Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Stable representation of a naturalistic movie emerges from episodic activity with gain variability

Ji Xia (xiaji@wustl.edu)
Washington University in St. Louis https://orcid.org/0000-0003-1349-9114
Tyler Marks
University of California, Santa Barbara
Michael Goard
University of California, Santa Barbara https://orcid.org/0000-0002-5366-8501
Ralf Wessel
Washington University in St. Louis

Article

Keywords: visual cortical responses, neuronal responses

Posted Date: January 7th, 2021

DOI: https://doi.org/10.21203/rs.3.rs-126977/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Version of Record: A version of this preprint was published at Nature Communications on August 27th, 2021. See the published version at https://doi.org/10.1038/s41467-021-25437-2.

Title: Stable representation of a naturalistic movie emerges from episodic activity with gain variability

Authors:
Ji Xia[1], Tyler D. Marks[2], Michael J. Goard[2,3,4*], Ralf Wessel[1*]
1 Department of Physics, Washington University in St. Louis
2 Neuroscience Research Institute, University of California, Santa Barbara
3 Department of Molecular, Cellular, and Developmental Biology, University of California, Santa Barbara
4 Department of Psychological & Brain Sciences, University of California, Santa Barbara
* These authors contributed equally

Abstract

Visual cortical responses are known to be highly variable across trials within an experimental session. However, the long-term stability of visual cortical responses is poorly understood. Chronic imaging experiments in V1 showed that neural responses to repeated natural movie clips were unstable across weeks. Single neuronal responses consisted of sparse episodic activity which were stable in time but unstable in spike rates across weeks. Further, we found that the individual episode, instead of neuron, served as the basic unit of the week-to-week fluctuation. To investigate how population activity encodes the stimulus, we extracted a stable one-dimensional representation of the time in the natural movie, using an unsupervised method. Moreover, most week-to-week fluctuation was perpendicular to the stimulus encoding direction, thus leaving the stimulus representation largely unaffected. We propose that precise episodic activity with coordinated gain changes are keys to maintain a stable stimulus representation in V1.

Introduction

Stimulus-driven activity is highly variable across repeated trials within a recording session [1–5]. Furthermore, in chronic recordings covering multiple stimulus sessions, session-to-session fluctuation tends to be qualitatively different from trial-to-trial variability within sessions [6–9]. Even without learning, the same neuron population responds unstably under the same environmental and behavioral conditions across days [10–14]. However, not all the brain areas share the same instability [9]. For example, PPC neurons [11] and hippocampal neurons [14] exhibit large changes across days, while HVC neural activity remains stable in long-term recordings [15].

How does stimulus-driven activity in V1 change across days under nominally constant condition? Recently, several studies shed light on how V1 stimulus-driven activity changes in the long term in responses to drifting gratings [16–18]. Even though day-to-day variations were larger than trial-to-trial variations [17], stable tuning over weeks were found in most tuned neurons [16,17]. Yet few reported on the long-term stability of neural responses to natural movies [18,19]. Natural movie responses are sparser and more precise than neural responses to artificial stimuli such as drifting gratings [20,21]. Moreover, responses to natural stimuli cannot be predicted from responses to drifting gratings [22,23]. Thus, the long-term stability of neural responses to natural movies are not necessarily the same as that to drifting gratings. Indeed, data from our group showed that single neural responses to natural movies were significantly more unstable than drifting grating responses [24].

This session-to-session fluctuation raises an important question: Is there a stable representation of natural stimuli hidden in the unstable neural activity in V1? Stable stimulus representation is possible when neural fluctuations reside in a space orthogonal to the stimulus encoding dimensions [25]. Intuitively, if one neuron's session-to-session fluctuation affected the encoding of stimulus, then the other neurons' fluctuation could compensate for its influence. Moreover, the stimulus could be encoded in a low-dimensional subspace of the high-dimensional population activity [26,27]. In that case, the random fluctuation in the high-dimensional neural space would likely be perpendicular to the low-dimensional subspace of stimulus encoding, often referred to as the stimulus encoding dimension. Clarification of these possibilities requires long-term recordings in response to repeated stimulation, identification of the stimulus encoding dimensions, and quantification of neural fluctuation within the high-dimensional population activity.

To address the question of stable stimulus representation in unstable neural activity, we analyzed a dataset from longitudinal two-photon calcium imaging of excitatory neurons in the primary visual cortex of awake, head-fixed mice during visual stimulation with repeated identical natural movie clips across weeks. We found that single neural responses consisted of episodic activity that were precise in time during the natural movie across weeks. However, firing rates during those spiking episodes were unstable across weeks. Moreover, within the same neuron, firing rates of different spiking episodes varied in distinct temporal patterns across weeks. By fitting a linear model, we found that episodic activity was the basic unit of the week-to-week fluctuation. Importantly, despite unstable episodic activity, we extracted a low-dimensional stable representation of time in the natural movie from neuronal population activity across weeks. We propose that precise episodic activity with coordinated gain changes are keys to maintain a stable stimulus representation in V1.

Results

Single neuron responses to natural movies are unstable across weeks

To investigate the long-term variability of cortical responses, we used a dataset that consisted of chronic GCaMP6s imaging of excitatory neurons in V1 L2/3 of awake, head-fixed mice (9 mice; 10 imaging fields) during visual stimulation with repeated natural movies (30 trials per session; one session per 7 ± 1 days; over 5-7 weeks) (Fig. 1a) [24]. Single neuron responses varied in largely stochastic manner across trials within a recording session (week) as described before [1–3], and, importantly, varied in a qualitatively different manner across weeks (Fig. 1b). We quantified this response variation across weeks in terms of the "similarity", defined as the correlation coefficient between trial-averaged neural responses (within a week) for a given neuron between pairs of weeks and averaged across all neurons. Similarity largely decreased over time using the first week of recording as the reference (Fig. 1c). Specifically, the similarities of the fifth week were significantly lower than the similarities of the second week (Mann-Whitney U test, $p < 0.01$). In a complementary analysis, to compare how single neuronal activity varied across weeks, we computed the difference of trial-averaged activity across weeks (Supplemental Fig. 1a). The change of trial-averaged $\Delta F/F$ across weeks was significantly higher than baseline variability within week (Supplemental Fig. 1b; Mann-Whitney U test, $p < 0.01$). In conclusion, consistent with

an earlier study [24], but using complementary analyses, we showed that single neuron responses to natural movies are unstable across weeks.

Single neuron responses consist of episodic activity with distinct episode-specific rate variations across weeks

The episodic nature of cortical neuron responses to naturalistic visual stimuli (Fig. 1b) [20,28–30] provides the unique opportunity to study neural variability with respect to episodic spiking. Neurons in visual cortex are known to respond to naturalistic movies sparsely with temporally precise, but stochastic, spiking within a few well-timed "spiking episodes" [20,21,31]. Is the change in single neuron spiking across weeks (Fig. 1c, Supplemental Fig. 1) dominated by changes of spike timing or by changes in spike counts? To address this question, we inferred spikes [32] and defined spiking episodes (Fig. 2a; see Methods) based on peaks in the smoothed peristimulus time histogram (PSTH). A neuron usually possessed multiple spiking episodes and episodes from different neurons overlapped (Fig. 2b). To quantify the precision of spike timing across weeks, we computed the durations of spiking episodes (Fig. 2c). The right skewed distribution of durations showed that most of the spiking episodes had short durations (median duration: 0.66 ± 0.17 sec, 10 imaging fields). Furthermore, compared with spiking episodes defined based on PSTH within weeks, the median of durations of spiking episodes defined based on PSTH across weeks only increased by at most 2 time steps (0.2 sec) for each imaging field (Fig. 2c, median duration of spiking episodes based on PSTH within weeks: 0.59 ± 0.074 sec, 10 imaging fields). The short durations of spiking episodes and small increase compared with data within weeks indicated that episodic activity had rather precise and stable spike timing across weeks. In contrast, the spike rates during those spiking episodes changed more from trial-to-trial across weeks than that within each week (Fig. 2d). Importantly, spike rates during each episode for a given neuron varied in different patterns from week to week (Fig. 2e). The diverse spike rate variation for different spiking episodes raised the question whether spike rates during spiking episodes within the same neuron change independently across weeks. We quantified the similarity of spike rate variability during different spiking episodes as mean correlation coefficient between mean spike rate across weeks (Fig. 2e). For most neurons, the similarity of spike rate changing patterns across spiking episodes was low, although significantly higher than the chance level (Fig. 2f, Mann-Whitney U test, $p < 0.0001$). This means that different spiking episodes within the same neuron have different, but not completely independent, spike rate variations across weeks. Moreover, the similarity of spike rate changing patterns was significantly lower than that expected from i.i.d. Poisson statistics (Fig. 2f, Mann-Whitney U test, $p < 0.0001$). Consequently, assuming spike trains of all the trials were independent Poisson spike trains, the spike rates of distinct spiking episodes within the same neuron followed significantly different variations across weeks. The difference in spike rate changing patterns of spiking episodes within the same neuron suggests that the basic unit of the week-to-week fluctuation is spiking episode instead of neuron.

Latent factors resembling episodic activity with gain changes capture the across-week fluctuations

To identify the basic unit of the week-to-week fluctuation in an unbiased fashion, we switched from single-neuron analysis (Fig. 1, 2) to population analysis (Fig. 3, Supplemental Fig. 2, 3), thus including the potential impact of coordinated activity. We decomposed population activity into

latent factors that can have independent gain changes across trials. For this purpose, we chose the recently introduced tensor component analysis (TCA) [3,33], which provides an unsupervised way to identify latent factors of the recorded population activity. Specifically, we organized neuronal responses into a 3-dimensional tensor (neuron x time x trials) and decomposed this tensor into R components, each consisting of a neuron factor, a temporal factor, and a trial factor (Fig. 3a). Thus, TCA achieves a simultaneous, interlocked dimensionality reduction across neurons, time, and trials. For each component, (i) the neuron factor indicates how the component is shared across neurons, (ii) the temporal factor reflects the component's temporal profile on every trial, and (iii) the trial factor enumerates how the component's gain changes across trials. Within this framework, a neuronal response can be approximated by the reconstructed response, which is a linear combination of these TCA components (Fig. 3b). As TCA components mainly capture correlated activity across neurons or trials [33], the reconstructed responses from TCA components can be viewed as denoised responses, i.e., the responses from which independent noise has been removed.

Within this unsupervised TCA dimensionality reduction method (Fig. 3c), the pronounced peaks in the temporal factors (Fig. 3c, center) revealed shared episodic activity across neurons (Fig. 2a, b). Importantly, the distribution of the temporal factors across all 40 TCA components (Fig. 3c, center) revealed the scattering of the episodic activity across the duration of a trial (Fig. 2b). For a given TCA component (of the chosen R = 40 components), the neurons with a high neuron factor value (Fig. 3c, left) had episodic activity timed near the peak in the temporal factor (Fig. 3c, center). Any given neuron tended to display high neuron factor values in multiple components (Fig. 3c, left), thus reflecting the occurrence of multiple activity episodes for any one neuron (Fig 2a). The co-activation of neurons within a given component (i.e., multiple neurons with a high neuron factor; Fig. 3c, left) revealed the temporal overlap between episodic activity from different neurons (Fig. 2b). Further, the diverse variation of the trial factor values (Fig. 3c, right) reflected the diverse gain variability of episodic activity (Fig. 2d), even for any given neuron (Fig. 2e, f).

In summary, the TCA dimensionality reduction confirmed in an unsupervised manner the episodic activity of single neurons (Fig. 2a), the temporal overlap of episodic activity from different neurons (Fig 2b), and the diversity of week-to-week fluctuations of episodic activity within a given neuron (Fig. 2d, e, f). In conclusion, the results from the TCA analysis (Fig. 3c) support the hypothesis that cortical coordination resides at the level of episodic activity, rather than at the level of neurons, as is commonly assumed [34].

Visual inspection of the trial factors across weeks indicated vastly diverse dynamics across weeks for different components. To illustrate this diversity of dynamics, we sorted the components by their trial factors using K-means clustering, choosing 5 or 6 clusters (Fig. 3c, Supplemental Fig. 3). Within each thus determined cluster, we further ordered the components by the time to peak in their temporal factors. This reorganization of the TCA analysis display revealed two important insights. First, trial factors changed in a distinctly different manner across weeks for different clusters of components. For instance, while the trial factors for the first cluster of components were largely homogeneous across weeks, the trial factors for the second cluster largely faded away after the second week. Of functional significance, with such vanishing trial factors, the

second cluster of components would contribute little to a stimulus representation in week 4 and beyond. We observed such diverse dynamics of trial factors across weeks for all imaging fields studied (Fig. 3d, e, Supplemental Fig. 3). Second, within each cluster of components, the pronounced peaks in the temporal factors were largely evenly distributed across the duration of the trial. Assuming that the peaks in the temporal factors (or equivalently the spiking episodes; see Fig. 2) contribute to cortical stimulus representation, the even distribution of these peaks suggests that every moment in the movie was evenly represented, however by different groups of neurons at different weeks. In conclusion, the diverse dynamics of trial factors across weeks for different components indicates a fluid long-term stimulus representation in visual cortex. Importantly, the fluid stimulus representation was structured at the level of episodic activity rather than the neuron.

Stable manifolds exist in unstable population activity

As expected from the interconnected nature of cortical circuits [35], we observed population-wide correlated neural fluctuations summarized by TCA components in the previous section (Fig. 3). Does a stable representation emerge from unstable population responses? To answer this question, we searched for a stable neural manifold using dimensionality reduction.

We mapped the high-dimensional denoised neuronal population responses (reconstructed responses; Fig. 3b) of episodic activity onto a low-dimensional space (manifold) and investigated the stability of the activity on this manifold (Fig. 4). For N recorded neurons, the denoised instantaneous population response $\Delta F/F$ is a point in an N-dimensional state space. In an attempt to preserve the manifold topology of neuronal population responses (Fig. 2a, b), we chose a mapping such that nearby points in the high-dimensional state space would also be adjacent in the resulting low-dimensional space. Since the structure of the presumed intrinsic manifold was not known a priori, we adopted the unsupervised algorithmic approach, Isomap, for the mapping (see Methods; [36]).

For visualization purposes, we plotted the mapped population responses in the first 3 Isomap dimensions (i.e., 3 eigenvectors with largest eigenvalues of the geodesic distance matrix; Fig. 4a). Each dot is a nonlinear projection of the instantaneous population activity into this three-dimensional space. Interestingly, most of the dots resided on a ring-shaped low-dimensional manifold, forming well-aligned trajectories of neural activity across trials (Fig. 4a, Supplemental Fig. 4a). To quantify the stability of these trajectories across trials, we projected all trajectories against a given Isomap dimension and compared projected trajectories across all trials (Fig. 4b). From the visual inspection of the projected trajectories in the first three Isomap dimensions we obtained a sense of the stability of these trajectories across trials and sessions. For further quantification, we used the average correlation coefficient of these projected trajectories from all pairs of projected trajectories as a measure of stability across trials (Fig. 4c). Stability was high for the first few Isomap dimensions, but beyond those decreased with increasing Isomap dimension.

In conclusion, this unsupervised analysis showed that stable low-dimensional latent variables exist in population activity consisting of unstable single neuronal responses (Fig. 1) that are sparse and temporally structured into episodic activity (Fig. 2, 3). This finding is likely to be of

functional significance. Even though the high-dimensional population vector contains considerable variability, there exists a stable low-dimensional subspace for potentially stable representation of visual stimuli. The discovery of a stable manifold set the stage for stable stimulus representation.

The manifold mediates a stable representation of the time within the movie clip.

To extract the stimulus representation potentially encoded in the stable neural manifold, we applied spline parameterization for unsupervised decoding (SPUD) [37] to population activity embedded in the first few Isomap dimensions. Here we only showed results for the first two Isomap dimensions for visualization, but the following results also hold for up to the first five dimensions (Supplemental Fig. 5a). The decoding process consisted of the following steps (Fig. 5a). First, we randomly split the 2-dimensional neural manifold into a training set (80%) and a testing set (20%). Second, we fit a 1-dimensional spline to the training set, and then assigned coordinates to the fitted spline. Third, we assigned each dot in the testing set a value according to the coordinate of its nearest point on the spline. Last, we circularly shifted or flipped the coordinates on the spline such that we achieved the best decoding performance (circular least mean squared error) for time in the movie of the testing dataset. We did this because when we assigned coordinates to the spline, the origin and direction of the coordinates were arbitrarily determined. To match the assigned coordinates with the actual time, we needed to determine the origin and direction of coordinates using the test set. The decoded time α closely traced the actual time t in the movie for population activity across weeks (Fig. 5b). We summarized the decoding error (circular absolute difference between t and α) from all the recorded imaging fields (Fig. 5c). As a comparison, the decoding error of SPUD was significantly lower than the decoding error from that of linear decoders (Supplemental Fig. 5b, Mann-Whitney U test, $p < 0.01$). In general, the decoding performance improved with an increasing number of recorded neurons in the imaging field (Fig. 5c).

One of the key reasons for the high decoding performance of population activity as analyzed by SPUD resides in the isometric representation [37]. Time in the movie was evenly represented along the fitted spline direction. In other words, equal amounts of population activity variations along the spline direction contributed to equal amounts of change across time in the trial. This isometric representation was related to the evenly distributed episodic activity across time in the trial, as shown by TCA components (Fig. 3c). If we removed episodic activity during a certain time window in the trial from the population activity, then the corresponding section in the manifold ring would collapse into the hyperplane perpendicular to the spline direction (Supplemental Fig 6).

Due to the high trial-to-trial variability in population activity, the ring-shaped neural manifold had many outlier dots. The outlier dots in the center of the ring corresponded to low amplitude of population activity, while outlier dots on outside of the ring corresponded to high amplitude of population activity (Supplemental Fig. 4b). The decoder failed at a few outlier dots. However, most of the neural variability seemed to be perpendicular to the direction of the fitted spline, thus, harmless to decoding. This observation gave us a hint about the mechanism that maintains stable neural correlates in the face of dynamical population activity.

Both week-to-week fluctuation and trial-to-trial variation within the week is restricted to non-coding directions.

In order to quantify to what extent neural variability influences the stimulus encoding, we calculated the variance of instantaneous population activity on the manifold along the direction parallel or perpendicular to the fitted spline. Specifically, we computed the parallel and perpendicular component of the instantaneous population activity variance employing the following steps. First, we reconstructed $\Delta F/F$population activity based on 40 TCA components (Fig. 6a). Second, we used Isomap to project population activity of all trials into a two-dimensional space. Third, as illustrated in Fig. 5, we separated the projected instantaneous population activity into training set and test set. Fourth, we calculated the fitted spline to the training set. Fifth, we computed the coordinates on the spline based on the test set data (Fig. 6b left panel). Sixth, for each time point in the movie, we calculated the variance of instantaneous population activity in the test set along the direction parallel or perpendicular to the spline (Fig. 6b right panel). Finally, we summarized the variance for all the time points in the movie (Fig. 6c). Variance of population activity along the spline direction was significantly smaller than that perpendicular to the spline direction. This observation held for 9 out of 10 imaging fields (Fig. 6d). In this computational framework, the spline direction signifies the stimulus coding direction. In conclusion, the comparatively small contribution of neural variability to stimulus encoding direction directly explains why the high neural variability we observed in spiking episodes (Fig. 2) did not harm the decoding performance of SPUD (Fig. 5).

The neural variability we measured here consisted of two portions: week-to-week fluctuations and trial-to-trial variability within each week. Are they both restricted to the non-coding direction? To answer this question, we quantified week-to-week variability and trial-to-trial variability within each week separately. For week-to-week fluctuations, first, we calculated the trial-averaged projected population activity in the two-dimensional space for each week (Fig. 6e). Second, we calculated the variance of those trial-averaged instantaneous population activity across weeks along the direction parallel or perpendicular to the spline. Finally, we summarized the variance for all the time points in the movie (Fig. 6f, left panel). The significantly larger week-to-week variance along the direction perpendicular to the spline compared with that of parallel direction suggested that the week-to-week fluctuation was also constrained to the non-coding direction. For trial-to-trial variability within each week, first, we calculated the variance of single-trial population activity for each week separately. Second, we summarized the variance for all the weeks and all the time points in the movie (Fig. 6f, right panel). The trial-to-trial variability within each week was larger along the direction perpendicular to the spline compared with that of parallel direction. Furthermore, the same observation held for most imaging fields (Fig. 6g). In conclusion, both week-to-week fluctuations and trial-to-trial variability within each week were restricted to the non-coding direction.

The precisely timed episodic activity constrains neural variability to non-coding directions

How is neural variability largely constrained to the direction perpendicular to stimulus coding direction? Is it caused by the reproducible timing of episodic activity, by the coordination between different episodes, or by the combination thereof? To answer these questions, we applied the previous analyses to shuffled reconstructed $\Delta F/F$ population activity.

