The Primate Hippocampus Constructs a Temporal Scaffold Anchored to Behavioral Events

Jon W Rueckemann[†*1,2], Yoni Browning[†3,4], Autumn J Mallory[1], Brian Kim[5], Adrienne L Fairhall[1], Elizabeth A Buffalo[*1,2]

University of Washington[1], Washington National Primate Center[2], UW Neuroscience Graduate Program[3], Allen Institute for Neural Dynamics[4], University of Michigan[5]

[†]Co-First Authors

*Corresponding Authors:

Jon W Rueckemann – jwrueck@uw.edu

Elizabeth A Buffalo – ebuffalo@uw.edu

Abstract

The hippocampus has been shown to support a variety of cognitive functions, ranging from spatial navigation[1,2] to anxiety regulation[3] to memory formation[4,5], but how its moment-to-moment neuronal activity contributes to cognition remains poorly understood. The activity of single hippocampal neurons is correlated with numerous perceptual features and task variables[6–14], raising the question of whether these response properties reflect distinct mechanisms or a single generalized computation. Here, we show that diverse hippocampal responses reflect unified event-driven dynamics in which population activity transitions between discrete ensemble states at behaviorally salient events. Recording from monkeys performing a virtual spatial alternation task, we found that population activity did not evolve smoothly over time but instead transitioned abruptly at each relevant event. These discontinuities segmented activity into distinct ensemble codes, effectively chunking conceptually defined task epochs. Notably, many neuronal responses persisted across visually distinct environments, demonstrating that these dynamics reflect abstract task structure rather

than specific sensory features. These results reveal that the hippocampus constructs a temporal scaffold anchored to relevant behavioral events, with each neural state tracking a distinct task phase. This organizational principle may explain the diverse correlates observed across studies: rather than individually encoding task features, hippocampal neurons collectively signal which phase of a behavioral sequence is currently active. Our findings suggest that the hippocampus parses experience into meaningful elements and tracks "position" within a learned behavioral structure.

Introduction

Damage to the hippocampus impairs memory[15–17] and cognition[18–21], yet we lack a fundamental account of what this structure computes. Much of what we know about its physiology comes from rodents navigating space, where neurons fire at specific locations and have been interpreted as encoding a map of the environment[1,2,22]. However, studies have also shown that hippocampal neurons respond to tones[19,20] and odors[22,23], track elapsed time during delays[11,23,24], and modulate their activity by behavioral context[12,14,25], revealing that the hippocampus encodes dimensions other than physical location. This diversity of responses has supported views in which individual hippocampal neurons act as conjunctive feature detectors that bind sensory, temporal, and contextual information within each moment[26,27].

A complementary perspective shifts the focus from representing individual moments to capturing the continuity that links them across time[28–32]. In this view, the hippocampus functions as a sequence generator whose activity evolves systematically across events[33], producing structured patterns that bridge the transitions of experience. In support of this view, studies have shown that neurons exhibit temporally extended firing across delays[34], predict upcoming actions[35–37], track iterations across trials[38], and link events into a holistic sequence[39–42], indicating that hippocampal activity conveys not only what is happening but how the current moment is embedded within the broader temporal progression of behavior. At the same time, other work highlights the hippocampus's acute sensitivity to salient boundaries[43–47], suggesting that behaviorally-relevant events play a pivotal role in structuring the unfolding sequence of experience. These perspectives converge on the idea that the hippocampus does more than encode discrete features—it tracks the evolving organization of experience, integrating the continuity of ongoing activity with boundaries sculpted both by external events and the organism's own interpretation of transitions.

Results

To examine how these principles manifest at the neuronal level in primates, we recorded neuronal activity from the hippocampus of three adult male rhesus monkeys (*Macaca mulatta*) during performance of a spatial alternation task in an immersive virtual environment (Fig. 1a). Subjects used a joystick to navigate Y-shaped mazes in which they were rewarded for alternating between left and right arms across successive trials (Fig. 1b). Subjects performed the task in multiple visually distinct Y-shaped mazes within a recording day, allowing us to dissociate responses tied to specific sensory features from those reflecting abstract task structure.

Chronically implanted microdrives with 124 independently-movable electrodes were used to record from the hippocampus in each animal (Fig. 1c). Recording locations were confirmed to sample the full longitudinal extent of the hippocampus through post-mortem histological reconstruction. To characterize how individual neurons responded to key features of the task, we aligned spiking activity to five behavioral events spanning all task phases: movement initiation from the central stem (Start), turn initiation at the choice point (Choice), turn initiation toward the reward location (Outcome), reward delivery (Reward), and onset of the inter-trial interval (ITI) (Fig. 1b). Individual neurons exhibited reliable responses aligned to these events (Fig 1d), and the response profiles included both sharp modulation time-locked to specific moments and sustained but irregular epoch-spanning activity.

Event-aligned responses across task phases

To systematically assess event-aligned responses, we analyzed activity within a 2-s window surrounding each event. We computed Shannon mutual information between instantaneous firing rate and peri-event time for each neuron, with significance assessed against a within-trial circular-shift null distribution ($P<0.01$). Across the population of 1659 uniquely recorded neurons, 566 neurons showed significant modulation to at least one event, with responses distributed across spatial and non-spatial task phases: Start (n=342; 60.4%), Choice (n=305; 53.9%), Outcome (n=348;

61.5%), Reward (n=350; 61.8%), and ITI (n=329; 58.1%). Strikingly, the vast majority of these neurons (n=485; 85.7%) were significantly responsive to multiple events (Fig. 1e), indicating that these neurons do not act as feature detectors that are narrowly selective for specific task variables or sensory stimuli. Accordingly, these findings suggest that hippocampal activity reflects a distributed code in which information emerges from ensemble dynamics across the neural population.

Events segment the task into distinct epochs

We next explored how neuronal responses contribute to population-level representations across the task. Spikes were first binned in non-overlapping 50ms intervals to compute instantaneous firing rates over the course of each trial. To compare trials of variable duration, the timeline of each trial was rescaled piecewise between the onset of task events (Movement Start, Choice Turn, Choice Arm, Outcome Arm, Reward, and Inter-Trial Interval; Fig. 2a). Each inter-event interval was linearly stretched or compressed to match its mean duration across the session, yielding a standardized trial lasting 17s that was sampled at 68 evenly spaced points (250ms resolution within rescaled time). Incorrect trials were excluded from further analysis. This temporal rescaling approach ensured that both movement periods and stationary epochs were represented (Fig. 2a). 588 neurons demonstrated significant Shannon mutual information ($P<0.01$) between instantaneous firing rate and task progression; these neurons composed the population used for subsequent analyses.

At the population level, these complex neuronal responses combined to form activity patterns that differentiated task phases across the trial. Rather than being uniformly dispersed across time, neuronal responses aligned to task events (Fig. 2b). Temporal cross-correlation of normalized population vectors[48,49] revealed sustained blocks of high within-epoch correlation separated by sharp discontinuities at event boundaries (Fig. 2c). These event-related discontinuities were not dependent on the specifics of the rescaling procedure: comparable block structure was observed when the navigation period was uniformly rescaled and with piecewise rescaling between alternative maze landmarks (Supp. Fig. 1). Together, these results demonstrate that neural populations

adopt distinct states within each task epoch, with transitions reliably aligned to behavioral events.

To verify these results were not dependent on temporal scaling, we rescaled trials by distance along the linearized maze. This approach similarly revealed within-epoch stability of the population response (Fig. 2d), confirming that hippocampal activity exhibits stable states within each phase of the task. While this spatial representation revealed within-epoch stability, it compresses temporally extended behavioral periods (e.g. turns) into single spatial coordinates. This binning increased off-diagonal correlations and emphasized the consistency within the long maze arms, but obscured the fine temporal structure evident in event-aligned data. Overall, these analyses reveal that hippocampal population activity is organized into discrete epochs delimited by behaviorally salient events, a structure that reflects the temporal architecture of the task itself.

Agnostic clustering of population activity matches behavioral epochs

To examine the variability of hippocampal dynamics on individual trials in light of this observed epoch structure, we analyzed pseudopopulations constructed from independently recorded neurons. Pseudopopulations provide a stringent test of whether behavioral epochs reflect generalizable organizational principles: if discrete clusters emerge from randomly sampled neurons, it indicates that the variability is sufficiently constrained to preserve distinct, epoch-specific population states. We generated large synthetic ensembles by sampling single trials from each neuron with replacement (1,000 left and 1,000 right pseudotrials), preserving each neuron's temporal profile and trial-to-trial variability while achieving dense sampling across the recorded population (588 neurons). At each moment in the task, activity across all neurons defined a population vector capturing the instantaneous network state. When projected into low-dimensional space, the vectors formed distinct clusters mirroring the behavioral epochs of the task (Fig. 2e).

We applied unsupervised clustering to the pseudopopulation activity, entirely blind to behavioral labels. Each pseudotrial yielded 68 population vectors (one per timepoint), and all vectors were pooled across trials and timepoints before clustering (136,000 total vectors from 2,000 pseudotrials). Clustering was performed in the native 588-dimensional neural space using Pearson correlation as the similarity metric. This approach was agnostic to trial identity, timepoint, and behavioral labels (Fig. 2f), providing a data-driven method to assess whether latent structure aligns with behavioral segmentation.

Unsupervised clustering revealed that pseudopopulation vectors organized into discrete, highly reproducible clusters. Two independent algorithms—Leiden community detection and k-means clustering—converged on highly similar partitions of the neural state space across parameter sweeps (Leiden: kNN $\in \{12, 24\}$, $\gamma \in \{0.2, 0.6, 1.0\}$; k-means: $k \in \{6–12\}$). Cluster assignments were consistent across repeated runs, with within-resolution Adjusted Rand Indices (ARIs) exceeding null distributions generated from comparing circularly shifted pseudopopulations (Mann-Whitney U, all $P<0.001$, Supp Fig. 2). Cross-parameter comparisons likewise yielded high ARIs (Mann-Whitney U, all $P<0.001$), indicating that similar cluster structure was preserved across parameter settings and that population organization was robust to changes in clustering resolution and algorithmic parameters. This convergence demonstrates that population vectors organize into discrete, stable clusters, revealing a structure that is intrinsic to hippocampal dynamics rather than an artifact of any particular clustering approach.

These clusters exhibited striking temporal organization. When mapped back onto the trial timeline, each cluster formed a nearly contiguous block rather than fragmenting randomly across time, defining discrete epochs that were reproducibly expressed across trials (Fig. 2g). Cluster transitions were remarkably consistent across clustering approaches and replications, with boundaries between clusters aligning at the same positions in the trial timeline (Supp Fig. 2). Comparing boundaries from each clustering resolution to those from circularly shifted pseudopopulations confirmed that the observed boundary configuration was not expected by chance (Benjamini–Hochberg–corrected $P < 0.01$ for ≥ 3 matching peaks across all resolutions). These transition

locations remained stable across all parameter sets and both clustering algorithms, even when the total number of clusters varied. While clustering resolution affected the number of clusters identified, higher resolutions split existing clusters rather than revealing new temporal boundaries, indicating that these transition points reflect fundamental organizing features of the population dynamics.

We next tested whether these cluster-derived boundaries align with behavioral structure by comparing boundary timing to the events that delimit the behavioral epochs (Fig 2a). Focusing on boundary timing rather than cluster identity allowed us to assess alignment independent of clustering granularity, since boundaries remained stable even when cluster identities and cluster counts differed across resolutions. We quantified alignment using two complementary metrics: F1 scores, which balance the rate of correctly identified behavioral epoch transitions against false detections, and simple hit rates, the proportion of behavioral epoch transitions that overlapped with a cluster boundary. Across all parameter sets, true boundaries showed markedly higher alignment with behavioral epochs than null distributions (Fig. 2h; Mann-Whitney tests: all $P<0.01$). Mid-level resolutions produced clusters that nearly matched behavioral epochs in both number and transition timing, resulting in the highest precision (F1 scores) and near-ceiling accuracy (hit rates). Both higher and lower resolutions identified behavioral boundaries accurately but with reduced precision due to cluster count mismatches. Together, these results demonstrate that hippocampal population activity adopts discrete states that transition precisely at behaviorally relevant events—forming a structure robust enough to emerge from wholly unsupervised clustering.

Variability across phases of the task reveals two modes of hippocampal processing

Having established epochal segmentation of hippocampal population activity, we next asked whether the fidelity of the population representation varies across the task. Specifically, we asked whether certain behavioral moments are encoded more reliably than others. We quantified how consistently population activity patterns were expressed by correlating each trial's activity against the leave-one-out population mean

at each time bin, yielding a trial-by-template correlation matrix (Fig. 3a). Extracting the diagonal of this matrix—which represents the similarity between corresponding time bins—provided a continuous measure of trial-to-template similarity across the trial (Fig. 3b). This measure revealed that similarity increased sharply at behaviorally-relevant events, indicating highly stereotyped population states, while the intervening epochs showed substantially lower similarity, indicating greater trial-to-trial variability. We constructed null distributions by disrupting temporal alignment through independent circular rotations of each neuron's activity (1,000 iterations). True trial-to-template similarity significantly exceeded null expectations in restricted task windows (permutation tests, Benjamini-Hochberg corrected, $P < 0.05$), confirming that instances of high reliability reflected coordinated alignment of the population with task structure.

To further examine how reliably population activity patterns map onto task progression, we asked whether task progression could be decoded from the population activity across individual trials. Using a Bayesian decoder with a Poisson-based framework, we constructed templates from half of the trials and decoded the remaining half. For each time bin in a held-out trial, the decoder produced a posterior probability distribution over all time bins in the standardized trial, representing the likelihood that the observed population activity corresponded to each moment in the task.

Posterior probability distributions revealed differential precision across the task (Fig. 3d). At behaviorally salient events, posteriors were tightly concentrated around the true time bin, indicating precise temporal decoding. By contrast, during the epochs between events, posteriors were far more diffuse, spreading across multiple bins within the epoch. Remarkably, while posteriors were diffuse within epochs, they respected epoch boundaries: the decoder reliably identified *which* epoch was occurring, even when uncertain about *when* within an epoch. This pattern reveals that hippocampal population activity provides precise temporal information at behavioral transitions but represents moments within epochs more categorically, reliably signaling epoch identity while being less informative about exact moment within that epoch.

