Representational drift in the mouse visual cortex

Representational drift in the mouse visual cortex

Daniel Deitch[1], Alon Rubin[1,2], Yaniv Ziv[1,2]*

[1]Department of Neurobiology, Weizmann Institute of Science, Rehovot 76100, Israel
[2] These authors equally contributed to this work
* Lead Contact and Corresponding Author: Yaniv Ziv (yaniv.ziv@weizmann.ac.il)

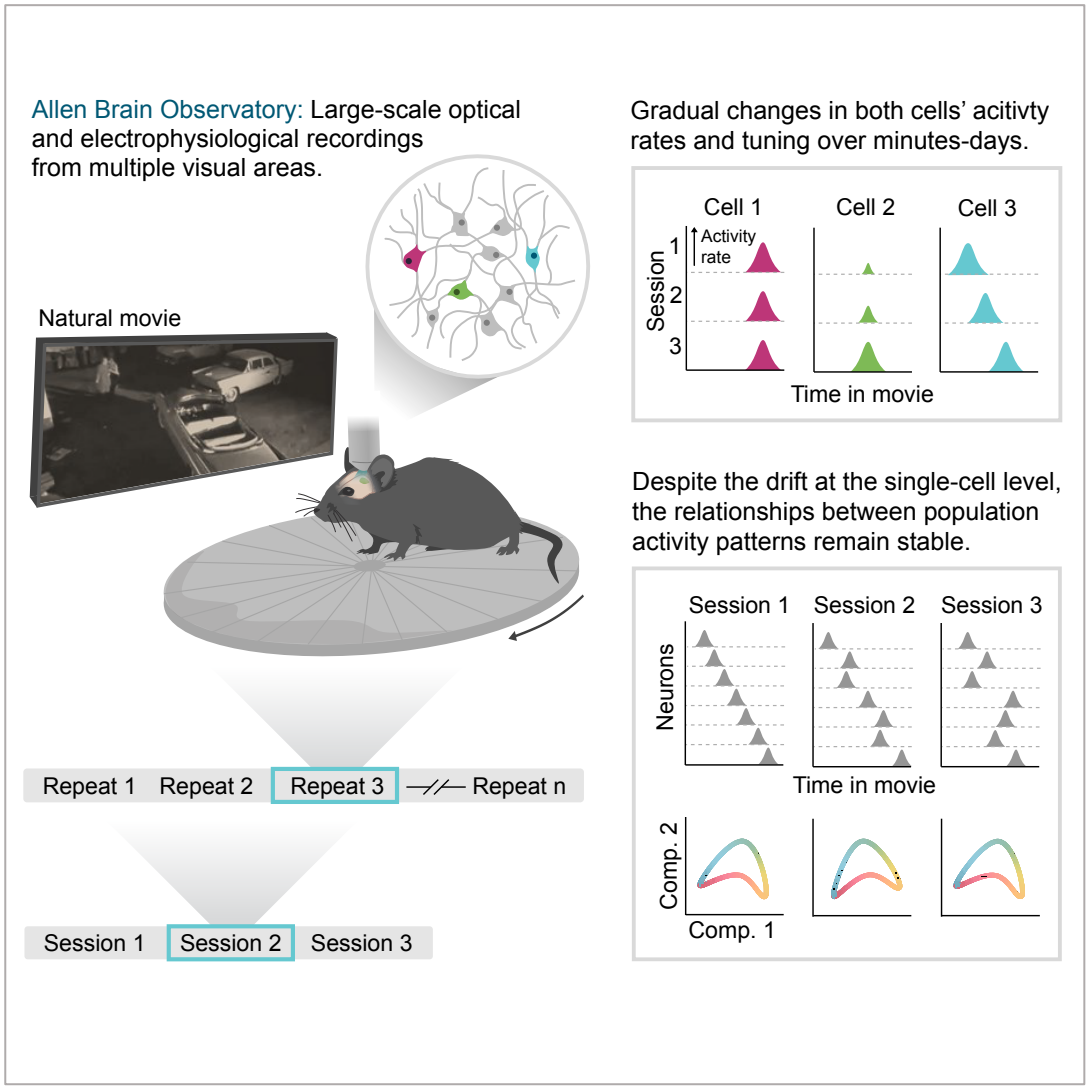

SUMMARY

Recent studies have shown that neuronal representations gradually change over time despite no changes in the stimulus, environment or behavior. However, such representational drift has been assumed to be a property of high-level brain structures, whereas earlier circuits, such as sensory cortices, have been assumed to stably encode information over time. Here, we analyzed large-scale optical and electrophysiological recordings from six visual cortical areas in behaving mice that were repeatedly presented with the same natural movies. Contrary to the prevailing notion, we found representational drift over timescales spanning minutes to days across multiple visual areas, cortical layers and cell types. Notably, neural-code stability did not reflect the hierarchy of information flow across areas. Although individual neurons showed time-dependent changes in their coding properties, the structure of the relationships between population activity patterns remained stable and stereotypic. Such population-level organization may underlie stable visual perception despite continuous changes in neuronal responses.

INTRODUCTION

One of the great marvels of the brain is that it achieves persistent functionality throughout adult life despite an extensive continuous turnover of its bio-molecular and cellular building blocks[1–5]. Recent advances in electrophysiology and optical imaging techniques enable to study in behaving animals the persistence over time of neuronal coding properties, such as the tuning of neurons to specific stimuli[6–15]. Some of these studies have exposed a substantial degree of variability in neuronal responses to the same stimuli over timescales spanning minutes to weeks, prompting neuroscientists to question the naïve assumption that stable neuronal codes are essential for stable brain functionality[4,6,8,12,14,16–30].

One example is the neuronal representations of space in the hippocampus and related brain areas, which gradually change over timescales of hours to days despite no apparent changes in the environment or behavior[6,26–28,31]. The finding of this so called 'representational drift'[32] was surprising, because classical models of memory consider the stability of the engram as the basis for the persistence of memory[33,34]. Notably, representational drift differs from mere variability in neuronal responses. In representational drift, the similarity between two representations of the same stimulus gradually decays with elapsed time, whereas variability in neuronal responsiveness does not lead to such gradual decay in the similarity between representations[21,32].

The specific mechanisms that underlie representational drift remain elusive, but it has been suggested that drift may be an inevitable outcome of the network dynamics in deep brain circuits that consist of multiple input and output loops[32]. Consistent with this logic, and given the need to support stable perception and motor outputs, it is plausible that brain circuits situated closer to the sensory input or to the motor output will display more stable neuronal representations than those of higher-order cortical areas[35]. While a direct examination of this hypothesis is still lacking, several recent studies of sensory cortices have found variability in neuronal responsiveness over days[9,11,16,36–39]. For instance, in the primary visual cortex (V1),

Rose et al. (2016) revealed session-to-session variability of neuronal visual tuning properties (e.g., ocular dominance), and Montijn et al. (2016) reported that neuronal responses to synthetic stimuli (drifting gratings) are variable across trials within the same day while showing modest gradual changes over days.

These studies provide clear indications that representations of visual stimuli in L2/3 neurons of V1 are variable over time. However, it remains unclear if and to what extent the visual cortex exhibits representational drift that is similar to that observed in deep circuits[6,26], in terms of the degree to which different aspects of cells' coding properties, such as tuning and activity rate, change over time. It is also unknown how the stability of neuronal coding properties differs across different cell-types and cortical layers within a given area.