First, we checked whether the neural manifold was an artifact of the method by applying Isomap to shuffled data. To remove both the reproducible timing of episodic activity and the coordination of episodic activity across neurons in the shuffled data, we circularly time shifted reconstructed $\Delta F/F$ responses by a random amount for every trial of each neuron independently (Fig. 7a). In other words, only the temporal statistics of $\Delta F/F$ responses were kept. As expected, neural trajectories from different trials were not aligned (Fig. 7b). However, trajectories were continuous instead of being a noisy point cloud. Such continuous trajectories arise from the smooth nature of shuffled reconstructed $\Delta F/F$ responses. This sanity check showed that the ring structure of the neural manifold (Fig. 6b) arose from the timing and coordination of the population activity and was not an artifact of the method.

Second, we checked whether the reproducible timing of episodic activity was sufficient to constrain the neural variability by applying Isomap to shuffled data with preserved trial structure. To merely remove the coordination between different episodes, but to maintain the amplitude of the covariance of neural activity, we chose to shuffle TCA factors instead of shuffling reconstructed population activity. In contrast, shuffling reconstructed population activity would decrease the covariance between neural activity across neurons, in addition to removing the coordination between episodic activity. For each TCA component, we randomly shuffled the neuron order in the neuron factor, and we circularly shifted the temporal factor and the trial factor by a random amount (Fig. 7c). Thus, by shuffling the factors for each component independently, we removed all the significant coordination between episodic activity. As expected, the removal of coordination between episodic activity resulted in a new manifold and a new spline (Fig. 7d). However, the variability of reconstructed population activity (based on shuffled TCA factors) continued to be largely constrained to the direction perpendicular to the spine (Fig. 7d). The smaller variability of population activity parallel to the spline is visible in the separation of dots of different color, where color indicates the time in the trial of the instantaneous population activity (Fig. 7d). Indeed, the quantification of variance showed that amplitude of neural variability along the spline was significantly smaller than that perpendicular to the spline (Fig. 7e). Moreover, the significant difference between variance along the direction parallel and perpendicular to the spline held for shuffled data with preserved trial structure from all the imaging fields (Fig. 7f). In conclusion, the fact that episodic activity is precise in time across trials (Fig. 2c) alone is sufficient for constraining neural variability to the direction perpendicular to the stimulus encoding direction. In contrast, the coordination among episodic activity plays no role in this constraint.

However, coordination between episodic activity is essential for uniquely representing time points in the trial. The neural manifold of shuffled data with preserved trial structure had a collapsed ring structure (Fig. 7d, Supplemental Fig. 7a) in contrast to the clear ring structure from original data (Fig. 5a, Fig. 6b). The collapsed ring structure would lead to ambiguous decoding due to the overlap between instantaneous population activity from different time points in the trial (Supplemental Fig. 7b). We quantified the shape of the neural manifold for original and shuffled data with preserved trial structure by calculating the distance from each dot representing instantaneous population activity to the center of the manifold (see Methods). For 8 out of 10 imaging fields, the neural manifold of shuffled data with preserved trial structure had a more

collapsed ring structure than the manifold of original data, as shown by a significantly smaller radius (Fig. 7g).

In summary, both the nature of precise episodic activity and the coordination between different activity episodes contributes to encode time in the natural movie. However, episodic activity reproducible in time alone is sufficient for restricting neural variability to non-coding directions.

Discussion

We showed that single neuronal responses to the natural movie in V1 consisted of episodic activity with variability in gain across weeks. Importantly, we found a stable low-dimensional subspace inside the highly variable high-dimensional neural space. Time in the movie was represented on a one-dimensional ring manifold isomorphically, where equivalent changes on the ring indicated equivalent changes in time. Moreover, we found that the limited influence of neural variability and week-to-week fluctuations on the stable representation of the natural movie was mediated by the fact that most of neural variability was constrained in the non-coding direction, augmenting the previous literature on population coding and neural variability [18,38–40]. Furthermore, we found that stable episodic activity was sufficient for restricting neural variability to non-coding directions independent of coordination between episodic activity.

To study the neural representation in V1, it is common practice in the field to measure tuning curves (trial-averaged single neuron activity) with respect to external variables [41–43] or decode external variables from neural activity with supervised methods, such as linear decoders [11,44,45]. In contrast, recent work introduced unsupervised methods in revealing the internal representation using neural data alone without reference to external variables [37,46]. Here, we identified an internal representation of time in the natural movie by parameterizing the neural manifold, without using any external information or prior assumptions.

There are several advantages in the dissociation of internal and external variables. First, such dissociation avoids the biases introduced by the chosen external variables. One caveat of interpreting the neural activity through the lens of the chosen external variable is that the encoded variable might be different but correlated with the chosen external variable. Thus, non-trivial tuning curves or supervised decoding results do not necessarily reveal the actual neural representation. Second, dissociation of internal and external variables permits discovering representation of cognitive variables. It is possible that the internal variable represents the animal's inference about an external variable. For example, as hypothesized by the sampling-based neural variability theory [47,48], neural variability in V1 might represent the perceptual uncertainty of certain visual features. In the future, it will be interesting to investigate whether the thickness of the ring manifold (Fig. 6) reflects the animal's perceptual uncertainty of certain scenes in the movie.

Even though population activity may never visit the same state in the high dimensional space, there exists a stable readout direction as indicated by the fitted spline (Fig. 5). The liquid state machine (LSM) [49], a computational paradigm for recurrent neural networks, describes a similar situation. Instead of viewing neural networks as "feature detectors", LSM views the network as liquid, continuously receiving external perturbations. Although the liquid neural trajectory keeps changing across time, we can get stable readout by training a linear readout unit. Note that our

work is different from LSM in the readout method, as we obtained stable readout in an unsupervised manner. LSM suggests that trial-to-trial variability reflects an accumulation of information instead of noise, as recurrent network activity implicitly contains the previous external perturbations. This recurrent-network perspective can be instructive for our future work. In our work, we found that trial-to-trial variability is mostly constrained in the direction perpendicular to spline direction (Fig. 6). However, we did not interpret the latent variables encoded in other directions except for the spline direction. Moreover, recent works suggest that V1 encodes various behavior and state variables besides visual-related variables [3,50,51]. New experimental design with behavior or state recordings might provide a more complete picture of internal representation in V1.

The low-dimensional internal representation offered us a better reference point to understand neural dynamics than the high-dimensional population activity [52]. As a promising future direction, it would be informative to study neural dynamics on or off the manifold with perturbations [53]. One way of perturbation is to modulate the visual stimulus [54,55]. For example, on some of the trials, we propose to overlay flash dots with some frames in the natural movie [56] and observe whether the neural trajectories first deviate from the ring manifold and then flow back. Another way of perturbation is to directly control neural activity with optogenetics [57–59]. As suggested by the TCA analysis, episodic activity shared across neurons were the building blocks for the ring manifold (Fig. 3). It will be interesting to see how the optogenetically mediated changes of spiking timing or amplitude of episodic activity impact population dynamics on or off the manifold.

Methods

Animals

For imaging visual cortical responses, a Emx1-Cre (Jax Stock #005628) x ROSA-LNL-tTA (Jax Stock #011008) x TITL-GCaMP6s (Jax Stock #024104) triple transgenic mouse line (n = 9) was bred to express GCaMP6s in cortical excitatory neurons [60]. Mice ranging in age from 6 - 20 weeks of both sexes (4 males and 5 females) were implanted with a head plate and cranial window and imaged starting >2 weeks after recovery from surgical procedures and up to 10 months after window implantation. The animals were housed on a 12 hr light/dark cycle in cages of up to 5 animals before the implants, and individually after the implants. All animal procedures were approved by the Institutional Animal Care and Use Committee at University of California, Santa Barbara.

Surgical procedures

All surgeries were conducted under isoflurane anesthesia (3.5% induction, 1.5 - 2.5% maintenance). Prior to incision, the scalp was infiltrated with lidocaine (5 mg/kg, subcutaneous) for analgesia and meloxicam (1-2 mg/kg, subcutaneous) was administered preoperatively to reduce inflammation. Once anesthetized, the scalp overlying the dorsal skull was sanitized and removed. The periosteum was removed with a scalpel and the skull was abraded with a drill burr to improve adhesion of dental acrylic. A 4 mm craniotomy was made over the visual cortex (centered at 4.0 mm posterior, 2.5 mm lateral to Bregma), leaving the dura intact. A cranial window was implanted over the craniotomy and sealed first with silicon elastomer (Kwik-Sil, World Precision Instruments) then with dental acrylic (C&B-Metabond, Parkell) mixed with black ink to

reduce light transmission. The cranial windows were made of two rounded pieces of coverglass (Warner Instruments) bonded with a UV-cured optical adhesive (Norland, NOA61). The bottom coverglass (4 mm) fit tightly inside the craniotomy while the top coverglass (5 mm) was bonded to the skull using dental acrylic. A custom-designed stainless steel head plate (eMachineShop.com) was then affixed using dental acrylic. After surgery, mice were administered carprofen (5-10 mg/kg, oral) every 24 hr for 3 days to reduce inflammation. The full specifications and designs for head fixation hardware can be found on the Goard lab website (https://goard.mcdb.ucsb.edu/resources).

Note that we performed glass prism implant surgeries to two of the mice [24] to record from L2-5 neurons in V1. In this work, we only performed analysis on L2/3 neurons.

Two-photon imaging

After >2 weeks' recovery from surgery, GCaMP6s fluorescence was imaged using a Prairie Investigator 2-photon microscopy system with a resonant galvo scanning module (Bruker). Prior to 2-photon imaging, epifluorescence imaging was used to identify the visual area being imaged by aligning to areal maps measured with widefield imaging. For fluorescence excitation, we used a Ti:Sapphire laser (Mai-Tai eHP, Newport) with dispersion compensation (Deep See, Newport) tuned to λ = 920 nm. For collection, we used GaAsP photomultiplier tubes (Hamamatsu). We used a 16x/0.8 NA microscope objective (Nikon) at 1x or 2x magnification, obtaining a square field of view with width ranging from 414 to 828 µm. Laser power ranged from 40–75 mW at the sample depending on GCaMP6s expression levels. Photobleaching was minimal (<1%/min) for all laser powers used. A custom stainless-steel light blocker (https://goard.mcdb.ucsb.edu/resources) was mounted to the head plate and interlocked with a tube around the objective to prevent light from the visual stimulus monitor from reaching the PMTs. During imaging experiments, the polypropylene tube supporting the mouse was suspended from the behavior platform with high tension springs to reduce movement artifacts.

For imaging across multiple weeks, imaging fields on a given recording session were manually aligned based on visual inspection of the average map from the reference session recording, guided by stable structural landmarks such as blood vessels and neurons with high baseline fluorescence. Physical controls were used to ensure precise placement of the head plate and the visual stimulus screen relative to the animal, and data acquisition settings were kept consistent across sessions. Recordings were taken once every 7 ± 1 days for 5-7 weeks. To acclimate to head fixation and visual stimulus presentation, mice were head-fixed and presented the full series of visual stimuli for 1 to 2 full sessions prior to the start of their experimental run.

2-Photon Post-processing

Images were acquired using PrairieView acquisition software and converted into TIF files. All subsequent analyses were performed in MATLAB (Mathworks) using custom code (https://goard.mcdb.ucsb.edu/resources). First, images were corrected for X-Y movement within each session by registration to a reference image (the pixel-wise mean of all frames) using 2-dimensional cross correlation. Next, to align recordings to the reference session, we used a semi-automated method similar to prior work [61,62]. First, anchor points were automatically generated

from matching image features between average projections detected by the 'Speeded-Up Robust Features' (SURF) algorithm (Computer Vision Toolbox, Mathworks), and were manually corrected and added through visual inspection when necessary. These anchor points defined a predicted displacement vector field that would be used to map coordinates from one session to the other. For each coordinate, the predicted vector was defined by the average (weighted inversely by distance) of the vectors for all defined anchor points. This vector field was then applied to every frame of the recording to warp the coordinates of each image to the reference coordinate plane.

To identify responsive neural somata, a pixel-wise activity map was calculated using a modified kurtosis measure. Neuron cell bodies were identified using local adaptive threshold and iterative segmentation. Automatically defined ROIs were then manually checked for proper segmentation in a graphical user interface (allowing comparison to raw fluorescence and activity map images). To ensure that the response of individual neurons was not due to local neuropil contamination of somatic signals, a corrected fluorescence measure was estimated according to:

$$F_{corrected}\,(n) \;=\; F_{soma}\,(n) \;-\alpha * F_{neuropil}(n)$$

where $F_{neuropil}$ was defined as the fluorescence in the region <30 μm from the ROI border (excluding other ROIs) for frame n and α was chosen from [0, 1] to minimize the Pearson's correlation coefficient between $F_{corrected}$ and $F_{neuropil}$. The $\Delta F/F$ for each neuron was then calculated as:

$$\Delta F/F \;=\; (F_n \;-\; F_0)\;/\;F_0$$

Where F_n is the corrected fluorescence ($F_{corrected}$) for frame n and F_0 defined as the mode of the corrected fluorescence density distribution across the entire time series.

Visual stimuli

All visual stimuli were generated with a Windows PC using MATLAB and the Psychophysics toolbox [63]. Stimuli used for two-photon imaging were presented on an LCD monitor (17.5 x 13 cm, 800 x 600 pixels, 60 Hz refresh rate) positioned 5 cm from the eye at a horizontal tilt of 30 deg to the right of the midline and vertical tilt of 18 deg downward, spanning 120 deg (azimuth) by 100 deg (elevation) of visual space in the right eye.

For drifting grating visual stimulation, 12 full-contrast sine wave gratings (spatial frequency: 0.05 cycles/deg; temporal frequency: 2 Hz) were presented full-field, ranging from 0 to 330 deg in 30 deg increments. We presented 8 repeats of the drifting grating stimulus; a single repeat of stimulus consisted of all 12 grating directions presented in order for 2 sec with a 4 sec inter-stimulus interval (gray screen).

For natural movie visual stimulation, we displayed a grayscale 30 sec clip from *Touch of Evil* (Orson Wells, Universal Pictures, 1958) containing a continuous visual scene with no cuts (https://observatory.brain-map.org/visualcoding/stimulus/natural_movies). The clip was contrast-normalized and presented at 30 frames per second. We presented 30 repeats of the natural movie

stimulus; each repeat started with 5 sec of gray screen, followed by the 30 sec of movie.

Spiking episodes

We first calculated deconvolved traces from $\Delta F/F$ using Suite-2p toolbox [32,20,17]. For every neuron, we binarized the deconvolved trace by thresholding at 3 standard deviation above 0 to get inferred spikes. To calculate peristimulus time histogram (PSTH) for a given neuron, we first summed the inferred spikes across trials and smoothed it using Bayesian adaptive regression splines [64]. Spiking episode in each neuron was defined in the following steps. First, we found peaks with a prominence larger than 3 in the smoothed PSTH. Second, the full width at half maximum (FWHM) of the peaks defined the duration of spiking episodes in most cases. When the FWHM of neighboring peaks overlapped, the duration was defined by the difference between the start of the first peak and end of the last peak.

Nonnegative Tensor Decomposition with missing data

We organized our data into a 3-way tensor χ ($N \times T \times K$) and let x_{ntk} represent the activity of neuron n at time t and trial k. Nonnegative tensor component analysis (TCA) decomposes χ into a sum of R rank-one tensors, where each rank-one tensor can be written as an outer product of 3 nonnegative vectors:

$$x_{ntk} \approx \sum_{r=1}^{R} w_b^r b_t^r a_k^t = \hat{x}_{ntk}$$

Nonnegative TCA with missing values were fit to minimize the squared reconstruction error: $|| M \star (\chi - \hat{\chi}) ||_F^2$ while $W \geq 0, B \geq 0, A \geq 0$

Here, $\hat{\chi}$ denotes the reconstructed data. $|| \cdot ||_F^2$ denotes the squared Frobenius norm of a tensor:

$$|| \chi ||_F^2 = \sum_{n=1}^{N} \sum_{t=1}^{T} \sum_{k=1}^{K} x_{ntk}^2$$

M denotes a masking tensor with the same shape as χ , and $\star$ denotes entrywise multiplication of two tensors. For fitting nonnegative TCA on $\Delta F/F$ data, we set $m_{ntk} = 0$ if $x_{ntk} < 0$, otherwise we set $m_{ntk} = 1$. Normalized reconstruction error is the squared reconstruction error normalized by $|| M \star \chi ||_F^2$.

Different from matrix decompositions, tensor decompositions are often unique [65]. However, when R is large or W, B, A have low rank, it could be difficult to optimize. To monitor this possibility, we calculated similarity between different TCA fitting results on the same dataset as described in [33]. We found that the similarity between fitting results is close to 1 for all the nonnegative TCA models reported in this work.

Preprocessing of $\Delta F/F$ data

$\Delta F/F$ data were normalized such that the averaged squared sum of $\Delta F/F$ traces over time equals to 1 for every neuron:

$$\sqrt{(\sum_{tk} x_{ntk}^2)/TK} = 1$$

This normalization step is crucial for ensuring TCA fitting is not biased by high firing rate neurons, since TCA is optimized to minimize the squared reconstruction error.

Choice of the number of components in TCA

We picked the number of TCA components such that they captured a significant amount of neural responses without over-fitting, checked with cross-validation as previously reported [33]. To perform cross-validation, we randomly masked out 50% of tensor entries in χ. The remaining data was training set and the masked-out data was test set. We trained nonnegative TCA with missing values to fit the training set. And then we used the trained TCA model to fit the test set. As we increase the number of components in TCA, if the normalized reconstruction error of the test set went up, the TCA model would overfit the training set. As previously reported [33], TCA is unlikely to overfit, even with up to 70 components. For this paper, we chose 40 components for TCA, given that 40 component TCA captured a significant amount of neural responses without over-fitting (Supplemental Fig. 2). Note that all the results in this paper were robust to changes in the number of TCA components (data not shown).

Isomap

The instantaneous (temporal frequency: 10 Hz) population response $\Delta F/F$ of N recorded neurons is a point in an N-dimensional state space. Each axis in this state space represents the activity of one neuron. A given trial of 35 sec duration generates a discrete sequence (temporal frequency: 10 Hz) of 350 such points. The population activity from all trials (30 trials per recording session and 6 sessions) forms a cloud of 63000 points in this N-dimensional state space. For the purpose of the unsupervised transformation of the high-dimensional point cloud to a low-dimensional space, we ignored the association of a point to a given trial and to the time within the trial. We computed the Euclidean distance between all points, irrespective of trial number and within-trial time. Based on the Euclidean distance we assigned 20 nearest neighbors to each point (choosing a higher number of nearest neighbors also works).

This step of nearest neighbor assignment is sensitive with respect to the existence of independent fluctuations of $\Delta F/F$ responses (i.e, independent noise). To discover meaningful structure in the population activity, we removed such independent noise. Rather than working with $\Delta F/F$ directly, we conducted the nearest neighbor assignment based on the "TCA-reconstructed $\Delta F/F$", from which the independent noise was removed.

By linking (edge) each point to its thus defined nearest neighbors, we translated the point cloud of population responses into a graph, i.e., a network of vertices (points) with edges (between a point and its nearest neighbor). The geodesic distance between two vertices in the graph is the distance of the shortest path connecting them. For our data set, the graph was described by the geodesic distance matrix of dimension 63000 x 63000.

Based on the pairwise geodesic distance between data points, we thus performed a transformation from the population responses in the N-dimensional state space to a space of lower dimensions. This isometric mapping method ("Isomap") was chosen to incorporate the presumed (but a priori unknown) manifold structure in the resulting transformation to a low-dimensional space. Isometric mapping preserves essential structure within the neuronal population responses. Note that the top k eigenvectors of the geodesic distance matrix represent the coordinates (Isomap dimensions) in the new k-dimensional Euclidean space.

With all 63000 data points successfully mapped into a state space of n dimensions, we recalled the assignment of each point to a given trial and to the time within the trial. This temporal sequence of data points formed the trajectory of population activity for a given trial in this low-dimensional space.

Shuffled data with circularly shifted responses across trials has much higher intrinsic dimensions than original data. Due to the curse of dimensionality [66] and the smoothness of shuffled responses, we need to define a larger neighborhood size for Isomap to reveal robust topology of the neural manifold in this case. Thus, we chose 100 nearest neighbors for Isomap for shuffled data.

Spline parameterization for unsupervised decoding (SPUD)

We used the SPUD algorithm described in [37]. We fitted the manifolds with piecewise linear curves. We chose to fit a curve $L(y)$ with 10 knots to the data points x_i embedded in the 2-dimensional spaces by Isomap. Initially, the positions of knots were determined by K-means clustering centroids of the data points. Each knot was connected to the other knot with the highest data point density in between to form the initial curve. Then, positions of the knots were iteratively optimized to minimize $(\sum_i \ ||(L(y) - x_i||)|L(y)|$, where $||(L(y) - x_i||$ is the Euclidean distance between the ith data point and the nearest point on the curve, and $|L(y)|$ is the length of the curve.

