Extended experience reveals distinct representational dynamics governing the CA1 spatial code

Alexandra T. Keinath[1,*], Coralie-Anne Mosser[1], Mark P. Brandon[1,2,*]

Affiliations:

[1] Department of Psychiatry, Douglas Hospital Research Centre, McGill University, 6875 Boulevard LaSalle, Verdun, QC H4H 1R3, Canada.

[2] Integrated Program in Neuroscience, Montreal Neurological Institute, McGill University, 3801 University Street, Montreal, Quebec H3A 2B4, Canada.

* Correspondence: atkeinath@gmail.com and mark.brandon@mcgill.ca

Summary

Hippocampal subregion CA1 is thought to support episodic memory by reinstating a stable spatial code. However, recent calcium imaging experiments have challenged this presumed function by demonstrated that this code is largely unstable on a timescale of days. This turnover may reflect homogenous drift within the population; alternatively, it may reflect distinct time-varying representational component(s) which coexists alongside other stable components. Here we characterized the mouse CA1 spatial code over more than a month of daily free exploration experience in an extended geometric morph paradigm. We find that this code is governed by distinct representational components with different long-term dynamics, including stable components representing the shape of space and prior experience. These components are mediated by separate neural ensembles with similar short-term spatial reliability and precision. Together, these results demonstrate that the long-term dynamics of the CA1 spatial code are defined by representational content, not homogenous drift.

Introduction

Hippocampal subregion CA1 represents a mixture of external and internal cues – including the shape of the environment[1], visual landmarks[2,3], objects[4], task-relevant information[5], and past experience[6–9] – through changes to the spatial tuning properties of its principal cells[10]. This code is hypothesized to support spatial and episodic memory by later reinstating its content. However, longitudinal experiments in mice have challenged this view[11,12], reporting that the vast majority of CA1 place cells change their spatial tuning properties more rapidly than expected (on the order of days), while only a subset of these cells (15-25%) provide a stable spatial representation on this timescale[12].

These long-term representational dynamics may reflect two possible phenomena, each carrying its own theoretical implications. On one hand, these dynamics may reflect homogenous population-wide representational drift, as could result from gradual but persistent plasticity mechanisms throughout the hippocampal circuit[13,14]. In this case, drift would be independent of the content of the representation, with complete representational turnover on sufficiently long timescales. Such drift would thus pose a strong challenge to the persistent reinstatement hypothesis, and would at the very least motivate a revision to include compensatory downstream mechanisms which might keep pace with, filter out, or otherwise overcome this drift[14]. Alternatively, these dynamics may instead reflect time-varying representational component(s) alongside other stable components, as might arise from the convergence of both stable and time-varying inputs onto the population. In this case, we would expect time-varying dynamics to be restricted to distinct component(s) of the representation at the level of the population, and possibly at the level of individual neurons. Such component-specific dynamics would pose a minimal challenge to the persistent reinstatement hypothesis, as downstream readers tuned to a stable representational component would be blind to any variability in other orthogonal components.

To adjudicate between these possibilities, one must assay multiple representational components over long timescales and determine whether drift depends on, or is independent of, representational content. To this end, we characterized the mouse CA1 spatial code over more than a month of daily experience in an extended version of the classic geometric morph paradigm. We find that the CA1 spatial code is defined by distinct representational components mediated by separate neural ensembles with different long-term dynamics, including stable components reflecting the shape of space and prior experiences.

Results

We recorded daily from CA1 with calcium imaging (Fig. 1A) in mice freely exploring open environments for 32 days in an extended version of a geometric morph paradigm (5 mice, 160 sessions total; Fig. 1B; Fig. S1). In this paradigm, mice were first familiarized with two geometrically-distinct environments, and later tested in deformed (morphed) versions of these environments spanning the shapespace between the two familiar environments. Previous acute versions of this paradigm have demonstrated that a subset of CA1 cells exhibit experience-dependent attractor-like properties in their spatial tuning, with an abrupt transition in their preferred firing locations near the midpoint of the morph sequence, while other cells fired at geometrically similar locations across the morph sequence[8,9]. By employing more geometrically-dissimilar familiar environments and extending this paradigm across a 32 day period, we can thus assay three potential components of the CA1 representation – a representation of the spatial geometry (*geometry*), a representation of experience-dependent attractor dynamics (*attractor*), and a representation which varies with time (*time*).

We began by first characterizing these data within the morph sequence akin to previous acute reports. Following motion correction[15], cells were segmented and calcium traces were extracted via constrained nonnegative matrix factorization[16,17] (Fig. S1). The likelihood of spiking events (henceforth *firing rate*) was inferred from deconvolution of the filtered calcium traces via a second-order autoregressive model[18]. Cell identity was tracked across recording sessions on the basis of extracted spatial footprints (n = 2745)[19]. We found that individual cells exhibited a variety of dynamics, including abrupt transitions in their preferred locations near the midpoint of the morph sequence, persistent firing at geometrically-similar locations across the morph sequence, and the appearance of punctate spatial tuning in a subset of environments (Fig. 1C). Across the population, the morph sequence produced sigmoidal transitions between the (partially-correlated) familiar environment rate maps (Fig. 1D; Fig. S2). Fitting these transitions with sigmoidal functions allowed us to quantify dynamic properties across sequences within each mouse, revealing qualitatively distinct patterns within each mouse and a trend toward increased decorrelation of maps across sequences (Fig. S2).

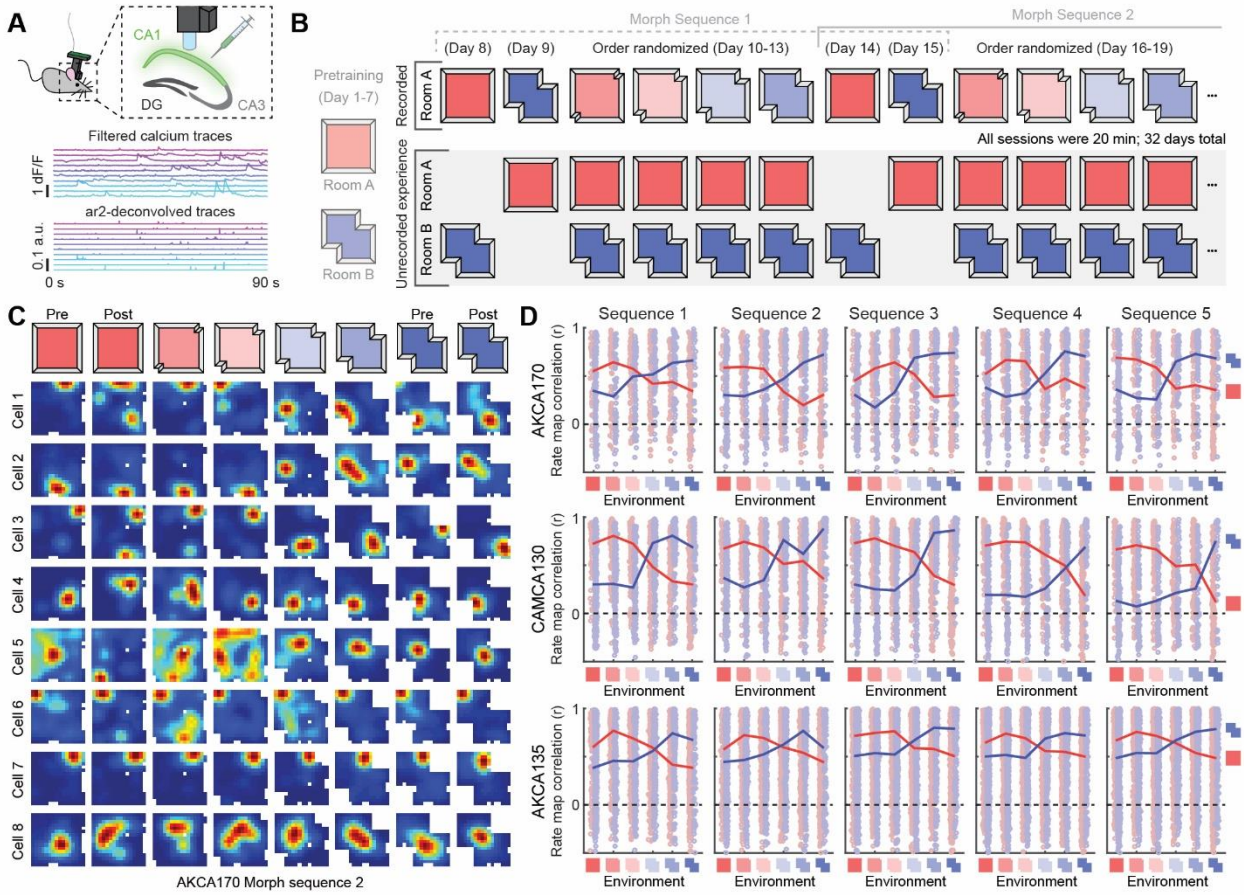

Figure 1. An adapted geometric morph paradigm characterizes the CA1 spatial code across extended experience. *(**A**) Schematic of the miniscope recording procedure (top), as well as resulting calcium traces and firing rates inferred from second order autoregressive deconvolution (bottom). (**B**) Schematic of the behavioral paradigm. Assignment of familiar environments to rooms and the order of familiar environment recordings were randomized across mice, and the order of unrecorded experience was randomized within mice. (**C**) Example of eight simultaneously recorded cells tracked across one morph sequence exhibiting a diversity of dynamics. Rate maps normalized from zero (blue) to the peak (red) within each session. (**D**) Rate map correlations between each environment and the familiar environments for each*

sequence. Bold lines indicate median across pairwise comparisons of cells (dots). Only comparisons between sessions where the cell had a within-session split-half rate map correlation exceeding the 99th percentile of the shuffled distribution in at least one session were included.

Traditional within-morph sequence analyses only characterize the representational similarity between subsets of sessions. To take advantage of the full representational structure, we explored population-level representational similarity across all 32 days. To this end, we first computed the median rate map similarity across all tracked cells for each pairwise comparison of sessions (Fig. 2A). We then reduced this representation to two dimensions via nonmetric multidimensional scaling, an unsupervised technique for visualizing high-dimensional structure in a digestible low-dimensional space that preserves the representational similarity between sessions as well as possible. The resultant embedding revealed two dimensions that strongly determined representational similarity in this paradigm: a spatial context component and a time-varying component (Fig. 2B). These components were pronounced in every mouse. Quantification of the embedded representations revealed that the contextual and time-varying components defined nearly orthogonal dimensions in this 2D subspace, with an absolute angular difference of 89.0° ± 5.8° (mean ± circular standard deviation) between context and time-varying dimensions (see *Methods*; all angles: [83.0°, 84.6°, 85.5°, 95.2°, 96.7°]; Rayleigh's test versus uniformity on the 0° to 180° range: p = 2.47e-3, z = 4.80). This reliable structure provides evidence that spatial context and time-varying components of the CA1 spatial code are representationally distinct at the level of the population.