Recently, the Allen Brain Institute published two large-scale, standardized physiological surveys of neuronal coding in the visual cortex (Allen Brain Observatory)[40,41]. These datasets consist of optical and electrophysiological recordings of tens of thousands of neurons from six different visual cortical areas in hundreds of awake behaving mice that were repeatedly presented with the same set of visual stimuli. Thus, they offer a unique opportunity to study coding stability across different areas of the visual cortex and over different timescales, from minutes to days. The fact that the same experiments were conducted using two different recording techniques (Neuropixels probes[42] and Ca^{2+} imaging) can help control for the limitations and biases associated with each technique. Furthermore, a specific set of stimuli – natural scene movies – were used in these experiments and on different days. This allows studying the stability of visual representations of complex stimuli that are more ethologically relevant than the synthetic stimuli traditionally used for longitudinal studies[43–45].

Using these datasets, we found that representational drift does occur across different visual areas, over timescales spanning minutes to days, and is characterized by both changes in the cells' activity rates and their tuning. We demonstrate that despite clear time-dependent changes in neuronal responsiveness to visual stimuli, the structure of relationships between neuronal population activity patterns remains stable, permitting the conservation of visual information over time.

RESULTS

We analyzed datasets from experiments that used two recording techniques: two-photon Ca^{2+} imaging[40] and electrophysiology via Neuropixels probes[41]. The Ca^{2+} imaging dataset comprises neuronal activity from nearly 60,000 neurons collected from six visual cortical areas, across different layers, from hundreds of adult mice that were presented with the same set of visual stimuli (Figure 1A-D). Each mouse was imaged from a single cortical area while performing three imaging sessions, separated by days. During each session, mice viewed a battery of natural and artificial stimuli (Figure 1C). The Neuropixels dataset comprises neuronal activity from nearly 100,000 single units collected from six visual areas, thalamic nuclei, and the hippocampus, from 58 adult mice (Figure 1E-H). Each mouse was implanted with multiple Neuropixels probes in different brain areas and underwent a single recording session while viewing a battery of natural and artificial stimuli (Figure 1G).

We focused our analysis on data recorded during the presentations of two natural movies because they were presented twice within the same recording session or in all imaging sessions across days. This enabled us to study the stability of neuronal representations on three different time scales: (1) Between movie repetitions within a single block across seconds-minutes; (2) Between different blocks within the same recording session across minutes-hours; and (3) Between sessions recorded on different days. In datasets from both recording techniques, we could readily identify neurons that displayed reliable and distinct tuning curves that were stable across different movie repeats, blocks, and days (Figure 1D, H).

Representational drift occurs across visual cortical areas over timescales of seconds-minutes

To study the stability of visual representations over timescales of seconds to minutes, we analyzed data recorded using Neuropixels probes during the presentations of 'Natural Movie 1'. We divided each movie repeat into equal time bins and constructed a population vector (PV) of neuronal activity for each time bin (Figure S1A and Methods). We then calculated the correlation across the PVs of all time bins of all movie repeats (Figure 2A). We found higher PV correlations between the same time bins across movie repeats than between different time bins, indicating distinct and stable representation of the movie sequence (Figure 2A inset). The average PV correlation values between the same time bins on two different movie repeats capture the stability of the ensemble representation between these repeats (Figure 2B). Calculating the mean PV correlation as a function of the interval between movie repeats showed a significant gradual decline, indicating representational drift in all studied visual areas (Figure 2C-E). We found similar drift using a decoder that was trained to infer the time bin associated with a given activity pattern across movie repeats (Figure S1K, top panel).

Changes in neuronal tuning and activity rates underlie drift in visual representations

What cellular properties could underlie the observed representational drift? Time-dependent decline in PV correlations may stem from changes in cellular excitability (Figure 2F) or from changes in the tuning of individual neurons to the presented stimuli (Figure 2G). To test the contribution of each of these factors to the observed changes in PV correlations over time, we used two complementary measures: (1) 'Ensemble rate correlation': For each movie repeat, we constructed a single vector constituting the overall activity rates of each cell

in the recorded population. We then calculated the correlations across all pairs of these vectors, which captured the changes in the cells' activity rates, irrespective of their tuning to different time points along the movie (Figure S1B). (2) 'Tuning curve correlation': For each neuron, at each movie repeat, we constructed a vector representing its responsiveness to each time bin in the presented movie (i.e., its tuning curve) and then correlated the tuning curves for the same neurons across different movie repeats (Figure S1C). These analyses revealed a significant decline in the ensemble rate correlation and a modest, yet significant, decline in the tuning curve correlation values as a function of elapsed time in all studied visual areas (Figure 2H,I). Notably, the changes in the cells' activity rates were largely independent of changes in their tuning (Figure S1D-J). Overall, changes in both the cells' activity rates and tuning contributed to drift in visual representations over seconds-minutes.

Representational drift cannot be explained by changes in arousal state, visual adaption or recording instability

Could the observed representational drift merely reflect changes in behavioral state or global fluctuation in neuronal activity levels? Indeed, we found a mild drop in running speed, pupil area, and global neuronal activity rates after the first few movie repeats, potentially reflecting changes in arousal[46–55] or visual adaptation[56,57] (Figure S2A-D). Repeating our analyses while removing the first several movie repeats or excluding cells that showed a significant decrease in their activity rates throughout the block, revealed a significant gradual decline in both ensemble rate correlation and tuning curve correlation values as a function of time (Figure S2E-H). Furthermore, the distribution of the differences in activity rates of the same individual neurons between the beginning and end of each block was centered around zero (Figure S2I). Together, these analyses suggest that representational drift is not driven by a systematic decline in firing rates, changes in the behavioral state or visual adaptation.

To minimize the contribution of recording instability to our observations, we restricted our analysis to cells whose tuning curves were highly correlated across different blocks, which increased the likelihood of tracking the same cells within a given block. Here too, we found gradual changes in visual representations in all studied cortical areas (Figure S2J-O). Notably, we obtained similar results in the Ca^{2+} imaging dataset, further substantiating that the observed drift is not due to recording instability (Figure S1K,L).

Representational drift is continuous over timescales of tens of minutes to hours

To determine the degree to which visual representations change over timescales of tens of minutes, we analyzed the stability within and across blocks of movie presentations. We found higher correlations within a given block compared to between blocks in all measurements, brain areas and datasets (Figure 3A-D and Figure S3A-D). Furthermore, the decline in ensemble rate correlations was gradual across blocks of different natural movies (Figure S3E-J). Thus, visual representations change over the course of tens of minutes.

Could the stability of visual representations be affected by the complexity of the stimulus? While here we found drift in the representations of natural movies, previous studies have demonstrated that tuning to moving grating are relatively stable[16,37]. Therefore, we repeated our analyses on visual representations of drifting gratings (Figure S4A-C), and found higher ensemble rate correlation values between two temporally proximal blocks relative to those of two temporally distal blocks in all brain areas and datasets (Figure S4D,F). In contrast, in

most visual areas we found no such significant difference in the tuning curve correlation values as a function of time (Figure S4E,G). Thus, representations of drifting gratings change over the course of tens of minutes, but these changes are characterized by changes in the cells' activity rates rather than in their tuning.