We picked a random origin on the curve and assigned coordinates from 0 to 1 to the point on the curve. The coordinate of each data point $x_i$was decoded as the coordinate of its nearest point on the curve. We shifted or flipped the coordinates of the data points to minimize the mean squared error between the decoded coordinates and the rescaled actual time in the movie (rescaled to (0,1]). The decoded time for a given data point was set to the resulting coordinate scaled up to (0, 35) seconds.

Note that the neural manifold for shuffled data often did not have a perfect ring structure (Supplemental Fig. 7a). The SPUD will fail without carefully choosing the positions of initial knots. For the purpose of force quantitative comparison between original and shuffled data (Supplemental Fig. 7b), we chose ten trial-averaged projected instantaneous population activity evenly distributed in time as the initial knots for the shuffled data analysis.

Radius of points on the manifold

We quantified the shape of the neural manifold for original and shuffled data by calculating the distance from each dot representing instantaneous population activity to the center of the manifold.

Center of the manifold was calculated as averaged coordinates across all the points. Empirically, the center was close to the origin.

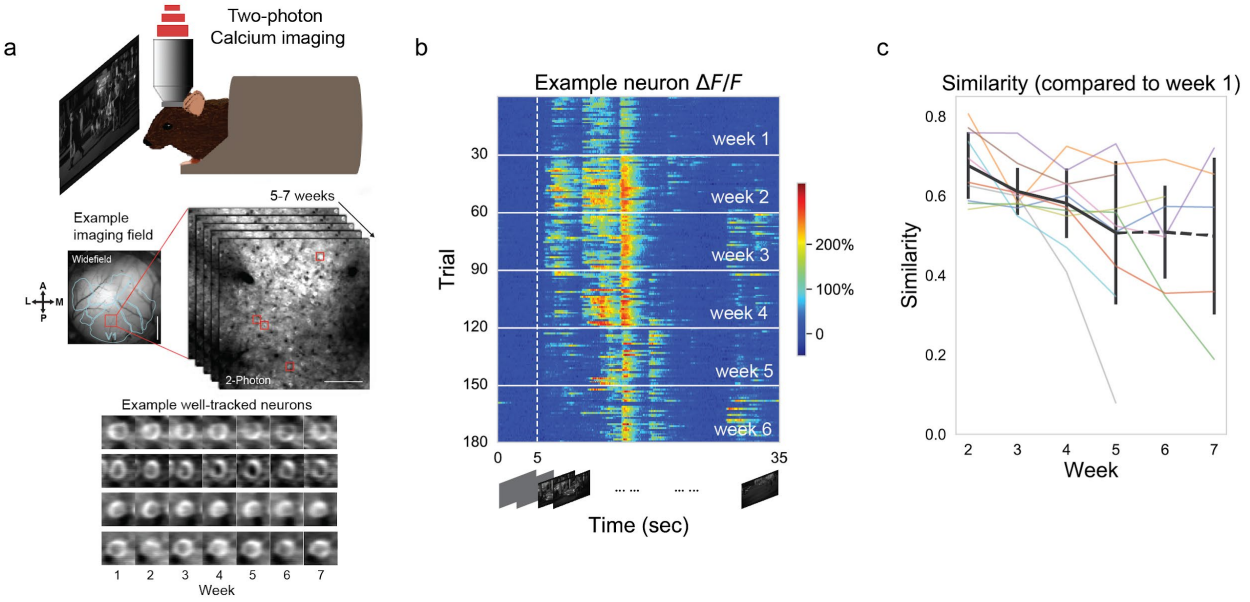

Figure 1 Single neuron responses to natural movies are unstable across weeks.

a. Experimental setup. We performed chronic calcium imaging of excitatory neurons in the primary visual cortex of awake, head-fixed mice during visual stimulation with repeated natural movies. Visual cortex (contralateral to visual stimulus delivery) is retinotopically mapped in Emx1-Cre::TITL-GCaMP6s mice. V1 fields are chosen from the region selective for the center of the presentation screen. Widefield scale bar = 1 mm; 2-photon scale bar = 100 μm. Average activity of four example well-tracked neurons across weeks are shown in the bottom panel.

b. $\Delta F/F$ responses of one example neuron during the same natural movie clip for 30 trials per experimental session for 6 weeks (movie starts at 5 sec and lasts for 30 sec duration). We recorded 1 experimental session per week.

c. Similarity (correlation coefficient between trial-averaged $\Delta F/F$) averaged over neurons during week 1 and that during other weeks are plotted for all the recorded imaging fields. Different imaging fields are denoted by different colors. The black curve with error bar denotes mean and standard deviation of similarity over imaging fields. Only a subset of imaging fields have recordings on week 6 (6 fields) and week 7 (5 fields). Specifically, the similarities of the fifth week were significantly lower than the similarities of the second week (Mann-Whitney U test, $p < 0.01$).

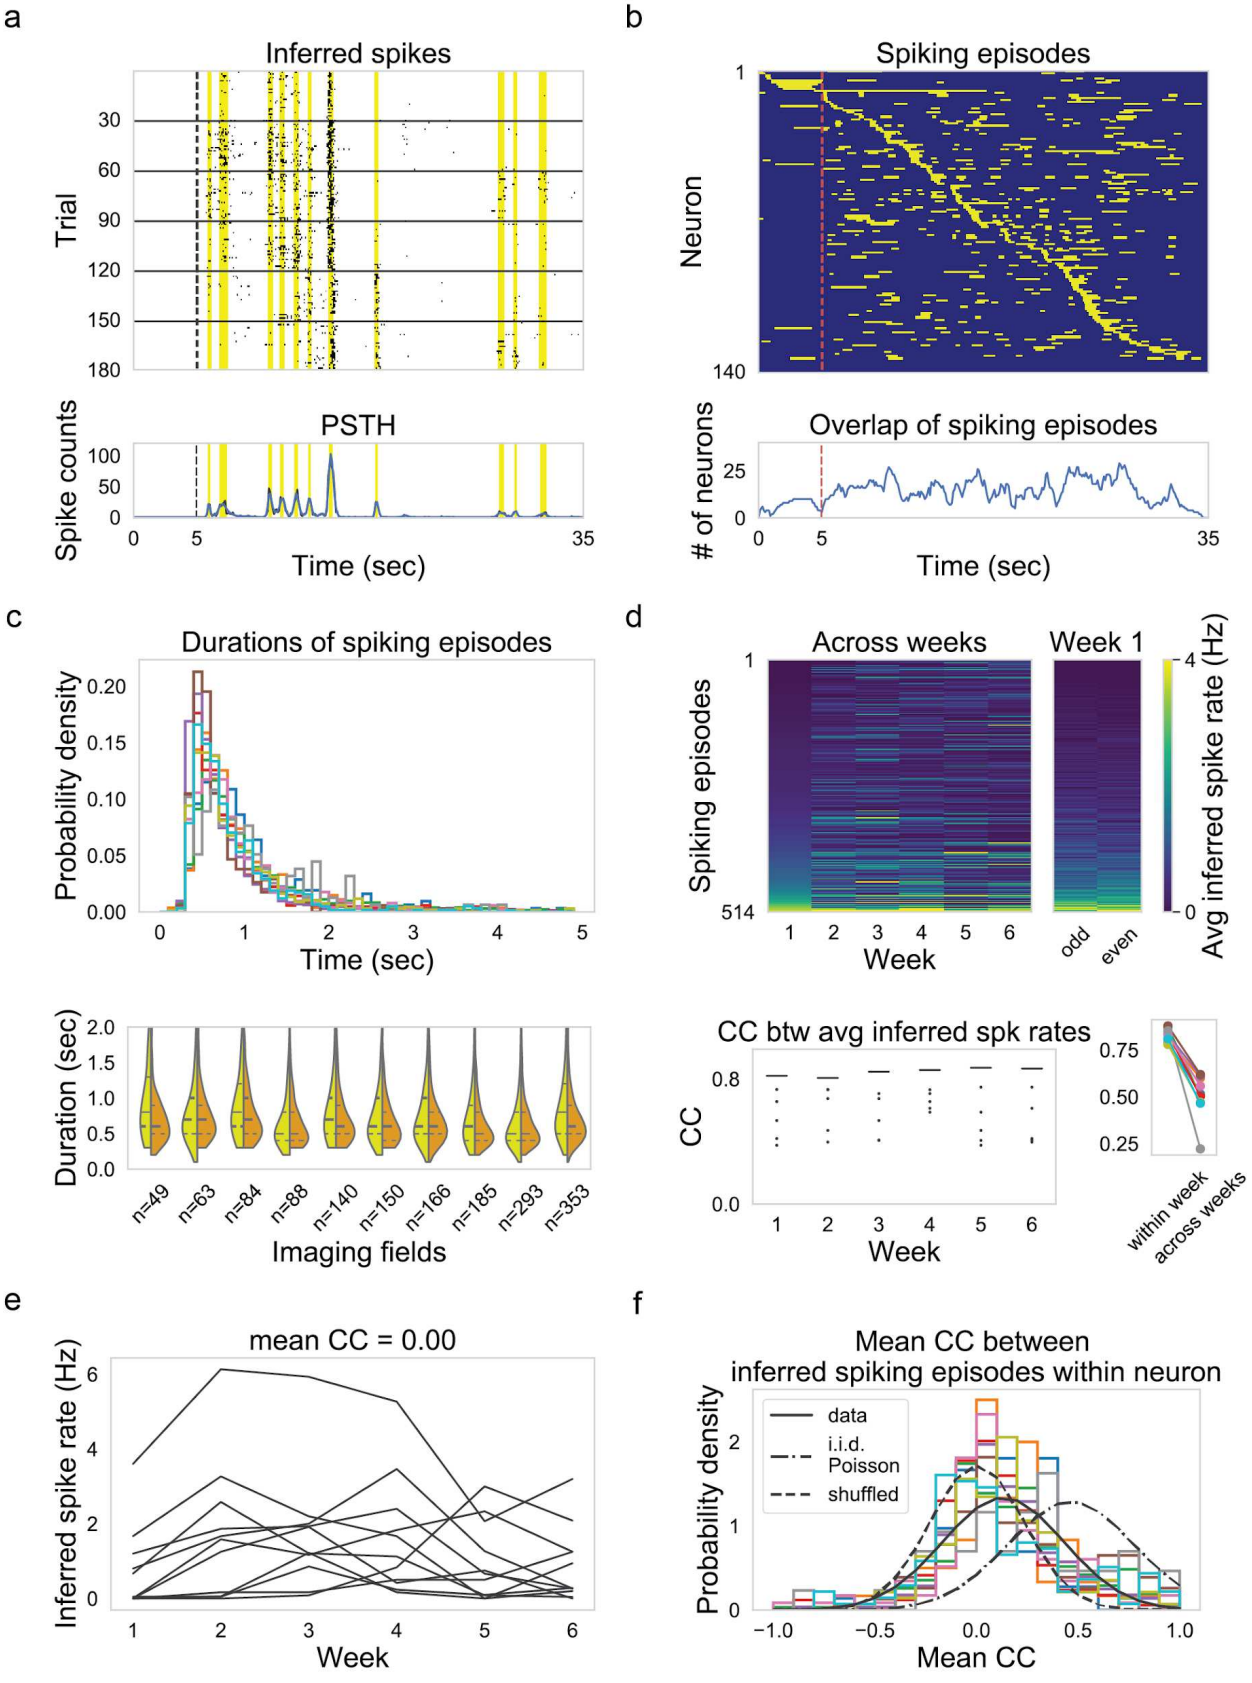
a
Inferred spikes
Trial
30
60
90
120
150
180
PSTH
Spike counts
100
50
0
0
5
35
Time (sec)
b
Spiking episodes
Neuron
1
140
Overlap of spiking episodes
# of neurons
25
0
0
5
35
Time (sec)
c
Durations of spiking episodes
Probability density
0.20
0.15
0.10
0.05
0.00
0
1
2
3
4
5
Time (sec)
Duration (sec)
2.0
1.5
1.0
0.5
0.0
n=49
n=63
n=84
n=88
n=140
n=150
n=166
n=185
n=293
n=353
Imaging fields
d
Across weeks
Week 1
Spiking episodes
1
514
1
2
3
4
5
6
Week
odd
even
4
0
Avg inferred spike rate (Hz)
CC btw avg inferred spk rates
CC
0.8
0.0
1
2
3
4
5
6
Week
0.75
0.50
0.25
within week
across weeks
e
mean CC = 0.00
Inferred spike rate (Hz)
6
4
2
0
1
2
3
4
5
6
Week
f
Mean CC between
inferred spiking episodes within neuron
Probability density
2
1
0
data
i.i.d.
Poisson
shuffled
−1.0
−0.5
0.0
0.5
1.0
Mean CC

Figure 2 Single neuron responses consist of episodic activity with distinct episode-specific rate variations across weeks.

a. Top: inferred spikes of the same neuron shown in Fig. 1b. Bottom: peristimulus time histogram (PSTH) (black) and smoothed PSTH (blue) of the same neuron. Shaded areas (yellow) denote spiking episodes for this neuron.
b. Top: spiking episodes for all the neurons in the example imaging field. Neurons are ordered by latency of their spiking episodes with the highest spiking rates. Bottom: number of neurons with overlapped spiking episodes.
c. Top: Distributions of durations of spiking episodes from all imaging fields. Different colors denote different imaging fields. Bottom: distribution of durations of spiking episodes defined from PSTH of trials across weeks (yellow) plotted against distribution of durations from PSTH of trials within weeks (orange).
d. Top: averaged spike rates over trials of all the spiking episodes in one example imaging field are plotted for different weeks and for even and odd trials in week 1. spiking episodes are ordered by their averaged spike rates during week 1. Bottom left: correlation coefficients (CC) between averaged spike rates of week pairs (dots) and even/odd trials within the week (lines) are shown for the example imaging field across weeks. Bottom right: CC within week averaged across weeks is plotted against CC across weeks averaged across all the week pairs (for all imaging fields, Mann-Whitney U test, $p < 0.005$). Different colors denote different imaging fields. Colormap maximum value is set to 4 Hz.
e. Mean spike rate during each spiking episode in the example neuron varies across weeks.
f. Histogram of mean CC between mean spike rates during spiking episodes within the same neuron. Different colors denote different imaging fields. The black solid line is a gaussian curve fitted to the distribution of mean CC from all the imaging fields (mean 0.13, s.t.d. 0.30). The black dash dotted line is a gaussian curve fitted to the distribution of mean CC between simulated independent and identically distributed Poisson spike trains with the firing rates of a randomly selected spiking episode for a given neuron (mean 0.46, s.t.d. 0.31). The black dashed line indicates the chance level, which is a gaussian curve fitted to the distribution of mean CC between spiking episodes with independently shuffled weeks (mean 0.0036, s.t.d. 0.23). Only neurons with more than one spiking episode were included in this analysis.

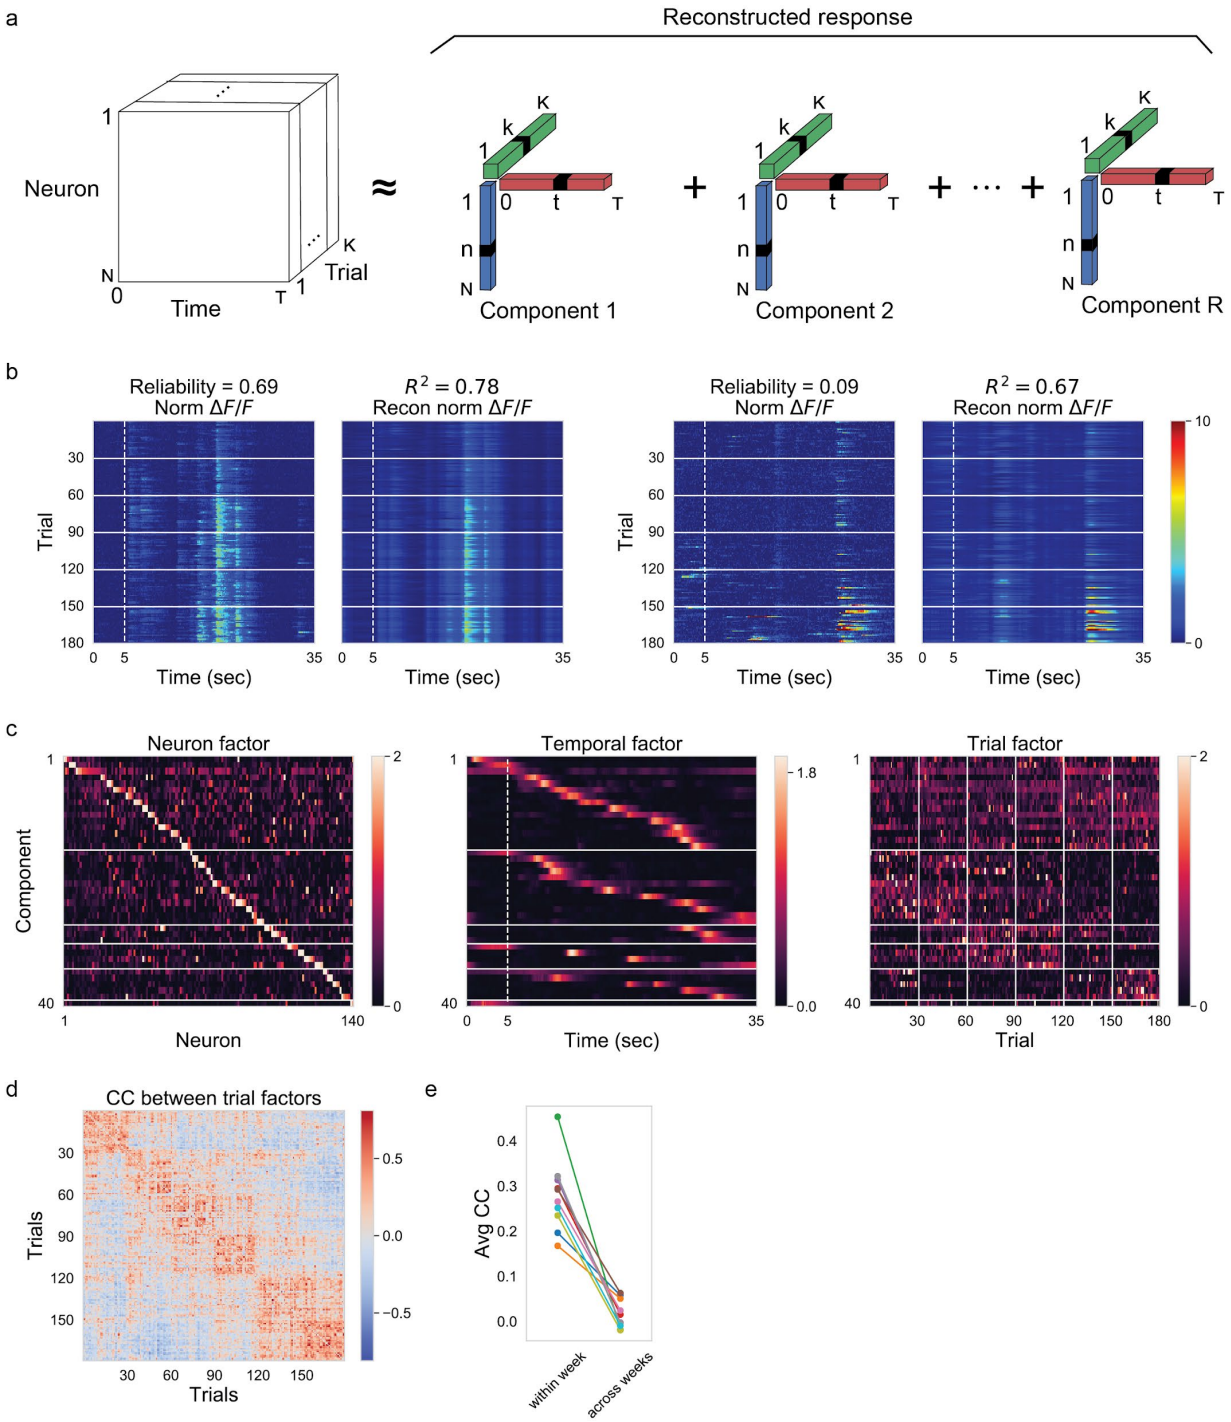

Figure 3 Latent factors resembling episodic activity with gain changes capture the across-week fluctuations

a. Schematic of Tensor Component Analysis (TCA). Neural activity ($\Delta F/F$) is organized into a third-order tensor with dimensions N x T x K. TCA approximates the data as a sum of outer products of three vectors from R components: neuron factors describe the weights of each neuron to that component, temporal factors describe the temporal dynamics of each component, and trial factors describe the modulation of the component across trials.

b. Normalized $\Delta F/F$ responses and reconstructed $\Delta F/F$ from 40 TCA components of two example neurons from the example imaging field. Reliability was defined as averaged correlation-coefficient between pairs of single-trial responses [3].
c. Neuron, temporal, and trial factors of nonnegative TCA with 40 components for the example imaging field. Colormap maximum values are set to 2 for neuron factors and trial factors. We ordered components according to the K-means clustering on their trial factors. Within each thus determined cluster, we further ordered the components by the time to peak in their temporal factors. We ordered neurons in the neuron factors by their dominant components.
d. Correlation coefficient (CC) between trial factors shown in c.
e. CC between trial factors averaged across trial pairs within week plotted against CC between trial factors averaged across trial pairs across weeks. Different color denotes different imaging fields. The week-to-week variability of trial factors was significantly larger than the corresponding trial-to-trial variability within each week (for all imaging fields, Mann-Whitney U test, $p < 0.0001$).