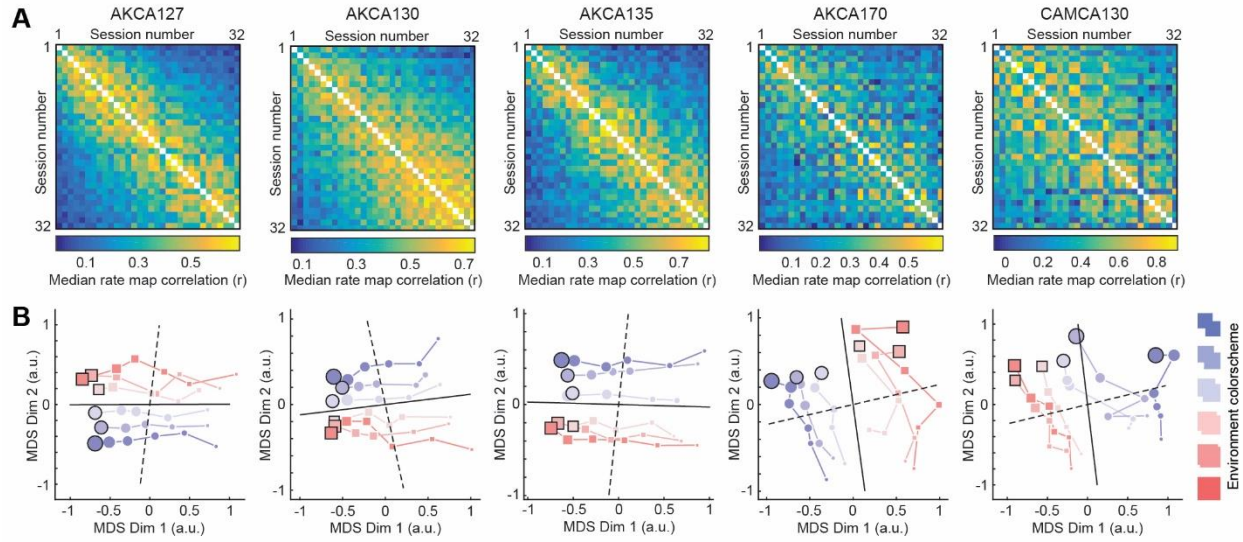

Figure 2. Dimensionality reduction reveals distinct temporal and contextual representational components in the CA1 spatial code. *(**A**) Population-level representational similarity matrices (RSMs) across all recording sessions for all five mice. Each RSM is computed by taking the median pairwise similarity across all tracked cells. (**B**) Session similarity structure when embedded in a two-dimensional space extracted via nonmetric multidimensional scaling. Dot size indicates session number, with earliest sessions indicated by the smallest dots and latest sessions indicated by the largest dots. Color indicates environmental geometry. Final sessions outlined in black. Note that estimated time (solid black line) and spatial context (dashed black line) dimensions are nearly orthogonal in this subspace (see Methods).*

The observation of distinct population-level representational components could be the product of a distributed code across neurons that exhibit mixed-selectivity. Conversely, it could also be a product of separate neural ensembles coding for individual representational components. To adjudicate between these possibilities, for each cell with at least 28 comparisons we created a representational similarity matrix (*RSM*) of the pairwise rate map correlations between all tracked sessions (Fig. 3A)[20]. (Twenty-eight was chosen as the number of pairwise comparisons within a single bookend morph sequence.) Next, we constructed predicted RSMs on the basis of similarity in time, spatial geometry, and contextual group as defined by the attractor dynamics (see *Methods*). Using these matrices, we computed the amount of variance in each (Fisher-transformed) cell RSM explained by a model containing the three predicted RSMs as well as all possible interactions (*total* r^2). We then determined the amount of total explained variance attributable to each factor by quantifying the reduction in explained variance when selectively excluding that factor (and its interactions).

The results of this analysis revealed remarkable selectivity in representational content and long-term dynamics across simultaneously-recorded cells (Fig. S3). Across all cells with a full model r^2 of at least 0.1915 (90th percentile of the shuffled distribution, n = 1316 of 2745, 47.9% of cells with sufficient comparisons, binomial test: p = ~0.0), explained variance was typically attributable to a single factor (Fig. 3B). Moreover, reorganization of RSMs for time-, attractor-, and geometry-loading cells revealed clear distinctions in their temporal and contextual dynamics (Fig. 3C; Fig S4). To quantify these observations, we computed null distributions by randomly shuffling RSM values within each comparison across cells. The resulting null distributions thus preserve the population-level statistics while destroying any within-cell cross-session effects. Relative to these distributions, the first-ranked factor was attributed slightly but significantly more variance, while the second- and third-ranked factors were attributed significantly less variance than would be expected given the variability within the population alone (Fig. 3D,E) – a pattern which held across a range of inclusion criteria and grew more pronounced with more stringent criteria (Fig. S5). These results provide evidence that simultaneously-recorded neural ensembles are selectively tuned to distinct representational components, governed by differing long-term dynamics.

Are differently-loading cells distinguished by more than their long-term dynamics and representational content? To address this possibility, we grouped cells according to primary factor loading and compared a number of properties as follows. We first examined whether groups differed by the short-term reliability and precision of their spatial codes. For cells in all groups, full model r^2 was correlated with median within-session split-half correlation (SHC; Fig. 3F), indicating that a better model fit was associated with a more reliable spatial code in the short-term regardless of group. Median within-session SHCs of cells meeting full model r^2 criteria exceeded those of cells not meeting criteria and were similarly high in all groups (Fig. 3G), indicating comparable spatial reliability on a timescale of minutes. Moreover, despite the differences in long-term dynamics between groups (Fig. 3H), within-session SHCs remained high across all sessions regardless of group (Fig. 3I). To quantify this observation, for each cell we computed the least squares linear fit of across-session map similarity and within-session SHC as a function of temporal lag between sessions. Time-loading cells exhibited a modest but reliable decrease in across-session rate map correlations for each day separating compared sessions, which was significantly greater than that observed in attractor-loading or geometry-loading cells (Fig. 3J). In contrast, time-loading cells exhibited significantly less of a decrease in within-session SHC as a function of lag, and differences between groups were relatively minor (Fig. 3J). Echoing these reliability results, the precision of the spatial code as measured by median spatial information content (SIC) was higher for cells meeting criteria and did not substantially differ between groups (Fig. 3K). Together, these results indicate that short-term

spatial reliability and precision is predictive of selectively coding some representational content, but is not predictive of the specific content itself.

