777, 2009. doi: 10.1523/JNEUROSCI.1319-09.2009.
- Guillermo Martín-Sánchez, Christian K Machens, and William F Podlaski. Three types of remapping with linear decoders: a population-geometric perspective. *PLOS Computational Biology*, 21(10):e1013545, 2025.
- Alexander Mathis, Pranav Mamidanna, Kevin M. Cury, Taiga Abe, Venkatesh N. Murthy, Mackenzie Weygandt Mathis, and Matthias Bethge. Deeplabcut: markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience*,

- 21(9):1281–1289, 2018. doi: 10.1038/s41593-018-0209-y. URL <https://doi.org/10.1038/s41593-018-0209-y>.
- Richard GM Morris, David G Amaral, Tim Bliss, Karen Duff, and John O’Keefe. *The hippocampus book*. Oxford university press, 2024.
- Edvard I. Moser, Emilio Kropff, and May-Britt Moser. Place cells, grid cells, and the brain’s spatial representation system. *Annual Review of Neuroscience*, 31:69–89, 2008. doi: 10.1146/annurev.neuro.31.060407.125601.
- Robert U. Muller and John 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.
- Andrea Navas-Olive, Liset M. de la Prida, et al. Cell-type-specific manifold analysis discloses independent geometric transformations in the hippocampal spatial code. *Neuron*, 2025. In press.
- John O’Keefe and Neil Burgess. Geometric determinants of the place fields of hippocampal neurons. *Nature*, 381(6581):425–428, 1996. doi: 10.1038/381425a0.
- 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.
- John O’Keefe and Lynn Nadel. *The hippocampus as a cognitive map*. Oxford university press, 1978.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, et al. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8026–8037, 2019. doi: 10.48550/arXiv.1912.01703. URL <https://doi.org/10.48550/arXiv.1912.01703>.
- Andy Peters. Neuropixels trajectory explorer. https://github.com/petersaj/neuropixels_trajectory_explorer, 2022. Neuropixels trajectory explorer with the Allen CCF mouse atlas.
- Adrien Peyrache, Natalie Schieferstein, and György Buzsáki. Internally organized mechanisms of the head direction sense. *Nature Neuroscience*, 20(12):1598–1607, 2017. doi: 10.1038/nn.4682.
- James B Ranck. Head direction cells in the deep cell layer of dorsolateral pre-subiculum in freely moving rats. *Electrical activity of the archicortex*, 1985.
- A David Redish. *Beyond the cognitive map: from place cells to episodic memory*. MIT press, 1999.

- Alon Rubin, Liron Sheintuch, Noa Brande-Eilat, Or Pinchasof, Yoav Rechavi, Nitzan Geva, and Yaniv Ziv. Revealing neural correlates of behavior without behavioral measurements. *Nature communications*, 10(1):4745, 2019.
- Michael E. Rule, Adrianna R. Loback, Dhruva V. Raman, Laura N. Driscoll, Christopher D. Harvey, and Timothy O’Leary. Stable task information from an unstable neural population. *eLife*, 9:e51121, 2020. doi: 10.7554/eLife.51121.
- Juliane I. Sanguinetti-Scheck and Michael Brecht. Home, head direction stability, and grid cell distortion. *eLife*, 9:e59869, 2020. doi: 10.7554/eLife.59869. URL <https://pmc.ncbi.nlm.nih.gov/articles/PMC7191526/>.
- Francesca Sargolini, Marianne Fyhn, Torkel Hafting, Bruce L. McNaughton, Menno P. Witter, May-Britt Moser, and Edvard I. Moser. Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science*, 312(5774):758–762, 2006. doi: 10.1126/science.1125572.
- Patricia E. Sharp. Subicular place cells generate the same “map” for different environments: Comparison with hippocampal cells. *Behavioural Brain Research*, 174(2):206–214, 2006. ISSN 0166-4328. doi: <https://doi.org/10.1016/j.bbr.2006.05.034>. URL <https://www.sciencedirect.com/science/article/pii/S0166432806003263>. The Mammalian Subiculum: Contrasting and complementary in vivo and in vitro approaches to subicular function.
- Patricia E. Sharp and Colin Green. Spatial correlates of firing patterns of single cells in the subiculum of the freely moving rat. *Journal of Neuroscience*, 14(4):2339–2356, 1994. doi: 10.1523/JNEUROSCI.14-04-02339.1994.
- David I. Shuman, Sunil K. Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine*, 30(3):83–98, 2013. doi: 10.1109/MSP.2012.2235192. URL <https://doi.org/10.1109/MSP.2012.2235192>.
- Trygve Solstad, Charlotte N. Boccara, Emilio Kropff, May-Britt Moser, and Edvard I. Moser. Representation of geometric borders in the entorhinal cortex. *Science*, 322(5909):1865–1868, 2008. doi: 10.1126/science.1166466.
- Nicholas A. Steinmetz, Cengiz Aydin, Anna Lebedeva, Michael Okun, Marius Pachitariu, Marius Bauza, Matthieu Beau, Jaspreet Bhagat, Carsten Bohm, Marc Broux, Shijie Chen, Jennifer Colonell, Richard J. Gardner, Bill Karsh, Fabian Kloosterman, Dimitar Kostadinov, Claudia Mora-Lopez, James O’Callaghan, Junchol Park, Jan Putzeys, Britton Sauerbrei, Ralf van Daal, Anders Z. Vollen, Shuxin Wang, Marnix Welkenhuysen, Zhe Ye, Joshua T. Dudman, Barundeb Dutta, Adam W.