We further quantified these observations by measuring epoch-level decoding accuracy (Fig. 3f). Trial-type-specific decoding—requiring identification of both epoch and left/right

trial type—showed high accuracy across epochs, although some epochs discriminated trial type better than others (all epochs >90% accurate, except ITI at 70%; χ^2 tests $P<0.001$ for each). When trial type was ignored, epoch identity was decoded with near-perfect accuracy (all >95% accurate; χ^2 tests $P<0.001$ for each). Thus, the population activity reliably signals the current phase of the task irrespective of its ability to disambiguate how much time has passed within an epoch.

Beyond accuracy, posterior entropy provides a complementary measure that quantifies uncertainty of the decoded state (Fig. 3e). Entropy was lowest at event boundaries and elevated within epochs, quantitatively confirming that the decoder was highly certain at the times of events but uncertain about precise temporal position within epochs (permutation tests, Benjamini-Hochberg corrected, $P < 0.05$). This matched the structure observed in the trial-to-template correlation analysis (Fig. 3b): both revealed high precision during behaviorally-relevant events, and elevated uncertainty during the intervening epochs, despite accurate decoding of epoch identity.

To probe the basis of this temporal structure, we next examined when individual neurons were most informative about task progression. We computed information-by-bin—a time-resolved estimate of mutual information between instantaneous firing rate and position in the task timeline. Summing across all recorded neurons yielded the total information capacity of the population at each time bin. To assess significance, we compared this to null distributions generated by independently circularly rotating each neuron's trials, recalculating information-by-bin, and summing across the population. The population's aggregate information showed prominent peaks aligned with behavioral events and epoch boundaries that significantly exceeded null expectations, with reduced information content in the intervening intervals (Fig. 3c). These information peaks mirrored the temporal structure revealed by the correlation and decoding analyses, with population activity most reliable at behavioral events and reduced information within intervening epochs (permutation tests, Benjamini–Hochberg corrected, $P<0.05$).

Across three independent analytical frameworks—trial-to-template correlation, Bayesian decoding, and information content—a consistent pattern emerged. Hippocampal

population activity exhibits two distinct representational modes that alternate across the task. At salient behavioral events, the population activity converges into precise, highly reliable states across trials that are informative about task progression. Between these events, activity is less reliable, exhibiting greater trial-to-trial variability and reduced temporal precision. Critically, this within-epoch variability does not reflect noise or decoupling from task structure: population activity remains confined to epoch-specific subspaces, reliably signaling which phase of the task is occurring even while being less informative about exact temporal position within that phase. This dual architecture points toward a fundamental organizing principle: the hippocampus does not encode experience as a continuous temporal stream but rather constructs a temporal scaffold anchored to behaviorally-relevant events, with the intervals between these events represented categorically. This architecture may explain how the hippocampus balances the need for precise temporal ordering of events with the flexibility to generalize across variable durations and trajectories.

Neuronal responses reflect abstract task structure

To test whether hippocampal activity reflects environment-specific features or the abstract structure of task phases, monkeys performed an A–B–A′ sequence of visually distinct but geometrically identical Y-mazes, with each maze run separated by >15 minutes of quiescence in dim lighting. The virtual environments were designed to maximize perceptual dissimilarity, with different landmark positions and stark contrasts in colors, surface textures, and decorations (Fig. 4a). Maze A′ repeated the visual appearance of maze A to assess within-session stability. We analyzed 306 neurons that were significantly responsive in maze A, maze B, or both across the full A–B–A′ regimen.

In rodent hippocampus, exposure to new environments typically triggers global remapping, where new place fields emerge and the active population changes[50]. By contrast, we found that many primate hippocampal neurons maintained stable responses across visually distinct mazes. Stable neurons exhibited significant mutual information with task progression in both mazes (P<0.01) and Pearson correlations of

trial-averaged firing rates greater than 0.7. Of the 91 neurons exhibiting stable responses between repeated presentations of maze A (A–A′), 50 (55%) preserved their response profiles between mazes A and B (Fig 4b,c). The unexpected, low A-A' stability (91/306) echoes the developing literature on representational drift[51] – albeit on a faster, within-day timescale[52]. Especially in this context, the substantial preservation of responses across visually distinct environments (A-B) demonstrates that hippocampal activity can encode task structure independent of visual features.

Event-specific analyses revealed differential stability across task phases. Restricting our analysis to neurons that were stable across A–A′, we assessed how many preserved their peri-event activity patterns (±2 seconds) when the environment changed (A–B). Stability remained high at Movement start (64%) and ITI start (71%) but was substantially lower at Choice (30%) and Outcome (35%) turns, and Reward start was intermediate (53%) (Fig. 4d; $\chi^2(4)= 31.30$, $P<0.001$). Post-hoc Fisher exact tests with Holm correction confirmed that Movement start and ITI start each differed significantly from Choice and Outcome events (adjusted $P<0.05$). These results indicate that the earliest and latest events of the trial carry the most task-generalizable coding, whereas the events where the subject must attend to make a choice and evaluate the outcome are less likely to be stable.

Stability across left and right trials followed a similar organizational principle. Only 14% of responsive neurons were stable across left versus right trials at the Choice turn and 13% at the Outcome turn, compared to substantially higher stability at Movement start (42%) and ITI start (35%) (Fig. 4e; $\chi^2(4)=125.38$, $P<0.001$). Post-hoc Fisher exact tests with Holm correction confirmed that Movement start and ITI start each differed significantly from Choice, Outcome, and Reward events (adjusted $P<0.05$). Together, these findings reveal a shared structure across maze and trial-type generalization: stable responses concentrate at the beginning and end of the trial, whereas decision and feedback epochs exhibit more labile coding where attention to task contingencies is necessary to be successful.

Population-level analyses confirmed that hippocampal representations generalize across visual environments. A Bayesian decoder trained on maze A accurately tracked

task progression when applied to maze B, with posterior probability patterns and entropy closely matching within-environment decoding (A′ from A) (Fig. 4f). Decoding entropy across the trial was indistinguishable between A→B and A→A′, indicating that overall uncertainty was not affected by environment identity (Wilcoxon rank-sum test with Benjamini-Hochberg correction, $P>0.05$ all bins; Fig. 4g). The characteristic two-mode structure (i.e. Fig 3d) persisted: posteriors were sharply concentrated at event boundaries and more diffuse within epochs, with diffusion constrained by epoch identity. Epoch-level decoding accuracy was higher in the repeated presentation of the same maze for most epochs than for the maze with a different appearance (χ^2 tests $p<0.01$; Fig. 4h). However, decoding epochs of maze B from maze A significantly exceeded chance (χ^2 tests $P<0.001$ for each), and several epochs were decoded as well as A→A′. Although decoding accuracy was modestly lower for A→B, epoch-level decoding was conserved across environments, indicating that primate hippocampal populations have the capacity to encode the ordinal structure of task events in a manner that abstracts over sensory context. Together, the convergence of single-neuron stability and population-level decoding demonstrates that a substantial portion of the hippocampal population responds to the abstract structure of the task rather than the specific visual environment.

Discussion

The hippocampus is central to memory and cognition, but the organizing principles of its ongoing activity have remained elusive. Here we show that hippocampal population dynamics are organized not by individual sensory or spatial features but by the temporal structure of the engaged task. Individual neurons responded to multiple behavioral events, indicating that they do not act as narrow feature detectors but participate in a distributed representation spanning the full task sequence. These population responses segmented the task into distinct phases, each defined by a unique ensemble code. Transitions were precisely aligned to behavioral events, revealing that the hippocampus partitions experience into meaningful elements. Many neurons exhibited similar activity patterns across visually distinct environments, showing that the underlying code generalizes beyond sensory context[53] to capture the abstract progression of task events. Together, these results reveal that hippocampal activity reflects the learned temporal organization of behavior: a structure that persists across contexts and provides a framework for linking elements within an episode.

A central and novel finding of this study is that hippocampal population activity alternates between two distinct modes of organization across the unfolding of experience. At behaviorally-relevant moments, hippocampal activity converged into tightly coordinated, low-variability configurations that were highly reproducible across trials. Between these events, activity patterns were more dispersed and variable, occupying broader regions of neural state space while remaining confined within the boundaries of each behavioral epoch. One possible explanation for this pattern is that externally driven input at behaviorally salient moments triggers transitions between internally stabilized states. In this framing, event-aligned inputs propel the system from one metastable configuration to the next, while intrinsic hippocampal dynamics support the persistence of each state during the inter-event interval. This dynamic could implement a temporal chunking mechanism in which extrinsic input creates anchors in time, and internally sustained activity bridges the intervals between them—together forming a scaffold for organizing experience. These dual modes suggest that the hippocampus toggles between moments of input-driven “transitions” and intrinsically

organized “maintenance” states, providing a mechanistic bridge that stitches a discontinuous series of goal-directed events into a continuous episode. The ability to link events across time may offer a unifying explanation for the hippocampus’s fundamental contribution to memory and cognition.

These dynamics offer a new lens for interpreting the functional role of the hippocampus. Traditional accounts have largely focused on the content of hippocampal representations, often casting the hippocampus as a conjunctive feature detector[26,27] or part of a circuit specialized for performing metric spatial calculations[2]. More recent frameworks emphasize the process by which these representations evolve over time to structure experience[28,29]. The present findings build on this history by revealing that hippocampal population activity is organized by transitions between externally anchored and intrinsically sustained states, a dynamic architecture that links discrete behavioral events into a coherent temporal progression. This organization resonates with theories of event segmentation and predictive processing: the hippocampus not only encodes the elements of experience but also tracks the ordinal relationships among them[54,55]. Individual events serve as nodes in a graph-like scaffold connected by bridging activity that captures their temporal relationships. By projecting the complex multidimensionality of experience onto a compact cognitive graph[56], the hippocampus provides a framework for organizing behavior and enabling inference across time. The capacity to bind separated events into unified representations may underlie the hippocampus’s broad role in memory and cognition[32,57]—establishing the computational foundation for linking episodes as coherent elements of experience.

Acknowledgements:

This work was supported by the NIH (BRAIN Initiative – U19NS107609 (EAB, ALF); NIMH – R01MH117777 (EAB, JR); NINDS – UF1NS126485 (EAB, ALF); NIDA – 5T90DA032436 (ALF); ORIP – P51OD010425), the McKnight Endowment Fund for Neuroscience (EAB), the Simons Foundation, Simons Collaboration for the Global Brain (EAB, ALF), and the Wayne E. Crill Endowment at the University of Washington (EAB). We are grateful to Megan Jutras for outstanding project support, Sierra Schleufer and Ian O'Leary for excellent technical assistance, and Drs. Charles Gray and Baldwin Goodell for microdrive development. We also thank Drs. Larry Squire, Lara Rangel, Andrew Alexander, and Gregory Horwitz for insightful discussions.

References:

1. O'Keefe, J. & Dostrovsky, J. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* **34**, 171–175 (1971).
2. McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I. & Moser, M.-B. Path integration and the neural basis of the "cognitive map." *Nat. Rev. Neurosci.* **7**, 663–678 (2006).
3. McNaughton, N. & Gray, J. A. Anxiolytic action on the behavioural inhibition system implies multiple types of arousal contribute to anxiety. *J. Affect. Disord.* **61**, 161–176 (2000).
4. Squire, L. R. Memory systems of the brain: a brief history and current perspective. *Neurobiol. Learn. Mem.* **82**, 171–177 (2004).
5. Squire, L. R. Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychol. Rev.* **99**, 195–231 (1992).
6. Aronov, D., Nevers, R. & Tank, D. W. Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit. *Nature* **543**, 719–722 (2017).
7. Danjo, T., Toyoizumi, T. & Fujisawa, S. Spatial representations of self and other in the hippocampus. *Science* **359**, 213–218 (2018).
8. Gauthier, J. L. & Tank, D. W. A dedicated population for reward coding in the hippocampus. *Neuron* **99**, 179-193.e7 (2018).
9. Radvansky, B. A. & Dombeck, D. A. An olfactory virtual reality system for mice. *Nat. Commun.* **9**, 839 (2018).

10. Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C. & Fried, I. Invariant visual representation by single neurons in the human brain. *Nature* **435**, 1102–1107 (2005).

11. Pastalkova, E., Itskov, V., Amarasingham, A. & Buzsáki, G. Internally generated cell assembly sequences in the rat hippocampus. *Science* **321**, 1322–1327 (2008).

12. Markus, E. J. *et al.* Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *J. Neurosci.* **15**, 7079–7094 (1995).

13. Wirth, S. *et al.* Single neurons in the monkey hippocampus and learning of new associations. *Science* **300**, 1578–1581 (2003).

14. Gulli, R. A. *et al.* Context-dependent representations of objects and space in the primate hippocampus during virtual navigation. *Nat. Neurosci.* **23**, 103–112 (2020).

15. Manns, J. R., Hopkins, R. O. & Squire, L. R. Semantic memory and the human hippocampus. *Neuron* **38**, 127–133 (2003).

16. Vargha-Khadem, F. *et al.* Differential effects of early hippocampal pathology on episodic and semantic memory. *Science* **277**, 376–380 (1997).

17. Zola, S. M. *et al.* Impaired recognition memory in monkeys after damage limited to the hippocampal region. *J. Neurosci.* **20**, 451–463 (2000).

18. Schapiro, A. C., Gregory, E., Landau, B., McCloskey, M. & Turk-Browne, N. B. The necessity of the medial temporal lobe for statistical learning. *J. Cogn. Neurosci.* **26**, 1736–1747 (2014).

19. Hassabis, D., Kumaran, D., Vann, S. D. & Maguire, E. A. Patients with hippocampal amnesia cannot imagine new experiences. *Proc. Natl. Acad. Sci. U. S. A.* **104**, 1726–1731 (2007).

20. Vikbladh, O. M. *et al.* Hippocampal contributions to model-based planning and spatial memory. *Neuron* **102**, 683-693.e4 (2019).

21. Moscovitch, M., Cabeza, R., Winocur, G. & Nadel, L. Episodic memory and beyond: The hippocampus and neocortex in transformation. *Annu. Rev. Psychol.* **67**, 105–134 (2016).

22. Hartley, T., Lever, C., Burgess, N. & O'Keefe, J. Space in the brain: how the hippocampal formation supports spatial cognition. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **369**, 20120510 (2014).

23. MacDonald, C. J., Lepage, K. Q., Eden, U. T. & Eichenbaum, H. Hippocampal "time cells" bridge the gap in memory for discontiguous events. *Neuron* **71**, 737–749 (2011).

24. Kraus, B. J., Robinson, R. J., 2nd, White, J. A., Eichenbaum, H. & Hasselmo, M. E. Hippocampal "time cells": time versus path integration. *Neuron* **78**, 1090–1101 (2013).