Representational drift persists over timescales of days and weeks

The Ca^{2+} imaging dataset contains three imaging sessions per mouse, spanning multiple days (Figure 3E), which allows assessing the long-term stability of neuronal representations[58]. We first took a conservative approach and restricted our analysis to cells that were active in both compared time points (either within a session or across sessions). Similarly to our observations within a given day, we found a gradual decrease in correlations in all measurements and brain areas (Figure 3F-I and Figure S5A), consistent with previous results in V1[37]. Repeating our analyses with the cells found active in at least one of the time points we compared revealed an even more pronounced decline in the difference between sessions (Figure S5B). Time-dependent decline in ensemble rate correlations was also evident during blocks of spontaneous activity (i.e., without visual stimulation), implying that gradual changes in excitability drive drift in cell activity rates (Figure S5C). While ensemble rate correlations between pairs of sessions significantly decreased as a function of the number days between sessions in all visual areas, the tuning curve correlations showed only a modest trend (Figure 3J,K). Overall, these results suggest that representational drift is continuous over days.

Notably, the distribution of the mean activity rates, number of active cells, running speed and pupil area were similar across sessions (Figure S5D-G), suggesting that the observed drift cannot be explained by gross changes in the population response or animal arousal across sessions. There was also no consistent time-dependent decay in the performance of a within-day decoder, within-day PV correlation values or PV correlation values between different pairs of subsequent sessions, indicating that representational drift over days is not a result of a gradual deterioration in neuronal activity or tuning (Figure S5H-J). Importantly, our results were robust to the specific choice of Ca^{2+} event detection method (Figure S5K,L) or cell registration algorithm[58] (Figure S6A-J).

Representational drift occurs throughout different cortical layers and cell types

Our analysis thus far has focused on excitatory cells in different visual areas, irrespective of cortical layers. Repeating our analysis while grouping the data based on the depth of each field of view revealed a significant gradual decrease in the PV correlation values in all cortical layers (Figure 4A-C). We did not find significant differences in the rate of the drift across cortical layers (Figure 4C), which is surprising given the differences in their connectivity and computational roles.

Next, we replicated our analyses using the data from SST, VIP and Pvalb inhibitory Cre lines. In all inhibitory Cre lines, we could identify neurons that displayed reliable and distinct tuning curves across different movies repeats occurring on different days (Figure 4D-F). Similarly to our analyses of data from excitatory Cre lines, we found significant representational drift in interneurons of different visual areas, across timescales spanning seconds to days (Figure 4G-I). Thus, representational drift is not intrinsically related to a specific subset of cells or cortical layers.

Neural-code stability does not follow the hierarchy of information flow across areas

To what extent does the hierarchy of information flow across visual areas affect the stability of visual representations? To address this issue, we compared the stability between pairs of thalamic (dorsal LGN and LP) and cortical areas (V1 and LM). Brain areas within these pairs are anatomically adjacent and show similar degree of tuning reliability to natural movies, but are distinct with respect to their level in the hierarchical structure of the visual system[41,59]. We found that V1 was consistently less stable than the downstream area LM across all measured timescales (Figure 5A-H). Likewise, LGN showed faster drift than the downstream LP (Figure 5A,B). Thus, our results do not support the hypothesis that lower visual areas are more stable than higher areas.

The internal structure of neuronal activity differs across visual brain areas

How could the visual system generate consistent perception despite representational drift? Recent studies in the hippocampus have shown that the structure of the relationships between neuronal population activity patterns remains stable over days[60], and may confer perceptual constancy in the face of changing coding properties of individual neurons[32]. Consistent with this notion, applying dimensionality reduction on the PVs of all time bins from all movie repeats uncovered a highly-organized internal structure of population activity patterns (Figure 6A).

If the internal structure of neuronal activity reflects the computational processes undertaken by the network, then it should differ across brain areas according to their distinct computational roles[60]. Thus, we next asked to determine the degree to which the internal structure of neuronal population activity is distinct for each visual area, stereotypic across individuals, and stable over time. We calculated for each area the PV for each time bin within a movie repeat and then calculated the correlations across all the PVs. This yielded a matrix (time by time; Figure 6B) that represented the structure of similarities between the neuronal population activity patterns at different time points of the presented movie (i.e. the internal structure of neuronal population activity). Applying dimensionality reduction to the similarity matrices from all movie repeats and from all visual areas revealed highly separated clusters that corresponded to the different visual areas (Figure 6C). Therefore, the neuronal population activity of each visual area forms a distinct internal structure.

We next explored the extent to which such an organization genuinely stems from the intrinsic functional properties of each brain area[61–64] and to what extent it is susceptible to biases in the analysis (e.g., due to incidental differences in the coding properties of the sampled neurons). We divided the dataset into two equal groups of mice, and then pooled together the data from each group to create two independent 'pseudo-mice', taking the same number of cells for each visual area in both pseudo-mice (see Methods). Hence, the resultant pseudo-mice have an equal number of randomly-sampled neurons for all visual areas, with an order of magnitude more neurons per visual area compared to individual mice. Applying dimensionality reduction to the data from two example pseudo-mice, revealed well-separated clusters (Figure 6D) that correspond to the different visual areas, similarly to what we found in individual mice (Figure 6C). Notably, the clusters of the same brain area across two pseudo-mice resided relatively close to each other in the reduced space, suggesting that the internal structures are not only distinct between visual areas but also stereotypic across

mice. The differences across visual areas and similarities between different pseudo-mice were also apparent when we constructed pseudo-mice using datasets from the two recording techniques (Figure 6E), or when comparing the low-dimensional latent structures (Figure 6F).

To quantitatively assess how stereotypical the representations are in different visual areas, we trained a decoder to classify the identity of the recorded visual area across pseudo-mice based solely on the internal structure of neuronal activity. This procedure revealed excellent classifications for all brain areas, which were higher than shuffled data and chance (Figure 6G). As expected, the decoder's performance increased with the number of cells included in the analysis (Figure 6H and Figure S7A-C), and was not different from chance when the activity of individual cells was temporally shuffled (Figure S7D). Further, neuronal responses to 'Shuffled natural movie 1' organized into less distinct internal structures compared to the responses to 'Natural movie 1' (Figure S7E-F), suggesting that the internal structure of neuronal activity is affected by the spatiotemporal coherence of the presented visual stimuli. Moreover, we found good classifications in all visual areas, even when using the neuronal activity in response to 'Full-field drifting gratings' (Figure S7G), indicating that the differences between the internal structures across visual areas are due to area-specific differences in visual field representation. Thus, the internal structure genuinely reflects coding properties that are unique to each visual area, and is not a trivial reflection of the structure of similarities between different frames in the movie.

The internal structure of neuronal activity is maintained over time despite drift in the coding of the neurons supporting it

Finally, we examined whether the internal structure of neuronal activity is more stable over time than the activity rates and tuning of individual neurons (Figure 7A). We calculated for each brain area the change in the correlations between the internal structures, and compared it to the change in the PV correlation over days (Figure 7B). While the PV correlations decayed with time, the correlations between the internal structures remained stable (Figure 7C and Figure S7H,I). The structure of the tuning curves' pairwise similarities ('signal correlations') also drifted with time, consistent with our finding that individual cells gradually change their tuning across days (Figure S7J,K). Notably, the stability over time of the internal structure also depended on the size of the neuronal population, as including more cells in the analysis resulted in a more stable structure. Conversely, the change in PV correlation or the structure of pairwise similarities did not depend on the number of cells, consistent with measurements that treat cells independently (Figure 7D and Figure S7L). Overall, our results suggest that the internal structure of neuronal population activity of each visual area is distinct, stereotypic, and stable across time despite drift in the activity rates and tuning of individual neurons.