- Hantman, Kenneth D. Harris, Albert K. Lee, Edvard I. Moser, John O’Keefe, Alfonso Renart, Karel Svoboda, Sebastian Haesler, Matteo Carandini, and Tim D. Harris. Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science*, 372(6539):eabf4588, 2021. doi: 10.1126/science.abf4588. URL <https://doi.org/10.1126/science.abf4588>.
- Hanne Stensola, Tor Stensola, Matthew B. Snow, May-Britt Moser, and Edvard I. Moser. The entorhinal grid map is discretized. *Nature*, 492(7427):72–78, 2012. doi: 10.1038/nature11649.
- Ilia Sucholutsky, Lukas Muttenthaler, Adrian Weller, Andi Peng, Andreea Bobu, Been Kim, Bradley C Love, Erin Grant, Iris Groen, Jascha Achterberg, et al. Getting aligned on representational alignment. *arXiv preprint arXiv:2310.13018*, 2023.
- Jeffrey S. Taube. The head direction signal: origins and sensory-motor integration. *Annual Review of Neuroscience*, 30:181–207, 2007. doi: 10.1146/annurev.neuro.29.051605.112854.
- Quanxin Wang, Song-Lin Ding, Yang Li, David Feng, Kuan Evans, Bhanu C. Ramamohan, Yan Chen, Mikio Mori, Aaron Szafer, Thomas Capiod, and et al. The allen mouse brain common coordinate framework: A 3d reference atlas. *Cell*, 181(4):936–953.e20, 2020. doi: 10.1016/j.cell.2020.04.007.
- Zhaoze Wang, Ronald W Di Tullio, Spencer Rooke, and Vijay Balasubramanian. Time makes space: Emergence of place fields in networks encoding temporally continuous sensory experiences. *Advances in Neural Information Processing Systems*, 37:37836–37864, 2024.
- Zhaoze Wang, Genela Morris, Dori Derdikman, Pratik Chaudhari, and Vijay Balasubramanian. Remi: Reconstructing episodic memory during internally driven path planning. *bioRxiv*, 2025.
- Tom J Wills, Colin Lever, Francesca Cacucci, Neil Burgess, and John O’Keefe. Attractor dynamics in the hippocampal representation of the local environment. *Science*, 308(5723):873–876, 2005. doi: 10.1126/science.1108905.
- Qiaorong S Yu, Zhaoze Wang, and Vijay Balasubramanian. When and where: A model hippocampal network unifies formation of time cells and place cells. *bioRxiv*, pages 2026–03, 2026.

