William Zheng ywz2@cornell.edu

Circuit mechanisms of representational drift and neural manifold collapse in Alzheimer's Disease

William Zheng

Abstract

The hippocampus supports generalization by organizing the coactivity dynamics of cell assemblies into disentangled neural manifolds, where spatial and temporal variables are encoded on orthogonal axes. Mechanistically, this geometry relies on upstream inputs to constrain global topology and local inhibition to drive representational drift. In Alzheimer's Disease (AD), spatial disorientation and impaired schema transfer suggest a collapse of this machinery, yet whether this results from upstream signal degradation or local circuit dysfunction remains undefined. Here, by combining dual-site, large-scale recordings with closed-loop optogenetics in 5XFAD mice, I aim to deconstruct the circuit mechanisms of this geometric failure. I hypothesize that 5XFAD mice exhibit "tangled" manifolds driven by the functional decoupling of CA3 inputs and the desynchronization of the local inhibitory signals. Specifically, I will determine if this decoupling arises from a "functional deafferentation" akin to CA3 silencing in wild-type mice, or an "aberrant noise model" driven by pathological CA3 hyperactivity where robust firing lacks predictive subspace alignment with CA1. Finally, I will test if stabilizing PV interneuron dynamics can causally rescue the orthogonality of the neural code, offering a unifying principle linking cellular pathology to the loss of cognitive flexibility.

Introduction

The hippocampus is essential for spatial navigation and episodic memory, forming an internal cognitive map. While traditionally characterized by the formation of distinct, uncorrelated representations for every environment, efficient navigation requires rapid generalization of structural knowledge to novel contexts. Recent geometric and topological analyses reveal that the hippocampus supports this flexibility by generating "disentangled" neural manifolds, where task variables such as spatial location, temporal progression, and behavioral choice are encoded on orthogonal axes. This geometry requires a precise circuit organization where CA3 inputs constrain spatial correlations across environments, while local parvalbumin (PV) inhibition drives the representational drift encoding temporal progression. However, whether this geometric machinery persists in neurodegenerative conditions, or how specific circuit pathologies disrupt these orthogonal codes to cause cognitive deficits remains unknown.

Alzheimer's Disease (AD) manifests clinically as early spatial disorientation and "cognitive rigidity" – an inability to transfer learned schemas to new contexts – suggesting a collapse of this geometric machinery. While molecular hallmarks such as amyloid plaques and PV interneuron dysfunction are well-documented, it remains unknown how these cellular deficits translate to system-level topological failures. Specifically, it is unclear whether the loss of generalization stems from a "tangling" of the neural manifold – where orthogonal axes collapse – or from a complete dissolution of topology into random noise. Furthermore, whether this geometric failure is driven by incoherent upstream inputs failing to anchor the spatial map, or by a loss of local inhibitory regulation disrupting the structure of representational drift, has not been mechanistically defined.

Here, we will combine large-scale and multi-region neural recordings in behaving 5XFAD and wild-type mice to explore the circuit mechanisms of this manifold collapse. In this study, we will

analyze hippocampal population dynamics in 5XFAD and wild-type mice to quantify geometric stability and manifold alignment across multiple environments. We will test the hypothesis that 5XFAD mice exhibit “tangled” manifolds driven by the functional decoupling of upstream spatial inputs and the desynchronization of local inhibitory signals. This work aims to causally resolve these mechanisms, establishing a unifying circuit principle linking cellular pathology to the degradation of cognitive flexibility.

Specific Aims

Aim 1: Quantify the topological collapse and loss of disentangled coding in hippocampal neural manifolds in 5XFAD mice

We will test the hypothesis that AD pathology disrupts the orthogonal encoding of task variables, resulting in manifolds that are geometrically “tangled” and driven by random drift rather than structured dynamics. We will record large-scale hippocampal population activity in 5XFAD and wild-type mice performing a continuous alternation task across multiple mazes. We aim to employ topological data analysis to assess whether 5XFAD mice fail to form the stable ring topology characteristic of abstract task learning by measuring the Lifespan Index. Furthermore, we will quantify the “tangling” of the neural code by measuring the orthogonality and principal angels of coding subspaces for space, time, and choice, and assessing their cross-condition generalization performance (CCGP) across environments. Finally, we will distinguish between unstructured noise and coherent remapping by quantifying the geometric alignment of manifolds across sessions. We predict that 5XFAD mice will exhibit reduced subspace orthogonality and poor manifold alignment, establishing a direct link between circuit-level geometric failure and behavioral rigidity.

Aim 2: Determine if the failure of spatial generalization arises from the functional decoupling of upstream CA3 inputs.