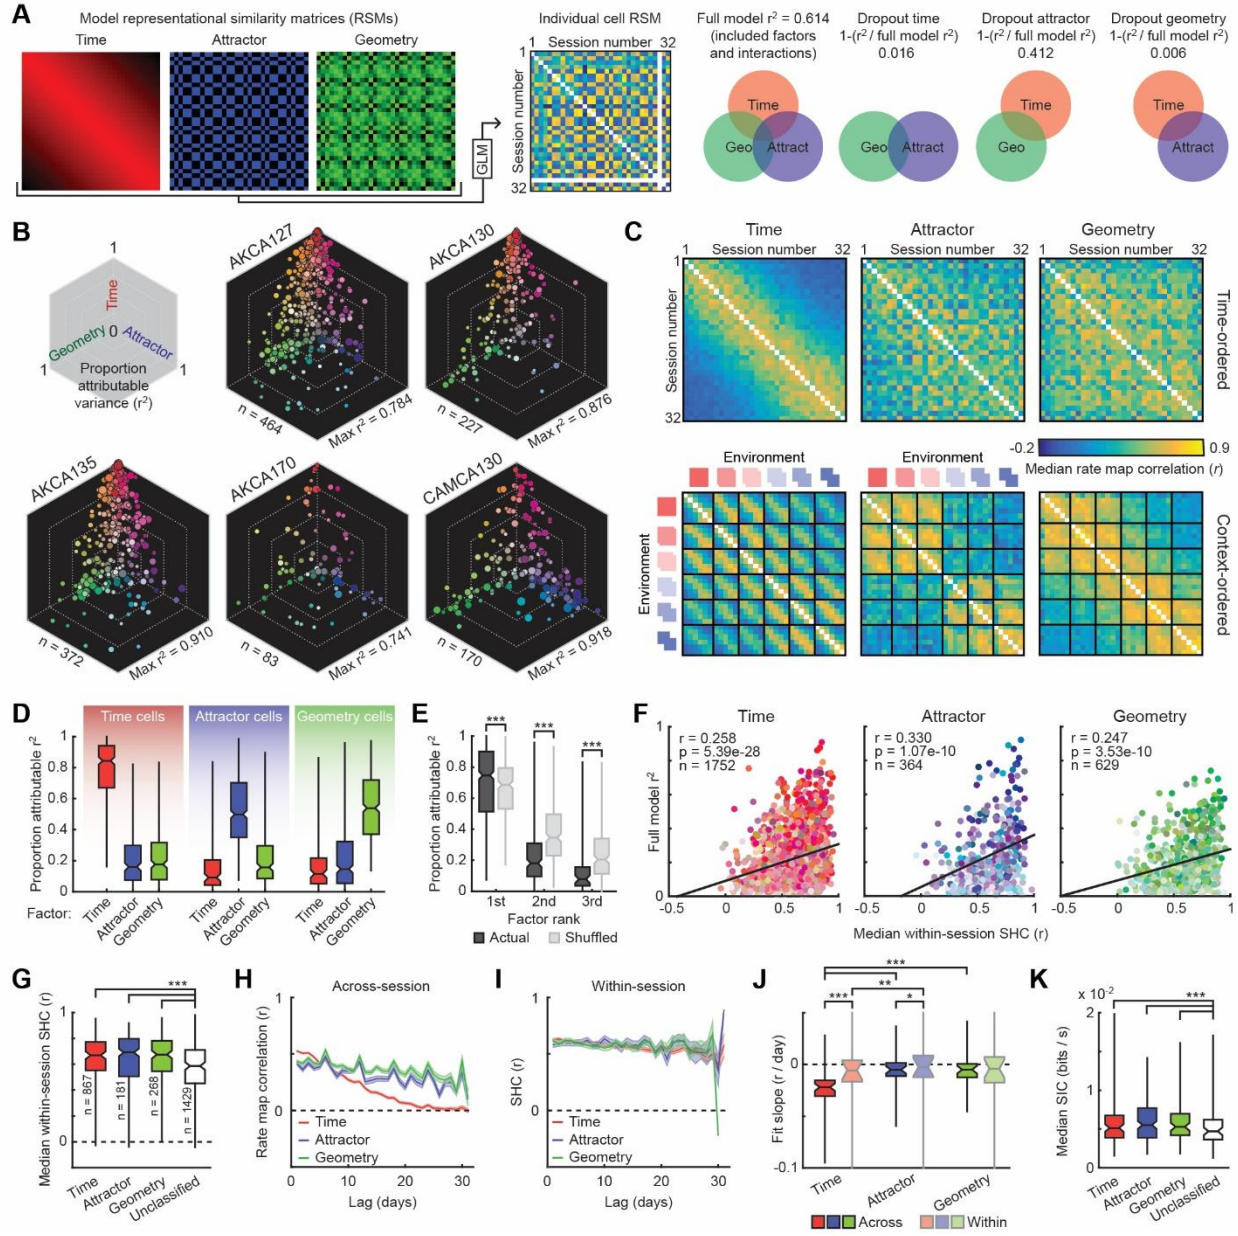

Figure 3. Characterizing individual cell RSM variance reveals representational selectivity at the single-cell level. *Exact outcomes for all statistical tests performed here are shown in Table S1. (**A**) Each cell RSM was fit with a general linear model consisting of time, attractor, and geometry components including interactions. Then the relative reduction in r^2 was computed when excluding each factor. Cell RSMs were Fisher-transformed before fitting. Note that model attractor and geometry RSMs tend to be correlated, limiting the amount of variance that can be attributed to one or the other. (**B**) Proportion attributable variance for cells meeting inclusion criteria. Note that cells tend to lie on the axes, indicating that most of their attributable variance is driven by a single factor. Dot size scaled to full model r^2. (**C**) RSMs of time, attractor,*

*and geometry-loading cells meeting criteria, ordered by time (top) and spatial context (bottom). Mouse CAMCA130 was excluded from this panel because of large changes in the attractor transition point across sequences. (**D**) Proportion of attributable variance for cells meeting inclusion criteria, grouped by primary factor. (**E**) Proportion of attributable variance when factors are ranked. (**F**) Full model r^2 and median within-session split-half rate map correlation (SHC) were similarly correlated across all cells regardless of the primary factor. (**G**) Median within-session SHC as a function of group for cells meeting criteria, and for cells not meeting criteria (unclassified). (**H**) Rate map similarity as a function of lag between sessions for cells meeting criteria (mean ± SEM). Jagged behavior of attractor and geometry cells reflects context-dependence in combination with structure in the behavioral paradigm. (**I**) SHC as a function of lag between sessions for cells meeting criteria (mean ± SEM). For each cell the session with the maximum SHC was taken as the reference session (i.e. lag = 0). (**J**) Change in across-session rate map correlations and within-session SHC per day for cells meeting criteria. (**J**) Median spatial information content (SIC) as a function of group and for cells not meeting criteria. *p < 0.05, **p < 0.01, ***p < 0.001*

We investigated whether other coding properties may distinguish these cell groups. Some theories posit that the most excitable cells are more likely to be recruited into a representation[13], and firing rate is known to correlate with spatial reliability and precision on short timescales[21]. In our data, cells meeting the full model r^2 criterion tended to have higher mean and peak firing rates than cells not meeting this criterion (Fig. 4A,B). Interestingly, attractor and geometry-loading cells exhibited higher firing rates than time-loading cells (Fig. 4A,B), possibly reflecting differences in the intrinsic properties and/or inputs specific to these ensembles. Biases in place field locations also differed by group (Fig. 4C). Time-loading cells tended to tile the navigable space more evenly with a bias toward corners, attractor cells tended to have fields near the center of the environment, and geometry-loading cells typically had fields near the deformed corners. Time-loading cells were tracked across more sessions than attractor- or geometry-loading cells (Fig. 4D), consistent with a lack of context-selectivity in these cells and indicating that variance attributed to time was not the result of cell dropout. Interestingly, all groups were tracked across fewer sessions than cells not meeting criteria, also consistent with the time- and context-dependence of grouped cells and suggesting that unclassified cells did not fail to reach criteria due to dropout. The proportion of each group within the population differentially predicted morph sequence properties: the proportions of geometry- and attractor-loading cells were positively related to the decorrelation between familiar environment representations, while the proportion of time-loading cells was negatively related to this decorrelation (Fig. S2). Signal-to-noise differences between groups were minor (Fig. S5). No clear anatomical organization according to primary factor was observed (Fig. S6). Thus, ensembles can be partially distinguished by properties such as firing rate, field location, and the number of tracked sessions.

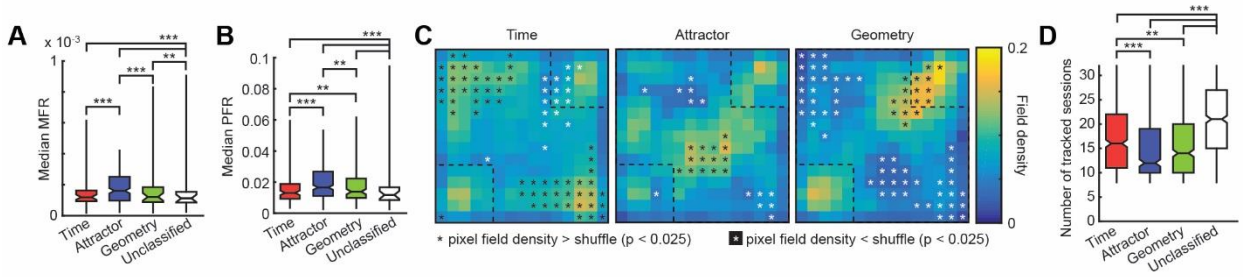

Figure 4. Component ensembles are partially distinguished by firing rate, field location, and number of tracked sessions. *Exact outcomes for all statistical tests performed here are*

*shown in Table S1. (**A**) Median mean firing rate (MFR) as a function of group for cells meeting criteria, and for cells not meeting criteria (unclassified). (**B**) Median peak firing rate (PFR) as a function of group and for cells not meeting criteria. (**C**) Field density within the environment separately for each group. Field density was computed by first averaging the maps from all tracked sessions within each cell. Next, density at a given pixel was defined as the proportion of cells with activity exceeding 75% of their average map PFR at that pixel. Shuffled density distributions were computed by randomly permuting group assignment 1000 times. (**D**) Number of sessions in which a cell was tracked as a function of group and for cells not meeting criteria. **p < 0.01, ***p < 0.001*

Discussion

Here we show that the CA1 spatial code is defined by distinct representational components with different long-term dynamics, including stable components representing spatial geometry and prior experience. Stable and time-varying components were distinct at the level of the population as well as individual cells. These components were not distinguishable by their short-term spatial stability or precision, though they were partially differentiated by other characteristics such as firing rate and field location.

These findings have several important implications. First, by assaying representational content across extended experience, these results go beyond prior work[12,22] to establish that differing long-term dynamics governing the CA1 spatial code are mapped to distinct representational components, including time, spatial geometry, and experience-dependent attractor dynamics. Components representing the spatial context, such as those coding spatial geometry and prior experience in similar environments, are persistently reinstated across a timescale of weeks. This mix of persistent and time-varying representational dynamics is consistent with many theories of episodic memory and hippocampal function, such as indexing theory[23] and temporal context models[24]. In such theories, a time-varying component allows for the encoding and memory of the temporal relationship between different events, while the persistent context-dependent component allows for linking events occurring within the same spatial context across time.

These results also demonstrate that long-term changes in the CA1 spatial code do not reflect homogenous representational drift, which could result from uniform plasticity at the level of the population[13]. Rather, these changes likely reflect either the selective influence of plasticity on a distinct subpopulation within CA1 and/or the influence of time-varying input(s) mapped onto this subpopulation, such as those found in the lateral entorhinal cortex[25,26] and hippocampal subregion CA2[27]. Such time-varying inputs may include surrounding soundscapes, scentscapes, motivations, and intervening experience that cannot be completely controlled in this paradigm, as well as a representation of time per se.

Our finding that these distinct components coexist within the population code adds to a growing literature that challenges the traditional view of a representationally homogenous CA1 spatial code[28–30], and in turn also challenges models which assume a single state space[31] or basis set[32] as the source of the CA1 spatial code. Notably, the short-term spatial reliability and precision of a cell was predictive of that cell coding some representational content but not the particular content itself – suggesting that such properties alone are generally insufficient for inferring the representational content of a cell. Finally, these results provide strong evidence of the heterogenous and malleable nature of hippocampal remapping[33,34], one that is linked to a dynamic recruitment and subsequent reinstatement of distinct neural ensembles[13].

References

1. O'Keefe, J. & Burgess, N. Geometric determinants of the place fields of hippocampal neurons. *Nature* **381**, 425–8 (1996).
2. Muller, R. U. & Kubie, J. L. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J. Neurosci.* **7**, 1951–68 (1987).
3. Leutgeb, S. *et al.* Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* **309**, 619–23 (2005).
4. Deshmukh, S. S. & Knierim, J. J. Influence of local objects on hippocampal representations: Landmark vectors and memory. *Hippocampus* **23**, 253–67 (2013).
5. 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**, 623–33 (2000).
6. Keinath, A. T., Nieto-Posadas, A., Robinson, J. C. & Brandon, M. P. DG–CA3 circuitry mediates hippocampal representations of latent information. *Nat. Commun.* **11**, 3026 (2020).
7. Ferbinteanu, J. & Shapiro, M. L. Prospective and retrospective memory coding in the hippocampus. *Neuron* **40**, 1227–39 (2003).
8. Wills, T. J., Lever, C., Cacucci, F., Burgess, N. & O'Keefe, J. Attractor dynamics in the hippocampal representation of the local environment. *Science* **308**, 873–6 (2005).
9. Colgin, L. L. *et al.* Attractor-map versus autoassociation based attractor dynamics in the hippocampal network. *J. Neurophysiol.* **104**, 35–50 (2010).
10. O'Keefe, J. & Dostrovsky, J. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* **34**, 171–5 (1971).
11. Kentros, C. G., Agnihotri, N. T., Streater, S., Hawkins, R. D. & Kandel, E. R. Increased attention to spatial context increases both place field stability and spatial memory. *Neuron* **42**, 283–95 (2004).
12. Ziv, Y. *et al.* Long-term dynamics of CA1 hippocampal place codes. *Nat. Neurosci.* **16**, 264–6 (2013).
13. Mau, W., Hasselmo, M. E. & Cai, D. J. The brain in motion: How ensemble fluidity drives memory-updating and flexibility. *Elife* **9**, (2020).
14. Chambers, A. R. & Rumpel, S. A stable brain from unstable components: Emerging concepts and implications for neural computation. *Neuroscience* **357**, 172–184 (2017).
15. Pnevmatikakis, E. A. & Giovannucci, A. NoRMCorre: An online algorithm for piecewise rigid motion correction of calcium imaging data. *J. Neurosci. Methods* **291**, 83–94 (2017).
16. Pnevmatikakis, E. A. *et al.* Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data. *Neuron* **89**, 285–99 (2016).
17. Zhou, P. *et al.* Efficient and accurate extraction of in vivo calcium signals from microendoscopic video data. *Elife* **7**, (2018).
18. Friedrich, J., Zhou, P. & Paninski, L. Fast online deconvolution of calcium imaging data. *PLoS Comput. Biol.* **13**, e1005423 (2017).
19. Sheintuch, L. *et al.* Tracking the Same Neurons across Multiple Days in Ca2+ Imaging Data. *Cell Rep.* **21**, 1102–1115 (2017).

