Probing compositional learning during spontaneous exploration of a reconfigurable 3D environment

Tzuhsuan Ma[1], Ann M Hermundstad[1], and Jakob Voigts[1]
[1]HHMI Janelia Research Campus, Ashburn, VA 20147; {mat, hermundstada, voigtsj}@janelia.hhmi.org;

Abstract

In natural settings, animals navigate richly-structured sensory surroundings and rapidly adapt to changes in these surroundings. While many studies have explored navigation in mazes and open arenas, relatively little is known about how animals navigate in terrain that lacks defined routes and is too complex to memorize. Here, we probe the structure of mouse behavior in a complex, reconfigurable 3D arena in darkness and without explicit reinforcement. Within the first several hours, mice quickly explore the whole arena and converge on a sparse set of running and jumping paths. Surprisingly, after this initial phase of exploration, mice continue to generate new long paths for several days. To capture this structure, we develop a hierarchical segmentation algorithm that compresses raw behavioral trajectories into a compact set of composable sub-paths, or "motifs". We find that the behavior is highly compressible, indicating that mice create long paths by combining reusable motifs, rather than through random exploration. To study the evolving dynamics of these behavioral compositions, we first show that mice combine motifs in a non-random manner, generating temporal structure that is not captured by a Markov-chain that preserves the average transition probabilities between motifs. Next, we examine different phases of behavior in generating novel compositions. We find diverse dynamics that involve the rapid creation and extinction of compositions, as well as slower and more subtle refinements such as morphing, short-cutting, streamlining, and reinforcing a composition. These results suggest that mice use diverse learning rules to configure compact behavioral trajectories through space.

Keywords: navigation, complex behavior, one-shot learning, latent learning, time series segmentation

Experimental setup and preliminary observation. We filmed mice in an arena of 1.8 m diameter made from 153 hexagonal tiles (Fig 1a) (Newman et al., 2023). Tile heights varied in steps of 2.5 cm to produce a variety of local shapes, from small steps to tall towers that could not be easily jumped. The arena was illuminated with infrared LEDs and otherwise dark, limiting sensory information to tactile and auditory cues about the local arena shape and olfactory cues left by the mice. Mice explored the arena in their dark period over four 8-hour sessions on consecutive days, and were housed in their home cage in between. We used an overhead camera and SLEAP (Pereira et al., 2022) to track the mice at ~ 25 Hz.

Upon initial placement, mice quickly explore the majority of the arena with a distinctive backtracking behavior—a slow and short forward path followed by an immediate fast backward path. After this initial exploration phase (1-2 hours), mice quickly develop preferences for specific paths that are used in different combinations to generate new sequences that continue to emerge over time (Fig 1b).

A tree segmentation compresses behavior into composable motifs. To study the temporal evolution of this behavior, we devise a tree segmentation that compresses raw trajectories into series of motif events (Fig 2a). The algorithm extracts composable motifs that are 1) repetitive (which defines a "leaf" occupancy) and 2) part of repetitive compositions (which defines a "parent" occupancy). The algorithm fits a tree to the branching structure of long paths by maximizing *compositionality = leaf occupancy + parent occupancy* (Fig. 2b).

Mouse behavior is highly compositional. To see how much compositional information is extracted from mouse behavior, we compare the tree segmentation to a random segmentation that preserves the distribution of motif lengths. We find that the tree segmentation achieves a 5x compression relative to a random segmentation, which indicates that 1) mouse behavior is highly compositional and therefore compressible, and 2) the tree segmentation is effective in extracting this information. We find that 75% of the behavior can be captured by a compact set of 233 motifs; 2381 motifs are required to cover the remaining 25% of behavior (Fig. 2c).

Mice exhibit non-random structure in their compositions. To understand whether mice simply compose motifs at random, we compare mouse behavior to a memoryless Markov chain (MC) that preserves the average transition probabilities between motifs (Fig. 3). We find that 1) mice acquire new motifs at a much slower rate, 2) mice sample the same motifs more frequently within a short interval, and consequently 3) exhibit longer pauses between executions of the same motif.

Exploring diverse compositional learning rules. To study such a non-randomness, we explore how novel compositions are created during different phases of behavior. We find that groups of similar compositions emerge in ways that are suggestive of diverse learning rules, including one-shot learning, reinforcement learning, latent learning, streamlining, and global reconfiguration (Fig. 4). The learning dynamics often deviate strongly from MCs (Fig. 4a, blue vs gray) when a handful of key motifs are either discovered or perturbed; this is not exhibited in a MC that simply samples those same motifs more frequently (Fig. 4a,c-d). After a local perturbation, mice rapidly augment their repertoire with 48% of the observed novel compositions (Fig. 4a), suggesting that mice exhibit latent learning (Fig. 4c,d) that arises from perturbing both local motifs and nonlocal compositions.