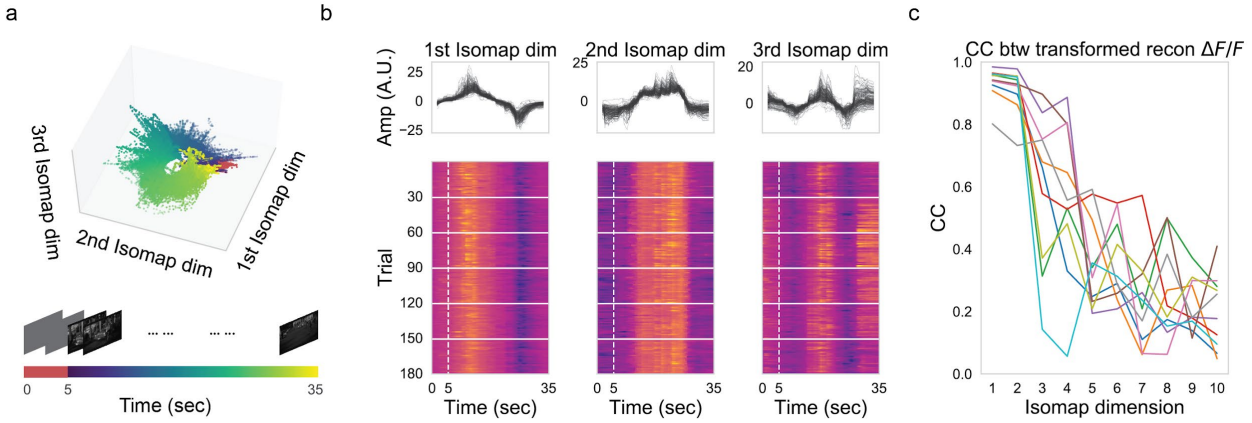

Figure 4 Stable manifolds exist in unstable population activity.

a. 3-dimensional neural trajectories extracted from reconstructed (denoised) $\Delta F/F$ populational activity across weeks from the example imaging field using Isomap. Each dot represents instantaneous population activity. Color of the dot indicates the corresponding time in the trial.

b. Neural trajectories along the first 3 Isomap dimensions (the same as shown in a) organized in trial by time matrices.

c. Correlation coefficients (CC) between transformed reconstructed $\Delta F/F$ (neural trajectories) across trials along each Isomap dimension are plotted for all 10 imaging fields. Different color denotes different imaging fields.

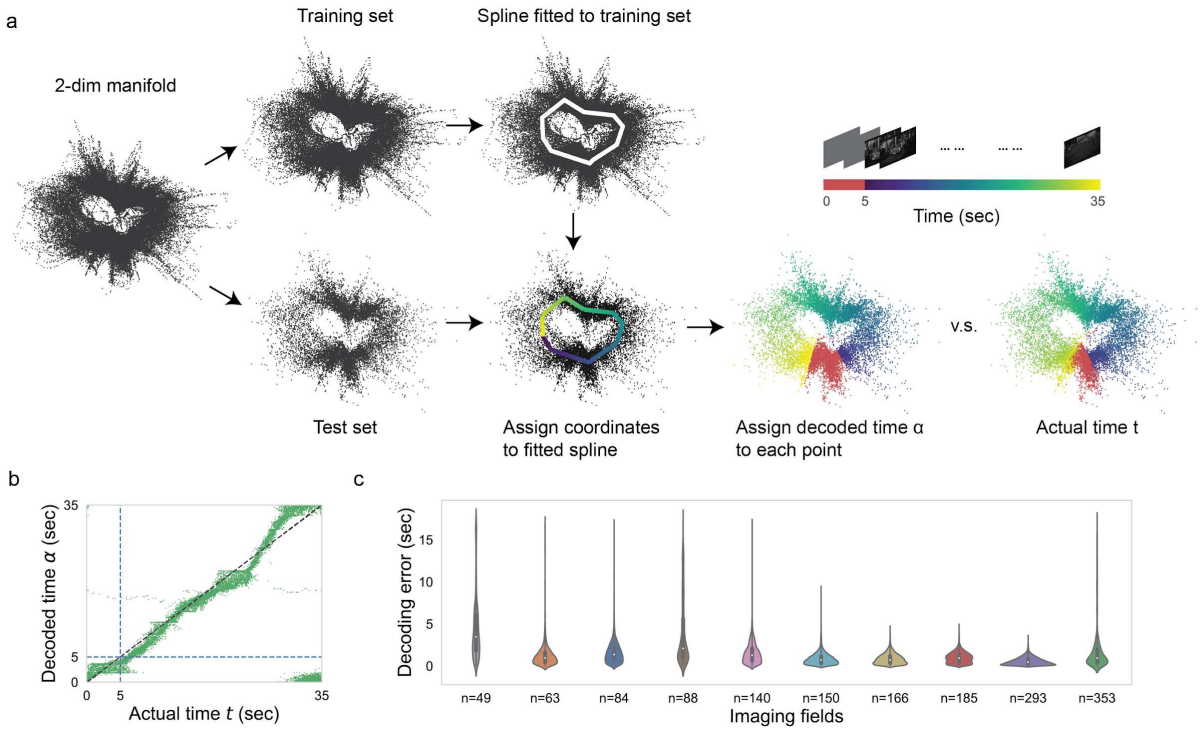

Figure 5 The manifold mediates a stable representation of the time within the movie clip.

a. Illustration of the unsupervised method with data from the example imaging field (n = 140): first, we projected reconstructed DF/F responses into the first two Isomap dimensions, each dot denotes instantaneous population activity; second, we randomly pick 80% of the instantaneous population activity as training set and rest of them as test set; third, we fitted a spline to the neural manifold of the training set and assigned coordinates with randomly picked origin to the fitted spline; finally, we shifted and flipped the coordinates on the fitted spline to match with the actual time and assigned decoded time to each point in the test set by its nearest coordinate on the spline.

b. Decoded time from the neural manifold plotted against actual time in the movie for the example imaging field (n = 140).

c. Violin plots showed decoding error (absolute circular difference between decoded time and actual time) for all the imaging fields. Imaging fields were ordered by the number of recorded neurons.

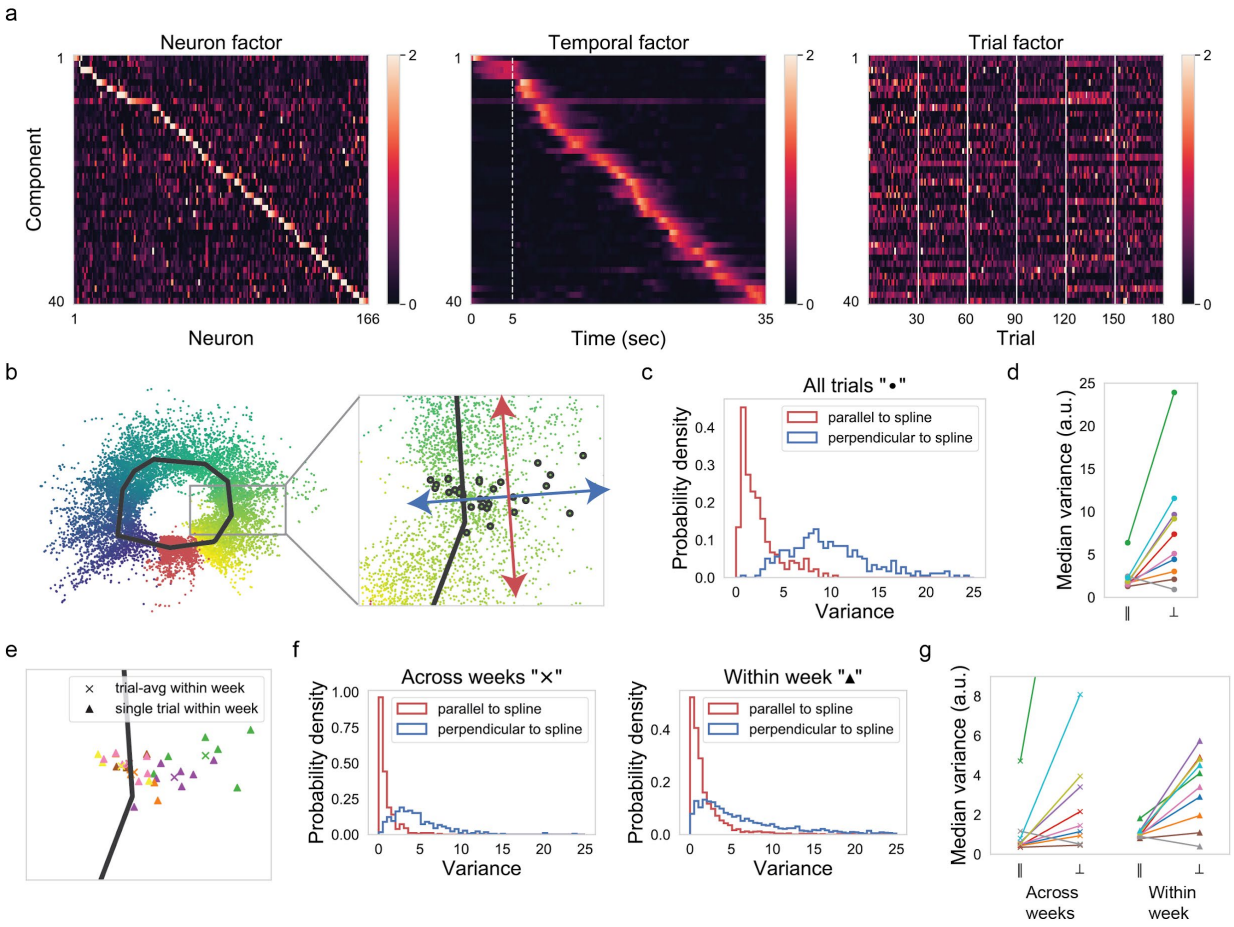

Figure 6 Both week-to-week fluctuation and trial-to-trial variation within the week is restricted to non-coding directions.

a. TCA components of one imaging field (n = 166). We ordered components by the time to peak in their temporal factors. We ordered neurons in the neuron factors by their dominant components. Colormap maximum values are set to 2 for all the factors.

b. Left: 2-dimensional neural manifold extracted from reconstructed (denoised) $\Delta F/F$ population activity (n = 166) across weeks using Isomap. Each dot represents instantaneous population activity in the test set. Color of the dot (Same colormap as Fig. 4a) indicates the corresponding time in the trial. Black line is the fitted spline to the training set. Right: Zoom-in view on the neural manifold. Instantaneous population activity corresponding to 28 s in the trial was highlighted with black shade. Blue arrow denotes the direction perpendicular to the spline, and red arrow denotes the direction parallel to the spline.

c. Histogram of the variance of population activity parallel or perpendicular to the spline for the imaging field shown in a&b.

d. Median variance of population activity parallel to the spline plotted against median variance of population activity perpendicular to the spline for all the imaging fields. Except for one imaging field (gray one), variance of population activity parallel to the spline was significantly smaller than the variance perpendicular to the spline (for all imaging fields, Mann-Whitney U test, $p < 0.0001$). The one outlier (gray line) is from an imaging field with

the least number of recorded neurons (n = 49), whose neural manifold didn't have a clear ring shape (Supplemental Fig. 5a).

e. The same zoom-in view on the neural manifold as shown in b (right). Each triangle represents instantaneous population activity within a week. Color of the triangle denotes different weeks. Each cross represents the trial-averaged instantaneous population activity within a week. Color of the cross also denotes different weeks.

f. Left: histogram of the variance of trial-averaged population activity within a week parallel or perpendicular to the spline for the imaging field. Right: histogram of the variance of single-trial population activity within a week parallel or perpendicular to the spline for the imaging field.

g. Median variance of trial-averaged population activity or single-trial population within a week parallel to the spline plotted against median variance of population activity perpendicular to the spline for all the imaging fields. Y axis is clippled at 9 for visualization. Except for one imaging field (gray one), variance of population activity parallel to the spline was significantly smaller than the variance perpendicular to the spline for both across weeks and within a week cases (for all imaging fields, Mann-Whitney U test, $p < 0.0001$). The one outlier (gray line) is from an imaging field with the least number of recorded neurons (n = 49).

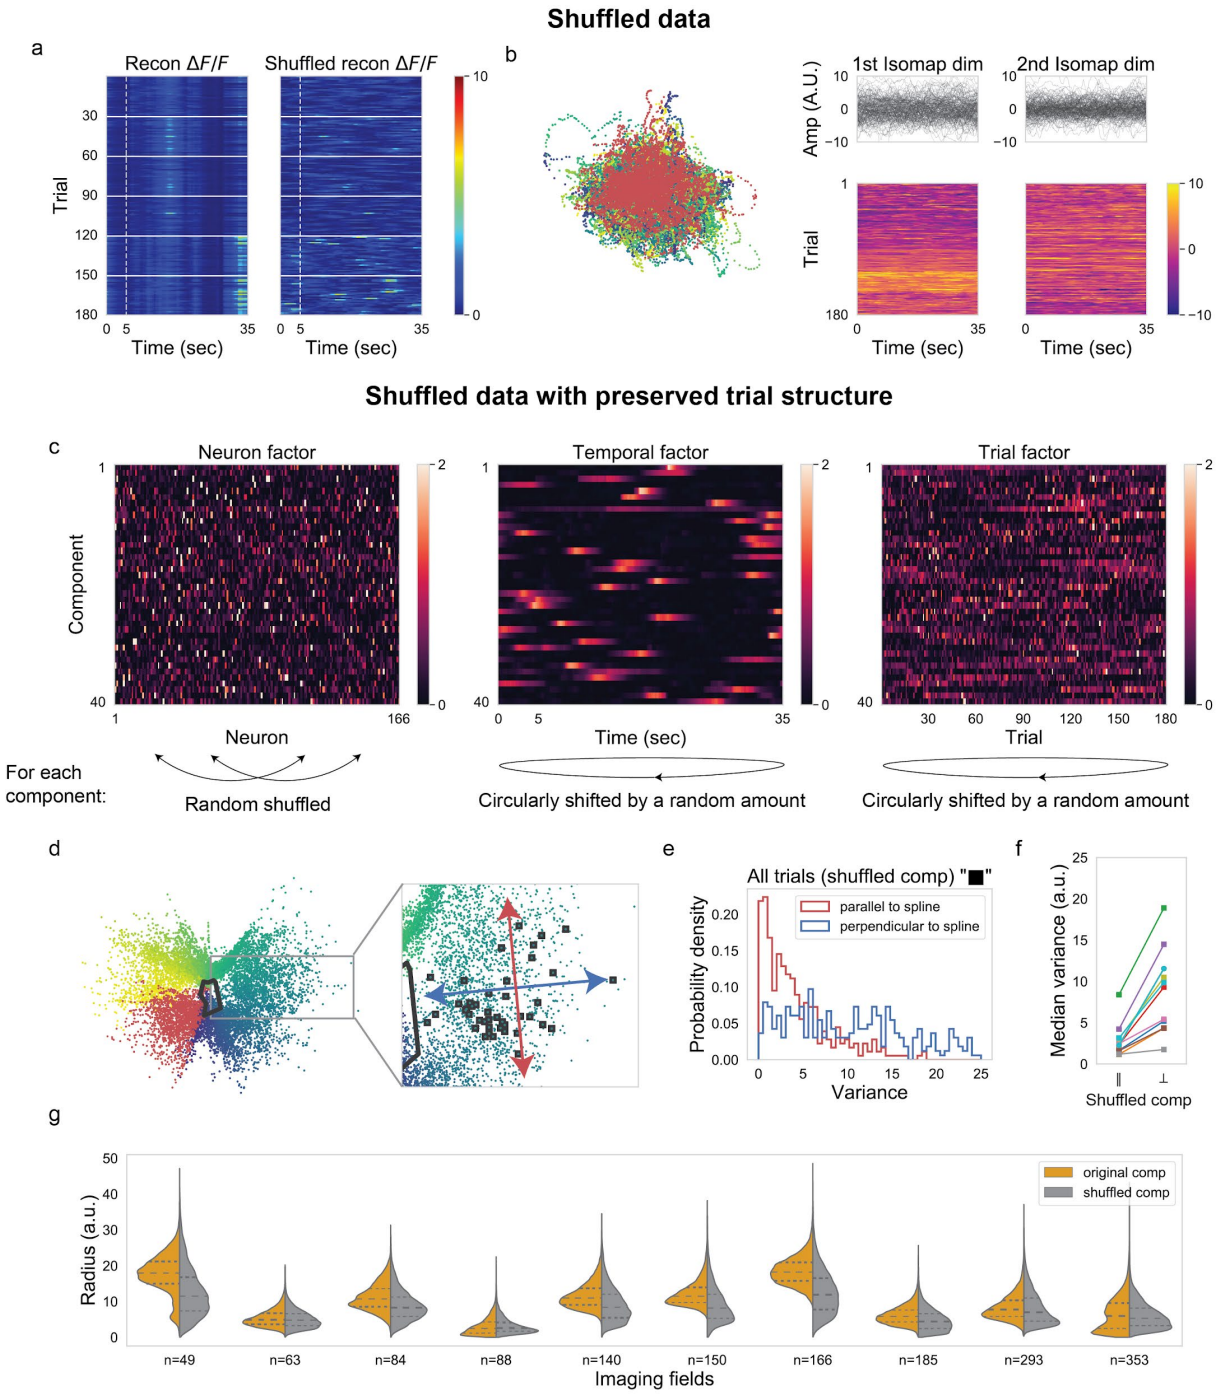

Figure 7 The precisely timed episodic activity constrains neural variability to non-coding directions.

a. Reconstructed $\Delta F/F$ responses and shuffled reconstructed $\Delta F/F$ responses of two example neurons from the imaging field (n = 166). The single neuron response was circularly shifted by a random amount independently for each trial for shuffling.

b. Left: 2-dimensional neural manifold extracted from shuffled reconstructed $\Delta F/F$ population activity (example single neuronal shuffled responses shown in a). The same

colormap was used as in Fig. 6b, left panel. Right: neural trajectories along the first 2 Isomap dimensions (the same as shown in the left panel) organized in trial by time matrices. Here we set the number of nearest neighbors of ISOMAP to be 100 (see Methods).

c. TCA components of the imaging field (n = 166) with shuffled factors. For each component, we independently shuffled neuron order in the neuron factor, circularly shifted the time factor and the trial factor by a random amount. Components with shuffled factors were ordered again in the same fashion as Fig.6a.

d. 2-dimensional neural manifold extracted from reconstructed (denoised) $\Delta F/F$ population activity (n = 166) from components with shuffled factors using Isomap. Each dot represents instantaneous population activity in the test set. Black line is the fitted spline to the training set. Instantaneous population activity corresponding to 16 s in the trial was highlighted with black shade.

e. Histogram of the variance of neural variability parallel or perpendicular to the spline for reconstructed (denoised) $\Delta F/F$ populational activity (n = 166) from components with shuffled factors (i.e., shuffled data with preserved trial structure).

f. Median variance of neural variability parallel to the spline plotted against median variance of neural variability perpendicular to the spline for all the imaging fields for reconstructed activity from components with shuffled factors. Variance of population activity parallel to the spline was significantly smaller than the variance perpendicular to the spline (Mann-Whitney U test, $p < 0.0001$) for all imaging fields.

g. Radius (distance to the center of the point cloud) distribution of points on the neural manifold from original TCA components plotted against radius distribution from TCA components with shuffled factors for all the imaging fields. Except for 3 imaging fields (n = 49, n = 63, n = 88), the radius of points on the neural manifold from original TCA components was significantly larger than the radius from TCA components with shuffled factors (Mann-Whitney U test, $p < 0.0001$).