20. Kriegeskorte, N. Representational similarity analysis – connecting the branches of systems neuroscience. *Front. Syst. Neurosci.* (2008). doi:10.3389/neuro.06.004.2008

21. Mizuseki, K. & Buzsáki, G. Preconfigured, skewed distribution of firing rates in the hippocampus and entorhinal cortex. *Cell Rep.* **4**, 1010–21 (2013).

22. Rubin, A., Geva, N., Sheintuch, L. & Ziv, Y. Hippocampal ensemble dynamics timestamp events in long-term memory. *Elife* **4**, (2015).

23. Teyler, T. J. & DiScenna, P. The hippocampal memory indexing theory. *Behav. Neurosci.* **100**, 147–154 (1986).

24. Howard, M. W. & Kahana, M. J. A Distributed Representation of Temporal Context. *J. Math. Psychol.* **46**, 269–299 (2002).

25. Tsao, A. *et al.* Integrating time from experience in the lateral entorhinal cortex. *Nature* **561**, 57–62 (2018).

26. Deshmukh, S. S. & Knierim, J. J. Representation of non-spatial and spatial information in the lateral entorhinal cortex. *Front. Behav. Neurosci.* **5**, 69 (2011).

27. Mankin, E. A., Diehl, G. W., Sparks, F. T., Leutgeb, S. & Leutgeb, J. K. Hippocampal CA2 Activity Patterns Change over Time to a Larger Extent than between Spatial Contexts. *Neuron* **85**, 190–201 (2015).

28. Sun, X. *et al.* Functionally Distinct Neuronal Ensembles within the Memory Engram. *Cell* **181**, 410-423.e17 (2020).

29. Tanaka, K. Z. *et al.* The hippocampal engram maps experience but not place. *Science (80-.).* **361**, 392–397 (2018).

30. Taxidis, J. *et al.* Differential Emergence and Stability of Sensory and Temporal Representations in Context-Specific Hippocampal Sequences. *Neuron* **108**, 984-998.e9 (2020).

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

32. Barry, C. *et al.* The boundary vector cell model of place cell firing and spatial memory. *Rev. Neurosci.* **17**, 71–97 (2006).

33. Sanders, H., Wilson, M. A. & Gershman, S. J. Hippocampal remapping as hidden state inference. *Elife* **9**, (2020).

34. Anderson, M. I. & Jeffery, K. J. Heterogeneous modulation of place cell firing by changes in context. *J. Neurosci.* **23**, 8827–35 (2003).

35. Giovannucci, A. *et al.* CaImAn an open source tool for scalable calcium imaging data analysis. *Elife* **8**, (2019).

36. Skaggs, W. E., McNaughton, B. L. & Gothard, K. M. An information-theoretic approach to deciphering the hippocampal code. in *Advances in neural information processing systems* 1030–1037 (1993).

Methods

Subjects

Naive male mice (C57Bl/6, Charles River) were housed in pairs on a 12-hour light/dark cycle at 22°C and 40% humidity with food and water ad libitum. All experiments were carried out during the light portion of the light/dark cycle, and in accordance with McGill University and Douglas Hospital Research Centre Animal Use and Care Committee (protocol #20157725) and with Canadian Institutes of Health Research guidelines.

Surgeries

During all surgeries, mice were anesthetized via inhalation of a combination of oxygen and 5% Isoflurane before being transferred to the stereotaxic frame (David Kopf Instruments), where anesthesia was maintained via inhalation of oxygen and 0.5-2.5% Isoflurane for the duration of the surgery. Body temperature was maintained with a heating pad and eyes were hydrated with gel (Optixcare). Carprofen (10 ml kg^{-1}) and saline (0.5 ml) were administered subcutaneously at the beginning of each surgery. Preparation for recordings involved three surgeries per mouse.

First, at the age of six to ten weeks, each mouse was transfected with a 350 nl injection of the calcium reporter GCaMP6f via the viral construct AAV5.CaMKII.GCaMP6f.WPRE.SV40 (CAMCA130) or AAV9.syn.GCaMP6f.WPRE.SV40 (all other mice). The original titre of the AAV9.syn.GCaMP6f.WPRE.SV40 construct, sourced from University of Pennsylvania Vector Core, was 3.26e14 GC-ml and was diluted in sterile PBS (1:1 AKCA170; 1:5 AKCA127; 1:10 AKCA130; 1:15 AKCA135) before surgical microinjection. The original titre of the AAV5.CaMKII.GCaMP6f.WPRE.SV40 construct, sourced from Addgene, was 2.3e13 GC-ml and was diluted in sterile PBS (1:3) before surgical microinjection.

One to three weeks post-injection, either a 1.8mm (AKCA127, AKCA130, AKCA135) or 0.5mm (AKCA170, CAMCA130) diameter gradient refractive index (GRIN) lens (Go!Foton) was implanted above dorsal CA1 (Referenced to bregma: ML = 2.0 mm, AP = -2.1 mm; Referenced to brain surface: DV = -1.35 mm). Implantation of the 1.8mm diameter GRIN lens required aspiration of intervening cortical tissue, while no aspiration was required for implantation of the 0.5mm diameter GRIN lens. Results observed using 1.8- or 0.5-mm diameter GRIN lenses were similar. In addition to the GRIN lens, two stainless steel screws were threaded into the skull above the contralateral hippocampus and prefrontal cortex to stabilize the implant. Dental cement (C&B Metabond) was applied to secure the GRIN lens and anchor screws to the skull. A silicone adhesive (Kwik-Sil, World Precision Instruments) was applied to protect the top surface of the GRIN lens until the next surgery.

One to three weeks after lens implantation, an aluminum baseplate was affixed via dental cement (C&B Metabond) to the skull of the mouse, which would later secure the miniaturized fluorescent endoscope (miniscope) in place during recording. The miniscope/baseplate was mounted to a stereotaxic arm for lowering above the implanted GRIN lens until the field of view contained visible cell segments and dental cement was applied to affix the baseplate to the skull. A polyoxymethylene cap with a metal nut weighing ~3 g was affixed to the baseplate when the mice were not being recorded, to protect the baseplate and lens, as well as to simulate the weight of the miniscope.

After surgery, animals were continuously monitored until they recovered. For the initial three days after surgery mice were provided with a soft diet supplemented with Carprofen for pain management (MediGel CPF, ~5mg kg^{-1} each day). Familiarization to both environments (while recording in the room A-associated environment to monitor imaging quality and habituate the mouse to recording) began 3 to 7 days following the baseplate surgery.

Data acquisition

In vivo calcium videos were recorded with a UCLA miniscope (v3; miniscope.org) containing a monochrome CMOS imaging sensor (MT9V032C12STM, ON Semiconductor) connected to a custom data acquisition (DAQ) box (miniscope.org) with a lightweight, flexible coaxial cable. The DAQ was connected to a PC with a USB 3.0 SuperSpeed cable and controlled with Miniscope custom acquisition software (miniscope.org). The outgoing excitation LED was set to between 2-8% (~0.05-0.2 mW), depending on the mouse to maximize signal quality with the minimum possible excitation light to mitigate the risk of photobleaching. Gain was adjusted to match the dynamic range of the recorded video to the fluctuations of the calcium signal for each recording to avoid saturation. Behavioral video data were recorded by a webcam mounted above the environment. Behavioral video recording parameters were adjusted such that only the red LED on the CMOS of the miniscope was visible. The DAQ simultaneously acquired behavioral and cellular imaging streams at 30 Hz as uncompressed avi files and all recorded frames were timestamped for post-hoc alignment.

All recording environments were constructed of a grey Lego base and black Lego bricks (Lego, Inc) according to the dimensions specified in the main text and supplemental figures. All external walls had a height of 22 cm. During recording, the environment was dimly lit by a nearby computer screen, which could serve as directional cue. During familiarization and all unrecorded experience the environments were well-lit by overhead room lighting. All sessions were 20 min, and only one session was recorded per day to avoid photobleaching. The mouse was always placed in the same corner at the start of the session and was allowed to explore the environment for 15 to 30 s prior to the start of data acquisition. Following each recording the environment was cleaned with disinfectant (Prevail).

Data preprocessing

Calcium imaging data were preprocessed prior to analyses via a pipeline of open source MATLAB (MathWorks; version R2015a) functions to correct for motion artifacts[15], segment cells and extract transients[16,17], and infer the likelihood of spiking events via deconvolution of the filtered transient trace through a second-order autoregressive model[18]. The motion-corrected calcium imaging data were manually inspected to ensure that motion correction was effective and did not introduce additional artifacts. Following this preprocessing pipeline, the spatial footprints of all cells were manually verified to remove lens artifacts. Position data were inferred from the onboard miniscope red LED offline following recording using a custom written MATLAB (MathWorks) script and were manually corrected if needed. Cells were tracked across sessions on the basis of their spatial footprints and/or centroids[19].

Data analysis

All analyses were conducted using the vector of likelihood of spiking events inferred via deconvolution, treating this vector as if it were the firing rate of the cell (henceforth *firing rate*) [18]. Similar results were observed when the likelihood of spiking was inferred as a binarized vector

indicating periods when calcium transients were rising, instead of deconvolution. Only periods of movement for which the mouse speed exceeded 2 cm s^{-1} were included in analyses, though no qualitative differences were observed if this criterion was relaxed.