We will test the hypothesis that topological collapse and generalization failure in AD are driven by the degradation of coherent upstream spatial signaling. We will perform simultaneous dual-site, large-scale recordings in the dorsal CA3 and CA1 of 5XFAD and wild-type mice across familiar (M1), remapping (M2), and transfer (M3) environments. We will compare the geometric alignment of 5XFAD manifolds against a “functional mimicry” control group of wild-type mice with chemogenetically silenced CA3 inputs. This experimental design aims to rigorously distinguish whether the AD deficit follows a “functional deafferentation model”, evidenced by random manifold rotations and geometric similarity to the CA3-silenced state, versus an “aberrant noise model,” characterized by pathological CA3 hyperactivity where robust firing lacks predictive subspace alignment with CA1. By quantifying cross-maze manifold alignment using Procrustes analysis and the CA3-CA1 communication subspace via Canonical Correlation Analysis (CCA), this aim will precisely identify the upstream circuit failure preventing the cross-context alignment of spatial maps in AD.

Aim 3: Causally rescue representational drift and temporal coding via optogenetic stabilization of local inhibition

Having identified upstream inputs as the driver of spatial collapse, we will test the hypothesis that the degradation of temporal structure arises from the desynchronization of local inhibitory circuits. We will implement a closed-loop optogenetic intervention in PV-Cre/5XFAD mice, delivering white-noise modulated stimulation to resynchronize interneuron activity during navigation. This design aims to causally dissociate the circuit mechanisms of AD, testing if

stabilizing the local inhibitory signal can rescue representational drift independent of upstream spatial inputs. We will quantify the rescue by comparing the Euclidean distance between adjacent trials, the orthogonality of the time-encoding subspace, and the accuracy of temporal decoding classifiers trained to predict trial number between stimulated and control sessions. We predict that restoring inhibitory precision will untangle the temporal manifold, significantly improving decoding accuracy and proving that local circuit dysfunction is the primary driver of temporal disintegration in the AD cognitive map.

Research Design

In this experiment, we will use 5XFAD transgenic mice maintained on a congenic C57BL/6J background. Before the surgery, mice will be accustomed to one maze 3-5 days. Following the implant in 5XFAD and wild-type littermates, targeting the dorsal CA1, and a subsequent recovery period, animals will be handled daily and accommodated to the experimenters and cables for a week. One day before training, mice will be placed under water restriction.

To define the baseline geometric deficit, we will train mice to run a figure-8 maze. Each trial begins at the center stem, requiring alternating choices between the left and right arms. The system introduces an initial 8-s delay by using automatic doors to keep the animal in the center stem. Each trail completion will trigger a small reward (30% sugar water), being pumped at the end of the side arms. Animals will run two sessions per day with 2-hour sleep session in between. With the recording data, we will visualize the intrinsic structure of the neural population activity by constructing 3D manifold embeddings (using UMAP) for both 5XFAD and wild-type mice. We hypothesize that 5XFAD mice will exhibit a “topological collapse,” manifesting as tangled trajectories caused by a fundamental misalignment between the time, choice, and location encoding subspaces. In contrast to the structured dynamics of wild-type mice – where these variables are orthogonally separated – the 5XFAD manifold will fail to independently factorize these dimensions. To further characterize this disintegration at the single-unit level, we will generate normalized place cell rate maps sorted by peak firing location across trials. This analysis is expected to demonstrate a clear dissociation: wild-type controls will display organized “temporal drift stripes” indicative of a healthy, evolving time code, whereas 5XFAD mice will exhibit disordered, diffusive firing patterns that confirm the degradation of spatiotemporal structure.

To determine if the cognitive deficit in AD arises from a failure to acquire the shared latent structure of the task, we will quantitatively compare hippocampal population geometry between wild-type and 5XFAD mice across the three environments (M1, M2, and M3). We will apply Topological Data Analysis (TDA) to measure the lifespan index of persistent homology features, specifically tracking the evolution of manifold topology from specific experiences to abstract schemas. We predict that wild-type mice will exhibit a topological transformation from a “two-loop” structure (representing distinct left/right trajectories) to a unified “single-loop” ring (representing the abstract alternation rule) as they generalize from M1 to M2, and crucially, will immediately instantiate this 1-ring topology in the novel M3. In contrast, 5XFAD mice will fail to manifest this topological convergence. We hypothesize their manifolds will persistently retain a specific “two-loop” or tangled structure across all environments, reflecting an inability to abstract the common task schema. Moreover, this topological stagnation will mirror their behavioral performance. Unlike wild-type littermates who exhibit “transfer learning” (immediate high performance in M3), 5XFAD mice are expected to show no significant performance gain

across transfers, treating each maze as a novel task. Mechanistically, we will confirm this using Multidimensional Scaling (MDS) and to calculate the principal angles, which we hypothesize a fundamental misalignment between spatial and temporal coding subspaces in 5XFAD mice, proving that the loss of coding orthogonality underlies their inability to generalize task structure.