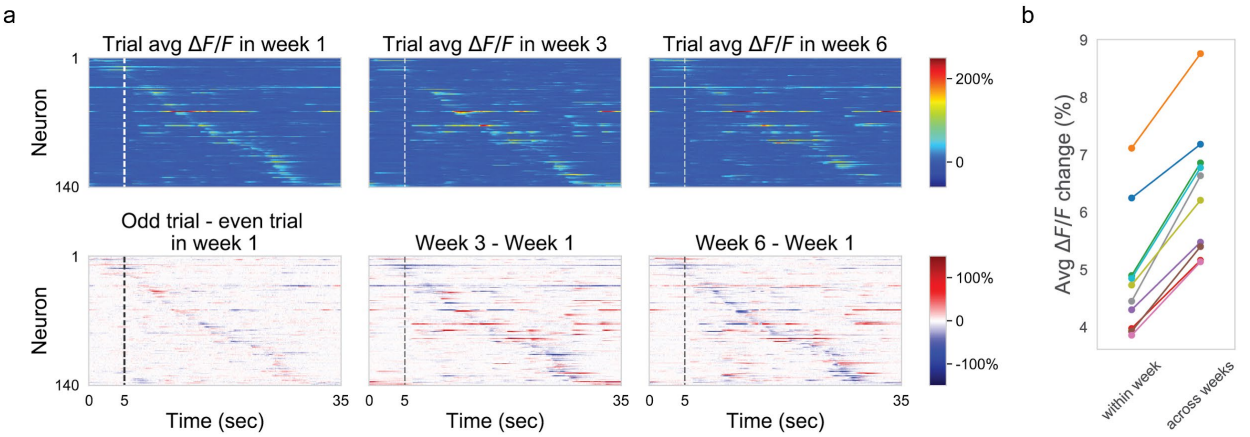

Supplemental Figure 1 Single neuron responses to natural movies are unstable across weeks.

a. Trial-averaged $\Delta F/F$ ordered by their peak time in week 1 are shown for different weeks for one example imaging field in the upper row. Difference between trial-averaged $\Delta F/F$ with the same neuron ordering are shown in the lower row.

b. $\Delta F/F$ change within week (L1 norm of difference between trial-averaged $\Delta F/F$ of even and odd trials) averaged across weeks is plotted against $\Delta F/F$ change across weeks (L1 norm of difference between trial-averaged $\Delta F/F$ of different weeks) averaged across all the week pairs (Mann-Whitney U test, $p < 0.01$). Different colors denote different imaging fields.

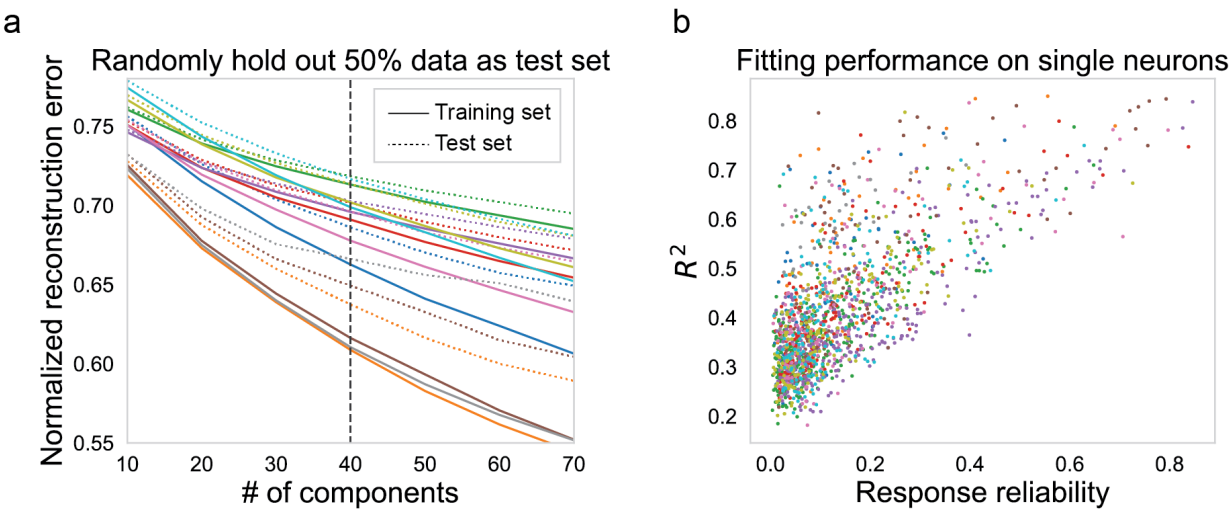

Supplemental Figure 2 Fitting performance of TCA.

a. Cross validation of TCA on one example dataset (545 neurons x 350 frames x 30 trials). Normalized reconstruction error plotted against the number of components of TCA for training set and test set for 10 imaging fields. Color denotes different imaging fields. Dashed line denotes the TCA model with 40 components.

b. Fitting performance R^2 plotted against response reliability for neurons pooled from 10 imaging fields. Each dot represents one neuron. Color denotes different imaging fields.

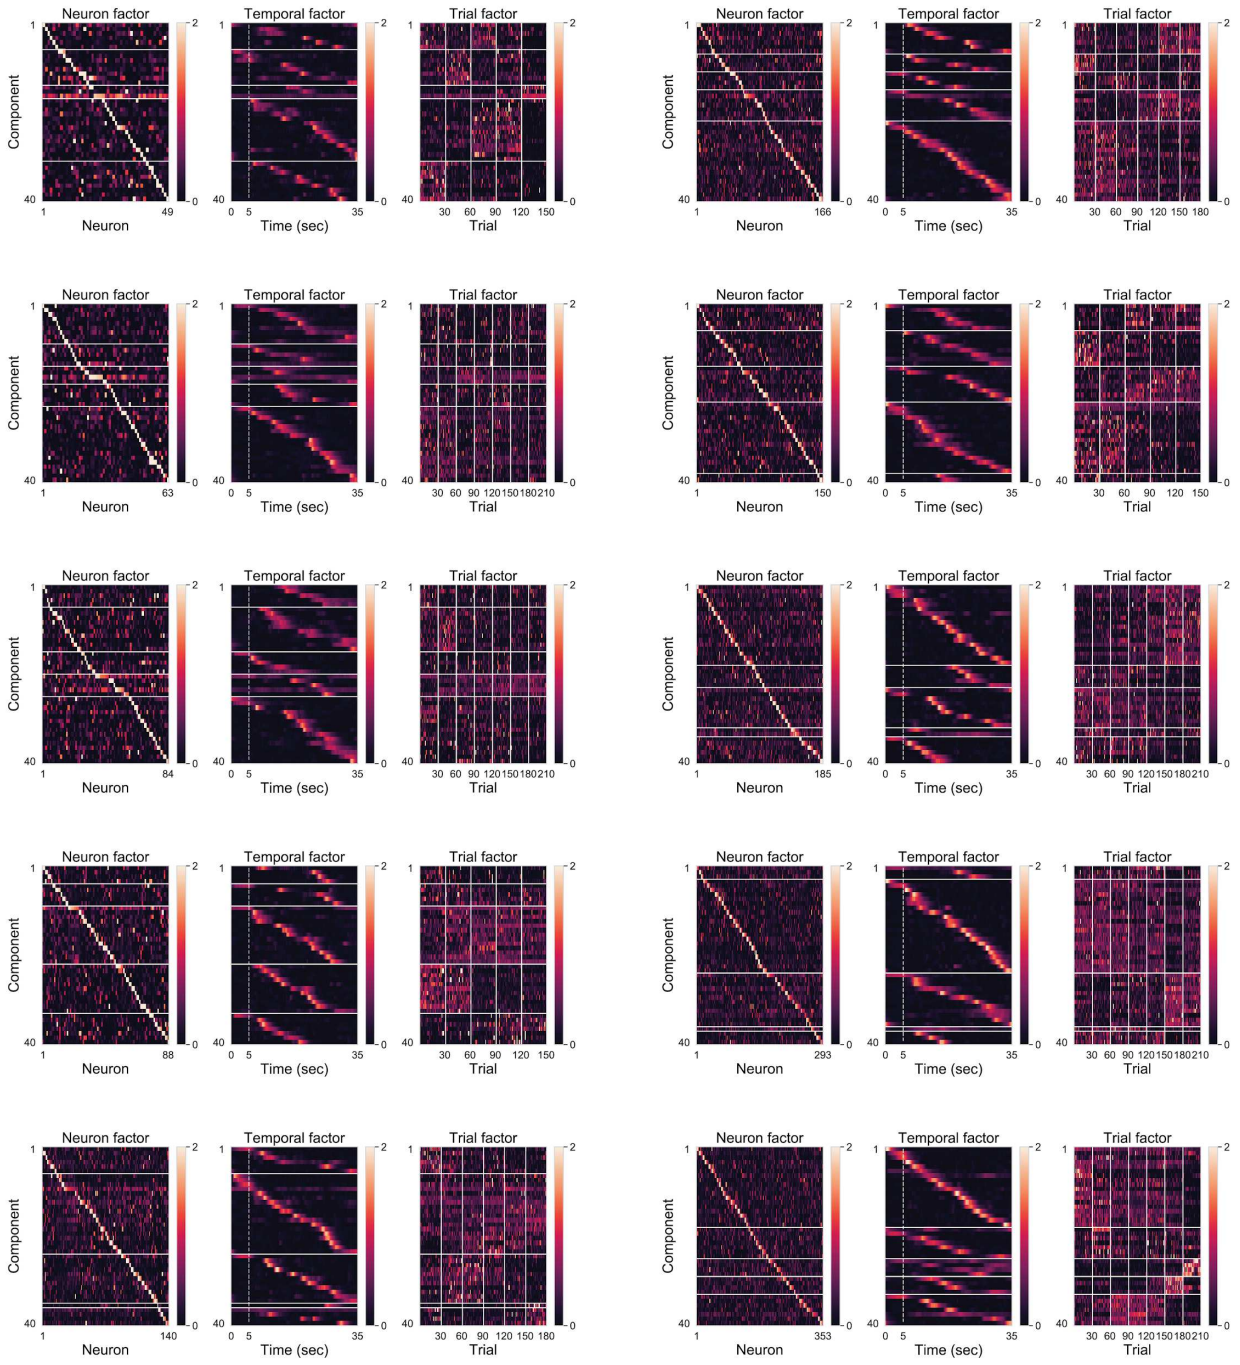

Supplemental Figure 3 Latent factors resembling episodic activity with gain changes capture the across-week variability.
Neuron, temporal, and trial factors of nonnegative TCA with 40 components for 10 imaging fields. We ordered neurons in the neuron factors by their dominant components. Colormap maximum values are set to 2. We ordered components according to the K-means clustering on their trial factors. Within each thus determined cluster, we further ordered the components by the time to peak in their temporal factors.

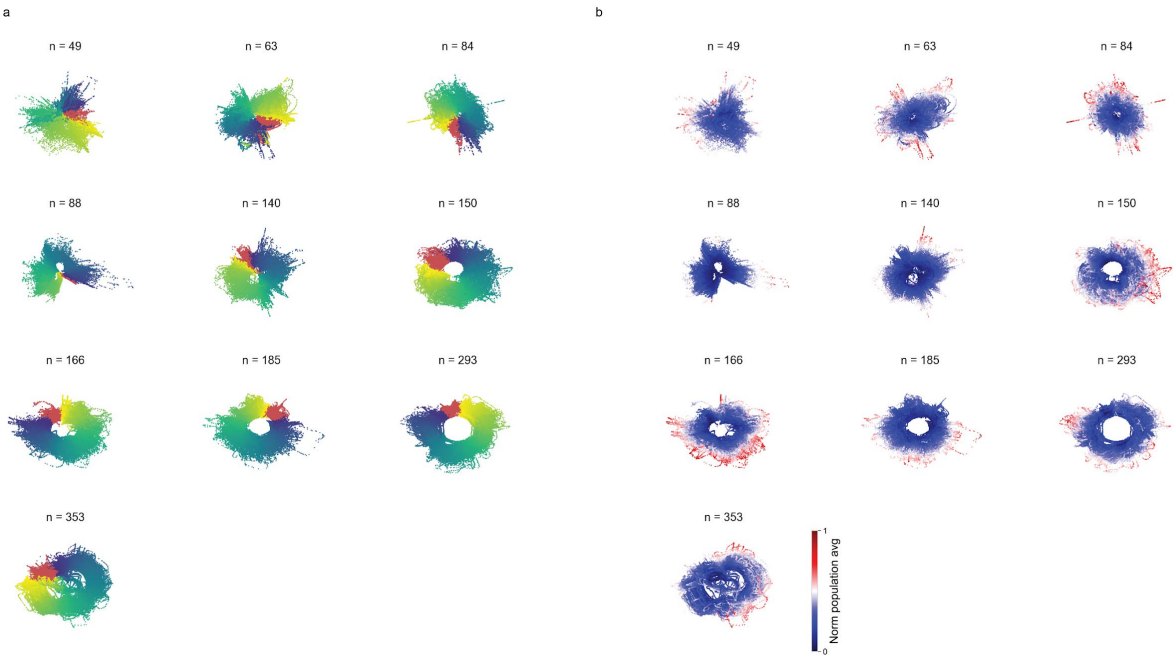

Supplemental Figure 4 Neural manifold for natural movie population responses in V1.

a. 2-dimensional neural manifold extracted from reconstructed (denoised) $\Delta F/F$ population activity across weeks using Isomap for 10 imaging fields. Each dot represents instantaneous population activity in the test set. Color of the dot (Same colormap as Fig. 4a) indicates the corresponding time in the trial.

b. Same neural manifold as in a. Color of the dot indicates the corresponding normalized reconstructed (denoised) $\Delta F/F$activity averaged over neurons in the imaging field.

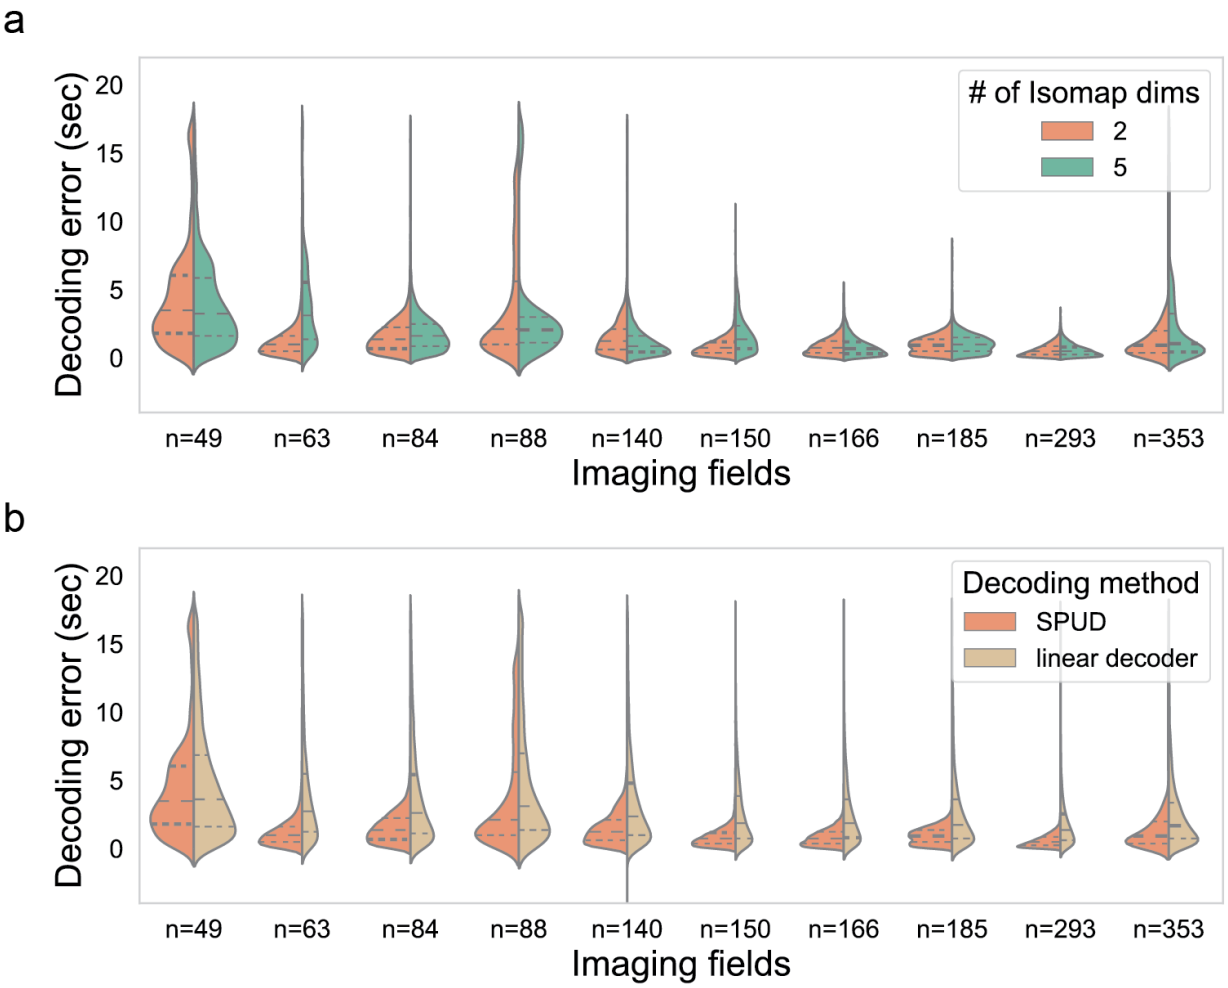

Supplemental Figure 5 Decoding performance of SPUD.

a. Decoding error (absolute circular difference between decoded time and actual time) of SPUD on 2-dimensional neural manifold plotted against decoding error of SPUD on 5-dimensional neural manifold for all the imaging fields. Imaging fields were ordered by the number of recorded neurons.

b. Decoding error (absolute circular difference between decoded time and actual time) of SPUD on 2-dimensional neural manifold plotted against decoding error of linear decoder using reconstructed (denoised) $\Delta F/F$ population activity for all the imaging fields. Imaging fields were ordered by the number of recorded neurons. Decoding error of SPUD was significantly smaller than decoding error of linear decoder for all the imaging fields (Mann-Whitney U test, $p < 0.01$)

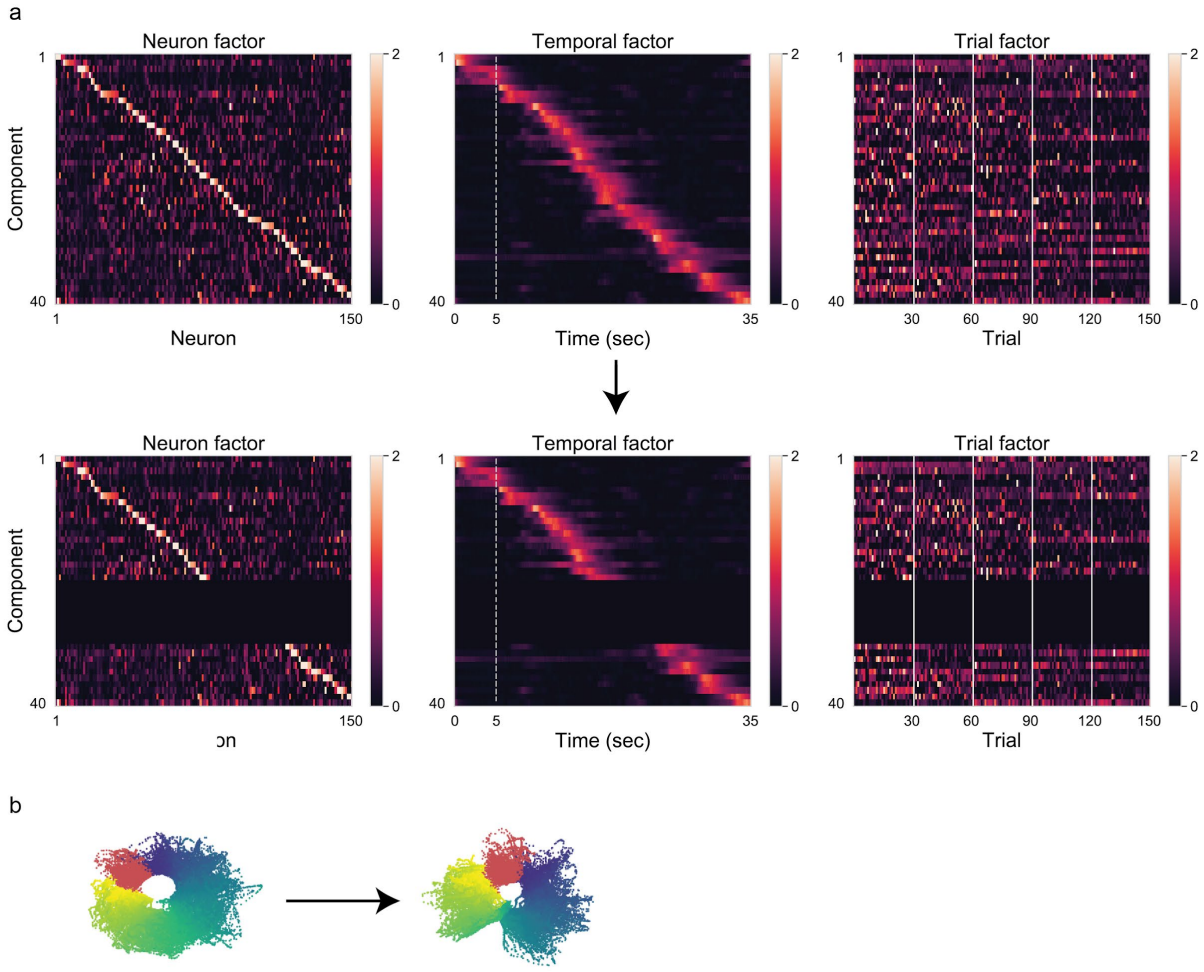

Supplemental Figure 6 Removing episodic activity during a certain time window leads to collapse of the ring manifold.