Rate maps were constructed by first binning the position data into pixels corresponding to a 2.5 cm x 2.5 cm grid of locations. Then the mean firing rate was computed for each pixel and then smoothed with a 4 cm standard deviation isometric Gaussian kernel. For all comparisons between rate maps, similarity was measured as the Pearson's correlation between corresponding pixels. Individual cell representational similarity matrices (RSMs) were thus computed as the Pearson's correlation between pairwise rate map comparisons for tracked sessions. Population-level RSMs were computed by taking the median across all cells for each pairwise comparison.

Within-sequence analyses summarized morph sequence dynamics with a transition plot, which captured the similarity of all 6 morph environment maps to the 2 familiar environment maps. To this end, rate map correlations between each environment and familiar rate maps at the beginning and end of the morph sequence were computed. Only comparisons between cells whose within-session split-half rate map correlation (SHC) exceed the 99th percentile of a shuffled control for at least one of the compared sessions were included. The shuffled control was computed for each cell by randomly circularly shifting its firing rate vector relative to the position data by at least 30 sec and recomputing the SHC 1000 times to create the null distribution. The median across all pairwise cell comparisons was taken as the measure of central tendency for transition plots. To characterize these plots, the Fisher-transformed median values comparing the morph environments to each of the two familiar environments were both fit to minimize mean squared error with a 4-parameter sigmoid of the form

$$f(x) = p_3 + \left(1 + \frac{1}{p_2 e^{(-x+p_1)}}\right) / p_4$$

where x is the position of the current environment in the shapespace (arbitrarily chosen to range from 1 to 6), and parameters p determine the shape of the sigmoid. The intersection of these two sigmoids was taken as the measure of the transition point between familiar maps. The maximum absolute difference between these two sigmoids within the sampled shapespace was taken as the measure of maximum map decorrelation. Sigmoidality was computed as the mean absolute difference between each sigmoid and a linear fit from endpoint to endpoint, with endpoints normalized to range from zero to one.

Two-dimensional embeddings of the population representation were computed and quantified as follows. First, the population RSM was computed by taking the median across cells for each comparison. Next, this RSM was computed to a distance matrix as 1-(RSM) with the diagonal set to a distance of 0. Finally, this distance matrix was reduced to a two-dimensional embedding via Kruskal's nonmetric multidimensional scaling, implemented by the MATLAB function mdscale with the default parameterizations. To estimate the context dimension within this embedding, the average difference separating neighboring-in-time familiar environment recordings was computed. To estimate the time dimension within this embedding, the average difference between the first and last recordings of each familiar environment was computed.

Characterization of individual cell RSM variance was computed as follows. Only cells tracked across at least 28 pairwise comparisons (8 sessions) were included. Model RSMs were

computed on the basis of temporal, geometric, and attractor group similarity, as schematized in Figure 3a. The time RSM was computed as the difference in session number between pairwise comparisons. The geometry RSM was the difference in environment shape between pairwise comparisons, assuming that consecutive environments mapped linear steps through the shapespace separating the two familiar environments. The attractor RSM was computed such that if a pairwise comparison was between two sessions on the same side of the within-sequence transition point then those environments were assigned a similarity of 1; otherwise, that comparison was assigned a similarity of 0. Models consisting of all factors and their interactions (full model) as well as models with each factor and their interactions removed were then fit to the Fisher-transform individual cell RSMs as described in the main text. Notably, geometry and attractor factors were often correlated (though the extent of their correlation differed depending on transition point dynamics and the particular subset of tracked sessions for each cell), often limiting the amount of explained variance that could be attributed to one or the other. The time RSMs tended to be uncorrelated with either the geometric or attractor RSMs.

Signal-to-noise (SNR) for each dF/F trace was computed as a measure of data quality prior to spike estimation via deconvolution as described in previously[35]. Briefly, because calcium transients around the baseline can only be positive, for each trace the variance of a normal noise distribution was estimated by scaling the standard deviation of trace values below baseline via a half-normal distribution, such that:

$$\mathrm{NOISE} = \frac{\mathrm{std}(\mathbf{t}(\mathbf{t} < 0))}{\sqrt{1 - \frac{2}{\pi}}}$$

where t is the detrended filtered calcium trace and $\mathbf{t}(\mathbf{t} < 0)$ are all filtered trace values below baseline. Next, we z-score the trace $\mathbf{t}$ such that $z(\mathbf{t}) = \mathbf{t}/\mathrm{NOISE}$. These z-scored values are next transformed into probabilities $p(\mathbf{t}) = \varphi(-z(\mathbf{t}))$, where φ denotes the cumulative distribution function of the standard normal distribution. Then we compute the lowest probability event over a time window of 0.4 s (N = 12 frames), the duration around which a GCamp6F transient is typically at its maximum, such that:

$$p_{\min} = \min_i \left(\prod_{j=0}^{N-1} \mathrm{p}(\mathbf{t}(i+j)) \right)^{1/N}$$

where $\min_i$ is the minimum value across all timepoints i from $i = 1$ to $i =$ the length of $\mathbf{t} - N + 1$. Finally, we define the average peak SNR as:

$$\mathrm{SNR} = -\varphi^{-1}(p_{\min})$$

where φ^{-1} is the quantile function of the standard normal distribution (logit function).

Spatial information content (SIC) was computed from the whole-session rate maps of each cell as described previously[36] via the equation:

$$\mathrm{SIC} = \sum_i \mathbf{s}_i \mathbf{r}_i \log_2 \left(\frac{\mathbf{r}_i}{\bar{\boldsymbol{r}}} \right)$$

where i is the rate map pixel index, $\mathbf{s}_i$ is the probability of sampling pixel i, $\mathbf{r}_i$ is the mean firing rate at pixel i, and $\bar{\mathbf{r}}$ is the mean firing rate across all pixels.

Histological validation of expression and recording targets

After experiments, animals were perfused to verify GRIN lens placement. Mice were deeply anesthetized and intracardially perfused with 4% paraformaldehyde in PBS. Brains were dissected and post-fixed with the same fixative. Coronal sections (50 μm) of the entire hippocampus were cut using a vibratome and sections were mounted directly on glass slides. Sections were split and half of all sections were stained for DAPI and mounted with Fluoromount-G (Southern Biotechnology) to localize GRIN lens placement and to evaluate virus expression. Due to the large imageable surface but restricted miniscope field of view (~0.5 mm x ~0.8 mm), we were unable to determine more specific localization of populations within the hippocampus for mice recorded with 1.8 mm lenses.

Statistics and reproducibility

All statistical tests are noted where the corresponding results are reported throughout the main text and supplement. All tests were uncorrected 2-tailed tests unless otherwise noted. Z-values for nonparametric Wilcoxon tests were not estimated or reported for comparisons with fewer than 15 datapoints. In all cases, the median served as the measure of central tendency unless otherwise noted. Box plots portray the minimum and maximum (whiskers), upper and lower quartiles (boxes), and median (cinch).

Code availability

All custom code written for reported analyses are publicly available at [insert Github link] or via request to the corresponding authors.

Data availability

The complete dataset for all experiments are publicly available at [insert DRYAD link] or via request to the corresponding authors.

Acknowledgements

We thank D. Aharoni for extensive guidance in using the UCLA miniscope. We thank J Quinn Lee for helpful feedback on prior versions of this manuscript. During this work ATK was supported by a McGill University Healthy Brains for Healthy Lives CFREF postdoctoral fellowship and a Natural Sciences and Engineering Research Council (NSERC) Banting postdoctoral fellowship. CAM was supported by a Fonds de Recherche du Québec – Santé (FRQS) postdoctoral fellowship. Funding was provided by the Canadian Institutes for Health Research (grants #367017 and #377074), the Natural Sciences and Engineering Research Council of Canada (Discovery grant #74105), the Canada Research Chairs Program, and the Brain Canada Foundation (Future Leaders in Canadian Brain Science).

Author Contributions

ATK contributed to experimental design, surgeries, recordings, analysis of data, as well as drafting and revising the manuscript. CAM contributed to surgeries and histology. MPB contributed to experimental design, analysis of data, as well as drafting and revising the manuscript.

Competing interests.

The authors declare no competing interests.

Supplementary information for:

Extended experience reveals distinct representational dynamics governing the CA1 spatial code

Keinath, et al.

