National Institute for Theory and Mathematics in Biology

Funded by the U.S. National Science Foundation and the Simons Foundation

Differential Geometry of Representations in the Mouse Brain

Yu Jin Oh*, Renan Costa*, Lucas Pinto‡, Ramon Nogueira‡

Humans and other animals show a remarkable ability to adapt their behavior in complex, high-dimensional, and noisy environments, yet the neural mechanisms underlying this flexibility remain unclear. We address this by characterizing the differential geometry of neural representations across the mouse cortex, during a complex visual accumulation task. We head-fixed mice in a VR environment and trained them to navigate a T-maze containing checkerboard stimuli with varying coherence, i.e. differing proportions of white and black squares (Fig. 1A). As mice progressed through the maze, they accumulated sensory evidence to infer whether the checkerboard favored white or black squares. Throughout this decision-making process, we monitored neural activity across the dorsal cortex. To analyze task-related neural responses, we fit artificial neural networks (ANNs) as encoding models to predict population-level activity from sensory and task parameters. This yielded a differentiable mapping between task variables and neural responses, allowing us to visualize the brain's activity as a smooth low-dimensional representational manifold (Fig. 1B).

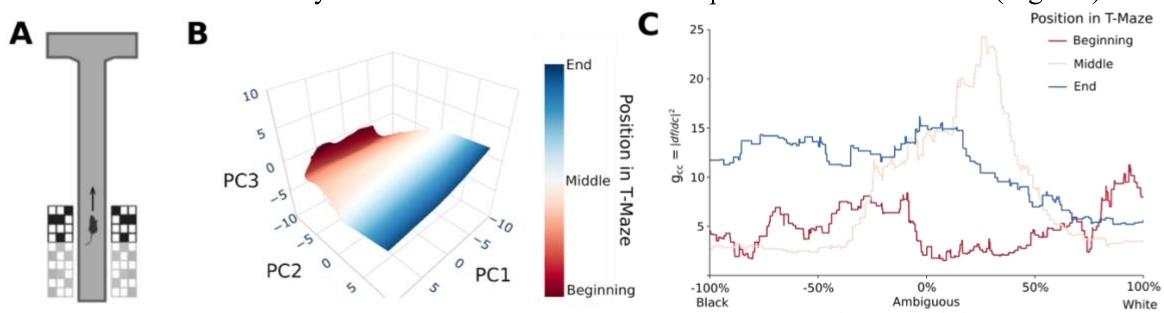

Figure 1: (**A**) Animals navigated a VR T-maze with checkerboards of varying coherence and had to make a turn at the end of the T-maze depending on whether the visual stimuli contained more white or black squares. (**B**) Neural networks fit to mouse V1 revealed a smooth representational manifold in which maze position and coherence were encoded in approximately orthogonal, low-dimensional subspaces. (**C**) The visual coherence component of the metric tensor (g_{cc}) is larger at the end of the maze (blue) than the beginning of the maze (red) and was expanded near the 0% coherence boundary in the middle of the maze (grey).

In the mouse primary visual cortex (V1), we found that this manifold represents both the mouse's location as well as the coherence of visual evidence. These two variables are encoded in approximately orthogonal subspaces, indicating largely independent representations and allowing the brain to track position and visual cues independently. Using powerful open-source computational tools (PyTorch), we calculated a "geometric ruler" (known as metric tensor, g_{ij}) to measure the exact shape of the manifold. Focusing on the coherence dimension, g_{cc}, we measured how much neural space the brain recruits as the coherence of the visual information changes (Fig 1C).

By the end of the stimulus presentation, neural representations of visual evidence were expanded near the decision boundary, potentially aiding evidence accumulation during decision-making. Future work will test whether similar geometric organization emerges in higher-level brain regions involved in decision-making and action. We hypothesize that such representational dilation will be most pronounced in posterior visual areas that are strongly engaged in evidence accumulation and closely linked to behavioral output.

This research was supported in part by grants from the NSF (DMS-2235451) and Simons Foundation (MP-TMPS-00005320).