הוא שהעתקות של הייצוגים המרחביים בחדרים שונים ניתנים ברוב המקרים לתיאור טוב באמצעות טרנספורמציות אפיניות ממימד נמוך. זאת על אף שהמודל בנוי בצורה המאפשרת תפיסה של דפוסים מורכבים מאוד. כלומר, המעבר בין ייצוג מרחבי של חדר אחד לייצוג של חדר אחר ניתן לתיאור באמצעות צירוף פשוט יחסית של סיבוב, שיקוף, מתיחה והזזה. יתרה מזאת, המיפויים שהתקבלו מהמודל היו חלקים ורציפים, ושמרו על המבנה הגאומטרי המקומי של המרחב. כאשר בוצעה השוואה לשיטה מסורתית המבוססת על התאמה בין דפוסי ירי של וקטורי אוכלוסייה, התקבלו תוצאות דומות, אך המודל החישובי הניב התאמות חלקות, מדויקות ועמידות יותר לרעש. העיקרון הפשוט שהתקבל נשמר באופן עקבי על פני גאומטריות שונות של חדרים, על פני אזורי המוח שתועדו, וגם על פני תת-אוכלוסיות עצביות שונות, לרבות כאלה בעלות קידוד מרחבי. ממצא זה מצביע על כך שהטרנספורמציה בין סביבות אינה תכונה של סוג תא בודד, אלא עיקרון כללי של הרשת המכילה קוד מרחבי. מנגד, כאשר נבחנו השינויים בכוונון של תאים בודדים, נמצאה הטרוגניות רבה, ולעיתים קרובות לא הייתה התאמה פשוטה בין שינויי התאים הבודדים לבין הטרנספורמציה הגלובלית של האוכלוסייה. בנוסף, במקרים שבהם הטרנספורמציה כללה שיקוף, נמצא כי צירי השיקוף היו מעוגנים לגאומטריה הפיזית של הסביבה, ובפרט לכיוון הקיר המשותף בין החדרים. לסיכום, עבודה זו מראה כי מיפוי מחדש מרחבי בהיפוקמפוס אינו תהליך אקראי בלבד, אלא נשלט ברמת האוכלוסייה על ידי עיקרון גאומטרי פשוט ומובנה. ממצא זה מציע מסגרת חדשה להבנת האופן שבו המוח שומר על קוהרנטיות מרחבית כאשר הוא עובר בין הקשרים שונים, ומבסס תשתית רעיונית וחישובית לפענוח מיקום גם בסביבות חדשות באמצעות כללי טרנספורמציה נלמדים.