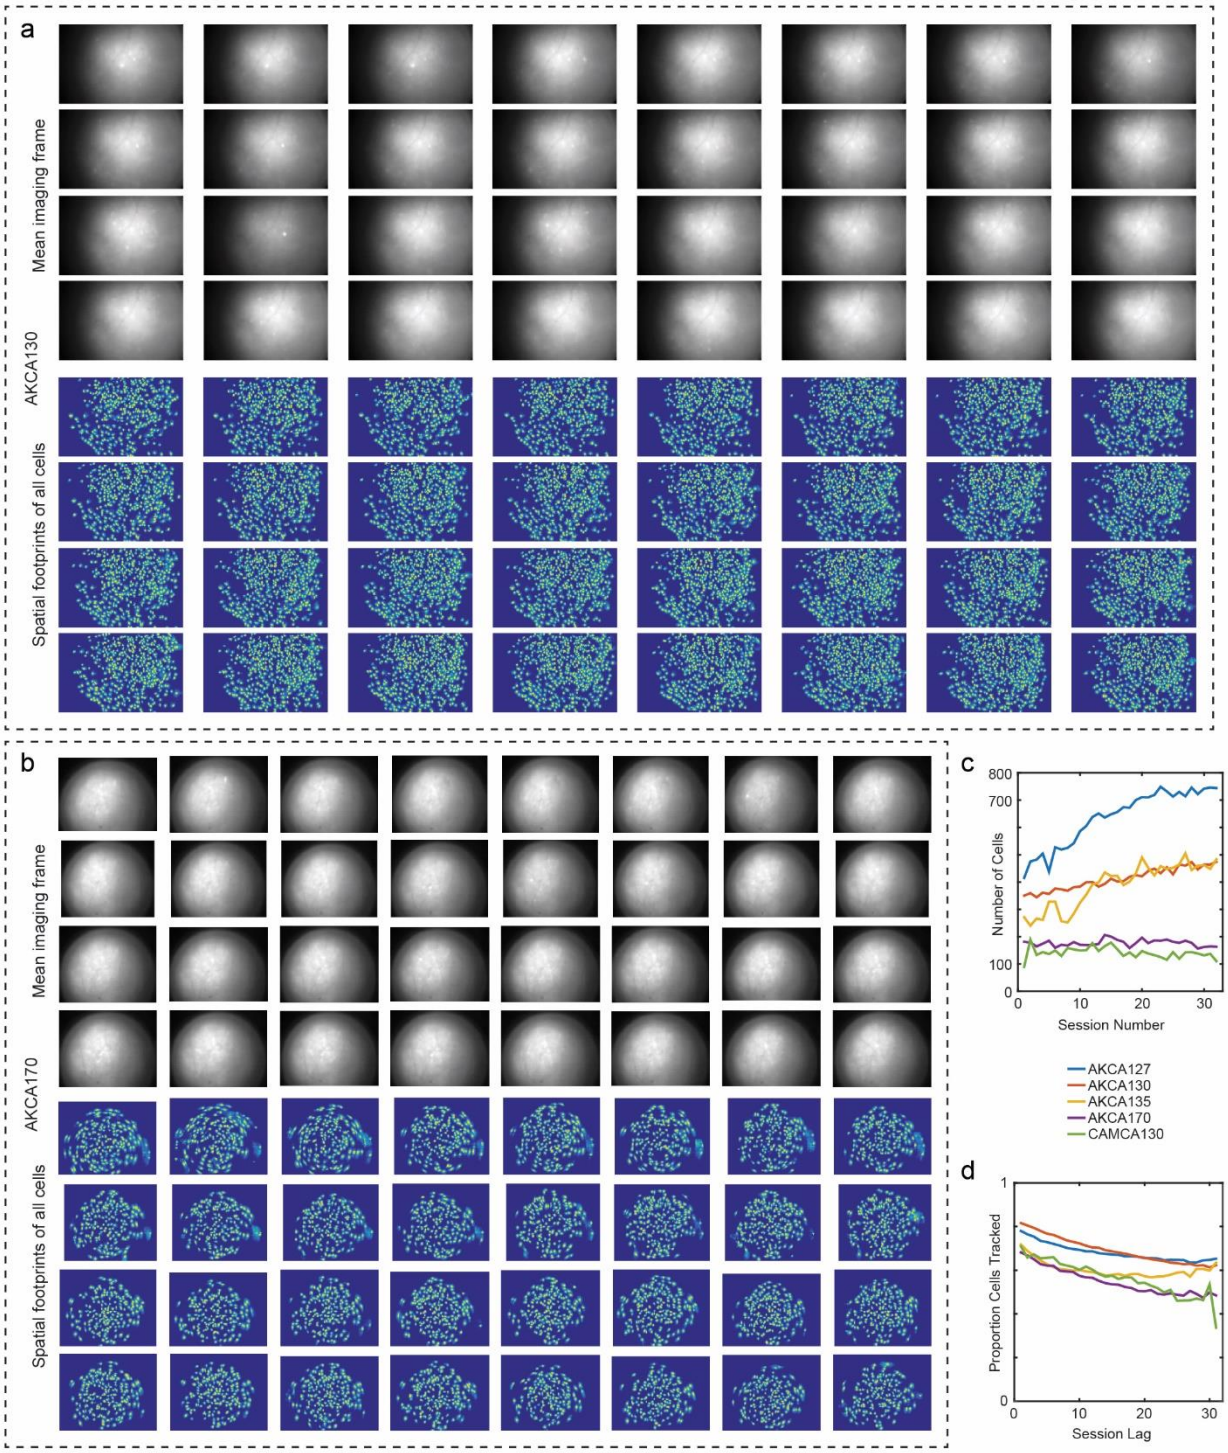

Supplementary Figure 1. Imaging across 32 days of recording. a,b) Arranged chronologically. Two example mice, AKCA130 (a) implanted with a 1.8 mm lens and AKCA170 (b) implanted with a 0.5 mm lens. Mean imaging frame normalized to maximum. Spatial footprints scaled and thresholded for ease of interpretation. c) Cell counts for each mouse across all 32 days. d) Proportion of the population registered as a function of lag between sessions for each mouse. Normalized to the smaller population for each pairwise comparison.

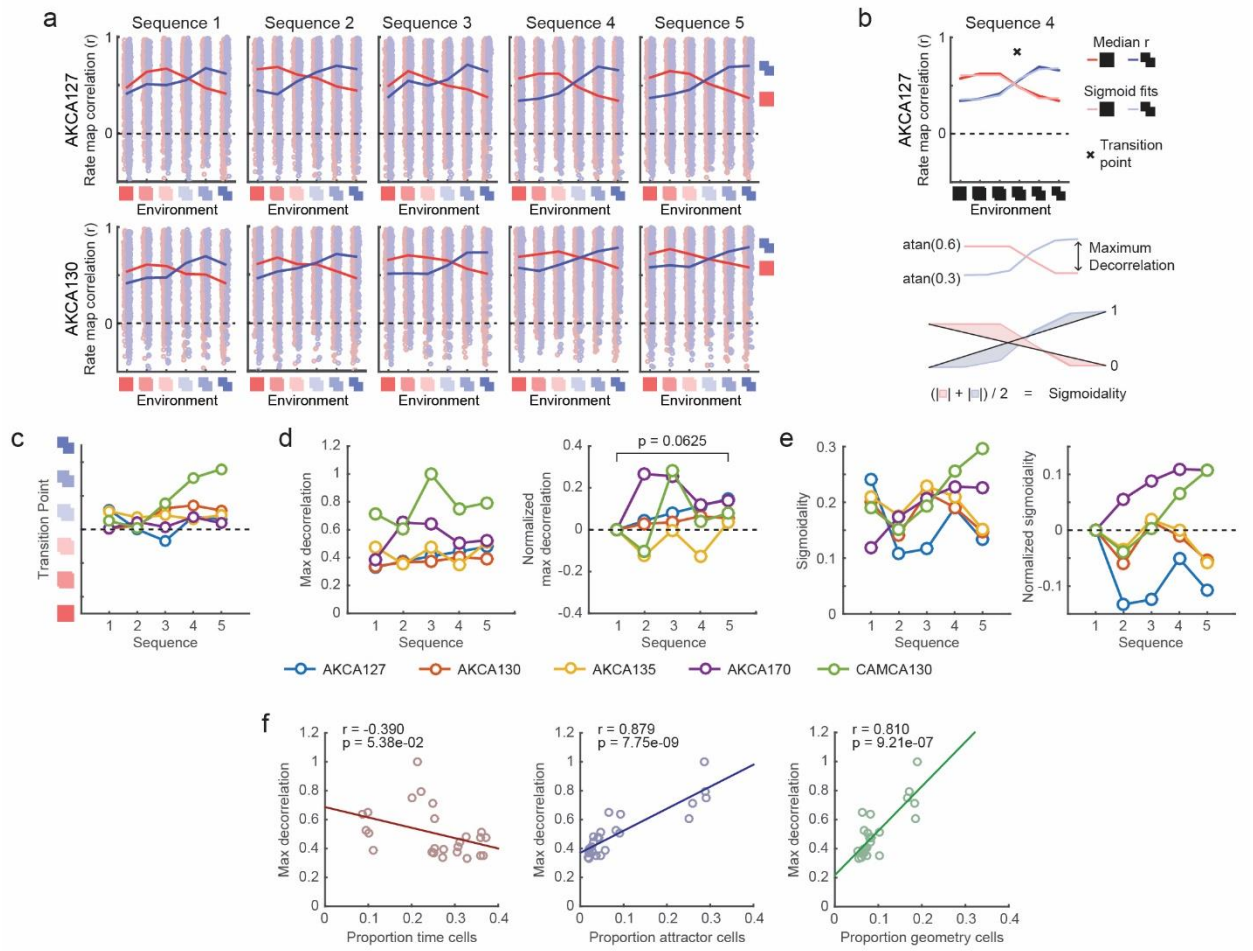

Supplementary Figure 2. Characterization of individual morph sequences. a) Rate map correlations between each environment and the familiar environments for each sequence as in Figure 1d for the two mice not shown there. b) To quantify transition characteristics, both median rate map correlations were fit with a sigmoid. Transition point was computed as the intersection of these sigmoids. Maximum decorrelation was computed as the maximum absolute difference between the Fisher transformation of these sigmoids. Sigmoidality was computed as the mean absolute difference between each sigmoid and a linear fit from endpoint to endpoint, with endpoints normalized to range from zero to one. c) Transition points across sequences for each mouse. Transition points tended to remain near the middle of the morph sequence except in the case of CAMCA130, which showed a progressive change in transition such that more morph environments resembled the square environment over time. Interestingly, the square environment was the room A-associated environment for this animal, where all test recordings took place. d) Max decorrelation by sequence and normalized by the first sequence. A trend toward increased decorrelation between the first and last sequences was observed (Wilcoxon rank sum test). e) Sigmoidality by sequence and normalized by the first sequence. No clear trend emerged. f) Relationship between maximum decorrelation and the mean proportion of cells meeting criteria for each primary factor within each sequence. A strong positive correlation was observed for attractor- and geometry-loading cells, and a marginal negative correlation for time-loading cells, consistent with a distinct contribution of each factor to representations of context and time, respectively.