25. Wood, E. R., Dudchenko, P. A., Robitsek, R. J. & Eichenbaum, H. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron* **27**, 623–633 (2000).

26. Eichenbaum, H. On the integration of space, time, and memory. *Neuron* **95**, 1007–1018 (2017).

27. Bellmund, J. L. S., Gärdenfors, P., Moser, E. I. & Doeller, C. F. Navigating cognition: Spatial codes for human thinking. *Science* **362**, eaat6766 (2018).

28. George, D. *et al.* Clone-structured graph representations enable flexible learning and vicarious evaluation of cognitive maps. *Nat. Commun.* **12**, 2392 (2021).

29. Whittington, J. C. R., McCaffary, D., Bakermans, J. J. W. & Behrens, T. E. J. How to build a cognitive map. *Nat. Neurosci.* **25**, 1257–1272 (2022).

30. Wallenstein, G. V., Eichenbaum, H. & Hasselmo, M. E. The hippocampus as an associator of discontiguous events. *Trends Neurosci.* **21**, 317–323 (1998).

31. Levy, W. B. A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus* vol. 6 579–590 Preprint at https://doi.org/10.1002/(sici)1098-1063(1996)6:6<579::aid-hipo3>3.0.co;2-c (1996).

32. Rueckemann, J. W. & Buffalo, E. A. Spatial Responses, Immediate Experience, and Memory in the Monkey Hippocampus. *Curr Opin Behav Sci* **17**, 155–160 (2017).

33. Buzsáki, G. & Tingley, D. Space and time: The hippocampus as a sequence generator. *Trends Cogn. Sci.* **22**, 853–869 (2018).

34. Shimbo, A., Izawa, E.-I. & Fujisawa, S. Scalable representation of time in the hippocampus. *Science Advances* vol. 7 eabd7013 Preprint at https://doi.org/10.1126/sciadv.abd7013 (2021).

35. Grieves, R. M., Wood, E. R. & Dudchenko, P. A. Place cells on a maze encode routes rather than destinations. *Elife* **5**, (2016).

36. Zutshi, I. *et al.* Hippocampal neuronal activity is aligned with action plans. *Nature* **639**, 153–161 (2025).

37. Crivelli-Decker, J. *et al.* Goal-oriented representations in the human hippocampus during planning and navigation. *Nat. Commun.* **14**, 2946 (2023).

38. Sun, C., Yang, W., Martin, J. & Tonegawa, S. Hippocampal neurons represent events as transferable units of experience. *Nat. Neurosci.* **23**, 651–663 (2020).

39. Terada, S., Sakurai, Y., Nakahara, H. & Fujisawa, S. Temporal and Rate Coding for Discrete Event Sequences in the Hippocampus. *Neuron* **94**, 1248-1262.e4 (2017).

40. Allen, T. A., Salz, D. M., McKenzie, S. & Fortin, N. J. Nonspatial sequence coding in CA1 neurons. *J. Neurosci.* **36**, 1547–1563 (2016).

41. Fortin, N. J., Agster, K. L. & Eichenbaum, H. B. Critical role of the hippocampus in memory for sequences of events. *Nat. Neurosci.* **5**, 458–462 (2002).

42. Paz, R. *et al.* A neural substrate in the human hippocampus for linking successive events. *Proc. Natl. Acad. Sci. U. S. A.* **107**, 6046–6051 (2010).

43. Barnett, A. J. *et al.* Hippocampal-cortical interactions during event boundaries support retention of complex narrative events. *Neuron* **112**, 319-330.e7 (2024).

44. Baldassano, C. *et al.* Discovering event structure in continuous narrative perception and memory. *Neuron* **95**, 709-721.e5 (2017).

45. Zheng, J. *et al.* Neurons detect cognitive boundaries to structure episodic memories in humans. *Nat. Neurosci.* **25**, 358–368 (2022).

46. Cohn-Sheehy, B. I. *et al.* The hippocampus constructs narrative memories across distant events. *Curr. Biol.* **31**, 4935-4945.e7 (2021).

47. Ezzyat, Y. & Davachi, L. Similarity breeds proximity: pattern similarity within and across contexts is related to later mnemonic judgments of temporal proximity. *Neuron* **81**, 1179–1189 (2014).

48. Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).

49. Yoo, S. B. M. & Hayden, B. Y. The Transition from Evaluation to Selection Involves Neural Subspace Reorganization in Core Reward Regions. *Neuron* **105**, 712-724.e4 (2020).

50. Colgin, L. L., Moser, E. I. & Moser, M.-B. Understanding memory through hippocampal remapping. *Trends Neurosci.* **31**, 469–477 (2008).

51. Driscoll, L. N., Duncker, L. & Harvey, C. D. Representational drift: Emerging theories for continual learning and experimental future directions. *Curr. Opin. Neurobiol.* **76**, 102609 (2022).

52. Mankin, E. A. *et al.* Neuronal code for extended time in the hippocampus. *Proc. Natl. Acad. Sci. U. S. A.* **109**, 19462–19467 (2012).

53. Baraduc, P., Duhamel, J.-R. & Wirth, S. Schema cells in the macaque hippocampus. *Science* **363**, 635–639 (2019).

54. Rueckemann, J. W., Sosa, M., Giocomo, L. M. & Buffalo, E. A. The grid code for ordered experience. *Nat. Rev. Neurosci.* **22**, 637–649 (2021).

55. Stachenfeld, K. L., Botvinick, M. M. & Gershman, S. J. The hippocampus as a predictive map. *Nat. Neurosci.* **20**, 1643–1653 (2017).

56. Muller, R. U., Stead, M. & Pach, J. The hippocampus as a cognitive graph. *J. Gen. Physiol.* **107**, 663–694 (1996).

57. Buffalo, E. A. Bridging the gap between spatial and mnemonic views of the hippocampal formation. *Hippocampus* **25**, 713–718 (2015).

Methods:

Animals. All experiments were carried out in accordance with National Institutes of Health guidelines and were approved by the University of Washington Institutional Animal Care and Use Committee. Three adult male rhesus monkeys (*Macaca mulatta*; M1: 9.3 kg, 7.5 years old; M2: 9.1 kg, 8.5 years old; M3: 12.3 kg, 8 years old) were used in the experiments.

Y-maze behavioural task. Before experimentation, monkeys were trained to manipulate a joystick mounted to their protective chair. Subjects first learned to center and approach a visual cue displayed on an immersive screen to obtain reward. After mastering joystick control, animals were trained to navigate simple linear paths in an immersive virtual environment. Virtual environments and task mechanics were developed in the Unity game engine.

Monkeys then performed a spatial alternation task in a Y-shaped virtual maze. The maze comprised three arms converging at a central choice point. Subjects navigated a semi-constrained path via joystick, which prevented backward movement or lateral deviation except at designated turn points.

Each trial began with the subject positioned at the base of the central stem. After traversing the stem, the subject reached a junction where it could select the left or right side arm. Reward contingencies required alternation between arms across successive trials, such that the correct choice on a given trial depended on remembering the outcome of the previous trial. At the end of a correct arm, subjects rounded a second corner to encounter a visual cue (a banana) marking the reward location. The cue was occluded until the corner was reached, ensuring that the outcome remained unknown until the side arm had been fully traversed. Contact with the banana triggered delivery of a slurry of banana and monkey chow (Purina).

On correct trials, reward was available for ~2s, followed by 1s of immobility before the screen went black for a 4s inter-trial interval (ITI). On incorrect trials, subjects were prevented from rounding the final corner, and the ITI was initiated immediately upon reaching the end of the chosen arm. Following the ITI, the subject was repositioned at the start of the central stem for the next trial.

Ten visually-distinct virtual environments were created for this experiment, and monkeys experienced multiple mazes in an A-B-A design, allowing for the determination of environment-specific responses in the neurons (i.e. remapping).

The events used to calibrate peri-event spiking activity (Figure 1) were (i) movement initiation on the center stem (Start), (ii) turn initiation at the choice point (Choice), (iii) turn initiation toward reward location (Outcome), (iv) reward delivery (Reward), and (v) onset of the inter-trial interval (ITI).

Implantation of chronic microdrive. A titanium headpost (Gray Matter Research) was implanted near the posterior vertex of the skull with screws. After acclimation and task

training, each monkey was implanted with a 124-channel microdrive (LS124; Gray Matter Research) targeting the full longitudinal axis of the right hippocampus. Chambers and drives were custom fit to each animal's skull and dural surface using preoperative MRI, ensuring electrode trajectories were directed toward hippocampus. Each channel carried an independently-movable tungsten electrode (100μm diameter, ~1 MΩ impedance; FHC) with more than 40 mm travel distance from the bottom of the microdrive. Implantation followed a three-stage procedure: 1) the chamber was implanted and sealed; 2) after healing and verification of sterility, a craniotomy was performed and sealed with a form-fitting plug; and 3) the sterile, custom microdrive was secured within the chamber and hermetically sealed.

Electrodes were advanced to ~10 mm below the brain surface on the day of implantation, then lowered over the next week to 2–3 mm above target. Each electrode was advanced via a fine-threaded screw (125μm per revolution).

Electrode positioning. During sedation, a sub-millimeter 3D surface scan of each animal's head and implant was obtained (Spyder; Artec). The surface mesh was registered with the preoperative MRI (Slicer 5.6.1) to ensure accurate alignment of the chamber relative to hippocampal anatomy and to generate precise trajectory projections for each electrode. Electrode depth was tracked through detailed turn records, allowing estimation of recording target positions. Adjustments were guided by audio monitoring of multi-unit spiking activity and visual inspection of local field potentials, which helped account for potential deviations in electrode trajectory. Electrodes were advanced only after completion of a recording session, allowing them to settle for at least a day before the next session to mitigate within-session drift. Not every electrode was moved each day; instead, advancing a subset was often sufficient to shift surrounding electrodes into fruitful recording positions.

Histological reconstruction. At the conclusion of experiments, 30μA of direct current was passed through each electrode to produce small electrolytic lesions at the electrode tips. After a minimum two-day survival period to allow gliosis, animals were transcardially perfused with 4% paraformaldehyde with electrodes in place. Following 48 hours of post-fixation, the microdrive was carefully withdrawn. The skull was positioned in a stereotax, and the brain was blocked in the coronal plane to facilitate matching of histological sections to the preoperative MRI. Tissue was cryoprotected by graded infiltration of glycerol and DMSO, before being rapidly frozen in −70°C isopentane[1]. Brains were coronally sectioned at 50μm, and every third section (150μm spacing) was Nissl-stained with cresyl violet.

A 10x micrograph was acquired for each Nissl section. Custom software used an iterative alignment procedure to register each section both to the preoperative MRI and to adjacent sections, resulting in a 3D Nissl volume stitched from the micrographs and coregistered to the MRI. A unique rescaling factor (1.75–2.50%, empirically determined for each animal) was applied to compensate for tissue shrinkage during cryoprotection.

The 3D Nissl volume and projected trajectories from the 3D surface scan were used to identify each electrode track in the tissue. This alignment enabled full reconstruction of the true path taken by each electrode to its terminal lesion. Daily recording positions were estimated from turn records as the relative distance along the electrode tract from the electrolytic lesion site. For electrodes that followed a straight trajectory, premortem position estimates were typically accurate within ~250μm; however, a notable subset showed deviations due to bending.

Electrophysiological recording. Neural signals were recorded with a Neuralynx Digital Lynx SX system (Neuralynx Inc., USA) using 128 channels sampled at 30 kHz with 24-bit A/D conversion. Broadband signals were acquired and grounded to the titanium recording chamber, with all channels referenced to ground.

Spike detection was performed offline. Continuous data were first high-pass filtered using a 4th-order Butterworth filter at 400Hz (zero-phase, 'filtfilt' MATLAB). Candidate spikes were identified by threshold crossing at 3.5 standard deviations below the mean of the filtered trace.

Spike sorting. Candidate units were clustered automatically with Isosplit (MountainSort[2]). Clusters were then triaged using custom visualization software that projected the spikes onto peak–valley and principal-component axes. This tool enabled rapid identification of (i) MountainSort-identified clusters consistent with noise for exclusion and (ii) channels without MountainSort clusters that contain putative units for rescue.

All putative clusters were subsequently curated manually in Offline Sorter (Plexon, USA). Sorting decisions were based on cluster density and separation in feature space. Units were accepted if they formed a stable and distinct cluster with consistent waveform morphology; clusters showing a lack of separation, infiltration by noise, or significant drift over the course of recording were excluded.

Temporal standardization of trials. To enable comparison across trials of variable durations, each trial was normalized to a standardized timeline sampled at 68 evenly spaced predictor bins. Because only the movement portion of the task varied in duration, four rescaling schemes were evaluated for this interval. The primary approach divided trials into six task epochs (Movement Start, Choice Turn, Choice Arm, Outcome Arm, Reward, and Inter-Trial Interval), with each epoch linearly rescaled to match its mean duration across sessions, yielding a 17-second trial with 250-ms bin resolution. Three alternative schemes served as controls: (1) uniform rescaling across the entire movement period, (2) piecewise rescaling between alternative landmarks (Movement Start, Center Stem Midpoint, Side Arm Midpoint, Reward, and Inter-Trial Interval), and (3) nonlinear rescaling based on distance in the linearized maze rather than elapsed time. Results throughout use the primary piecewise method; all schemes yielded consistent findings.

Spikes were first binned at 50-ms resolution to preserve fine temporal structure. These 50-ms rates were then mapped to the nearest predictor bin and averaged when multiple samples fell within the same bin. For population analyses, these averaged rates provided one value per predictor bin per trial. For mutual information calculations, individual 50-ms samples were retained to capture within-bin variability.

Neuron inclusion criteria. Units were considered eligible for analysis if they fired at least 100 spikes within at least one maze during a given session, provided that maze also contained ≥50 correct trials. In addition, the unit was required to fire at least once on ≥25% of those trials.

To be classified as task-responsive, a neuron was required to show both elevated information content and consistent trial-level activity patterns. For each neuron, trial activity was represented as binned spiking across the full trial timeline (predictor details below). Three statistics were computed: (i) Shannon mutual information between binned spikes and trial progression, (ii) correlation of trial-averaged activity between the first and second halves of trials, and (iii) correlation between odd- and even-numbered trials of the same trial type (left–left, right–right). Significance for all three measures was determined relative to the same null distribution, generated by independently circularly shifting spikes within each trial. A neuron was classified as task-responsive if its observed statistic exceeded the 99th percentile of this null.