תקציר

תופעת המיפוי מחדש של ייצוגים מרחביים בהיפוקמפוס בעת מעבר בין סביבות שונות היא אחת התופעות המרכזיות בחקר הניווט והזיכרון המרחבי. מחקרים רבים תיארו כיצד שדות הירי (איזורים במרחב בהם יש פעילות עצבית) של נוירונים בודדים משתנים בין סביבות, אך העיקרון המארגן העומד בבסיס שינוי זה עדיין אינו ברור. בעוד שבאזורי היפוקמפוס מסוימים, ובעיקר באזור CA1, המיפוי מחדש נראה לעיתים אקראי, אזורים אחרים כגון הסוביקולום וקליפת המוח האנטוריןלית התיכונה מציגים דפוסים מרחביים יותר מובנים. מכאן עולה השאלה האם גם כאשר תאים בודדים משנים את דפוס פעילותם באופן הטרוגני, ייתכן שברמת האוכלוסייה נשמרת חוקיות כללית ופשוטה יותר. שיטות הניתוח המסורתיות של מיפוי מחדש נשענות בדרך כלל על השוואת מפות קצב ירי של תאים בודדים בין סביבות. גישה זו תרמה רבות להבנת תאים בעלי קידוד מרחבי מובהק, אך היא מוגבלת בשני מובנים מרכזיים. ראשית, היא פחות יעילה עבור נוירונים שאינם מציגים כונון מרחבי חד. שנית, והיא החשובה ביותר, היא מתייחסת לנוירונים כאל יחידות עצמאיות ומתעלמת מהפעילות ההדדית ביניהם ברמת הרשת. לפיכך, ייתכן שעקרונות כלליים לגבי הייצוג המרחבי חומקים מן העין כאשר הניתוח נעשה ברמת התא הבודד בלבד. בחיבור זה אנו מציגים מסגרת חישובית יצירתית לחקר המיפוי מחדש ברמת אוכלוסיית הנוירונים. השיטה שלנו מבוססת על מודל לא ליניארי ותלוי זמן, שמאומן ישירות מהנתונים שהוקלטו ואינו מניח מראש צורה מסוימת של העתקה בין הסביבות. מטרת המודל הייתה לחזות את מיקום העכבר מתוך הפעילות העצבית בסביבות שונות (חדרים). לצורך כך נעשה שימוש בארכיטקטורה המורכבת ממקודד זמני משותף לכל החדרים, הממפה מקטע מהפעילות העצבית של האוכלוסייה למרחב לטנטי משותף, וממפנענחים נפרדים לכל חדר, אשר ממפים את הייצוג הלטנטי לקואורדינטות מרחביות במערכת הייחוס של אותו חדר. מבנה זה מאפשר לבחון כיצד פעילות עצבית שנמדדה בסביבה אחת הייתה מתפרשת אילו נקראה באמצעות המפענח של סביבה אחרת, וכך להסיק ישירות את יחסי ההתאמה בין הייצוגים המרחביים של שתי הסביבות. הניתוח יושם על מספר רב של תאים שהוקלטו מעכברים חופשיים שנעו בזירה המורכבת משני חדרים נפרדים, החדרים היו בצורות מגוונות בעלי גאומטריות שונות. בכל ניסוי העכבר הסתובב תחילה בחדר אחד למשך 20 דקות, לאחר מכן בחדר השני למשך 20 דקות, ולבסוף חזר לחדר הראשון ל-20 דקות נוספות. ההקלטות בוצעו באמצעות אלקטרודות נוירופיקסל (neuropixels), ובסך הכול בוצעו 71 ניסויים משישה בעלי חיים, שכללו אלפי נוירונים מקליפת המוח האנטוריןלית התיכונה, סוביקולום והפרה-סוביקולום. בנוסף, בחנו גם תתי-אוכלוסיות פונקציונליות שונות של תאים, ובהן תאי גבול אלוצנטריים, תאי גבול אגוצנטריים, תאי כיוון ראש, תאים בעלי מידע מרחבי גבוה ותאי מהירות. תחילה הראינו כי המודל מסוגל לחזות באופן אמין את מיקום העכבר מתוך הפעילות העצבית של האוכלוסייה. מצאנו כי שימוש בחלון זמן של מספר שניות משפר משמעותית את ביצועי החיזוי ביחס למודלים פשוטים יותר שחוזים רק על סמך הפעילות העצבית הרגעית. הממצא המרכזי של המחקר

המחקר נעשה בפקולטה לרפואה בהנחיית ד"ר הדס בניסטי ופרופ' דורי דרדיקמן.

מחבר חיבור זה מצהיר כי המחקר, כולל איסוף הנתונים, עיבודם והצגתם, התייחסות והשוואה למחקרים קודמים וכו', נעשה כולו בצורה ישרה, כמצופה ממחקר מדעי המבוצע לפי אמות המידה האתיות של העולם האקדמי. כמו כן, הדיווח על המחקר ותוצאותיו בחיבור זה נעשה בצורה ישרה ומלאה, לפי אותן אמות מידה.

אני מודה לטכניון על התמיכה הכספית הנדיבה בהשתלמותי.

עקרון גאומטרי של העתקה מרחבית בהיפוקמפוס

חיבור על מחקר

לשם מילוי חלקי של הדרישות לקבלת התואר
מגיסטר למדעים בבמדעי הרפואה

דניאל צור

הוגש לסנט הטכניון – מכון טכנולוגי לישראל
ניסן תשפ"ו חיפה אפריל 2026

**עקרון גאומטרי
של העתקה מרחבית
בהיפוקמפוס**

דניאל צור