a. TCA components of one imaging field (n = 150) and the modified TCA components. For modified TCA components, we set the components with episodic activity during a certain time window to zeros.

b. 2-dimensional neural manifold extracted from reconstructed $\Delta F/F$ population activity from original TCA components and neural manifold extracted from reconstructed $\Delta F/F$ population activity from modified TCA components.

a

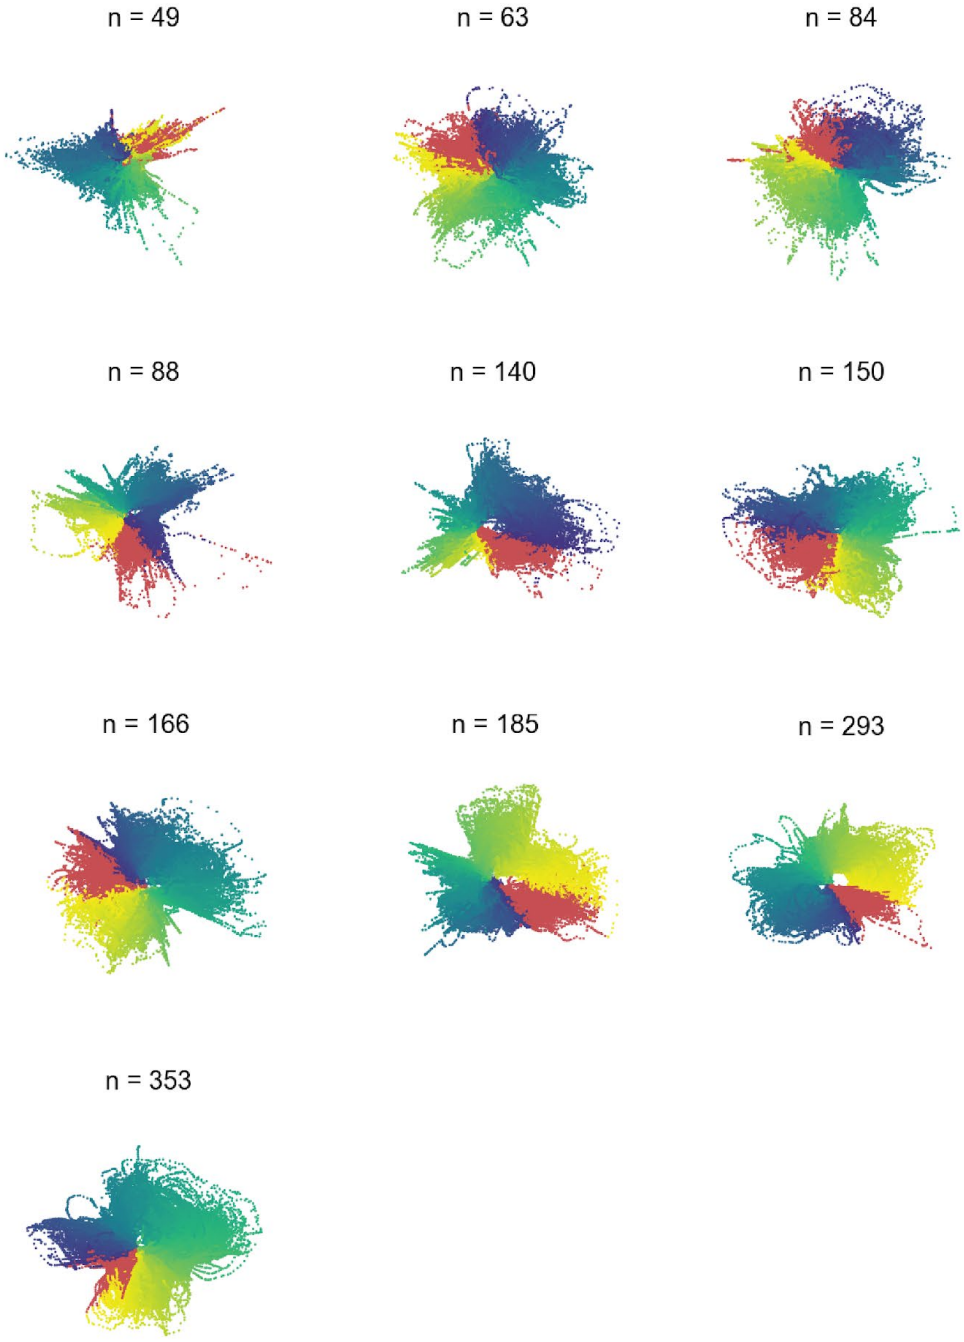
n = 49
n = 63
n = 84
n = 88
n = 140
n = 150
n = 166
n = 185
n = 293
n = 353

b

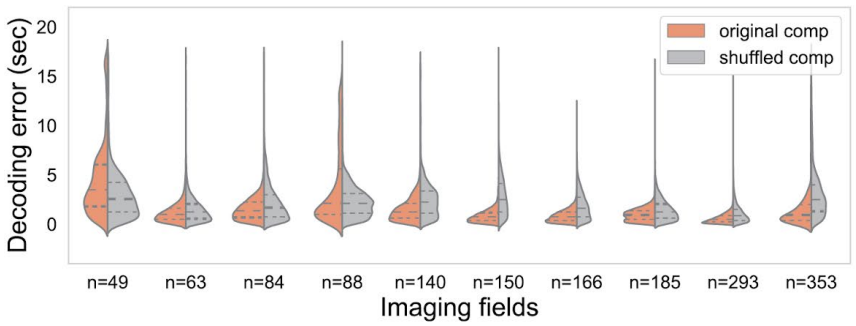
Decoding error (sec)
20
15
10
5
0
original comp
shuffled comp
n=49
n=63
n=84
n=88
n=140
n=150
n=166
n=185
n=293
n=353
Imaging fields

Supplemental Figure 7 Neural manifold of reconstructed $\Delta F/F$ population activity from TCA components with shuffled factors

a. 2-dimensional neural manifold extracted from reconstructed $\Delta F/F$ population activity from TCA components with shuffled factors as described in Fig. 7c using Isomap for 10 imaging fields. Each dot represents instantaneous population activity in the test set. Color of the dot (Same colormap as Fig. 4a) indicates the corresponding time in the trial.

b. Decoding error (absolute circular difference between decoded time and actual time) of SPUD on 2-dimensional neural manifold from reconstructed $\Delta F/F$ population activity from original TCA components plotted against decoding error of SPUD on 2-dimensional neural manifold from reconstructed $\Delta F/F$ population activity from TCA components with shuffled factors as described in Fig. 7c for all the imaging fields. Imaging fields were ordered by the number of recorded neurons. Except for 2 imaging fields (n = 49, n = 88), decoding error of reconstructed $\Delta F/F$ population activity from original TCA components was significantly smaller than decoding error of reconstructed $\Delta F/F$ population activity from TCA components with shuffled factors (Mann-Whitney U test, $p < 0.0001$). Note that for TCA components with shuffled factors, we chose ten trial-averaged projected instantaneous population activity evenly distributed in time as the initial knots for the SPUD to enable force quantitative comparison (see Methods).

Reference

1. Softky, W. R. & Koch, C. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.* **13**, 334–350 (1993).
2. Tomko, G. J. & Crapper, D. R. Neuronal variability: non-stationary responses to identical visual stimuli. *Brain Res.* **79**, 405–418 (1974).
3. Xia, J., Marks, T. D., Goard, M. J. & Wessel, R. Diverse co-active neurons encode stimulus-driven and stimulus-independent variables. *J. Neurophysiol.* (2020) doi:10.1152/jn.00431.2020.
4. Wright, N. C., Hoseini, M. S. & Wessel, R. Adaptation modulates correlated subthreshold response variability in visual cortex. *J. Neurophysiol.* **118**, 1257–1269 (2017).
5. Hoseini, M. S. *et al.* Dynamics and sources of response variability and its coordination in visual cortex. *Vis. Neurosci.* **36**, E012 (2019).
6. Chambers, A. R. & Rumpel, S. A stable brain from unstable components: Emerging concepts and implications for neural computation. *Neuroscience* **357**, 172–184 (2017).
7. LeMessurier, A. M. & Feldman, D. E. Plasticity of population coding in primary sensory cortex. *Curr. Opin. Neurobiol.* **53**, 50–56 (2018).
8. Rule, M. E., O'Leary, T. & Harvey, C. D. Causes and consequences of representational drift. *Curr. Opin. Neurobiol.* **58**, 141–147 (2019).
9. Clopath, C., Bonhoeffer, T., Hübener, M. & Rose, T. Variance and invariance of neuronal long-term representations. *Philosophical Transactions of the Royal Society B: Biological Sciences* vol. 372 20160161 (2017).
10. Rokni, U., Richardson, A. G., Bizzi, E. & Seung, H. S. Motor learning with unstable neural representations. *Neuron* **54**, 653–666 (2007).
11. Driscoll, L. N., Pettit, N. L., Minderer, M., Chettih, S. N. & Harvey, C. D. Dynamic Reorganization of Neuronal Activity Patterns in Parietal Cortex. *Cell* vol. 170 986–999.e16 (2017).

12. Mankin, E. A. *et al.* Neuronal code for extended time in the hippocampus. *Proc. Natl. Acad. Sci. U. S. A.* **109**, 19462–19467 (2012).
13. Betzel, R., Wood, K. C., Angeloni, C., Geffen, M. N. & Bassett, D. S. Stability of spontaneous, correlated activity in mouse auditory cortex. doi:10.1101/491936.
14. Ziv, Y. *et al.* Long-term dynamics of CA1 hippocampal place codes. *Nat. Neurosci.* **16**, 264–266 (2013).
15. Katlowitz, K. A., Picardo, M. A. & Long, M. A. Stable Sequential Activity Underlying the Maintenance of a Precisely Executed Skilled Behavior. *Neuron* vol. 98 1133–1140.e3 (2018).
16. Jeon, B. B., Swain, A. D., Good, J. T., Chase, S. M. & Kuhlman, S. J. Feature selectivity is stable in primary visual cortex across a range of spatial frequencies. *Scientific Reports* vol. 8 (2018).
17. Rose, T., Jaepel, J., Hubener, M. & Bonhoeffer, T. Cell-specific restoration of stimulus preference after monocular deprivation in the visual cortex. *Science* vol. 352 1319–1322 (2016).
18. Montijn, J. S., Meijer, G. T., Lansink, C. S. & Pennartz, C. M. A. Population-Level Neural Codes Are Robust to Single-Neuron Variability from a Multidimensional Coding Perspective. *Cell Rep.* **16**, 2486–2498 (2016).
19. Deitch, D., Rubin, A. & Ziv, Y. Representational drift in the mouse visual cortex. *BioRxiv* (2020) doi:10.1101/2020.10.05.327049.
20. Vinje, W. E. & Gallant, J. L. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* **287**, 1273–1276 (2000).
21. Baudot, P. *et al.* Animation of natural scene by virtual eye-movements evokes high precision and low noise in V1 neurons. *Front. Neural Circuits* **7**, 206 (2013).
22. Olshausen, B. A. & Field, D. J. How close are we to understanding v1? *Neural Comput.* **17**, 1665–1699 (2005).

23. David, S. V. Natural Stimulus Statistics Alter the Receptive Field Structure of V1 Neurons. *Journal of Neuroscience* vol. 24 6991–7006 (2004).
24. Marks, T. & Goard, M. Stimulus-dependent representational drift in the primary visual cortex. *BioRxiv* (2020).
25. Druckmann, S. & Chklovskii, D. B. Neuronal Circuits Underlying Persistent Representations Despite Time Varying Activity. *Current Biology* vol. 22 2095–2103 (2012).
26. Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience* vol. 17 1500–1509 (2014).
27. Gao, P. & Ganguli, S. On simplicity and complexity in the brave new world of large-scale neuroscience. *Curr. Opin. Neurobiol.* **32**, 148–155 (2015).
28. Buracas, G. T., Zador, A. M., DeWeese, M. R. & Albright, T. D. Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* **20**, 959–969 (1998).
29. Kumbhani, R. D., Nolt, M. J. & Palmer, L. A. Precision, reliability, and information-theoretic analysis of visual thalamocortical neurons. *J. Neurophysiol.* **98**, 2647–2663 (2007).
30. Spacek, M. A. & Swindale, N. V. Cortical state and natural movie responses in cat visual cortex. doi:10.1101/031765.
31. Herikstad, R., Baker, J., Lachaux, J.-P., Gray, C. M. & Yen, S.-C. Natural movies evoke spike trains with low spike time variability in cat primary visual cortex. *J. Neurosci.* **31**, 15844–15860 (2011).
32. Pachitariu, M. *et al.* Suite2p: beyond 10,000 neurons with standard two-photon microscopy. doi:10.1101/061507.
33. Williams, A. H. *et al.* Unsupervised Discovery of Demixed, Low-Dimensional Neural Dynamics across Multiple Timescales through Tensor Component Analysis. *Neuron* **98**, 1099–1115.e8 (2018).
34. Yuste, R. From the neuron doctrine to neural networks. *Nat. Rev. Neurosci.* **16**, 487–497

(2015).

35. Lichtman, J. W., Pfister, H. & Shavit, N. The big data challenges of connectomics. *Nat. Neurosci.* **17**, 1448–1454 (2014).
36. Tenenbaum, J. B. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science* vol. 290 2319–2323 (2000).
37. Chaudhuri, R., Gerçek, B., Pandey, B., Peyrache, A. & Fiete, I. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nat. Neurosci.* **22**, 1512–1520 (2019).
38. Averbeck, B. B., Latham, P. E. & Pouget, A. Neural correlations, population coding and computation. *Nature Reviews Neuroscience* vol. 7 358–366 (2006).
39. Zylberberg, J., Cafaro, J., Turner, M. H., Shea-Brown, E. & Rieke, F. Direction-Selective Circuits Shape Noise to Ensure a Precise Population Code. *Neuron* **89**, 369–383 (2016).
40. Rumyantsev, O. I. *et al.* Fundamental bounds on the fidelity of sensory cortical coding. *Nature* **580**, 100–105 (2020).
41. Hubel, D. H. & Wiesel, T. N. Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology* vol. 148 574–591 (1959).
42. O'Keefe, J. & Dostrovsky, J. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research* vol. 34 171–175 (1971).
43. Fyhn, M. Spatial Representation in the Entorhinal Cortex. *Science* vol. 305 1258–1264 (2004).
44. Rule, M. E. *et al.* Stable task information from an unstable neural population. *Elife* **9**, (2020).
45. Stringer, C., Michaelos, M. & Pachitariu, M. High precision coding in visual cortex. doi:10.1101/679324.
46. Rubin, A. *et al.* Revealing neural correlates of behavior without behavioral measurements. *Nat. Commun.* **10**, 4745 (2019).

47. Orbán, G., Berkes, P., Fiser, J. & Lengyel, M. Neural Variability and Sampling-Based Probabilistic Representations in the Visual Cortex. *Neuron* **92**, 530–543 (2016).
48. Berkes, P., Orbán, G., Lengyel, M. & Fiser, J. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* **331**, 83–87 (2011).
49. Maass, W., Natschläger, T. & Markram, H. Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput.* **14**, 2531–2560 (2002).
50. Niell, C. M. & Stryker, M. P. Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron* **65**, 472–479 (2010).
51. Stringer, C. *et al.* Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, 255 (2019).
52. Jazayeri, M. & Afraz, A. Navigating the Neural Space in Search of the Neural Code. *Neuron* vol. 93 1003–1014 (2017).
53. Sadtler, P. T. *et al.* Neural constraints on learning. *Nature* **512**, 423–426 (2014).
54. Rikhye, R. V. & Sur, M. Spatial Correlations in Natural Scenes Modulate Response Reliability in Mouse Visual Cortex. *Journal of Neuroscience* vol. 35 14661–14680 (2015).
55. Ponce, C. R. *et al.* Evolving Images for Visual Neurons Using a Deep Generative Network Reveals Coding Principles and Neuronal Preferences. *Cell* vol. 177 999–1009.e10 (2019).
56. Clawson, W. P., Wright, N. C., Wessel, R. & Shew, W. L. Adaptation towards scale-free dynamics improves cortical stimulus discrimination at the cost of reduced detection. *PLOS Computational Biology* vol. 13 e1005574 (2017).
57. Carrillo-Reid, L. & Yuste, R. Playing the piano with the cortex: role of neuronal ensembles and pattern completion in perception and behavior. *Curr. Opin. Neurobiol.* **64**, 89–95 (2020).
58. Carrillo-Reid, L., Han, S., Yang, W., Akrouh, A. & Yuste, R. Controlling Visually Guided Behavior by Holographic Recalling of Cortical Ensembles. *Cell* **178**, 447–457.e5 (2019).

59. Marshel, J. H. *et al.* Cortical layer-specific critical dynamics triggering perception. *Science* **365**, (2019).
60. Madisen, L. *et al.* Transgenic Mice for Intersectional Targeting of Neural Sensors and Effectors with High Specificity and Performance. *Neuron* vol. 85 942–958 (2015).
61. Pho, G. N., Goard, M. J., Woodson, J., Crawford, B. & Sur, M. Task-dependent representations of stimulus and choice in mouse parietal cortex. *Nat. Commun.* **9**, 2596 (2018).
62. Huber, D. *et al.* Multiple dynamic representations in the motor cortex during sensorimotor learning. *Nature* **484**, 473–478 (2012).
63. Brainard, D. H. The Psychophysics Toolbox. *Spatial Vision* vol. 10 433–436 (1997).
64. Dimatteo, I., Genovese, C. R. & Kass, R. E. Bayesian curve-fitting with free-knot splines. *Biometrika* vol. 88 1055–1071 (2001).
65. Kruskal, J. B. Three-way arrays: rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics. *Linear Algebra and its Applications* vol. 18 95–138 (1977).
66. Bellman, R. *Adaptive Control Processes: A Guided Tour*. (1961).

Figures

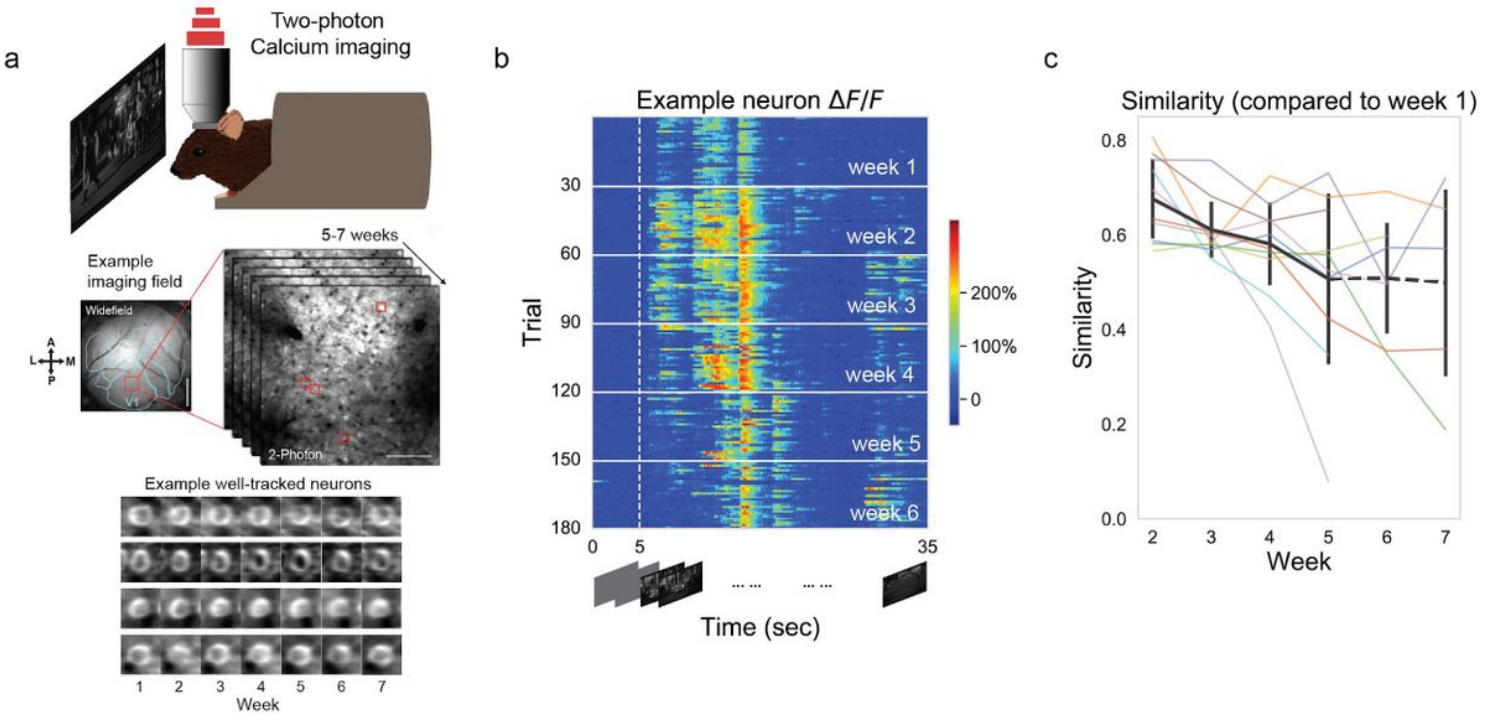

Figure 1

Single neuron responses to natural movies are unstable across weeks.a. Experimental setup. We performed chronic calcium imaging of excitatory neurons in the primary visual cortex of awake, head-fixed mice during visual stimulation with repeated natural movies. Visual cortex (contralateral to visual stimulus delivery) is retinotopically mapped in Emx1-Cre::TITL-GCaMP6s mice. V1 fields are chosen from the region selective for the center of the presentation screen. Widefield scale bar = 1 mm; 2-photon scale bar = 100 μm. Average activity of four example well-tracked neurons across weeks are shown in the bottom panel. b. □□□□/□□ responses of one example neuron during the same natural movie clip for 30 trials per experimental session for 6 weeks (movie starts at 5 sec and lasts for 30 sec duration). We recorded 1 experimental session per week. c. Similarity (correlation coefficient between trial-averaged □□□□/□□) averaged over neurons during week 1 and that during other weeks are plotted for all the recorded imaging fields. Different imaging fields are denoted by different colors. The black curve with error bar denotes mean and standard deviation of similarity over imaging fields. Only a subset of imaging fields have recordings on week 6 (6 fields) and week 7 (5 fields). Specifically, the similarities of the fifth week were significantly lower than the similarities of the second week (Mann-Whitney U test, $p < 0.01$).