To mitigate the possibility of counting the same neuron across days, multiple criteria were applied to define unit uniqueness. A unit was considered distinct if (i) the electrode from which it was recorded had been advanced, (ii) more than 5 days had elapsed since the last included recording on that channel, or (iii) the Pearson correlation of its trial-averaged activity with that of the previous day was <0.5. In addition, neurons were conservatively excluded if their mean activity subjectively appeared similar to nearby recordings during manual curation.

Shannon mutual information. Mutual information between instantaneous firing rate and bin identity was computed from all 50-ms rate samples. Let *r* denote the discretized firing rate and *x* the predictor bin (68 x 2 trial-type). Information was given by:

The marginal and conditional probabilities were estimated as $P(r)$and $P(r \mid x)$, and information was given by:

$$I_{\text{Shannon}} = \sum_{x} P(x) \sum_{r} P(r \mid x) \log_2 \frac{P(r \mid x)}{P(r)}$$

where P(x) is the probability of occupying bin x, P(r|x) is the conditional probability of observing rate r in bin x, and P(r) is the probability of observing rate r. Spike counts were quantized in two-spike increments to keep response categories compact in high-firing neurons, maintaining a manageable dynamic range and ensuring comparability of information estimates across neurons.

Significance was assessed against a within-trial circular-shift null (1000 iterations) in which each neuron's spike train was independently circularly shifted within each trial.

Skaggs information. Information per spike was computed according to Skaggs et al. (1993), using the trial-average rate curve, r(x) obtained by averaging 50-ms rate samples within each predictor bin. Trial-average rate curves were Gaussian smoothed with a kernel of 3 bins.

$$I_{\text{Skaggs}} = \frac{\sum_x P(x)\, r(x) \log_2 \left(\frac{r(x)}{\bar{r}}\right)}{\bar{r}}$$

where $\bar{\boldsymbol{r}} = \sum_x P(x)\, r(x)$. This normalization expresses information in bits per spike, facilitating comparison across neurons with different overall firing rates.

Sparsity. Tuning compactness of the smoothed trial-averaged rate curves r(x) was quantified by probability-weighted sparsity (Skaggs et al., 1996):

$$Sparsity = \frac{\left(\sum_x P(x)\, r(x)\right)^2}{\sum_x P(x)\, r(x)^2}$$

Population analysis

Population activity across the standardized trial is visualized as an NxB matrix of mean firing rates, where N is the number of neurons and B is the number of predictor bins (68 * 2 trial types). For each neuron, firing rates were averaged across all correct trials and linearly rescaled from 0 to 1 by dividing by the neuron's maximum rate. For visualization, neurons were grouped by the similarity of their mean activity profiles using hierarchical clustering, and the resulting clusters were reordered to reflect the general temporal progression of activity across the trial (Fig. 2b).

Rate soft-normalization. To balance the contribution of neurons with differing overall firing rates, each neuron's binned activity was soft-normalized[3,4]. For each neuron, the full set of predictor-binned firing rates—concatenated across all trials—was treated as a single time series. The transformed rate was given by

$$r_i'(t) = \frac{r_i(t)}{(\max\, r_i - \min\, r_i) + 0.2 \max\, r_i},$$

where $r_i(t)$denotes the binned firing rate of neuron i at time t.

This transformation preserves each neuron's relative modulation while reducing disparities in scale across the population. The soft-normalized firing rates were used in subsequent population-level analyses, including temporal correlations, agnostic clustering, and dimensionality analyses, but not Bayesian decoding.

Temporal cross-correlations. To quantify the temporal organization of ensemble activity, we computed pairwise correlations between population vectors across all predictor bins. For each bin, the population vector comprised the soft-normalized firing

rates of all neurons. The resulting BxB matrix captured the similarity of ensemble activity patterns across trial progression, with blocks of high correlation indicating stable population states and sharp boundaries marking event-aligned transitions.

Behavior agnostic clustering

To identify latent structure in population dynamics independent of task labels, we applied unsupervised clustering to neural activity patterns without incorporating behavioral epoch boundaries or trial-type information. Because neurons were recorded across different sessions, we constructed pseudopopulations that combined neurons while preserving each cell's temporal dynamics and trial-to-trial variability.

Pseudopopulation construction. For each trial type, pseudotrials were generated by randomly sampling one trial from each neuron and assembling their binned, soft-normalized firing rates into a common population vector. Each pseudotrial spanned 68 time bins across the standardized trial. This procedure was repeated 1,000 times per trial type to yield sufficient pseudotrials for clustering analyses. Pseudopopulations included all task-responsive neurons and maintained each neuron's trial-level variability while enabling analyses across the full recorded population.

Null distributions. To assess statistical significance of clustering structure, null pseudopopulations were created by disrupting temporal coordination across neurons while preserving each neuron's firing rate distribution. For each neuron, the binned rate time series was circularly shifted by a random offset independently across trials, eliminating temporal alignment across the population but maintaining within-trial firing statistics. Pseudotrials were then assembled from these permuted data using the same sampling procedure, repeated 1,000 times to produce null ensembles matched in size and composition to the empirical sets. Test statistics from empirical pseudopopulations were compared against distributions generated from these null ensembles.

Clusters methods. Population vectors from all time bins and trials were pooled and clustered directly in high-dimensional space using Pearson correlation distance. Dimensionality reduction via PCA was used only for visualization, not for clustering. Two complementary unsupervised methods were applied:

Leiden community detection. The Leiden algorithm identifies clusters by optimizing community structure in a k-nearest-neighbor (kNN) graph. We constructed kNN graphs with $k = 12$ (corresponding to the connectivity threshold $\ln(n)$ for the dataset size) and $k = 24$ (as a higher-connectivity comparator). For each graph, clustering used the Reichardt–Bornholdt configuration model with three resolution parameters ($\gamma = 0.2, 0.6, 1.0$), which control the propensity to split communities. Each parameter combination was replicated 100 times to assess partition stability.

k-Means clustering. As an alternative approach, k-means clustering was applied with correlation distance for $k = 6–12$ clusters. Each fit was optimized across 100 random

initializations, retaining the solution with the lowest within-cluster sum of squares. This process was repeated ten times for each value of k to demonstrate consistency.

Cluster stability assessment. Clustering consistency across clustering runs was quantified using Adjusted Rand Index (ARI). The Rand Index (RI) measures the proportion of sample pairs on which two labelings agree:

$$\mathrm{RI} = \frac{\mathrm{TP} + \mathrm{TN}}{\mathrm{TP} + \mathrm{FP} + \mathrm{FN} + \mathrm{TN}},$$

where TP and TN are correctly grouped or separated pairs, and FP and FN are errors. The Adjusted Rand Index (ARI) corrects for chance agreement,

$$\mathrm{ARI} = \frac{\mathrm{RI} - E[\mathrm{RI}]}{1 - E[\mathrm{RI}]},$$

where $E[\mathrm{RI}]$is the expected RI under random labeling.

For each parameter set (e.g., Leiden: kNN = 12, γ = 0.6), metrics were computed across all pairs of replicate runs to assess within-parameter stability. Metrics were also compared across different parameter sets to assess consistency of boundary placement across connectivity and resolution choices.

Metrics were computed both directly (unmatched) and after Hungarian matching of cluster identities. Unmatched comparisons penalize both differences in cluster assignments and differences in the number of recovered clusters. Matched comparisons compensate for hierarchical splits or merges and isolate whether partitions preserve consistent boundary placement irrespective of cluster count. This dual approach distinguished instability in the number of clusters from instability in boundary placement, allowing focus on temporal precision of state transitions rather than categorical labels.

Cluster boundary detection

To isolate the temporal precision of population state transitions from differences in cluster granularity, we focused analyses on when population activity changes occurred rather than on the categorical identity of states. This boundary-based approach asks when transitions happen, independent of what the cluster labels are or how many clusters are recovered. The high trial-to-trial consistency of clustering outcomes enabled meaningful comparison of state transition timing across trials and clustering runs.

Template construction. For each clustering run, we identified the most common cluster label at each time bin across trials. These modal labels defined a cluster template representing the dominant segmentation pattern for that run. The template thus captured the typical sequence of population states without requiring that every trial follow this exact pattern.

Boundary identification on individual trials. Trial-specific boundaries were identified as time bins where the cluster label changed relative to the preceding bin. To exclude

brief, single-bin fluctuations within otherwise stable epochs, a boundary was considered legitimate only if the transition occurred within ±3 bins of the corresponding template boundary. This criterion ensured that detected boundaries represented sustained state transitions consistent with the majority pattern across trials.

Aggregation across trials and runs. For each clustering run, the probability of observing a legitimate boundary at each bin was calculated across trials, yielding a boundary likelihood profile. These profiles were then aggregated across the 100 replicate runs within each parameter set (e.g., Leiden: kNN = 12, γ = 0.6), producing a run-by-bin matrix of boundary likelihoods. This matrix served as the foundation for analyses of boundary prevalence and behavioral alignment. Averaging this matrix yielded a boundary prevalence curve (values 0–1) representing the probability of a boundary at each bin for that parameterization. Peaks in the prevalence curve indicated bins where boundaries consistently recurred across replicates.

Equivalent prevalence curves were computed for null pseudopopulations generated by independent circular shifts of neuronal activity.

Statistical significance. To assess significance, we quantified the probability that null boundaries occurred within ±1 bin of each true prevalence peak. Multiple comparisons were controlled using the Benjamini–Hochberg procedure, identifying bins where observed boundary recurrence exceeded null expectations. These boundary prevalence curves were subsequently used to assess alignment with behavioral epoch boundaries.

Clustering alignment to behavioral events. Comparisons were made on boundary locations rather than labels to test whether population transitions aligned with behavioral transitions independent of cluster count or identity. To test alignment with behavior, cluster-derived boundaries were compared to behavioral epoch boundaries using ±1 bin tolerance. The F1 score was used to quantify the accuracy of overlap for solutions across clustering resolutions by accounting for the true and false hit rates. True positives (TP) were cluster boundaries within ±1 bins of an epoch boundary; false positives (FP) were boundaries outside this window; false negatives (FN) were epoch boundaries without a nearby cluster boundary. Precision and recall were calculated as

$$P = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FP}}, \qquad R = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}},$$

and combined as

$$F_1 = 2\,\frac{P\,R}{P + R}.$$

Distributions of F1 scores from true vs. null clusterings were compared with Mann–Whitney tests. Effect sizes between true and null distributions were quantified using Cliff's delta (δ), a nonparametric measure of stochastic dominance defined as

$$\delta = \frac{n_{>} - n_{<}}{n_X n_Y},$$

where $n_{>}$and $n_{<}$are the number of pairwise comparisons in which a true F1 value exceeds or falls below a null value, respectively, and $n_X$and $n_Y$are the sample sizes of

the two groups. Values range from –1 (all null > true) to +1 (all true > null), with 0 indicating complete overlap.

We also used the simple hit rate to determine if the overlap of cluster boundaries with behavior boundaries exceed chance as quantified through Mann–Whitney tests and Cliff's δ effect size.

Trial-specific analyses

Trial-wise correlation. To quantify cross-trial consistency, we compared each trial's binned population vector to templates derived from leave-one-out (LOO) means. Separate templates were constructed for left and right trials by averaging across all other trials, excluding the current one. For each trial, the Pearson correlation was computed between its activity vector at each bin and every bin of the templates, yielding a cross-temporal similarity matrix. Rows of this matrix capture the similarity of each moment in the trial to the entire template trajectory, providing a measure of both moment-to-moment alignment with the expected template and cross-temporal structure that reflects variability in representational entropy.

To summarize alignment over time, we extracted the diagonal of the cross-temporal matrix, representing the similarity between each trial bin and the corresponding bin of the template. These diagonal values were averaged across trials to yield a mean similarity curve, with confidence intervals obtained by repeating the pseudopopulation sampling procedure 1,000 times.

Null distributions were generated by independently circularly rotating each neuron's activity across bins, preserving marginal firing statistics while disrupting coordinated temporal structure. The LOO procedure was repeated 1,000 times on these null populations, producing null similarity matrices and curves. Statistical significance was assessed binwise by comparing the true curve to the null distribution via empirical p-values, with multiple comparisons controlled using the false discovery rate ($q < 0.05$).

Bayesian decoding. Neuronal population vectors were decoded under a Bayesian framework with a Poisson firing model. Pseudopopulation trials were split 50/50: templates were constructed from the mean firing rate of each neuron at each bin in the training set, and the test set was decoded by computing posterior probabilities over bins.

The posterior probability of bin *b* given observed spike counts $\mathbf{n} = (n_1, n_2, \dots, n_N)$ from N neurons was computed via Bayes' rule:

$$P(b \mid \mathbf{n}) = \frac{P(\mathbf{n} \mid b)\, P(b)}{\sum_{b'} P(\mathbf{n} \mid b')\, P(b')}.$$

Here $P(b)$is the prior probability (assumed uniform across bins), and $P(\mathbf{n} \mid b)$ is the likelihood under a Poisson model:

$$P(\mathbf{n} \mid b) = \prod_{i=1}^{N} \frac{[f_i(b)\,\Delta t]^{n_i}\, e^{-f_i(b)\,\Delta t}}{n_i!},$$

where $f_i(b)$ is the mean firing rate of neuron *i* in bin *b* (from the training template) and Δt is the bin duration.

This cross-validation was repeated 1,000 times with random splits to construct confidence intervals for the entropy and accuracy quantifications below.

Decoding uncertainty. Entropy of the posterior probability across trials was used to assess decoding uncertainty. Entropy at each task progression bin was quantified as:

$$H = -\sum P(t) \log\,(P(t))$$

where *t* indexes the bin of population template (templates for left and right trials are concatenated). The mean entropy curve across cross-validations was used for statistical comparison to the null distribution.

Decoding accuracy. Decoding accuracy was defined as the cumulative posterior probability assigned to the correct behavioral epoch. This border of the behavioral epoch was given a ±1 bin tolerance to account for slight timing jitter (~250ms). Accuracy was computed both with trial-type separation (left vs. right) and after collapsing across trial types to account for trial-type specificity versus generalization.

Chance baselines were defined by the fraction of bins belonging to each epoch relative to total template length, adjusted for tolerance windows. Bin-wise accuracies were compared against chance using exact Wilcoxon signed-rank tests (right-tailed), with rank-biserial correlation as effect size. In complementary analyses, mean epoch accuracy was tested relative to chance with a binomial test. These approaches produced consistent results.

Information-by-Bin. To localize how information content varied across the standardized trial structure, we computed information-by-bin, a time-resolved adaptation of the "information-by-position" metric introduced by Olypher and Fenton[5]. Spike counts were quantized in two-spike increments to keep response categories compact in high-firing neurons.