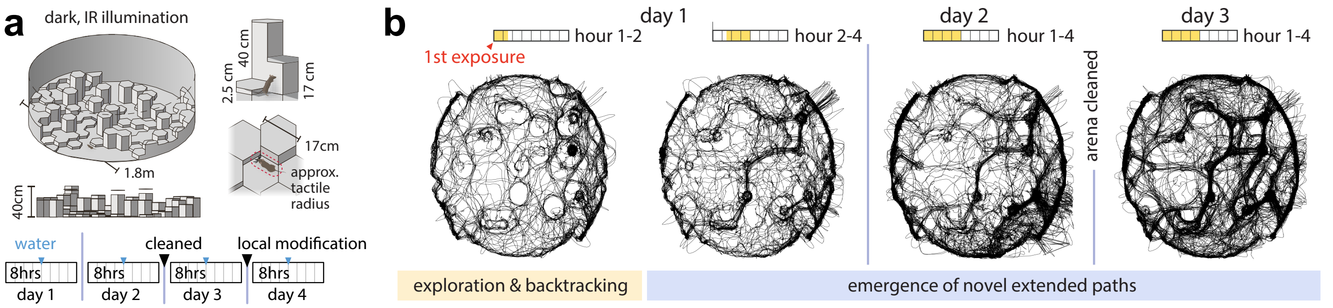

Figure 1: Behavior evolves over time during exploration. a) Exerimental setup. **b)** During exploration (yellow), mice explore new paths and backtrack along previously run paths. This exploration leads to a rapid convergence of the occupancy map, but novel paths continue to emerge over the course of multiple days.

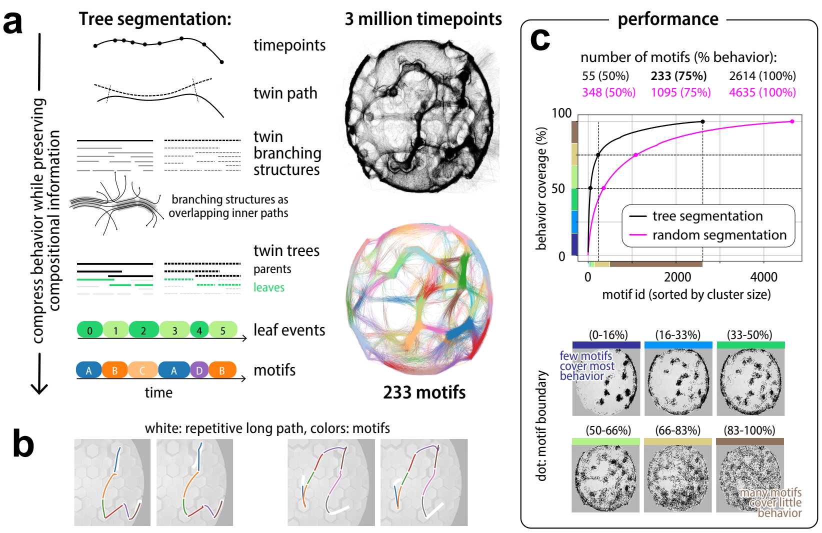

Figure 2: We compress mouse behavior into composable motifs. a) A tree segmentation algorithm extracts motifs by searching for 1) repetitive paths and 2) repetitive compositions—we find motifs that maximize both. **b)** Example segmentations. **c)** The tree segmentation achieves a ~5x higher compression rate compared to a random segmentation that preserves the distribution of motif lengths. Lower: tightly clustered boundaries show how a small subset of extracted motifs capture most behavior; the remaining majority of motifs cover a small fraction of the behavior.

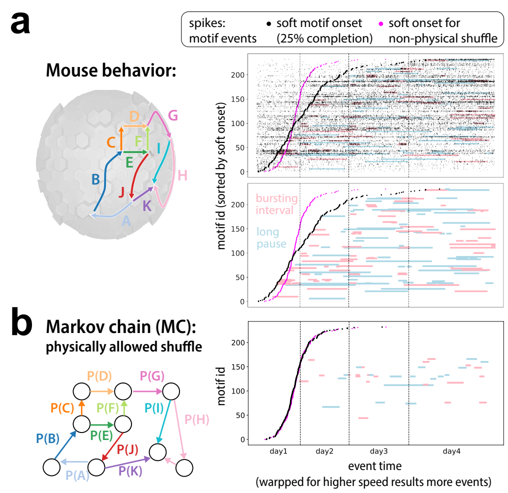

Figure 3: Mice do more than randomly compose. a) Upper: mouse behavior as a trajectory of motif events (spikes). Lower: Same plot without spikes. **b)** A Markov chain shows much closer statistics to a nonphysical random shuffle than to mice, as marked by its narrower spread of motif onsets, fewer bursts of motifs, and fewer long pauses. A burst or long pause event is defined as a 1/1000 outlier from a Poisson process (random shuffle).