DISCUSSION

We found representational drift over timescales of minutes to days across the visual system. Surprisingly, our analysis does not support the hypothesis that primary (or lower) sensory areas display more stable coding than downstream (higher) areas[35]. If anything, our analysis shows that the coding stability of some cortical (V1 and LM) and subcortical (LGN and LP) areas exhibit an opposite trend with respect to their hierarchy. We further show that the

structure of the relationship between neuronal population activity patterns is stereotypic across mice and stable over time, pointing to a possible network mechanism that can reliably preserve visual information despite drift in the coding properties of individual neurons.

Our work joins a number of longitudinal studies that quantified coding stability in the visual cortex and adds to these studies in several aspects[16,36–38]. While most previous work focused on excitatory L2/3 neurons in V1 and emphasized variability in neuronal responses to synthetic stimuli[16,36,38], our analysis encompasses multiple different visual areas, cortical layers, and cell types and focuses on changes in neuronal representations of natural movies. Consistent with our results, a recent study in V1 found more pronounced changes in cells' tuning to natural movies relative to that of drifting gratings[65].

In the experiments analyzed here, neuronal responses to the exact same stimuli were recorded using both electrophysiology and Ca^{2+} imaging, which allowed us to validate the results and control for biases specific for each technique. The majority of our analyses showed highly similar results across the two datasets. However, small differences were found in few of the analyses (e.g. Figure 3D compared to Figure S3C, and Figure S4E compared to Figure S4G), which likely reflect the differences in sensitivity and resolution of the two technologies[66–69].

Our findings that changes in activity rates and tuning are largely independent (Figure S1D-J) and that ensemble rate correlations can gradually change over time even in the absence of a visual stimulation (Figure S5J) suggest that different mechanisms may govern distinct aspect of neuronal function, such as excitability[70–73] and synaptic connectivity[74]. Additionally, the existence of drift across minutes to days raises the possibility that different mechanisms drive drift on different timescales.

To generate consistent perception, the visual system must cope with changes in the coding of visual information[75,76]. It has been suggested that a system that carries a high-dimensional distributed code may maintain its functionality under representational drift by either confining the drift to the null space of the code, or via a compensatory plasticity of the downstream reader[29,32]. In both cases, the similarities across representations of different stimuli are expected to be somewhat conserved over time, even under a significant change in the representations themselves. Here, we demonstrate that the relationships between representations are stable over time, consistent with recent studies showing that a stable structure (manifold) of population activity resides in a variable or drifting high-dimensional neural activity space and may underlie a stable behavior[60,77–80].

Measuring coding stability is challenging because various factors could affect longitudinal recordings in a way that could lead to the appearance of drift, even if the neuronal activity itself is stable. For these reasons, we performed several control analyses which together suggest that our results are not due to recording instability (Figure S2J-O, Figure S3F,I and Figure S6). Importantly, we found that drift occurs even in the absence of an overt sign of changes in the behavioral state, and has different properties from habituation or adaptation[56,57]. There are of course behavioral variables that were not recorded in the experiments we analyzed that could affect neuronal responses in the visual cortex[46,81]. However, as long as such variables do not gradually change with time, they should not affect the stability of the visual representations.

Overall, taken together with other findings of drift in the hippocampus and other cortical areas[82], our results imply that representational drift is an inherent property of neural networks, and that population-level organization of information could contribute to robust, time-invariant representations despite drifting or variable coding at the level of individual neurons.

ACKNOWLEDGMENTS

Y.Z. is the incumbent of the Daniel E. Koshland Sr. Career Development Chair. Y.Z. is supported by grants from the Abraham and Sonia Rochlin Foundation, Hymen T. Milgrom Trust, Israel Science Foundation (grant 2113/19), Human Frontier Science Program, and European Research Council (ERC-CoG 101001226). We thank Timothy O'Leary, Ofer Yizhar, Rafi Malach, Ivo Spiegel, Michal Rivlin, Jerome Lecoq, Michael Rule, Meytar Zemer, Liron Sheintuch, Eyal Bitton, Maya Salomon, Alice Eldar, Ofer Givton, and Noa Eren for helpful advice and comments on the manuscript.

AUTHOR CONTRIBUTIONS

Conceptualization, D.D., A.R. and Y.Z.; Software & Formal Analysis, D.D.; Writing, D.D., A.R., and Y.Z.; Review & Editing, D.D., A.R., and Y.Z.; Supervision, A.R. and Y.Z.

DECLARATION OF INTERESTS

The authors declare no competing interests.