For each neuron, the probability distribution of discretized spike counts was estimated within each predictor bin (x) and compared to the overall spike-count distribution across all bins.
The resulting information-by-bin value,

$$I_{\text{bin}}(x) = \sum_{k} P(k \mid x) \log_2 \frac{P(k \mid x)}{P(k)}$$

indicates how much the spiking statistics in bin x differ from those observed overall—i.e., how informative that bin of the trial is about the neuron's firing pattern.

The total mutual information for the neuron is the weighted sum of these binwise contributions,

$$I_{\text{Shannon}} = \sum_{x} P(x)\, I_{\text{bin}}(x)$$

which expresses the overall reduction in uncertainty about trial time provided by the neuron's firing activity.

Cross-maze analyses

Cross-maze consistency and decoding analysis. Neuronal activity was recorded as animals performed an A–B–A′ sequence of visually distinct but geometrically identical virtual mazes. Maze A′ repeated the original visual appearance to provide a control for within-session stability. To be considered for analysis, a neuron must be significantly responsive in one of the mazes, and each of the other mazes must contain ≥50 correct trials, have at least 100 spikes and the unit was required to fire at least once on ≥25% of trials on each maze. A total of 306 of the 588 significantly-responsive neurons were recorded under the full A-B-A'.

Individual neuron analysis. For each neuron, the average firing rate profile across the full trial timeline was compared pairwise across mazes (A–B, A–A′). A neuron's response was classified as 'stable' between two mazes if: (1) it showed significant peri-event Shannon information in both mazes (as defined in the neuron inclusion criteria above), and (2) the Pearson correlation of averaged firing rates between mazes was >0.7.

For event-specific analyses, peri-event responses were quantified in a ±2 second window centered on each task event (Movement Start, Choice Turn, Outcome Turn, Reward onset, ITI onset). Spike counts were binned at 50-ms resolution (identical to all other MI and trial-level analyses) and averaged across trials within each maze to generate peri-event rate vectors. The same stability criteria used for full-trial responses (significant mutual information in both mazes; Pearson correlation of peri-event rate vectors >0.7) were applied to classify each neuron as stable vs. unstable at each event.

To assess whether stability generalized across trial type rather than across maze appearance, the identical peri-event procedure was repeated for left vs. right trials within the same maze. Event-specific stability rates were therefore computed both for maze identity generalization (A–B) and trial-type generalization (left–right).

Differences in event-wise stability proportions across the five events were evaluated with χ^2 tests (df=4). Planned post-hoc pairwise comparisons were performed using Fisher exact tests with Holm correction for multiple comparisons.

Population decoding. A Bayesian decoder trained on maze A activity was used to decode trial progression in maze B and maze A′ pseudopopulations. Posterior probability matrices were compared across conditions. Similar to within-maze Bayesian decoding above, precision was quantified as the entropy of the posterior distribution, and accuracy as the likelihood of correctly decoding the behavioral epoch.

Decoding uncertainty (posterior entropy) was compared *per time bin* between cross-maze conditions (A→B vs. A→A′) using Wilcoxon rank-sum tests, with Benjamini–Hochberg false discovery rate control across bins ($p=0.01$).

Epoch-level decoding accuracy was evaluated in two parallel ways: (1) accuracy defined relative to the correct epoch *within trial-type* (left and right analyzed separately), and (2) accuracy defined relative to the correct epoch *after collapsing across trial types* to assess structure that generalizes beyond lateralization. For each definition, accuracy proportions were compared between cross-maze conditions (A→B vs. A→A′) using χ^2 tests conducted independently within each epoch. Chance baselines were determined by the fraction of template bins belonging to each epoch (±1 bin tolerance), and accuracy for each condition was compared against chance using χ^2 tests. Multiple epoch-wise comparisons were corrected using Holm adjustment.

Methods References:

1. Rosene, D. L., Roy, N. J. & Davis, B. J. A cryoprotection method that facilitates cutting frozen sections of whole monkey brains for histological and histochemical processing without freezing artifact. *J. Histochem. Cytochem.* **34**, 1301–1315 (1986).
2. Chung, J. E. *et al.* A Fully Automated Approach to Spike Sorting. *Neuron* **95**, 1381-1394.e6 (2017).
3. Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).
4. Yoo, S. B. M. & Hayden, B. Y. The Transition from Evaluation to Selection Involves Neural Subspace Reorganization in Core Reward Regions. *Neuron* **105**, 712-724.e4 (2020).
5. Olypher, A. V., Lánský, P., Muller, R. U. & Fenton, A. A. Quantifying location-specific information in the discharge of rat hippocampal place cells. *J. Neurosci. Methods* **127**, 123–135 (2003).

Figure 1 | Event-aligned responses of hippocampal neurons during a virtual alternation task

a, Immersive virtual Y-maze in which monkeys navigated via joystick to alternate between left and right arms across successive trials.
b, Two-dimensional maze schematic showing the five behavioral events used for alignment: movement initiation (Start), turn initiation at the choice point (Choice), turn toward the reward (Outcome), reward delivery (Reward). Not shown: onset of the inter-trial interval (ITI – off-maze, black screen).
c, Rendering of a chronically implanted microdrive targeting the hippocampus, based on MRI-derived 3D models of one subject. The lateral view depicts electrode trajectories entering the hippocampus from the implanted chamber.
d, Histogram over the number of events that neurons respond to, demonstrating that most neurons responded to multiple task events.
e, Example neurons illustrating diverse peri-event firing profiles aligned to the five task events from **b**. Trials ordered by length for clarity. Maze schematics show firing rates during movement on left vs. right trials. Time insets show firing rates following reward onset (purple) and ITI onset (orange).

Figure 2 | Population activity segments the task into discrete event-aligned states

a, Standardized trial timeline produced by piecewise temporal rescaling between task events (left) with corresponding epochs mapped onto the 2D maze (right).
b, Population heatmap showing mean firing rates (neurons × bins) across the standardized trial, revealing event-aligned structure.
c, Temporal cross-correlation matrix of soft-normalized population vectors, showing strong within-epoch consistency separated by discontinuities at epoch transitions.
d, Equivalent analysis after rescaling trials by distance along the linearized maze, confirming consistent epochal segmentation under a spatial alignment framework.
e, Three-dimensional PCA projection of pseudopopulation responses from 1,000 trials (Left trials shown), color-coded by behavioral epoch as in a.
f, Same projection as e, colored by behavior-agnostic cluster identities; clusters correspond closely to the epochs in a.
g, (Top) Cluster labels mapped back to time for 100 Left and 100 Right pseudotrials, showing contiguous epoch-like blocks. (Bottom) Probability of each cluster label across bins for Left and Right trials plotted as continuous curves, revealing reproducible boundaries aligned to behavioral events.
h, Alignment of cluster-derived boundaries with behavioral transitions, quantified by F1 scores and Hit Rates relative to null distributions, demonstrating strong correspondence between unsupervised segmentation and task epochs. Significance bars reflect Mann–Whitney U tests, $P < 0.001$.

Figure 3 | Dual representational modes alternate across behavioral epochs

a, Mean temporal cross-correlation between each trial's population vector and a leave-one-out mean template (mean across trials) showing that events defining epoch boundaries correspond to moments of high reliability across trials.
b, Mean trial-to-template similarity (diagonal of a) across the standardized trial, showing reproducible peaks at behaviorally-relevant events and reduced similarity within epochs.
c, Aggregate information-by-bin for the population, revealing transient peaks in information content aligned to behavioral events.
d, Bayesian decoding posteriors showing precise decoding at epoch boundaries and diffuse distributions within epochs.
e, Entropy of the posterior probability as a measure of decoding uncertainty, lowest at event boundaries and elevated between them.
f, Epoch-level decoding accuracy across the standardized trial. Colored bars indicate decoding that is both epoch- and trial-type-specific; black bars indicate decoding that is epoch-accurate irrespective of trial type.

Dots denote bins significantly different from null via permutation test ($P<0.01$).

Figure 4 | Hippocampal representations generalize across visually distinct environments

a, Visually distinct virtual Y-mazes (A, B, A') differing in landmarks, colors, and textures; maze A' is a repeat of maze A.
b, Counts of neurons that remain stable or remap between mazes A and B among those stable in A–A'.
c, Correlation of full-trial firing patterns across mazes A–A' (x-axis) and A–B (y-axis); dashed line marks the r = 0.7 stability threshold.
d, Event-wise stability across mazes A and B for neurons stable in A–A', showing variable preservation across task phases (bar differences $P<0.001$, χ^2 test).
e, Event-wise stability across Left and Right trial types, illustrating heightened trial-type stability at Movement Start and the ITI initiation (bar differences $p<0.001$, χ^2 test).
f, Cross-maze (A–B) Bayesian decoding posteriors for Left (top) and Right (bottom) trials, showing accurate tracking of task progression.
g, Decoding entropy across the trial for A–B (color) and A–A' (gray) conditions, plotted separately for Left (top) and Right (bottom) trials.
h, Epoch-level decoding accuracy for A–B (solid color) and A–A' (striped, same hue). Colored portions denote trial-type-specific decoding; black portions denote epoch-accurate decoding without regard for trial type (gray portions and black portions each indicate significant differences, $P<0.01$, χ^2 test). Top, Left trials; bottom, Right trials.

Supplementary Figure 1 | Epoch structure is preserved across alternative alignment schemes.

Columns compare four normalization approaches: **a**, event-based temporal rescaling, (**b**), midpoint-based temporal rescaling, **c**, linearized-distance alignment, and **d**, holistic temporal rescaling. All analyses use the same 588 task-responsive neurons.

Top row: population heatmaps (neurons × bins) for each alignment method showing mean firing rates.

Middle row (Full Trial Progression): temporal cross-correlation matrices of soft-normalized population vectors (50-ms binned spikes; piecewise mapped to standardized bins), revealing sustained within-epoch blocks under all four alignment methods.

Bottom row (Rescaled Segment Only): zoomed cross-correlation matrices restricted to the portion of the trial that is warped under each alignment method, demonstrating that the pattern of state transitions is robust to the specific warping scheme used.

Supplementary Figure 2 | Convergent boundary structure across clustering algorithms and parameterizations.

a, Boundary likelihood for Leiden clustering as a function of task progression. Each row reflects one clustering run (100 replicates per parameter set), ordered by resolution $\gamma \in \{0.2, 0.6, 1.0\}$ and kNN $\in \{12, 24\}$. Boundary likelihood reflects the proportion of trials with a cluster transition at each bin (±3-bin tolerance).

b, Same analysis as (a) for k-means clustering (k = 6–12; 100 replicate initializations per k).

c, Aggregate boundary prevalence curve for Leiden clustering (across all resolutions and kNN parameters), showing consistent peaks at common temporal locations.

d, Aggregate boundary prevalence curve for k-means clustering.

e, Boundary prevalence curves for kNN = 12 Leiden clustering shown separately for coarse, intermediate, and fine resolutions (γ = 0.2, 0.6, 1.0). Dotted lines indicate the distribution of peak likelihood across null pseudopopulations (circularly shifted pseudotrials).

f, Within-resolution similarity of cluster solutions for Leiden clustering (kNN = 12). Violin plots show pairwise Adjusted Rand Index (ARI) across replicate runs for true vs. null pseudopopulations, for γ = 0.2, 0.6, and 1.0. Left set of plots shows ARI without compensating for cluster count mismatches; right set shows ARI after greedy Hungarian matching of cluster identities. Significance bars reflect Mann–Whitney U tests, $P < 0.001$.

g, Cross-resolution similarity for Leiden clustering comparing γ = 0.2 vs 0.6, 0.2 vs 1.0, and 0.6 vs 1.0. Violin plots show ARI after greedy Hungarian matching; each comparison includes one true distribution and two null distributions (directional asymmetry accounted for by swapping true/null label sets). Significance bars reflect Mann–Whitney U tests, $P < 0.001$.