Furthermore, we will leverage this geometric framework to track the granular trajectory of pathology, hypothesizing that the neural manifold will exhibit a progressive topological disintegration that correlates with disease duration. Specifically, we anticipate that the degree of subspace misalignment will quantifiably worsen over time, allowing us to distinguish early-stage geometric perturbations detectable at two weeks from the irreversible manifold collapse observed after one month of progression.

To identify the upstream driver of the generalization deficit, we will conduct simultaneous dual-site, large-scale recordings targeting the dorsal CA3 and CA1 in 5XFAD and wild-type mice navigating the familiar, remapping, and transfer environments (M1, M2, M3). We will compare 5XFAD neural dynamics against a "functional mimicry" control group of wild-type mice with chemogenetically silenced CA3 inputs. This design enables us to rigorously distinguish between two competing circuit models: a "functional deafferentation model," where 5XFAD manifolds exhibit the random cross-context rotations and geometric disintegration characteristic of the CA3-silenced state in wild-type mice; versus an "aberrant noise model," where CA3 exhibits pathological hyperactivity that fails to propagate informative signals. We will quantify these relationships using Procrustes analysis to measure the alignment of manifolds across mazes and Canonical Correlation Analysis (CCA) to calculate the strength of the communication subspace. We predict that 5XFAD mice will exhibit a "disjointed" geometry indistinguishable from the CA3-silenced controls – characterized by low CCA scores despite potentially high firing rates – suggesting that AD pathology functionally uncouples the hippocampus from the upstream information required for spatial generalization.

Finally, to determine if the temporal fragmentation of the 5XFAD manifold is driven by the failure of local inhibitory circuits to pace the ensemble, we will implement a closed-loop optogenetic rescue experiment. We will target local inhibitory circuits by injecting a Cre-dependent channelrhodopsin vector into the dorsal CA1 of PV-Cre/5XFAD mice, allowing for the selective optogenetic control of PV interneurons. During the alternation task, running in a figure-8 maze, we will deliver white-noise modulated photostimulation designed to resynchronize PV interneuron activity, effectively stabilizing the local inhibitory signal that regulates population drift. We will evaluate the success of this intervention by quantifying the smoothness of representational drift, calculated as the Euclidean distance between the neural population states of adjacent trials. We hypothesize that while non-stimulated 5XFAD epochs will exhibit erratic transitions (large, variable Euclidean jumps), the optogenetically stabilized epochs will restore a smooth, continuous trajectory comparable to wild-type dynamics. Furthermore, we will train temporal decoding classifiers to predict the specific trial number from the neural population activity. We expect that re-establishing inhibitory precision will significantly increase the decoding accuracy of these classifiers, confirming that the intervention successfully rescues the temporal fidelity of the cognitive map.

William Zheng ywz2@cornell.edu

Figure Plans

Figure 1: The experimental schematic of the mazes and visualization of the topological collapse of 3D embedding of neural manifolds in 5XFAD mice, and structured neural dynamics in wild-type mice. The visualization of normalized place cell rate map of 5XFAD mice which reveal disordered, diffusive firing patterns compared to the structured temporal drift stripes of wild-type controls.

Figure 2: Comparison of hippocampal population geometry between wild-type and 5XFAD mice, reflecting their different acquisition of shared task structure. Quantification of the geometric structure of neural manifold by measuring the lifespan index, population activity similarity, revealing a collapsed manifold structure in 5XFAD mice. Quantification of the principal angles be between spatial and temporal subspaces to demonstrate a significant reduction in coding orthogonality in 5XFAD mice.

Figure 3: Schematic of simultaneous dual-site, large-scale recordings in dorsal CA3 and CA1 to interrogate upstream spatial inputs. Visualization of the Communication Subspace via Canonical Correlation Analysis (CCA) reveals that robust wild-type alignment is functionally decoupled in 5XFAD mice. Procrustes analysis of cross-maze manifold alignment comparing 5XFAD geometry to "functional mimicry" controls (CA3-silenced wild-types) will distinguish between a "functional deafferentation model" characterized by geometric similarity to the silenced state, and an "aberrant noise model" characterized by pathological CA3 hyperactivity. This comparison will definitively confirm whether the spatial generalization deficit arises from the complete loss of upstream drive or the corruption of signals by robust, non-informative noise.

Figure 4:

Schematic of the closed-loop white-noise optogenetic intervention targeting PV interneurons in 5XFAD mice to restore inhibitory pacing. Visualization of drift dynamics quantifying the Euclidean distance between neural states of adjacent trials, illustrating a transition from erratic 5XFAD trajectories to smooth, continuous wild-type dynamics upon stimulation. This stabilization is predicted to restore subspace orthogonality and significantly increase the accuracy of temporal decoding classifiers (predicting trial number), confirming the causal rescue of the cognitive map's temporal fidelity.