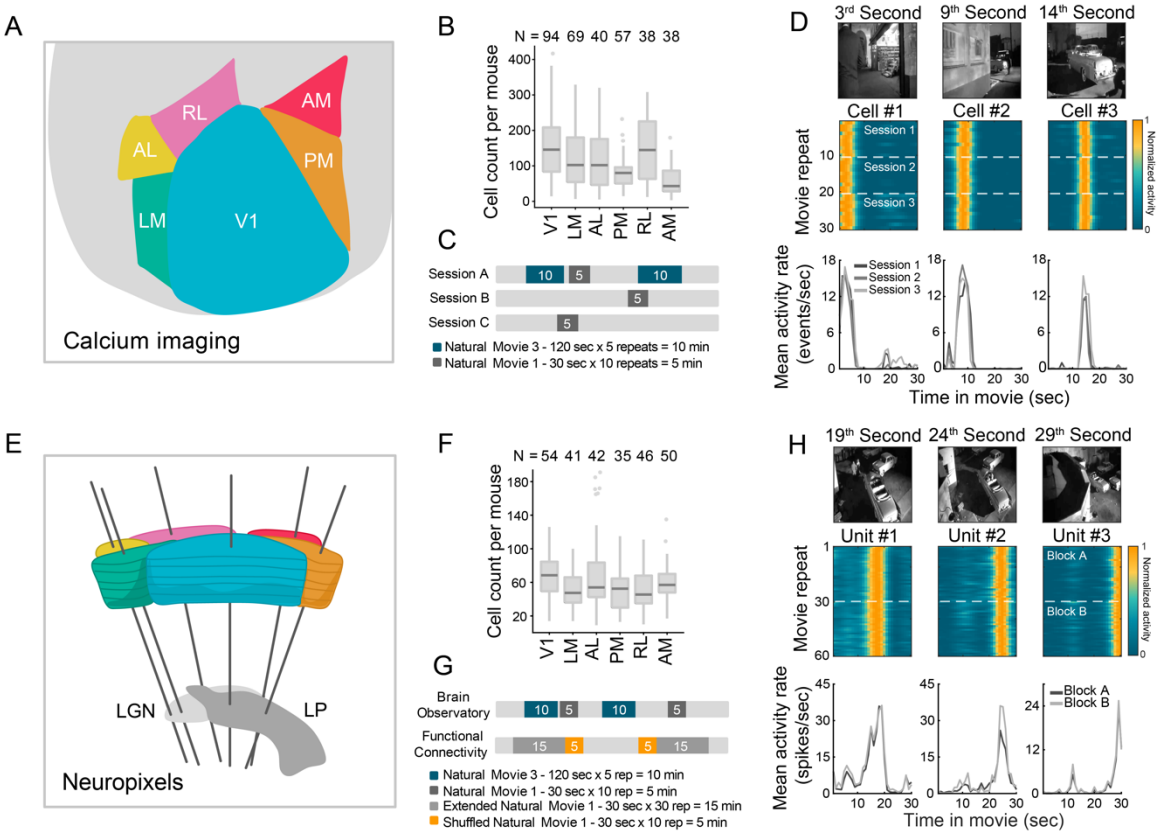

Figure 1. Neurons recorded from various visual cortical areas show reliable tuning to natural movies. (A-D) Ca^{2+} imaging dataset. (A) Schematic of the different brain areas imaged using two-photon Ca^{2+} imaging. V1 - primary visual area, LM – lateral-medial visual area, AL - anterolateral visual area, PM - posteromedial visual area, RL - rostrolateral visual area, AM - anteromedial visual area. (B) Distribution of cell counts per mouse across brain areas for the Ca^{2+} imaging dataset. (C) Experimental design. Each mouse performed three sessions in a random order, separated by a different number of days. Indicated stimuli ('Natural movie 1' and 'Natural movie 3') were used in our main analyses (see Methods). (D) Responses of three cells across different 'Natural movie 1' repeats spanning three sessions. (E-H) Neuropixels dataset. (E) Schematic (adapted from Siegle et al. (2021)) of the different brain area recordings using Neuropixels probes. LGN – lateral geniculate nucleus. LP – lateral parietal nucleus. (F) Cell counts per mouse across brain areas for the Neuropixels dataset. (G) Experimental design. Thirty-two of the mice performed the 'Brain Observatory' battery and 26 performed the 'Functional Connectivity' battery. Indicated stimuli ('Natural movie 1', 'Natural movie 3', and 'Shuffled natural movie 1') were used in our analyses. (H) Responses of three cells across different 'Natural movie 1' repeats spanning two blocks within the same session. Box plots in panels B and F show the data range (whiskers), 25_{th} and 75_{th} percentiles (box), and median (dark line). Outliers are marked by gray dots.

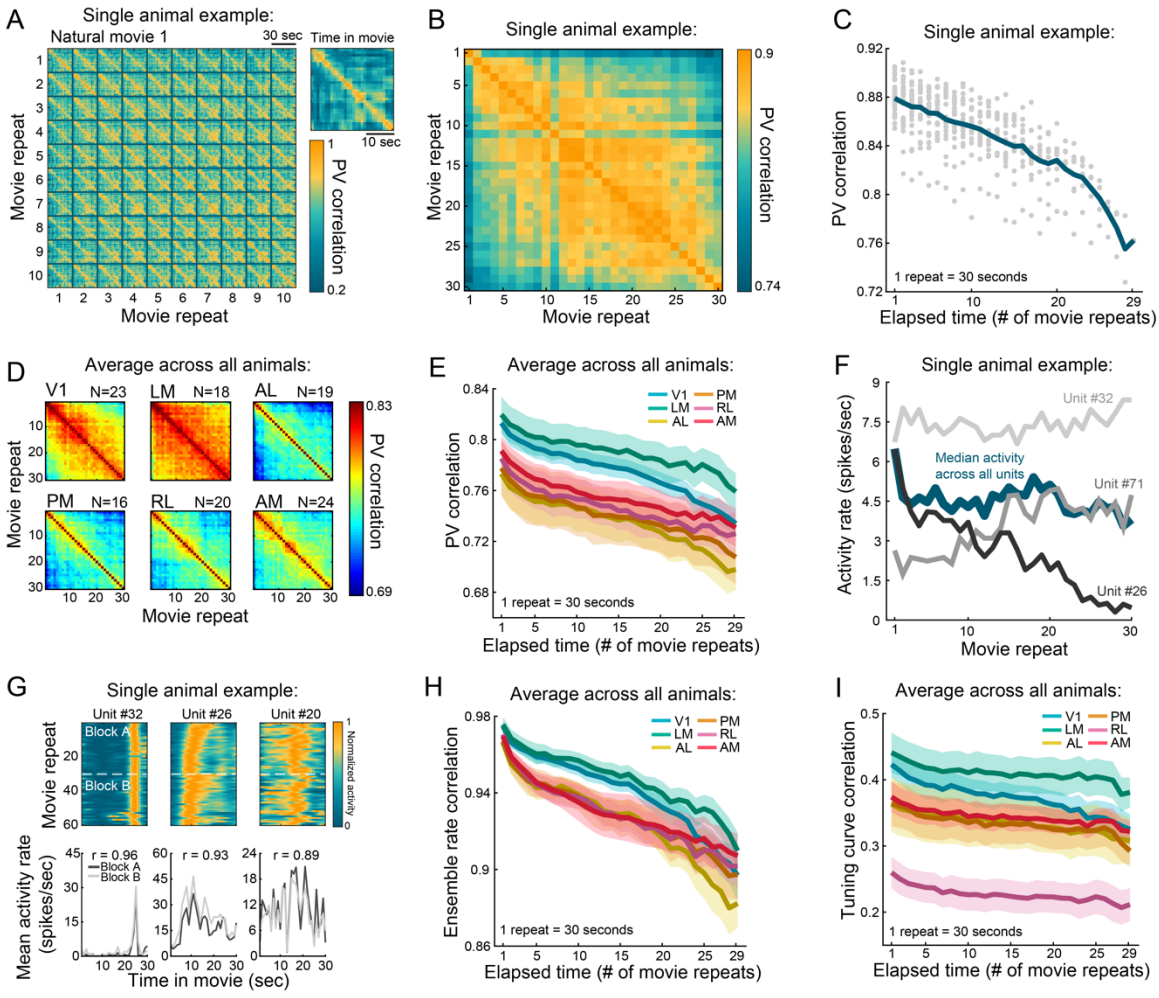

Figure 2. The visual cortex exhibits representational drift across subsequent presentations of a natural stimulus over timescales of seconds-minutes.
(A-I) Analyses using data from the Neuropixels 'Functional Connectivity' group during the presentation of 'Natural movie 1'. (A) PV correlation between the first 10 (out of 30) movie repeats of the first block, recorded from area PM of a representative mouse. Inset: average PV correlation over all pairs across different movie repeats. (B) Mean PV correlation for each pair of movie repeats from the same mouse shown in A. For visualization, the diagonal was set to the maximal value. (C) Mean PV correlation as a function of time. Each data point represents the mean PV correlation value for a single pair of movie repeats from B. (D) Mean PV correlation between movie repeats across animals and brain areas. (E) PV correlation as a function of time. (F) Mean activity rates for three units from area PM of the same representative mouse across movie repeats. (G) Responses of three V1 cells from the same mouse across different movie repeats, spanning two blocks within the same session. (H-I) Ensemble rate (H) and tuning curve (I) correlation across animals as a function of time. All visual areas showed a significant decrease in PV, ensemble rate and tuning curve correlations as a function of time ($\chi^2_{(28)} \geq 92.33$, $p<10^{-3}$, Friedman's tests with Holm–Bonferroni correction). Data in E, H and I are mean ± SEM across mice. See also Figures S1 and S2.