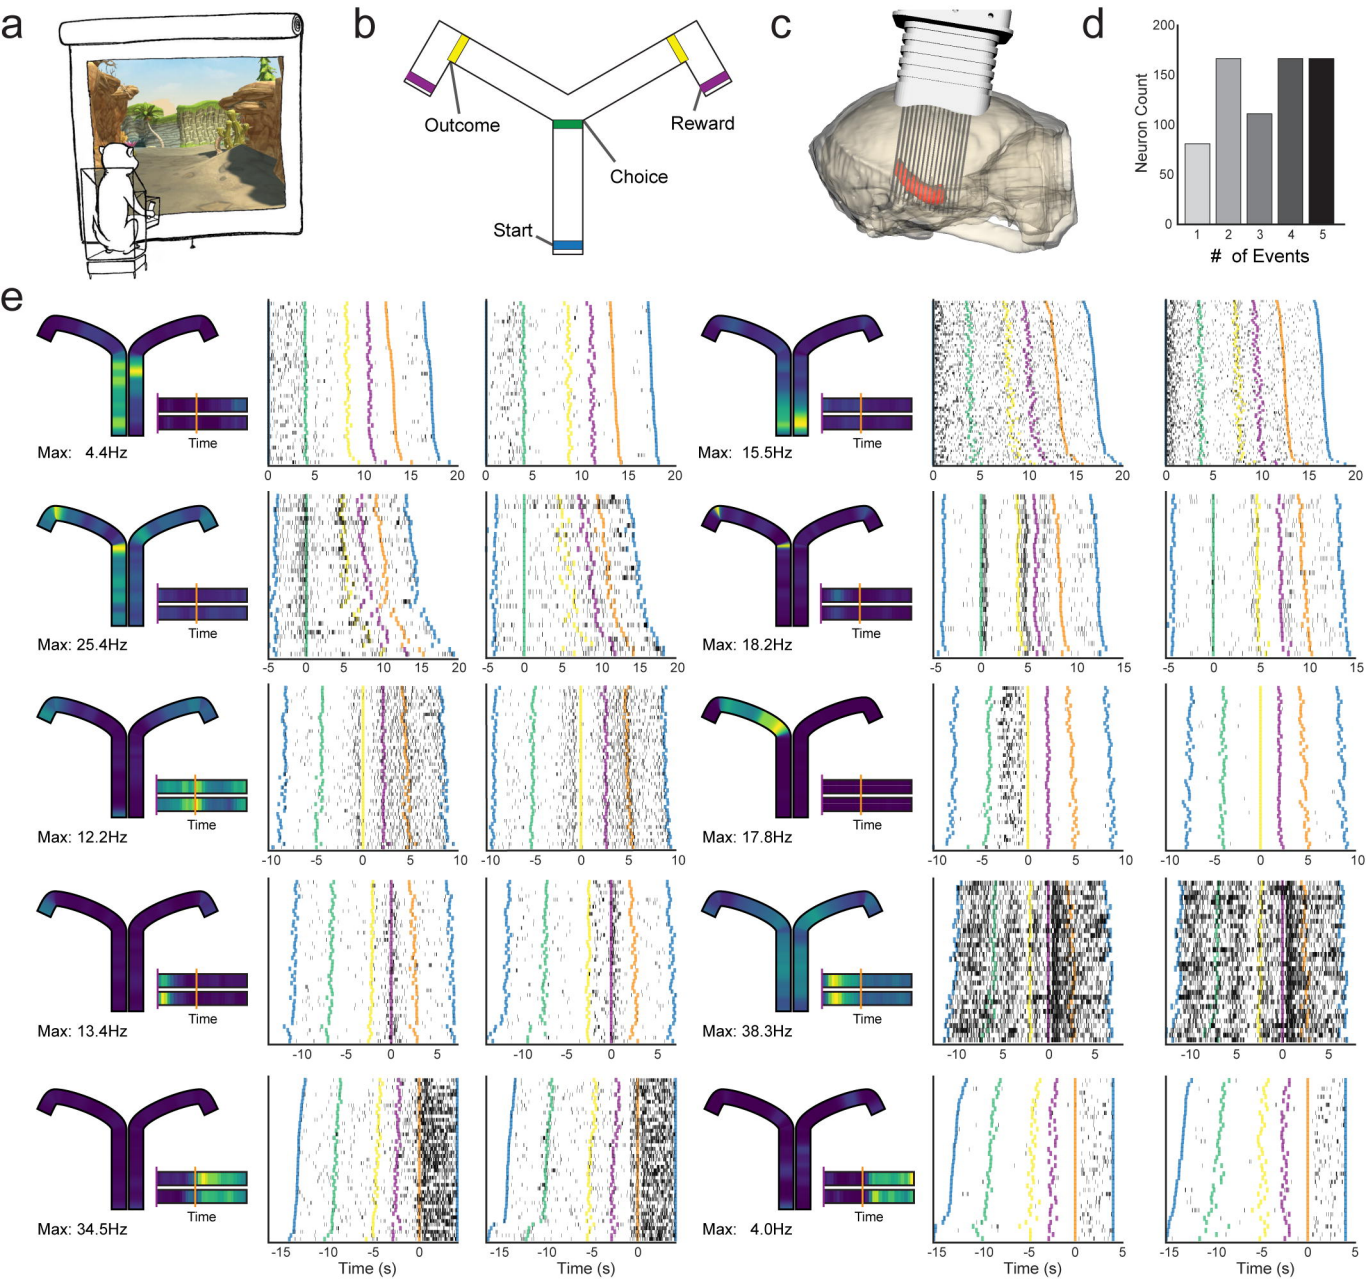
a
b
Outcome
Reward
Choice
Start
c
d
Neuron Count
# of Events
e
Max: 4.4Hz
Max: 15.5Hz
Max: 25.4Hz
Max: 18.2Hz
Max: 12.2Hz
Max: 17.8Hz
Max: 13.4Hz
Max: 38.3Hz
Max: 34.5Hz
Max: 4.0Hz
Time
Time (s)

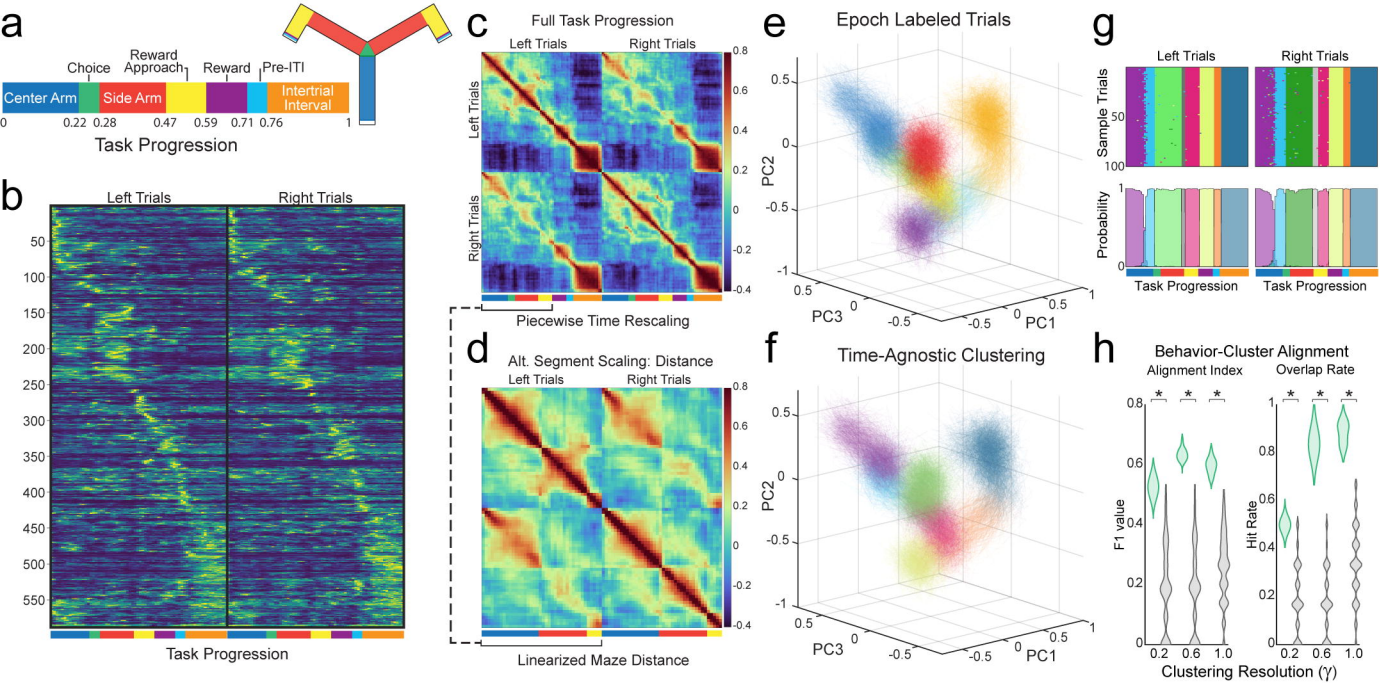
a
Choice
Reward Approach
Reward
Pre-ITI
Center Arm
Side Arm
Intertrial Interval
0
0.22
0.28
0.47
0.59
0.71
0.76
1
Task Progression
b
Left Trials
Right Trials
Task Progression
c
Full Task Progression
Left Trials
Right Trials
Piecewise Time Rescaling
d
Alt. Segment Scaling: Distance
Left Trials
Right Trials
Linearized Maze Distance
e
Epoch Labeled Trials
PC1
PC2
PC3
f
Time-Agnostic Clustering
PC1
PC2
PC3
g
Left Trials
Right Trials
Sample Trials
Probability
Task Progression
Task Progression
h
Behavior-Cluster Alignment
Alignment Index
Overlap Rate
F1 value
Hit Rate
Clustering Resolution (γ)

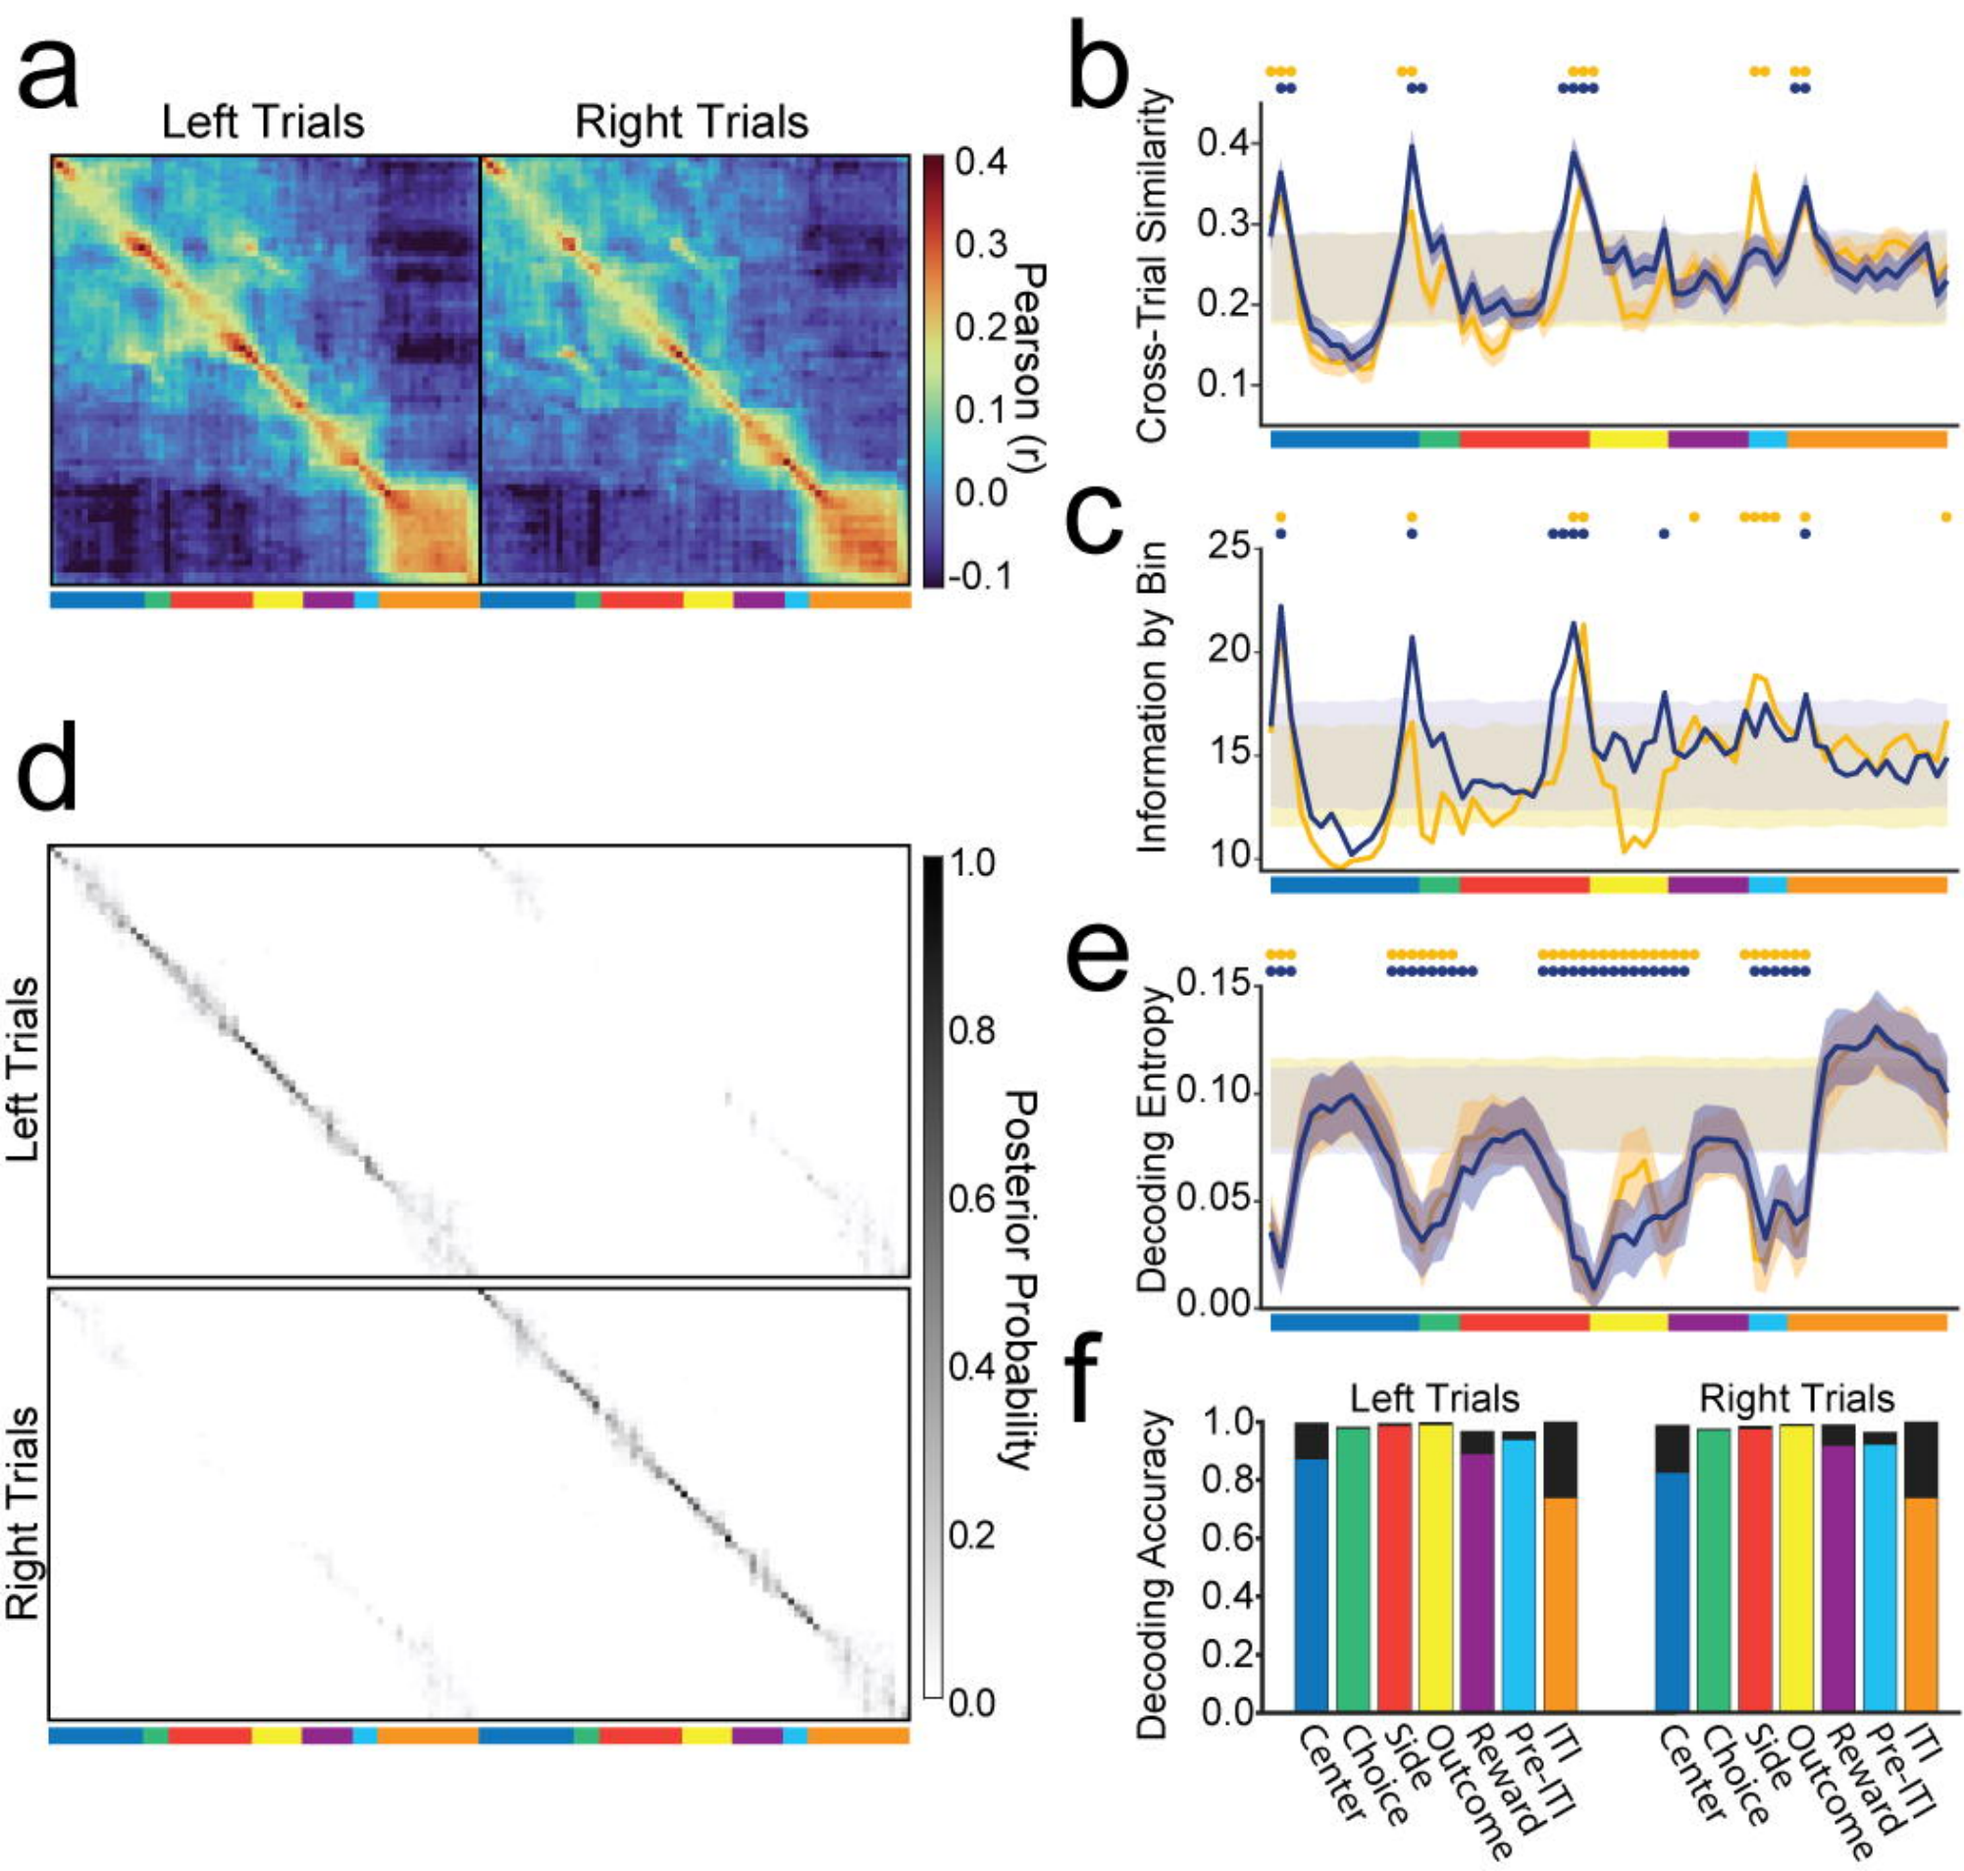
a
Left Trials
Right Trials
Pearson (r)
0.4
0.3
0.2
0.1
0.0
-0.1
b
Cross-Trial Similarity
0.4
0.3
0.2
0.1
c
Information by Bin
25
20
15
10
d
Left Trials
Right Trials
Posterior Probability
1.0
0.8
0.6
0.4
0.2
0.0
e
Decoding Entropy
0.15
0.10
0.05
0.00
f
Decoding Accuracy
Left Trials
Right Trials
1.0
0.8
0.6
0.4
0.2
0.0
Center
Choice
Side
Outcome
Reward
Pre-ITI
ITI