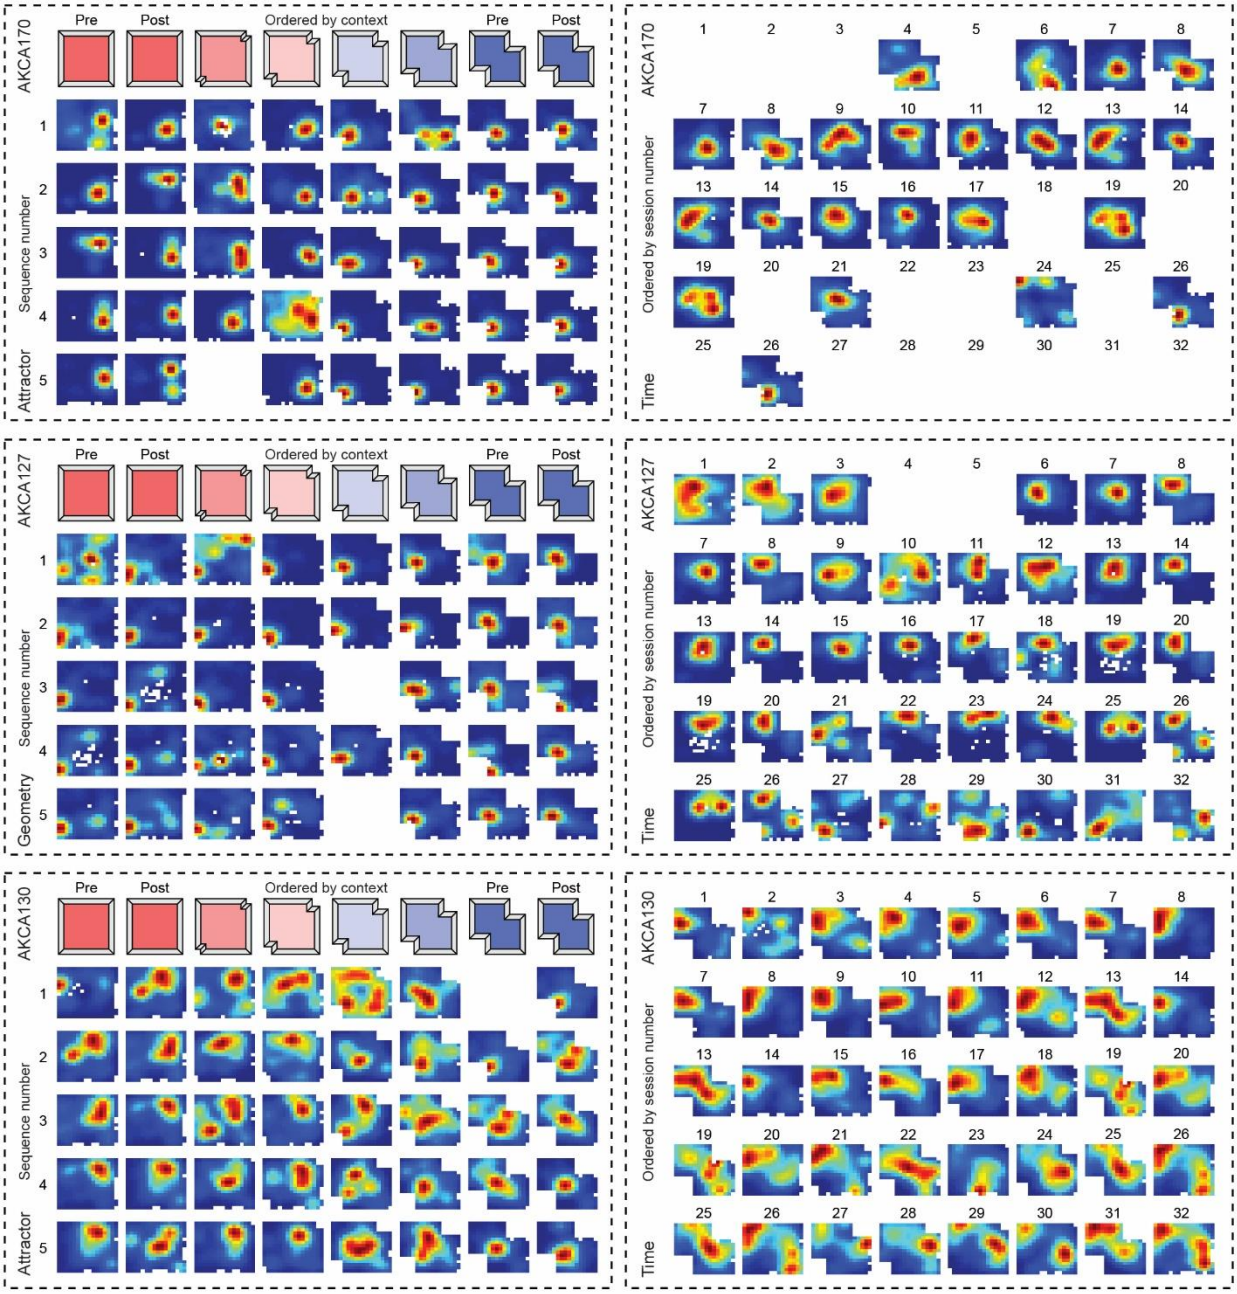

Supplementary Figure 3. Examples of individual cells across all sessions. Two simultaneously-recorded cells each from 3 mice. All cells met full model r^2 inclusion criteria, with their primary factors identified in the lower left of each panel. For ease of interpretation, context-loading cells are organized by context while time-loading cells are organized by session number. The attractor cell from AKCA170 is the example cell whose RSM is shown in Figure 3a.

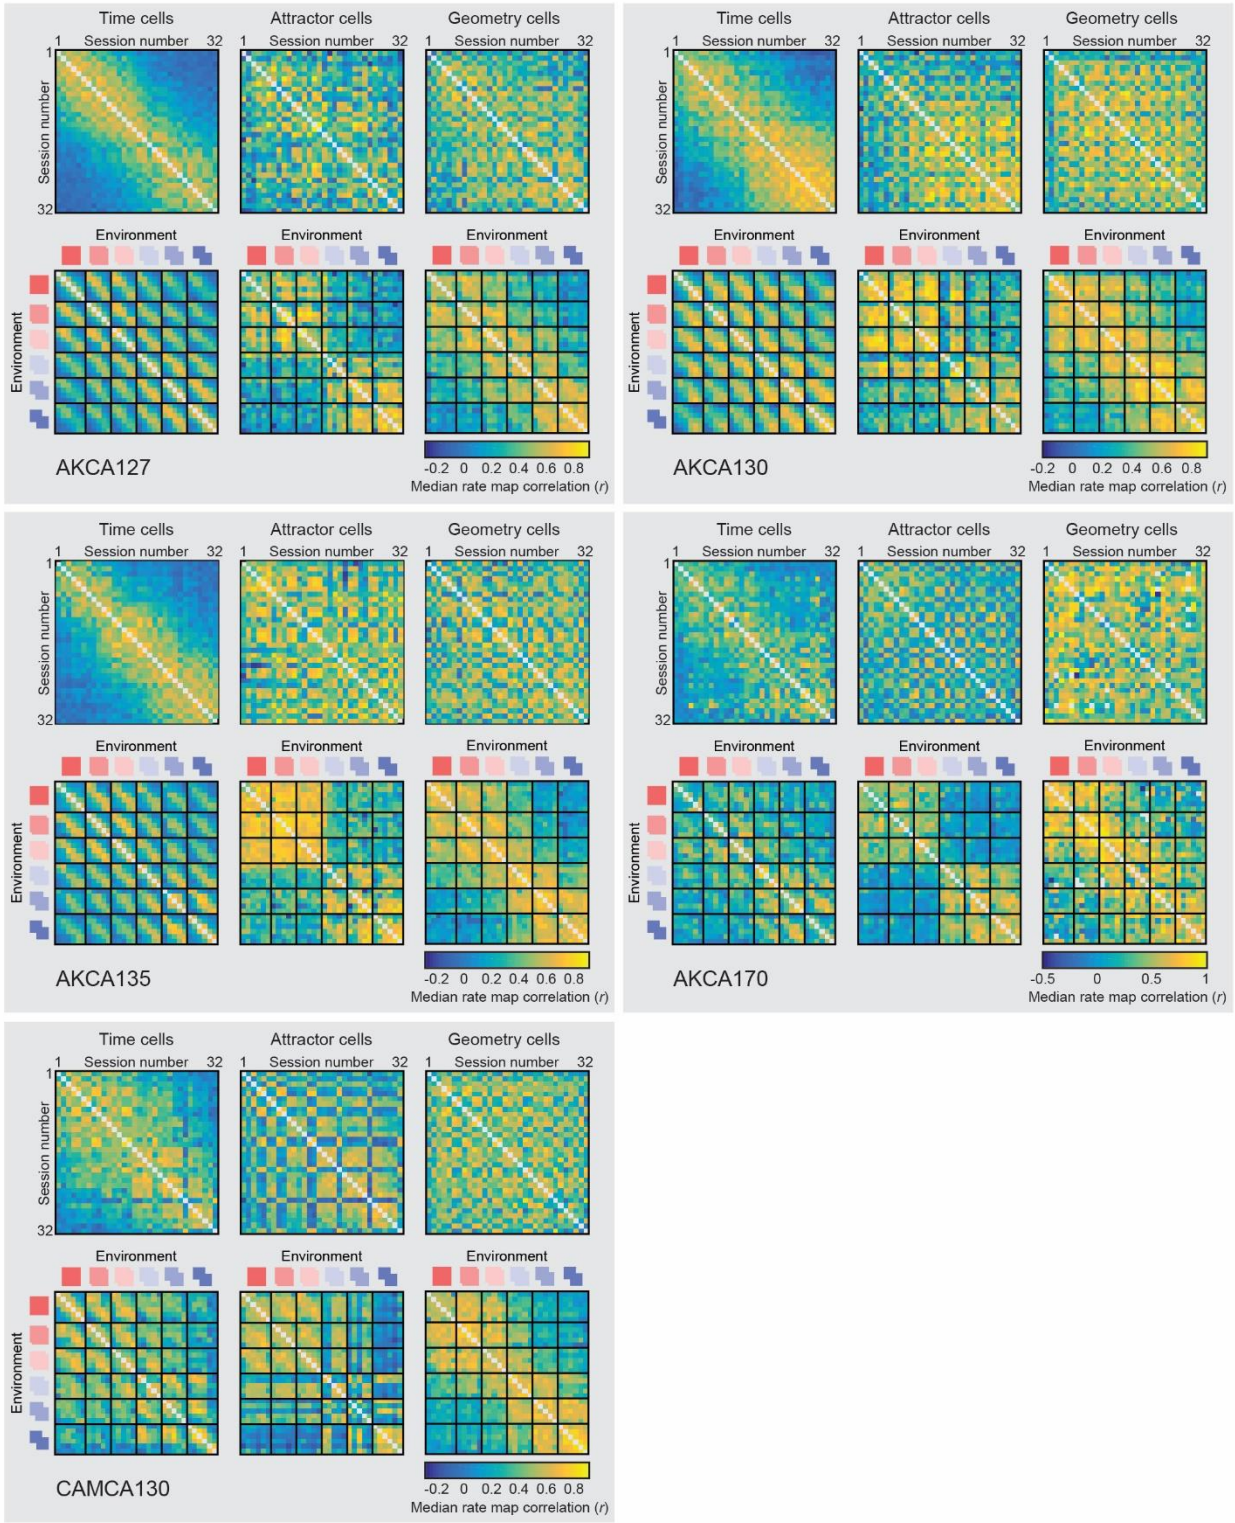

Supplementary Figure 4. Population RSMs for each cell group separated by mouse. RSMs of time, attractor, and geometry-loading cells meeting criteria, ordered by time (top) and context (bottom), separated by mouse.