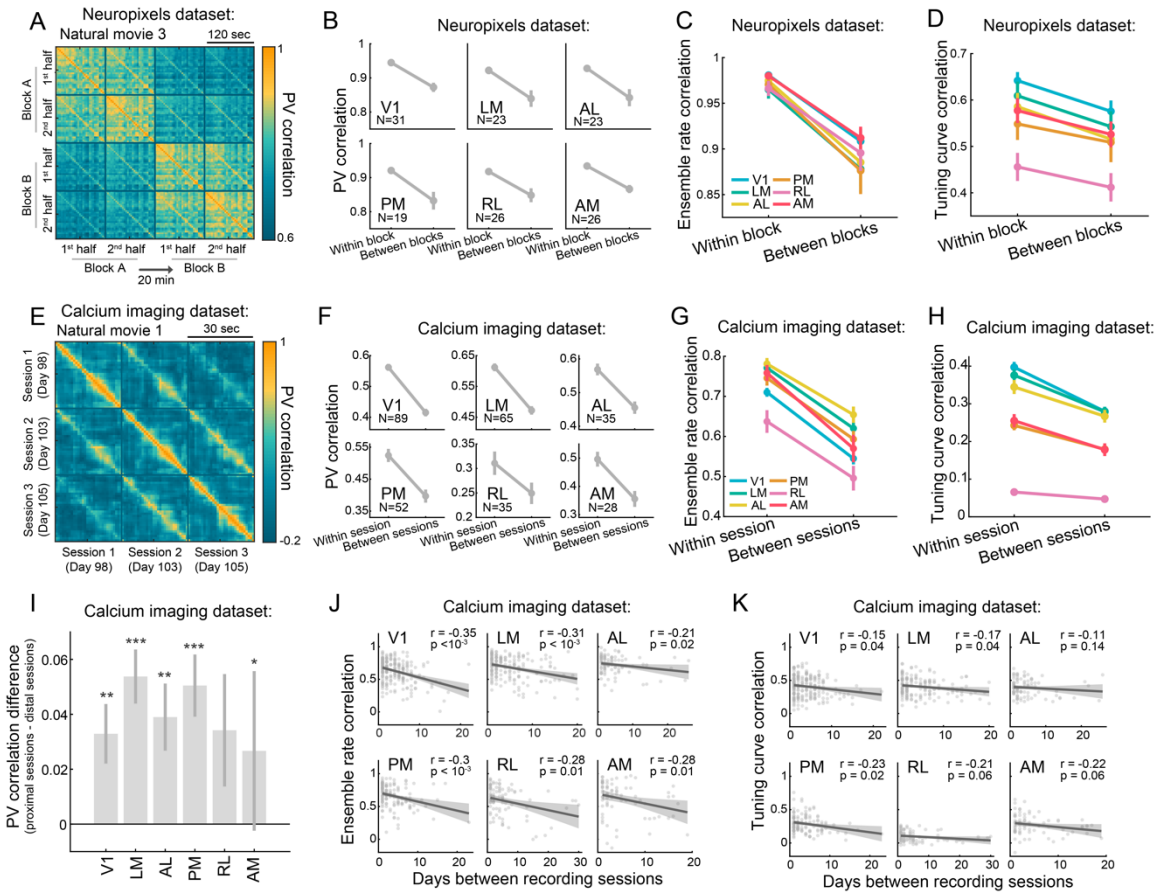

Figure 3. Visual representations gradually change over timescales of minutes-days.
(A-D) Data from the Neuropixels 'Brain Observatory' group during the presentation of 'Natural movie 3'. (A) PV correlation between the 1st (repeats 1-2) and 2nd (repeats 3-5) halves of two different blocks of 'Natural movie 3' in a single visual area. The presented example is the average matrix across mice in area LM. (B-D) PV (B), ensemble rate (C) and tuning curve (D) correlations between the two halves of the same block and between halves of different blocks. For all measurements and areas p≤0.002, two-tailed Wilcoxon signed-rank test with Holm–Bonferroni correction. (E-K) Ca^{2+} imaging dataset during the presentation of 'Natural movie 1'. (E) PV correlation between three different sessions from a representative mouse recorded in V1. The age of the mouse (in days) is indicated in parenthesis. (F-H) PV (F), ensemble rate (G) and tuning curve (H) correlations between the two halves of the same session and between halves of different sessions. For all measurements and areas $p<10^{-3}$, two-tailed Wilcoxon signed-rank test with Holm–Bonferroni correction. (I) The difference between the PV correlation of two temporally proximal sessions and that of two temporally distal sessions (V1($Z=3.35$,$p=0.001$), LM($Z=4.64$,$p<10^{-3}$), AL($Z=2.85$,$p=0.006$), PM($Z=3.92$,$p<10^{-3}$), RL($Z=1.38$,$p=0.083$), AM($Z=1.99$,$p=0.046$), one-tailed Wilcoxon signed-rank test with Holm–Bonferroni correction; *p<0.05, **p<0.01, ***p<0.001). (J-K) Ensemble rate (J) and tuning curve (K) correlation as a function of the number of days between sessions. Each mouse is represented by 2-3 data points, corresponding to different intervals between sessions, with a regression line of ± CI of 95% (one-tailed Pearson's correlation with Holm–Bonferroni correction). Data in B-D and F-I are mean ± SEM across mice. See also Figures S3-S6.