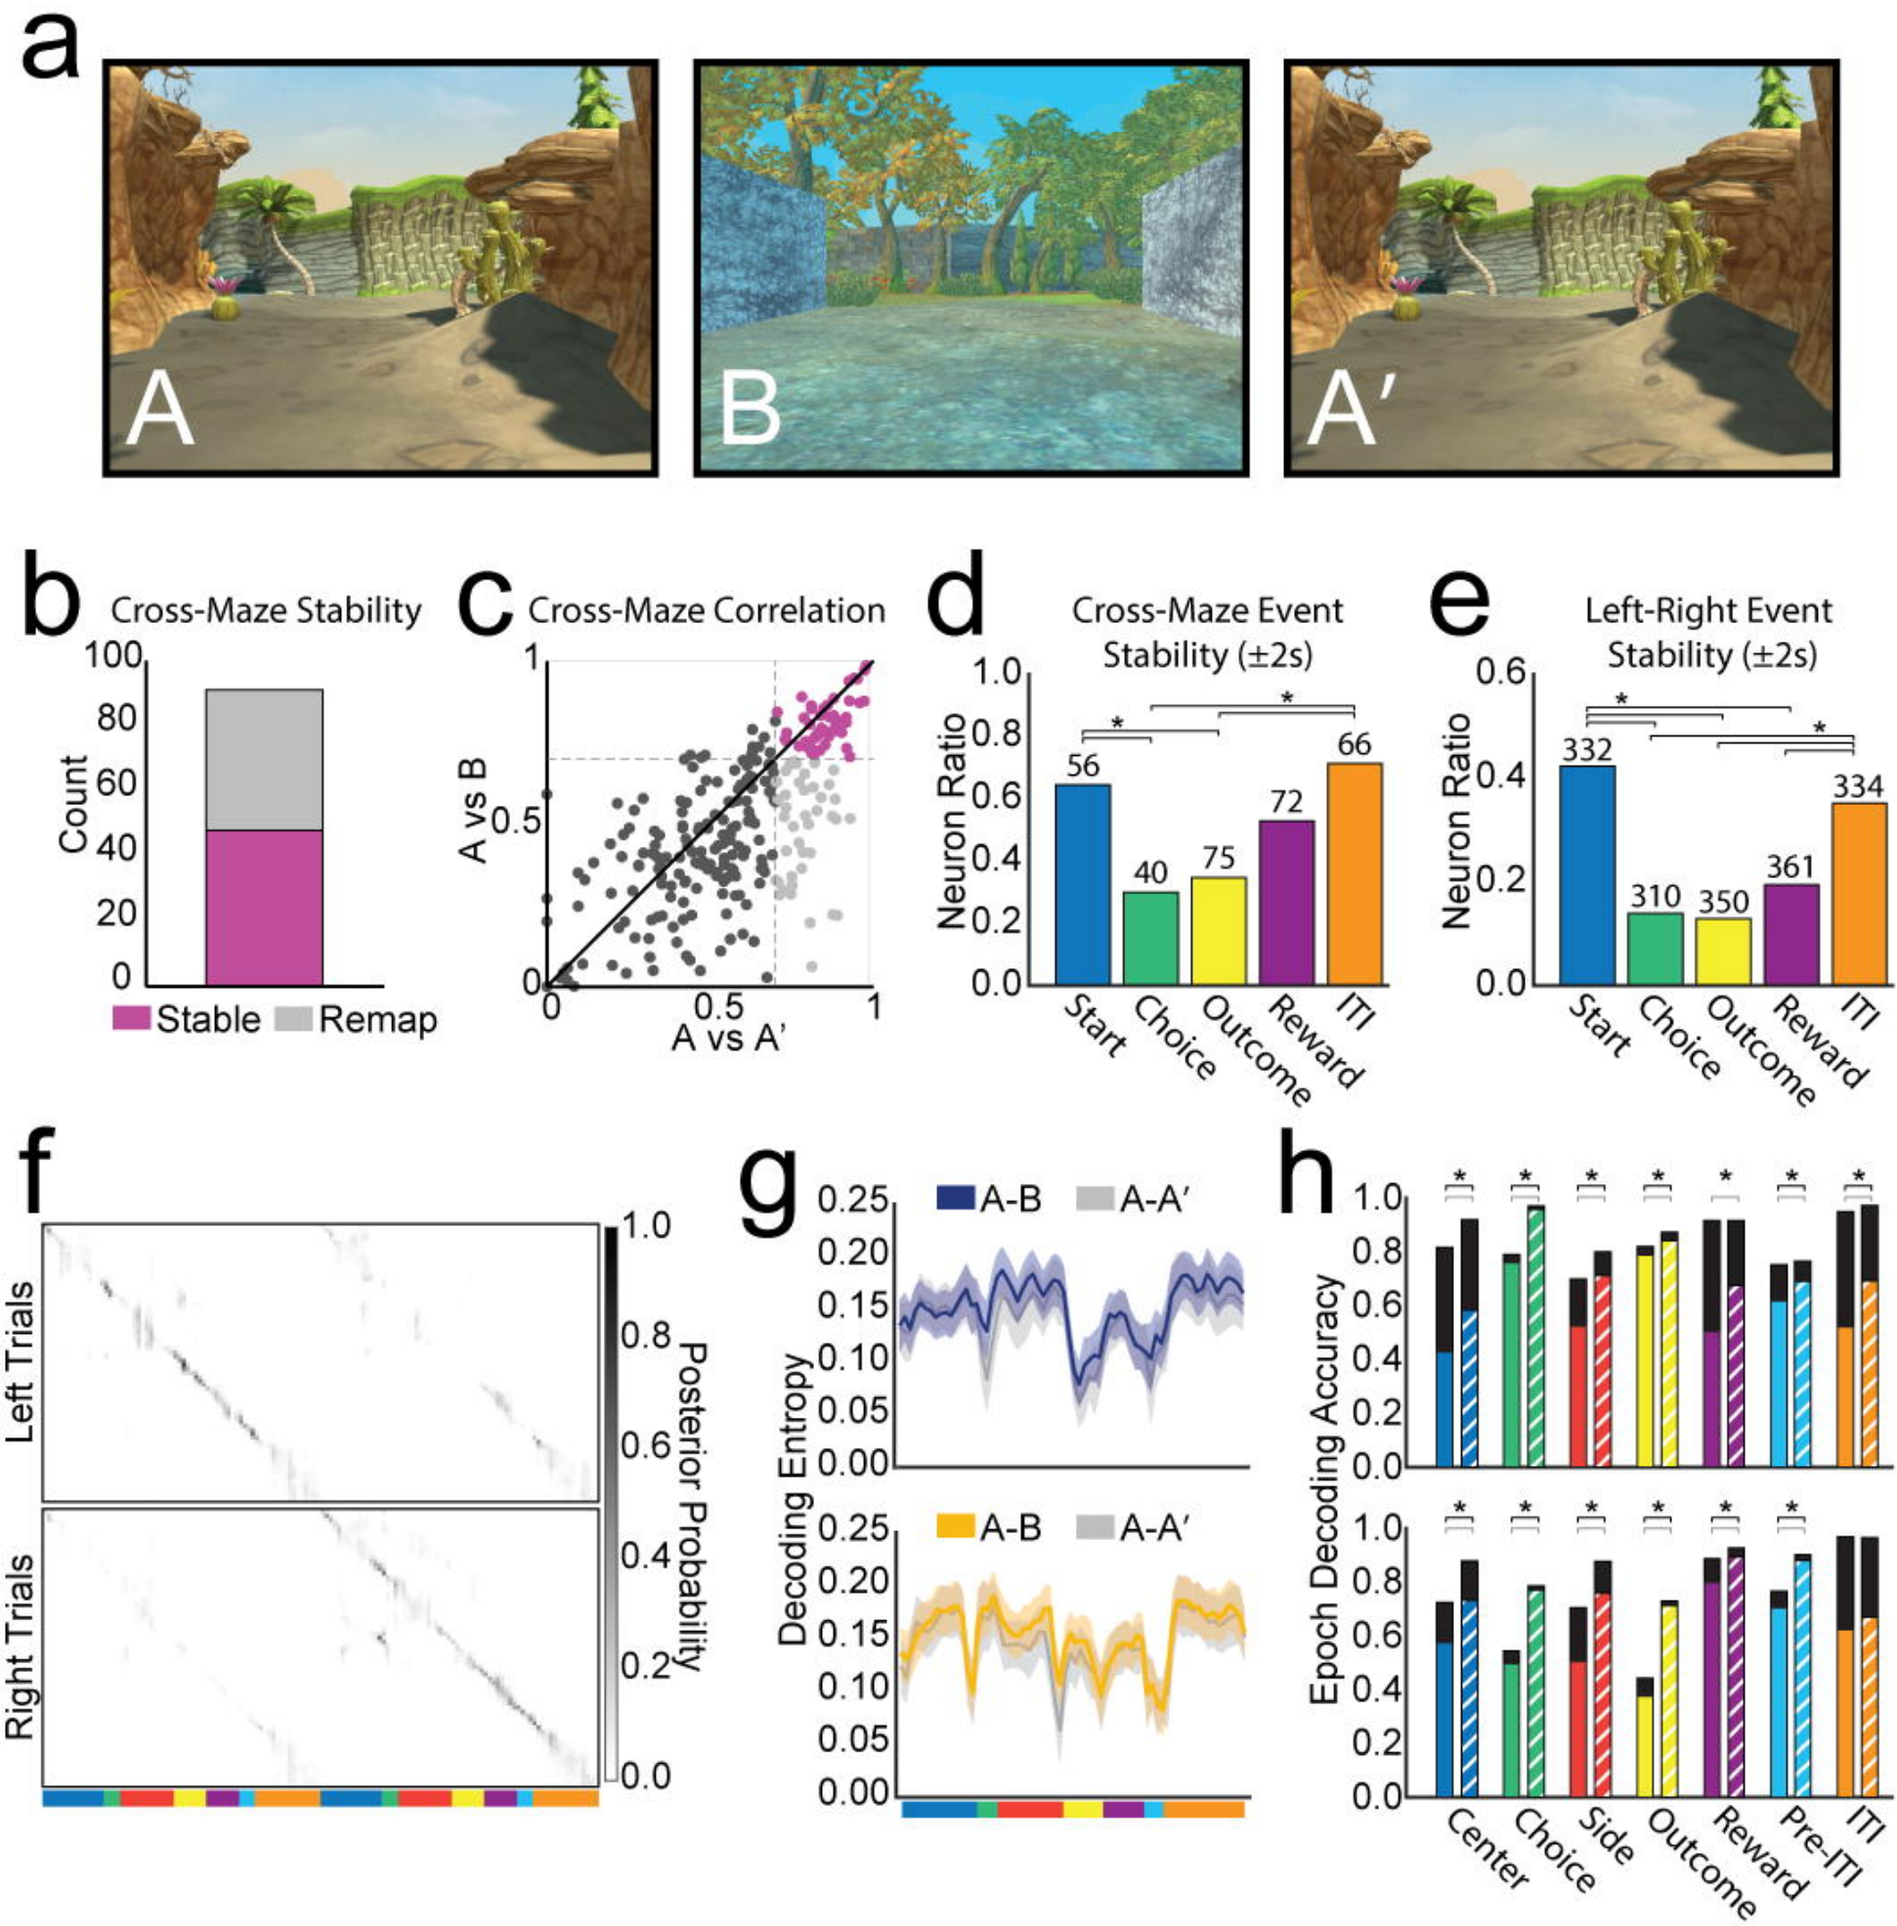
a
A
B
A′
b
Cross-Maze Stability
Count
Stable
Remap
c
Cross-Maze Correlation
A vs B
A vs A'
d
Cross-Maze Event Stability (±2s)
Neuron Ratio
56
40
75
72
66
Start
Choice
Outcome
Reward
ITI
e
Left-Right Event Stability (±2s)
Neuron Ratio
332
310
350
361
334
Start
Choice
Outcome
Reward
ITI
f
Left Trials
Right Trials
Posterior Probability
g
A-B
A-A′
Decoding Entropy
h
Epoch Decoding Accuracy
Center
Choice
Side
Outcome
Reward
Pre-ITI
ITI