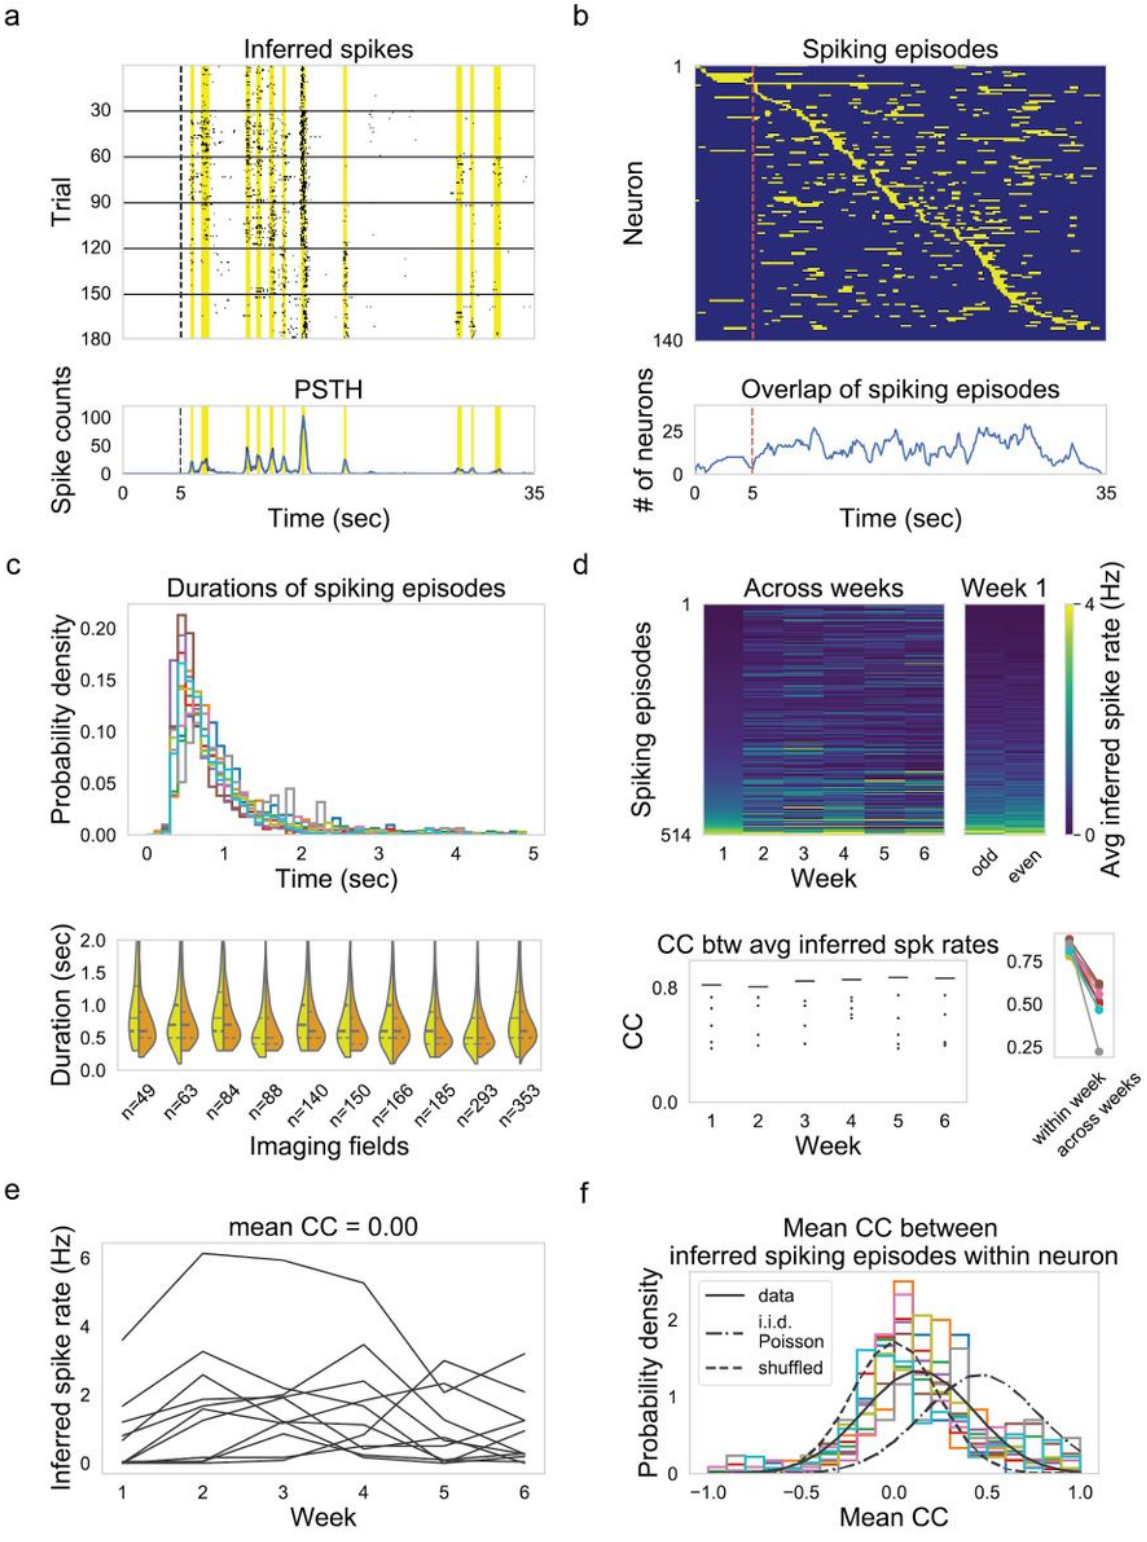

Figure 2

Single neuron responses consist of episodic activity with distinct episode-specific rate variations across weeks. a. Top: inferred spikes of the same neuron shown in Fig. 1b. Bottom: peristimulus time histogram (PSTH) (black) and smoothed PSTH (blue) of the same neuron. Shaded areas (yellow) denote spiking episodes for this neuron. b. Top: spiking episodes for all the neurons in the example imaging field. Neurons are ordered by latency of their spiking episodes with the highest spiking rates. Bottom: number

of neurons with overlapped spiking episodes.c. Top: Distributions of durations of spiking episodes from all imaging fields. Different colors denote different imaging fields. Bottom: distribution of durations of spiking episodes defined from PSTH of trials across weeks (yellow) plotted against distribution of durations from PSTH of trials within weeks (orange). d. Top: averaged spike rates over trials of all the spiking episodes in one example imaging field are plotted for different weeks and for even and odd trials in week 1. spiking episodes are ordered by their averaged spike rates during week 1. Bottom left: correlation coefficients (CC) between averaged spike rates of week pairs (dots) and even/odd trials within the week (lines) are shown for the example imaging field across weeks. Bottom right: CC within week averaged across weeks is plotted against CC across weeks averaged across all the week pairs (for all imaging fields, Mann-Whitney U test, $p < 0.005$). Different colors denote different imaging fields. Colormap maximum value is set to 4 Hz. e. Mean spike rate during each spiking episode in the example neuron varies across weeks. f. Histogram of mean CC between mean spike rates during spiking episodes within the same neuron. Different colors denote different imaging fields. The black solid line is a gaussian curve fitted to the distribution of mean CC from all the imaging fields (mean 0.13, s.t.d. 0.30). The black dash dotted line is a gaussian curve fitted to the distribution of mean CC between simulated independent and identically distributed Poisson spike trains with the firing rates of a randomly selected spiking episode for a given neuron (mean 0.46, s.t.d. 0.31). The black dashed line indicates the chance level, which is a gaussian curve fitted to the distribution of mean CC between spiking episodes with independently shuffled weeks (mean 0.0036, s.t.d. 0.23). Only neurons with more than one spiking episode were included in this analysis.

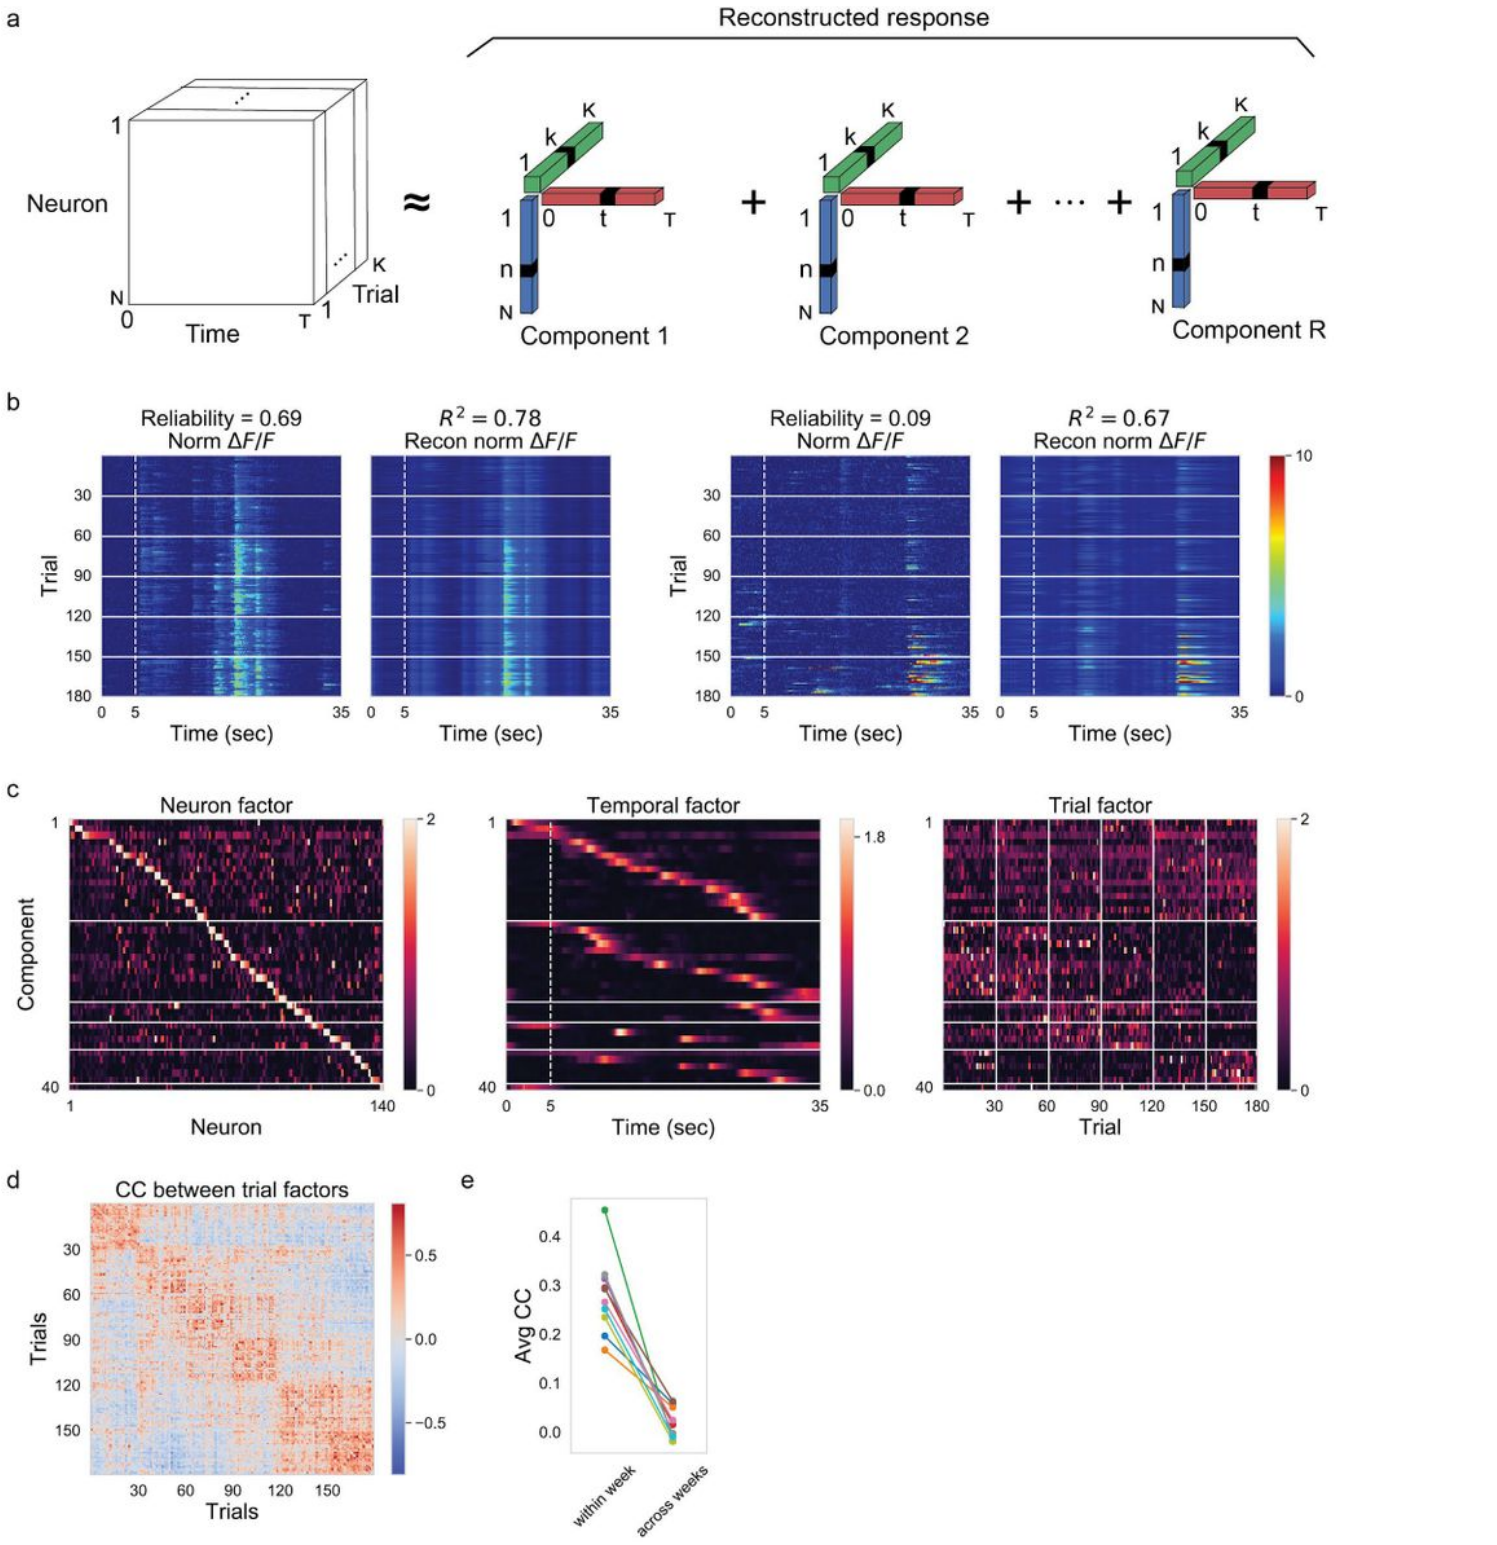

Figure 3

Latent factors resembling episodic activity with gain changes capture the across-week fluctuations a. Schematic of Tensor Component Analysis (TCA). Neural activity (□□□□/□□) is organized into a third-order tensor with dimensions N x T x K. TCA approximates the data as a sum of outer products of three vectors from R components: neuron factors describe the weights of each neuron to that component, temporal factors describe the temporal dynamics of each component, and trial factors describe the modulation of

the component across trials.b. Normalized □□□□/□□ responses and reconstructed □□□□/□□from 40 TCA components of two example neurons from the example imaging field. Reliability was defined as averaged correlation-coefficient between pairs of single-trial responses 3.c. Neuron, temporal, and trial factors of nonnegative TCA with 40 components for the example imaging field. Colormap maximum values are set to 2 for neuron factors and trial factors. We ordered components according to the K-means clustering on their trial factors. Within each thus determined cluster, we further ordered the components by the time to peak in their temporal factors. We ordered neurons in the neuron factors by their dominant components.d. Correlation coefficient (CC) between trial factors shown in c. e. CC between trial factors averaged across trial pairs within week plotted against CC between trial factors averaged across trial pairs across weeks. Different color denotes different imaging fields. The week-to-week variability of trial factors was significantly larger than the corresponding trial-to-trial variability within each week (for all imaging fields, Mann-Whitney U test, $p < 0.0001$).

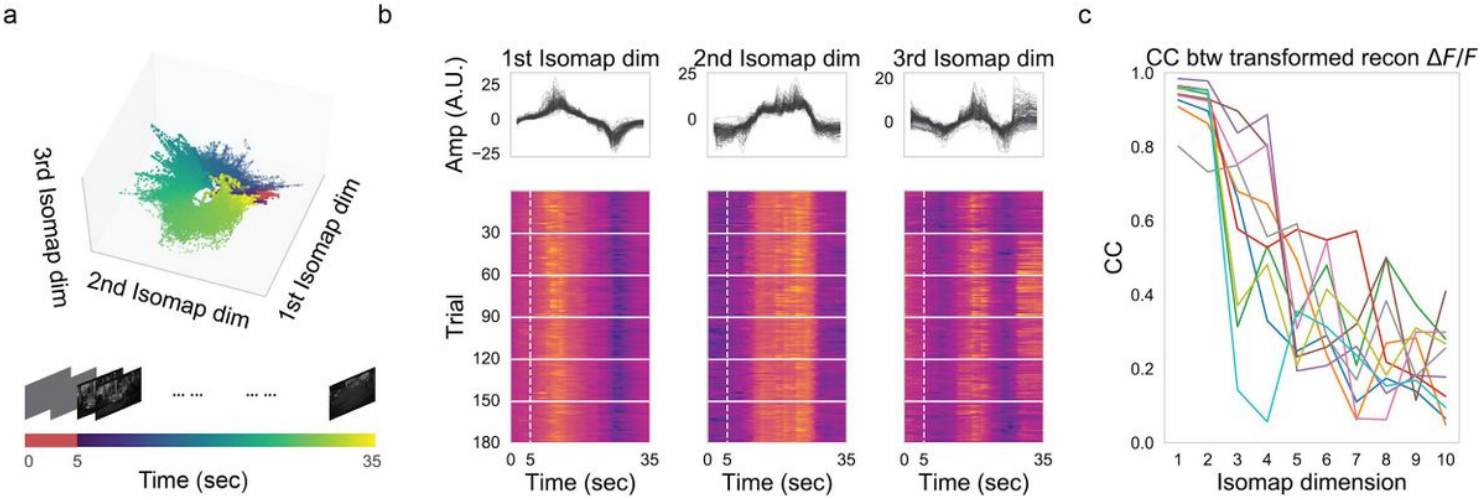

Figure 4

Stable manifolds exist in unstable population activity.a. 3-dimensional neural trajectories extracted from reconstructed (denoised) □□□□/□□ populational activity across weeks from the example imaging field using Isomap. Each dot represents instantaneous population activity. Color of the dot indicates the corresponding time in the trial.b. Neural trajectories along the first 3 Isomap dimensions (the same as shown in a) organized in trial by time matrices.c. Correlation coefficients (CC) between transformed reconstructed □□□□/□□ (neural trajectories) across trials along each Isomap dimension are plotted for all 10 imaging fields. 7Different color denotes different imaging fields.