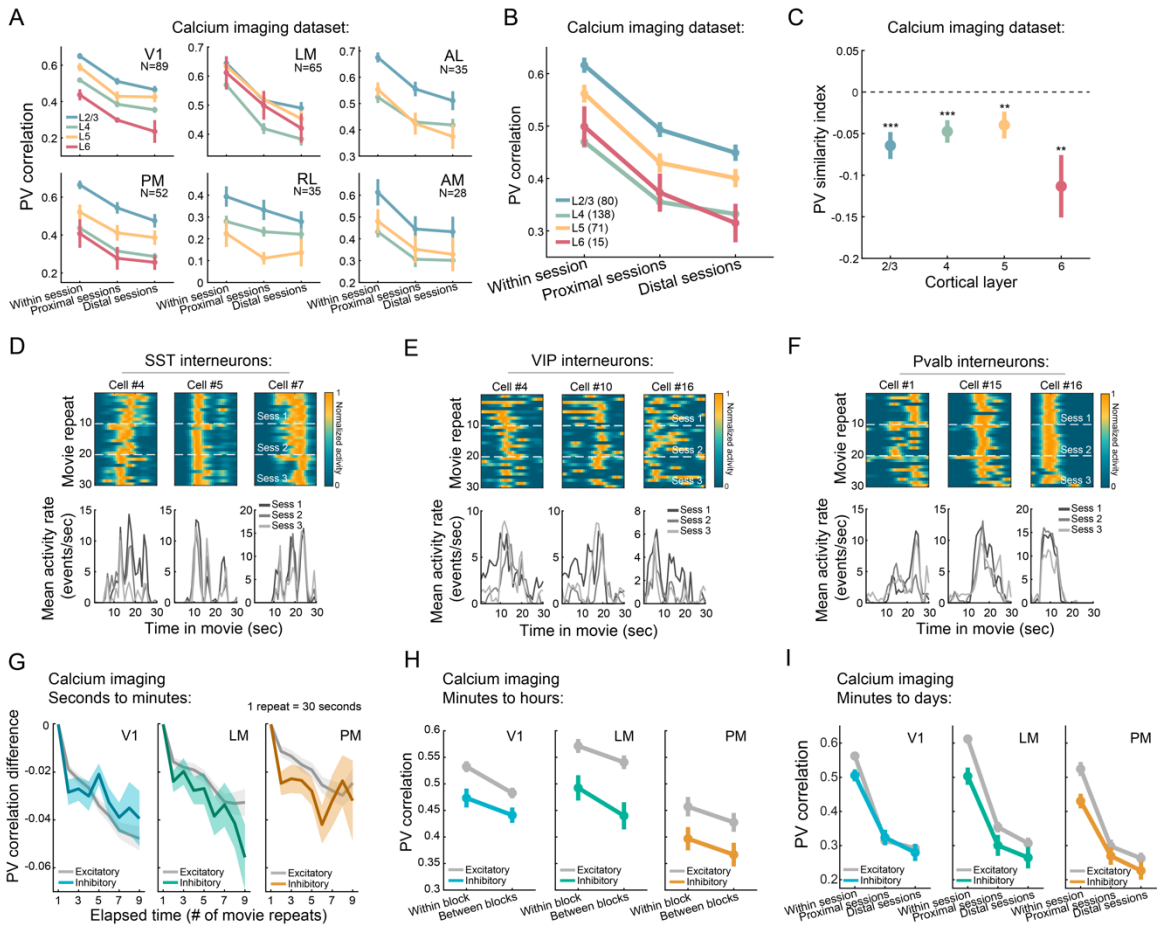

Figure 4. Representational drift across multiple cortical layers and cell types.
(A-C) Data from the Ca^{2+} imaging dataset during the presentation of 'Natural movie 1'. (A-B) PV correlation between the two halves of the same session, between halves of two temporally proximal sessions and between halves of two temporally distal sessions grouped based on cortical layers for each of six visual areas (A) or after pooling across all mice and visual areas (B). Colors indicate different cortical layers. (C) Normalized difference between the PV correlation of proximal sessions and distal sessions for all layers. The difference in PV correlations between proximal sessions and distal sessions was significant for all layers ($p \leq 0.0062$, one-tailed Wilcoxon signed-rank test with Holm–Bonferroni correction). There was no significant difference in the PV similarity indices between the different layers (two-tailed Mann-Whitney rank sum tests with Holm–Bonferroni correction). (D-I) Analyses were done using both the excitatory and inhibitory Cre lines from the Ca^{2+} imaging dataset during the presentation of 'Natural movie 1'. (D-F) Responses of three SST (D), VIP (panel E), and Pvalb (F) example cells from area V1 across different movie repeats spanning three sessions. (G) Difference in PV correlation as a function of time for the inhibitory (colored) and excitatory (gray) Cre lines imaged from areas V1, LM and PM; All areas in the inhibitory Cre lines showed a significant decrease in PV correlations as function of time ($\chi^2_{(8)} \geq 24.19, p \leq 0.002$, Friedman's tests with Holm–Bonferroni correction). (H) PV correlation between the two halves of the same block and between halves of different blocks for the inhibitory (colored) and excitatory (gray) Cre lines; $p \leq 0.005$ for all areas in the inhibitory Cre lines, two-tailed Wilcoxon signed-rank test with Holm–Bonferroni correction. (I) PV correlation between the two halves of the same session, between halves of two temporally proximal sessions and between halves of two temporally

distal sessions for the inhibitory (colored) and excitatory (gray) Cre lines. The difference in PV correlations between proximal sessions and distal sessions was significant for all areas in the inhibitory Cre lines ($p \leq 0.01$, one-tailed Wilcoxon signed-rank test with Holm–Bonferroni correction). Data in A-C and G-I are mean ± SEM across mice.

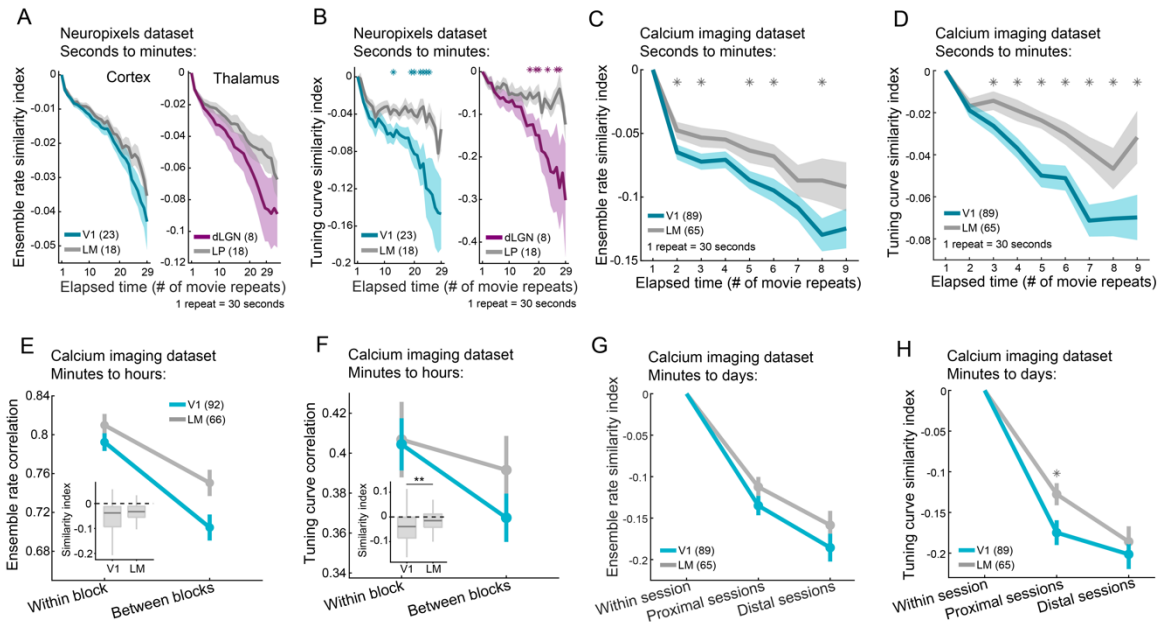

Figure 5. Comparison of representational drift between lower and higher visual areas. (A-B) Data from the Neuropixels 'Functional Connectivity' group during the presentation of 'Natural movie 1'. Normalized difference in ensemble rate (A) and tuning curve correlations (B) as a function of time for dLGN, LP, V1 and LM; significant differences in similarity indices between the different pairs of visual areas are indicated by colored asterisks (purple for dLGN compared to LP, and blue for V1 compared to LM; *p≤0.05, two-tailed Mann–Whitney rank-sum tests). (C-H) Data from the Ca^{2+} imaging dataset during the presentation of 'Natural movie 1' (C,D, G, H) or 'Natural movie 3' (E,F). (C-D) Ensemble rate (C) and tuning curve (D) similarity index as a function of time for V1 and LM. The difference in similarity indices between V1 and LM; *p≤0.05, two-tailed Mann–Whitney rank-sum tests. (E-F) Ensemble rate (E) and tuning curve (F) correlation between the two halves of the same block and between halves of different blocks. Inset: distribution of normalized difference in correlations between 'within block' and 'between blocks' for V1 compared to LM; two-tailed Mann–Whitney rank-sum test for ensemble rate (Z=1.63,p=0.101) and tuning curve (Z=2.6,p=0.009) correlations. (G-H) Normalized difference in ensemble rate (G) and tuning curve (H) correlations between the two halves of the same session, between halves of two temporally proximal sessions and between halves of two temporally distal sessions. The difference in similarity index between V1 and LM; *p≤0.05, two-tailed Mann–Whitney rank-sum tests. Data are mean ± SEM across mice. The number of mice is indicated in parentheses.