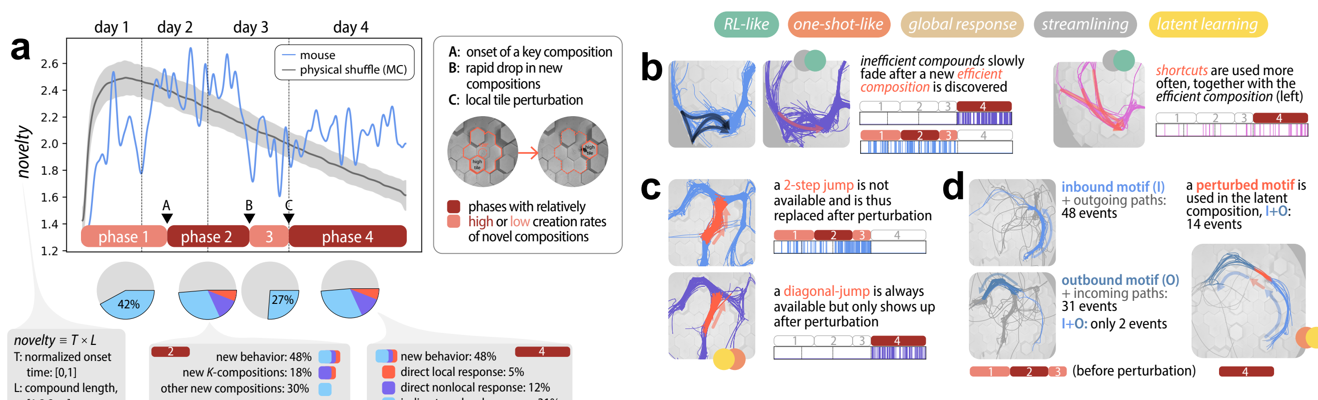

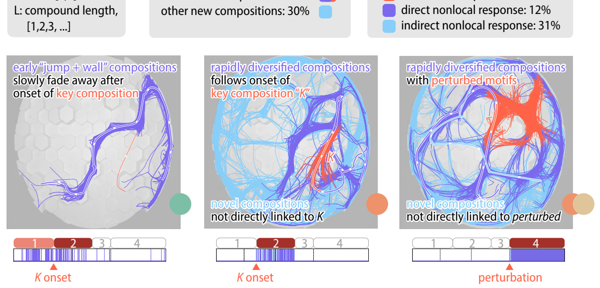

Figure 4: Rich compositional dynamics suggest diverse learning rules . a) A novelty trace is computed to measure the emergence of novel compositions. Different learning phases are heuristically defined (right panel). Phase 2 and 4, which show high creation rates, are discussed in detail in the lower panel. The results imply one-shot-like (orange sticker) learning dynamics after the acquisition of a key composition or in response to a perturbation on day 4. Note that such a global response to a local perturbation implies that it is mostly compositions--rather than local motifs--that are perturbed. **b)** Certain streamlined compositions display very different, reinforcement-learning-like (RL-like), dynamics. **c-d)** Mice also show latent learning in which a novel motif or composition suddenly appears in response to the perturbation, despite the fact that the motif or composition was previously available before the perturbation.

Acknowledgments

This work is funded by the Howard Hughes Medical Institute. We thank Nishan Shettigar, Emily Dennis and Roian Egnor for helpful discussions, and Angela Jacobs for husbandry and experimental support.

References

Newman, J. P., Zhang, J., Cuevas-Lopez, A., Miller, N. J., Honda, T., van der Goes, M.-S. H., ... others (2023). A unified open-source platform for multimodal neural recording and perturbation during naturalistic behavior. *bioRxiv*.

Pereira, T. D., Tabris, N., Matsliah, A., Turner, D. M., Li, J., Ravindranath, S., ... others (2022). Sleap: A deep learning system for multi-animal pose tracking. *Nature methods*, *19*(4), 486–495.