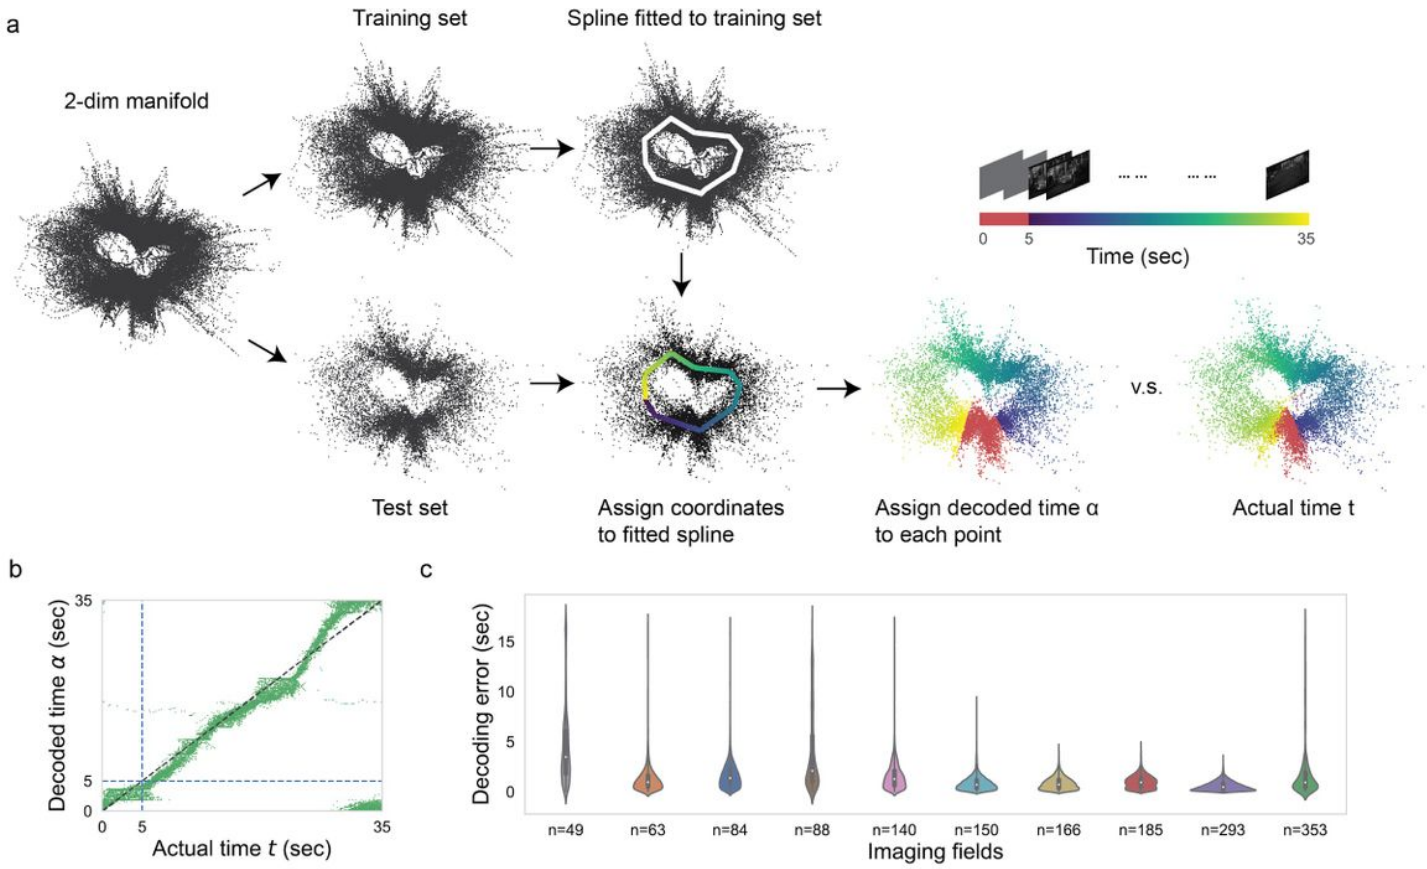

Figure 5

The manifold mediates a stable representation of the time within the movie clip. a. Illustration of the unsupervised method with data from the example imaging field (n = 140): first, we projected reconstructed DF/F responses into the first two Isomap dimensions, each dot denotes instantaneous population activity; second, we randomly pick 80% of the instantaneous population activity as training set and rest of them as test set; third, we fitted a spline to the neural manifold of the training set and assigned coordinates with randomly picked origin to the fitted spline; finally, we shifted and flipped the coordinates on the fitted spline to match with the actual time and assigned decoded time to each point in the test set by its nearest coordinate on the spline.b. Decoded time from the neural manifold plotted against actual time in the movie for the example imaging field (n = 140).c. Violin plots showed decoding error (absolute circular difference between decoded time and actual time) for all the imaging fields. Imaging fields were ordered by the number of recorded neurons.

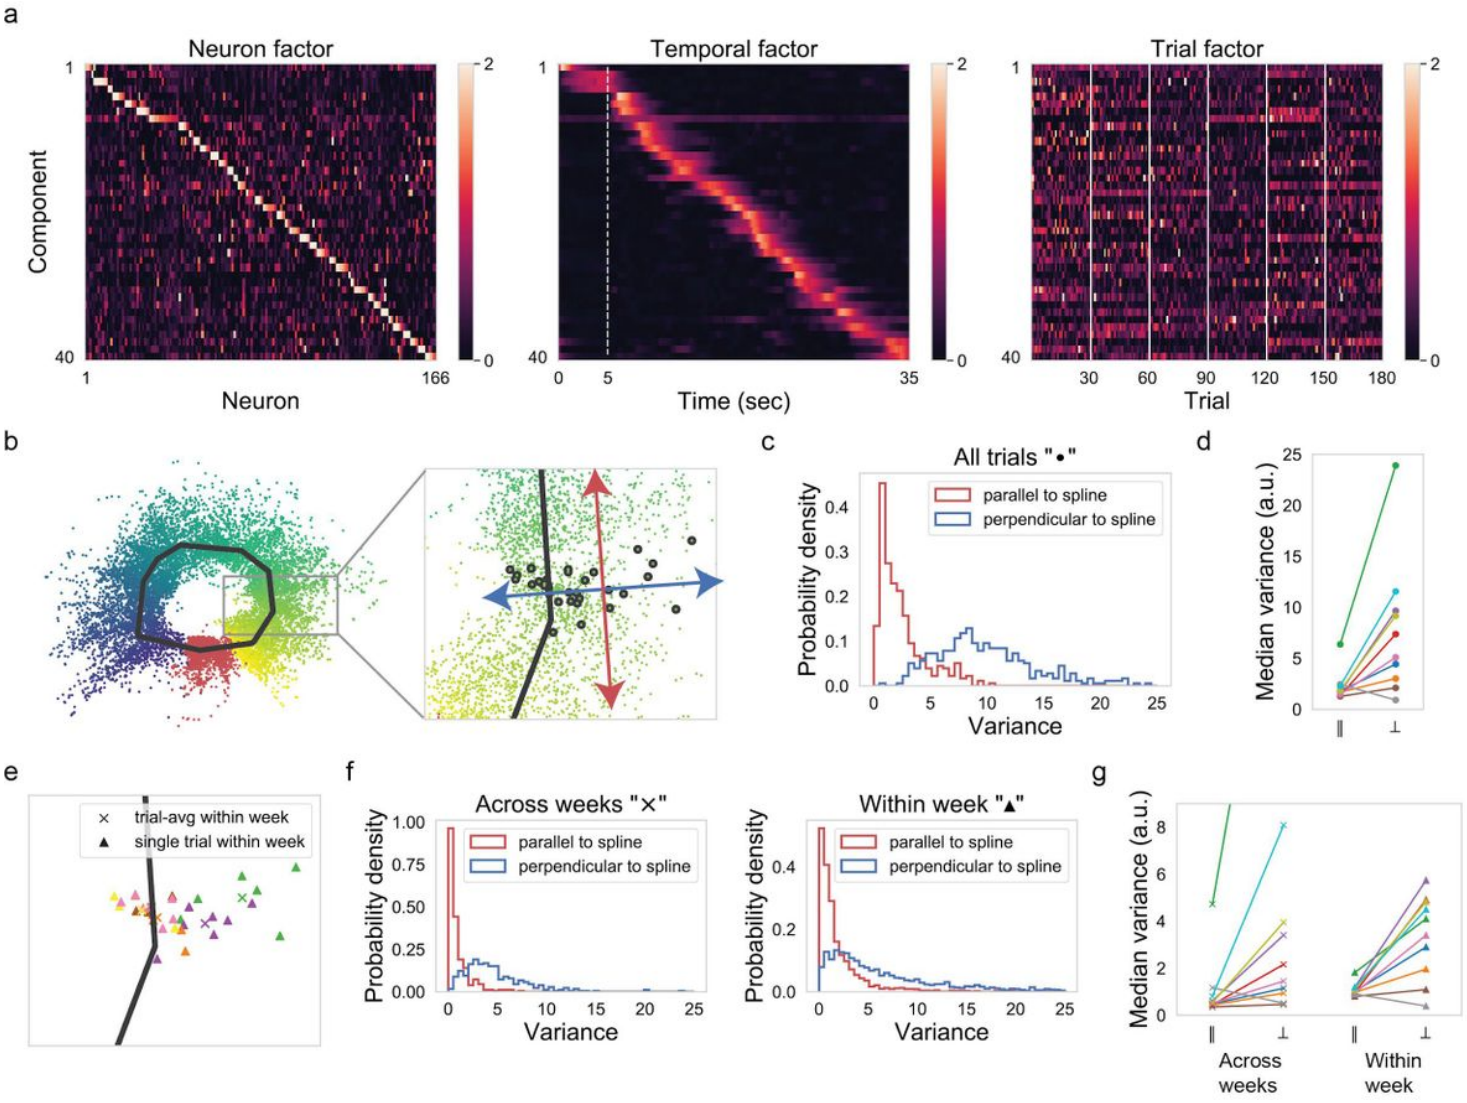

Figure 6

Both week-to-week fluctuation and trial-to-trial variation within the week is restricted to non-coding directions.a. TCA components of one imaging field (n = 166). We ordered components by the time to peak in their temporal factors. We ordered neurons in the neuron factors by their dominant components. Colormap maximum values are set to 2 for all the factors.b. Left: 2-dimensional neural manifold extracted from reconstructed (denoised) □□□□/□□ population activity (n = 166) across weeks using Isomap. Each dot represents instantaneous population activity in the test set. Color of the dot (Same colormap as Fig. 4a) indicates the corresponding time in the trial. Black line is the fitted spline to the training set. Right: Zoom-in view on the neural manifold. Instantaneous population activity corresponding to 28 s in the trial was highlighted with black shade. Blue arrow denotes the direction perpendicular to the spline, and red arrow denotes the direction parallel to the spline. c. Histogram of the variance of population activity parallel or perpendicular to the spline for the imaging field shown in a&b.d. Median variance of population activity parallel to the spline plotted against median variance of population activity perpendicular to the spline for all the imaging fields. Except for one imaging field (gray one), variance of population activity parallel to the spline was significantly smaller than the variance

perpendicular to the spline (for all imaging fields, Mann-Whitney U test, $p < 0.0001$). The one outlier (gray line) is from an imaging field with the least number of recorded neurons (n = 49), whose neural manifold didn't have a clear ring shape (Supplemental Fig. 5a).e. The same zoom-in view on the neural manifold as shown in b (right). Each triangle represents instantaneous population activity within a week. Color of the triangle denotes different weeks. Each cross represents the trial-averaged instantaneous population activity within a week. Color of the cross also denotes different weeks. f. Left: histogram of the variance of trial-averaged population activity within a week parallel or perpendicular to the spline for the imaging field. Right: histogram of the variance of single-trial population activity within a week parallel or perpendicular to the spline for the imaging field.g. Median variance of trial-averaged population activity or single-trial population within a week parallel to the spline plotted against median variance of population activity perpendicular to the spline for all the imaging fields. Y axis is clippled at 9 for visualization. Except for one imaging field (gray one), variance of population activity parallel to the spline was significantly smaller than the variance perpendicular to the spline for both across weeks and within a week cases (for all imaging fields, Mann-Whitney U test, $p < 0.0001$). The one outlier (gray line) is from an imaging field with the least number of recorded neurons (n = 49).

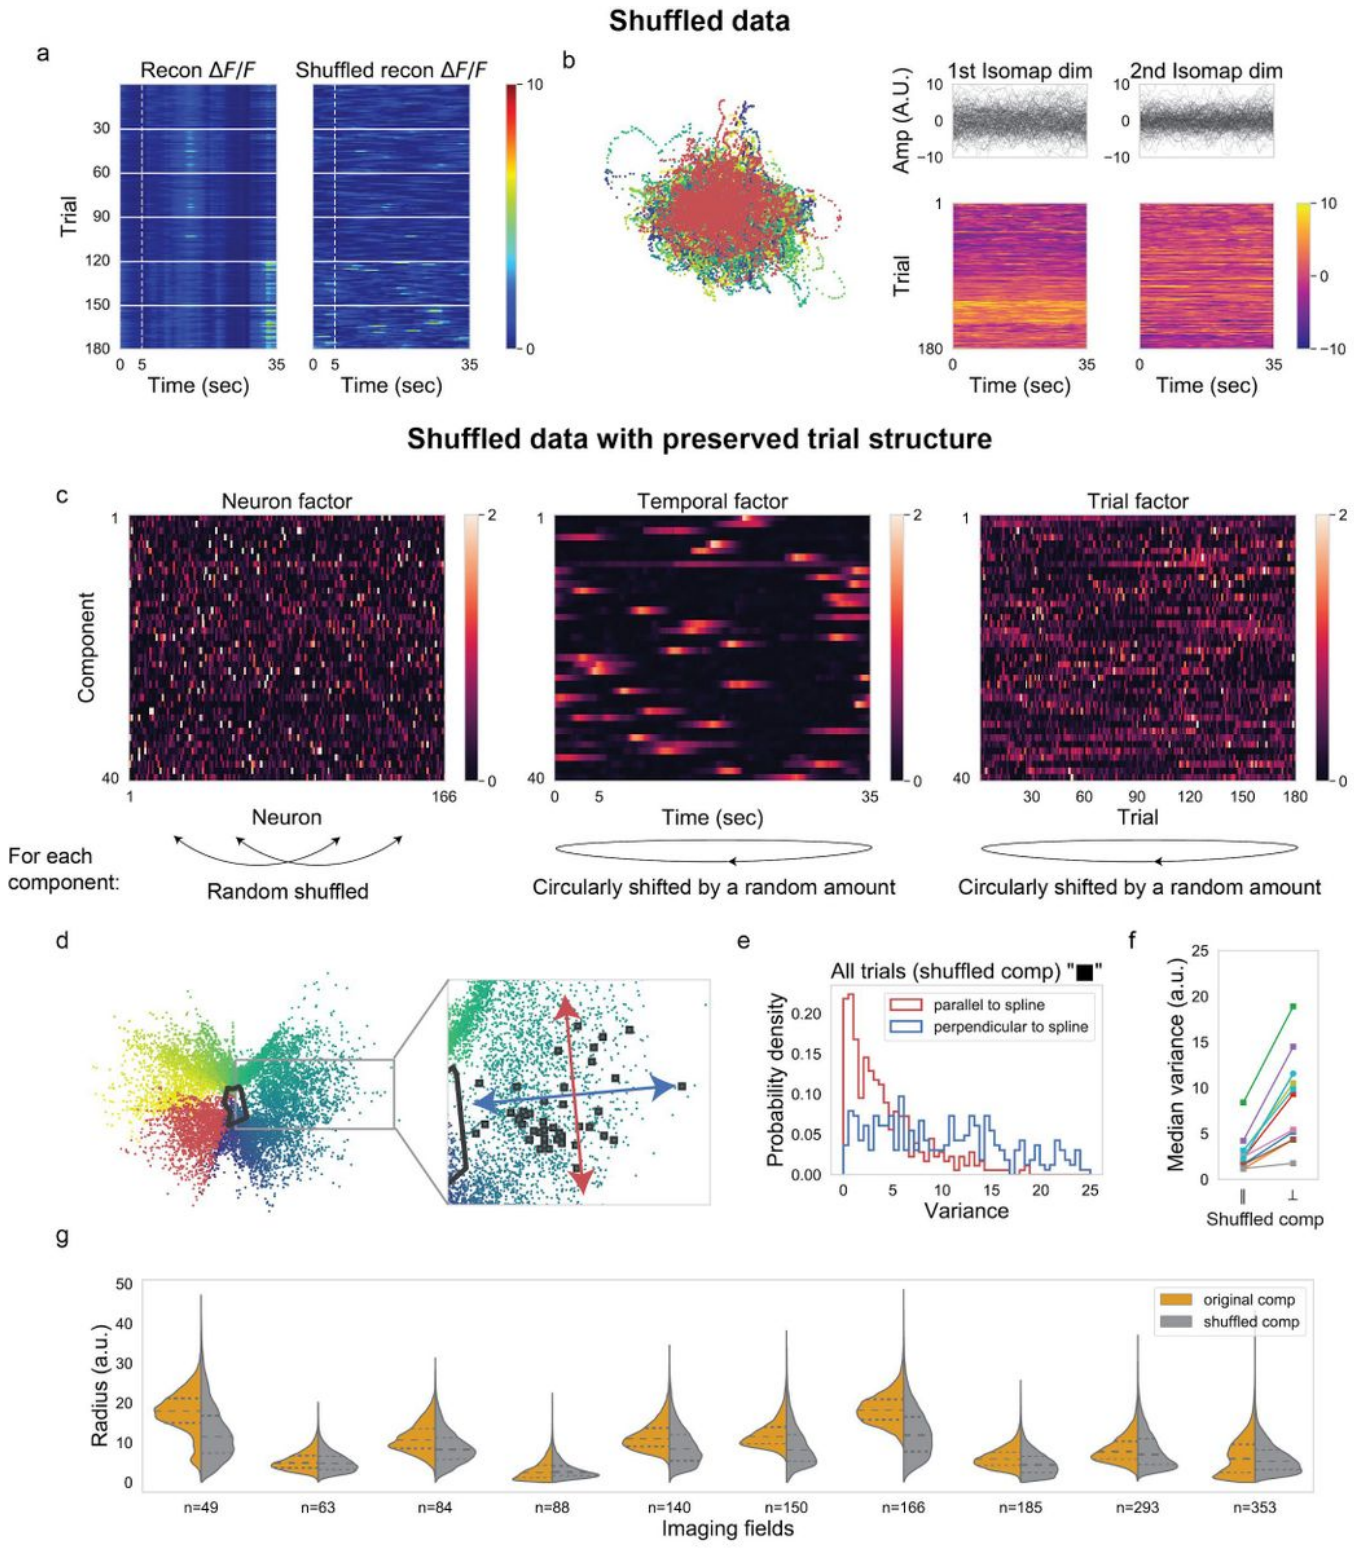

Figure 7

The precisely timed episodic activity constrains neural variability to non-coding directions.a. Reconstructed ΔF/F responses and shuffled reconstructed ΔF/F responses of two example neurons from the imaging field (n = 166). The single neuron response was circularly shifted by a random amount independently for each trial for shuffling.b. Left: 2-dimensional neural manifold extracted from shuffled reconstructed ΔF/F population activity (example single neuronal shuffled responses

shown in a). The same colormap was used as in Fig. 6b, left panel. Right: neural trajectories along the first 2 Isomap dimensions (the same as shown in the left panel) organized in trial by time matrices. Here we set the number of nearest neighbors of ISOMAP to be 100 (see Methods).c. TCA components of the imaging field (n = 166) with shuffled factors. For each component,we independently shuffled neuron order in the neuron factor, circularly shifted the time factor and the trial factor by a random amount. Components with shuffled factors were ordered again in the same fashion as Fig.6a. d. 2-dimensional neural manifold extracted from reconstructed (denoised) □□□□/□□ population activity (n = 166) from components with shuffled factors using Isomap. Each dot represents instantaneous population activity in the test set. Black line is the fitted spline to the training set. Instantaneous population activity corresponding to 16 s in the trial was highlighted with black shade.e. Histogram of the variance of neural variability parallel or perpendicular to the spline for reconstructed (denoised) □□□□/□□ populational activity (n = 166) from components with shuffled factors (i.e., shuffled data with preserved trial structure). f. Median variance of neural variability parallel to the spline plotted against median variance of neural variability perpendicular to the spline for all the imaging fields for reconstructed activity from components with shuffled factors. Variance of population activity parallel to the spline was significantly smaller than the variance perpendicular to the spline (Mann- Whitney U test, $p < 0.0001$) for all imaging fields.g. Radius (distance to the center of the point cloud) distribution of points on the neural manifold from original TCA components plotted against radius distribution from TCA components with shuffled factors for all the imaging fields. Except for 3 imaging fields (n = 49, n = 63, n = 88), the radius of points on the neural manifold from original TCA components was significantly larger than the radius from TCA components with shuffled factors (Mann-Whitney U test, $p < 0.0001$).