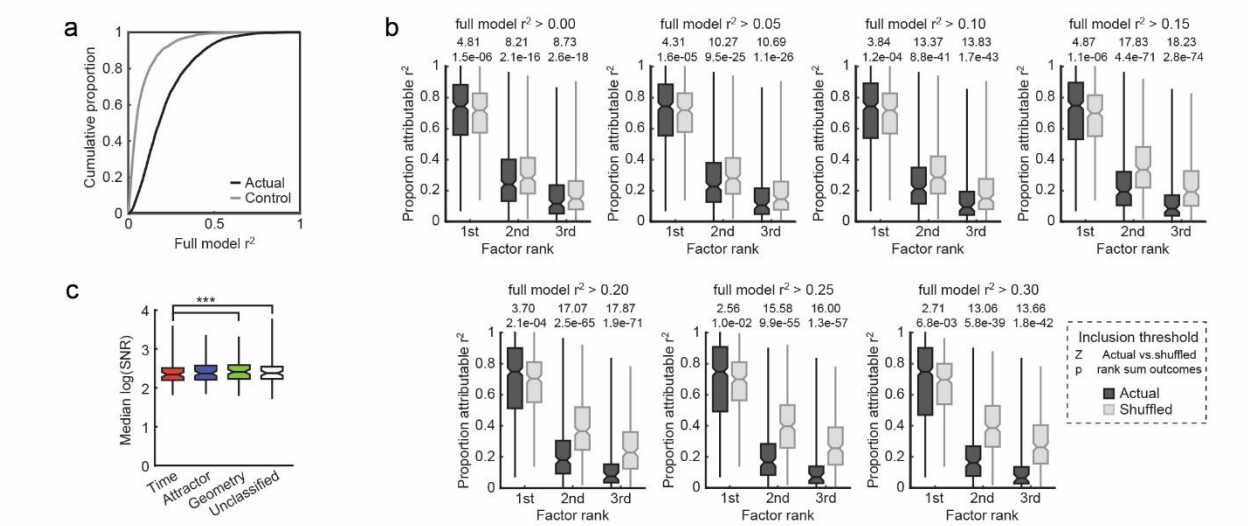

Supplementary Figure 5. Reported ranked-factor results are robust to a variety of inclusion criteria. All statistics denote the outcome of Wilcoxon rank sum tests between actual and control distributions. a) Cumulative distributions of full model r^2 for the actual and shuffled control populations (n = 2745). b) Proportion of attributable variance when factors are ranked for a variety of minimum full model r^2 thresholds. In all cases, the primary factor was attributed more variance, while the secondary and ternary factors were attributed less variance than would be expected by chance given population statistics. c) Signal-to-noise (SNR) as a function of group for cells meeting criteria and for cells not meeting criteria (Time vs Attractor, Z = -1.88, p = 6.01e-02; Time vs Geometry, Z = -3.60, p = 3.19e-04; Time vs Unclassified, Z = -3.42, p = 6.22e-04; Attractor vs Geometry, Z = -0.97, p = 3.31e-01; Attractor vs Unclassified, Z = 0.25, p = 8.02e-01; Geometry vs Unclassified, Z = 1.80, p = 7.19e-02).

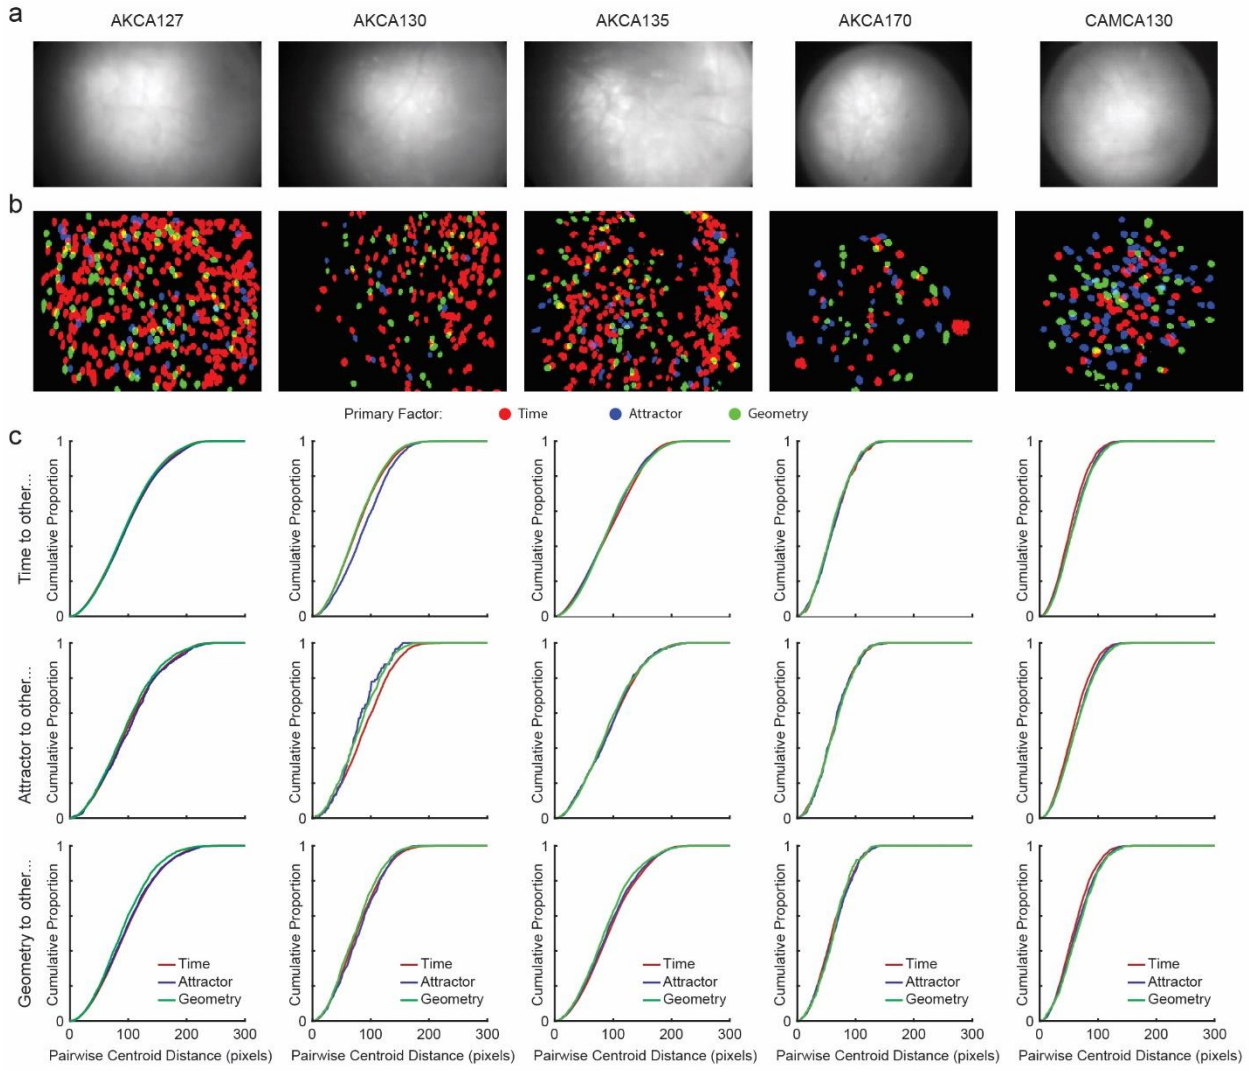

Supplementary Figure 6. Anatomical organization of cells as a function of primary factors. a) Mean imaging frame for each mouse. b) Distribution of spatial footprints (SFPs) of all cells meeting full model r^2 inclusion criterion, color-coded by primary factor. c) Cumulative distribution of pairwise distances between SFP centroids for all combinations of cell groups separately for each animal. No evidence of reliably different anatomical distributions as a function of primary factor loading was observed.

Figure	Panel	Groups	Outcome
3	E	Rank 1st vs shuffled	Z = 4.27, p = 1.99e-05
		Rank 2nd vs shuffled	Z = -16.15, p = 1.19e-58
		Rank 3rd vs shuffled	Z = -17.48, p = 1.89e-68
	G	Time vs Attractor	Z = -0.82, p = 4.12e-01
		Time vs Geometry	Z = -0.16, p = 8.74e-01
		Time vs Unclassified	Z = 9.96, p = 2.20e-23
		Attractor vs Geometry	Z = 0.56, p = 5.76e-01
		Attractor vs Unclassified	Z = 5.44, p = 5.27e-08
		Geometry vs Unclassified	Z = 6.06, p = 1.36e-09
	J	Across: Time vs Attractor	Z = -15.40, p = 1.62e-53
		Across: Time vs Geometry	Z = -18.65, p = 1.16e-77
		Across: Attractor vs Geometry	Z = 0.14, p = 8.91e-01
		Within: Time vs Attractor	Z = -2.89, p = 3.88e-03
		Within: Time vs Geometry	Z = -1.95, p = 5.11e-02
		Within: Attractor vs Geometry	Z = 0.85, p = 3.96e-01
		Across Time vs Within Time	Z = -21.59, p = 2.02e-103
		Across Attractor vs Within Attractor	Z = -2.17, p = 3.01e-02
		Across Geometry vs Within Geometry	Z = -1.36, p = 1.74e-01
	K	Time vs Attractor	Z = -1.88, p = 5.95e-02
		Time vs Geometry	Z = -1.46, p = 1.43e-01
		Time vs Unclassified	Z = 4.68, p = 2.86e-06
		Attractor vs Geometry	Z = 0.59, p = 5.54e-01
		Attractor vs Unclassified	Z = 4.21, p = 2.60e-05
		Geometry vs Unclassified	Z = 4.50, p = 6.77e-06
4	A	Time vs Attractor	Z = -4.87, p = 1.12e-06
		Time vs Geometry	Z = -0.55, p = 5.85e-01
		Time vs Unclassified	Z = 4.08, p = 4.44e-05
		Attractor vs Geometry	Z = 3.53, p = 4.11e-04
		Attractor vs Unclassified	Z = 6.61, p = 3.83e-11
		Geometry vs Unclassified	Z = 3.04, p = 2.38e-03
	B	Time vs Attractor	Z = -5.22, p = 1.80e-07
		Time vs Geometry	Z = -2.31, p = 2.11e-02
		Time vs Unclassified	Z = 4.74, p = 2.18e-06
		Attractor vs Geometry	Z = 2.82, p = 4.80e-03
		Attractor vs Unclassified	Z = 7.49, p = 6.84e-14
		Geometry vs Unclassified	Z = 5.34, p = 9.18e-08
	D	Time vs Attractor	Z = 4.21, p = 2.51e-05
		Time vs Geometry	Z = 2.85, p = 4.33e-03
		Time vs Unclassified	Z = -12.50, p = 7.10e-36
		Attractor vs Geometry	Z = -1.71, p = 8.65e-02
		Attractor vs Unclassified	Z = -10.22, p = 1.54e-24
		Geometry vs Unclassified	Z = -10.70, p = 1.06e-26

Supplementary Table 1. Exact statistical outcomes for all tests not reported in main text figures. All tests are Wilcoxon rank sum tests. P-values are uncorrected.