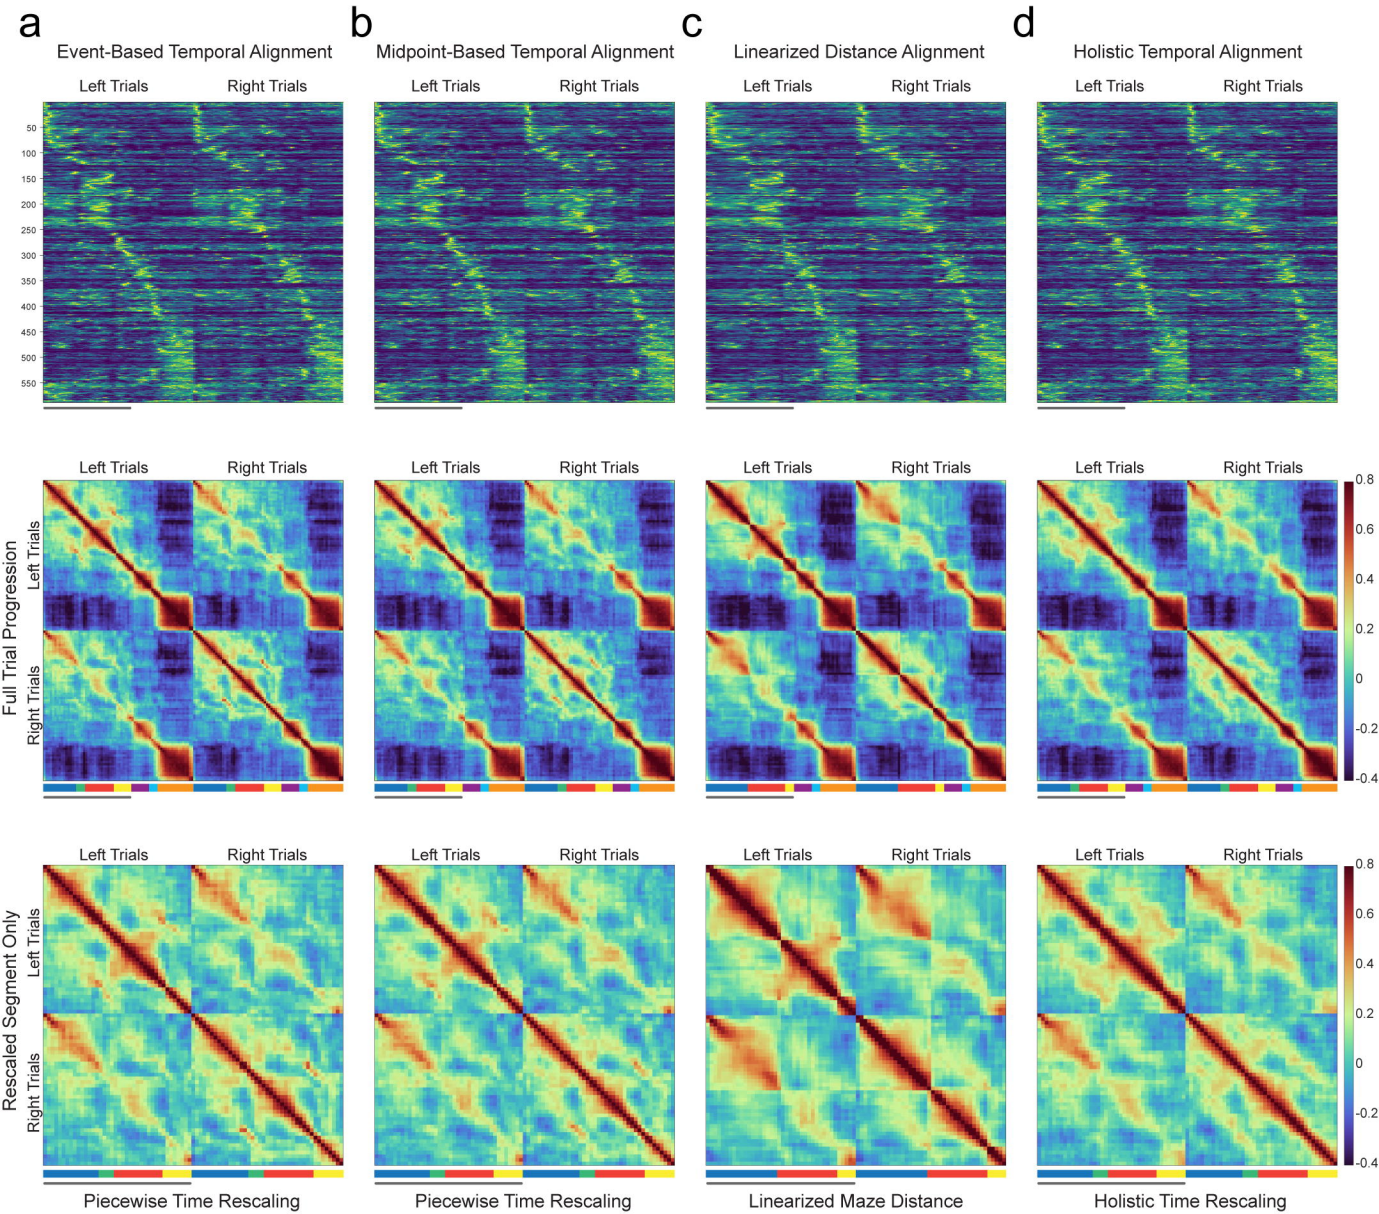
a
Event-Based Temporal Alignment
b
Midpoint-Based Temporal Alignment
c
Linearized Distance Alignment
d
Holistic Temporal Alignment
Left Trials
Right Trials
50
100
150
200
250
300
350
400
450
500
550
Full Trial Progression
Rescaled Segment Only
0.8
0.6
0.4
0.2
0
-0.2
-0.4
Piecewise Time Rescaling
Piecewise Time Rescaling
Linearized Maze Distance
Holistic Time Rescaling

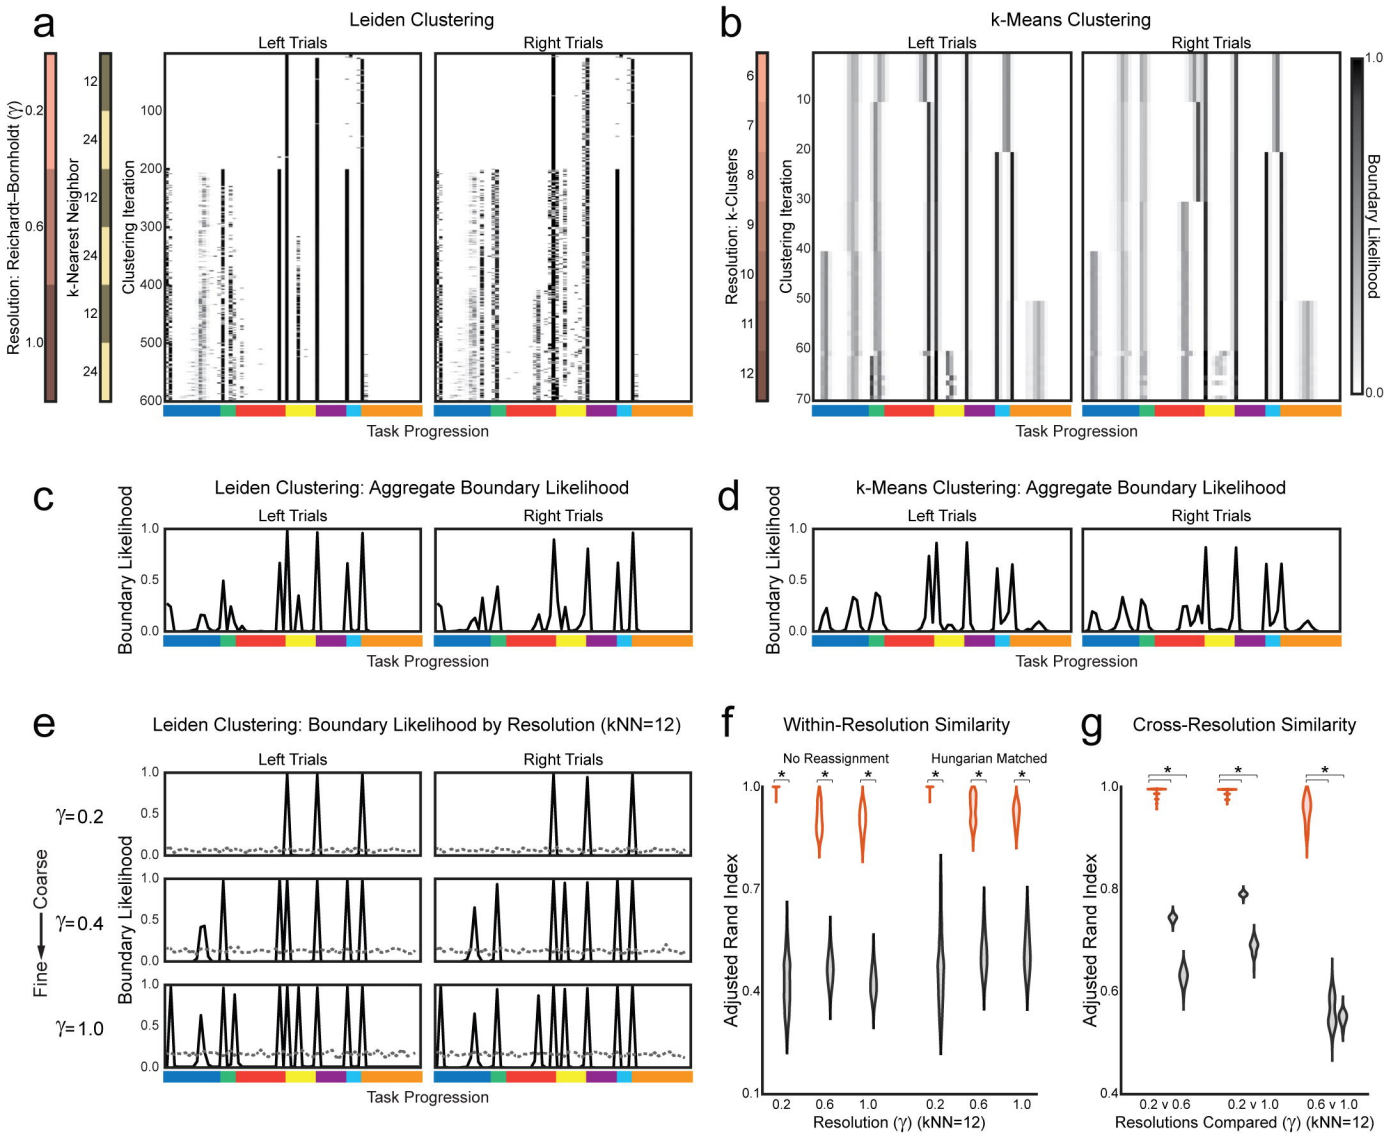
a
Leiden Clustering
Left Trials
Right Trials
Resolution: Reichardt–Bornholdt (γ)
k-Nearest Neighbor
Clustering Iteration
Task Progression
b
k-Means Clustering
Left Trials
Right Trials
Resolution: k-Clusters
Clustering Iteration
Boundary Likelihood
Task Progression
c
Leiden Clustering: Aggregate Boundary Likelihood
Left Trials
Right Trials
Boundary Likelihood
Task Progression
d
k-Means Clustering: Aggregate Boundary Likelihood
Left Trials
Right Trials
Boundary Likelihood
Task Progression
e
Leiden Clustering: Boundary Likelihood by Resolution (kNN=12)
Left Trials
Right Trials
γ=0.2
γ=0.4
γ=1.0
Coarse
Fine
Boundary Likelihood
Task Progression
f
Within-Resolution Similarity
No Reassignment
Hungarian Matched
Adjusted Rand Index
Resolution (γ) (kNN=12)
g
Cross-Resolution Similarity
Adjusted Rand Index
0.2 v 0.6
0.2 v 1.0
0.6 v 1.0
Resolutions Compared (γ) (kNN=12)