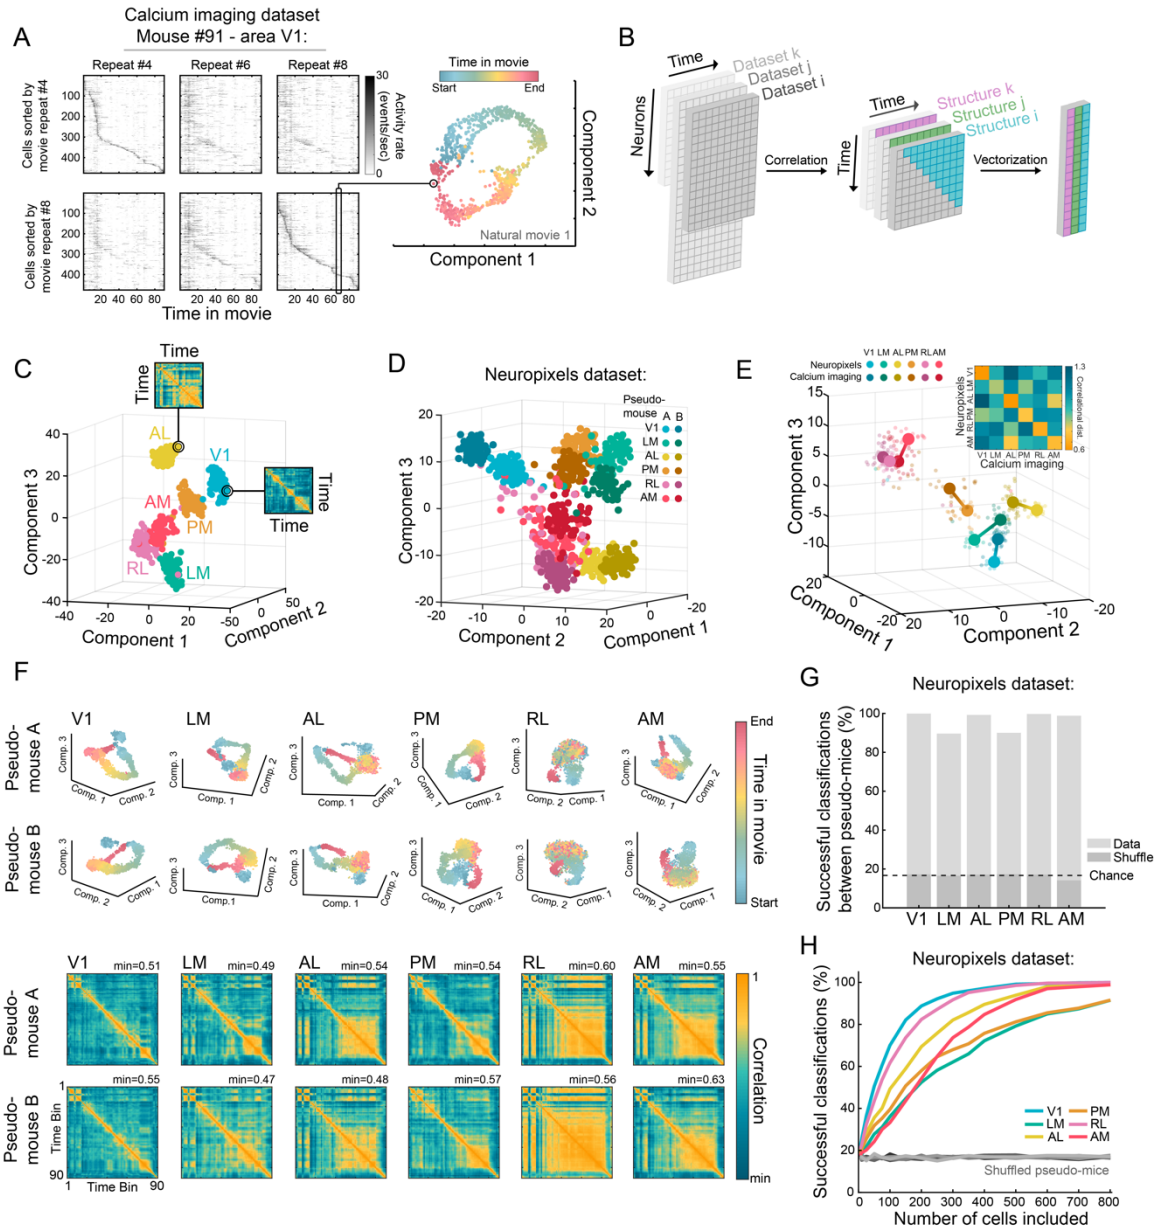

Figure 6. The internal structure of neuronal activity of each visual area is distinct and stereotypic across mice. (A) Dimensionality reduction (tSNE) of population activity of a single mouse recorded from area V1 (left) recovers a low-dimensional structure (right). Each point represents a single time-point of population activity of a single movie repeat, colored according to time in the presented movie. (B) Workflow for the extraction of the internal structure from the neuronal population responses. (C) Dimensionality reduction applied to the internal structures of different visual areas from a single representative mouse recorded via Neuropixels. Each data point corresponds to an internal structure of a single 'Natural movie 1' repeat. Insets: example of internal structures from area V1 (blue; repeat #49) and area AL (yellow; repeat #37). (D) Example of a dimensionality reduction on the internal structures of 'Natural movie 1' produced from two Neuropixels 'pseudo-mice'. Each data point corresponds to an internal structure of a single movie repeat. (E) Dimensionality reduction applied to the internal structures from different brain areas of two 'pseudo-mice' created using data from all

the mice of each dataset (Neuropixels and Ca^{2+} imaging). Each data point corresponds to an internal structure of a single repeat of 'Natural movie 1'. Large circles indicate the centers of mass of the data points for each area, with a line connecting between Neuropixels and Ca^{2+} imaging datasets. Inset: a correlation-distance matrix between the internal structure of each area across recording techniques. (F) Top: Dimensionality reduction on the population activity of two example pseudo-mice recorded using Neuropixels probes recovers a distinct low-dimensional structure for each visual area. Each point represents a single time-point of population activity of a single 'Natural movie 1' repeat. Bottom: Structure of similarities between trial averaged population activity of each visual area for the same example pseudo-mice shown in the top panels. (G-H) Percentage of successful classifications of the internal structures to their corresponding visual areas across pairs of Neuropixels pseudo-mice (G; data are mean across 1000 pairs of pseudo-mice), and as a function of the number of cells included in the analysis (H; data are mean across n=2000 iterations). See also Figure S7.

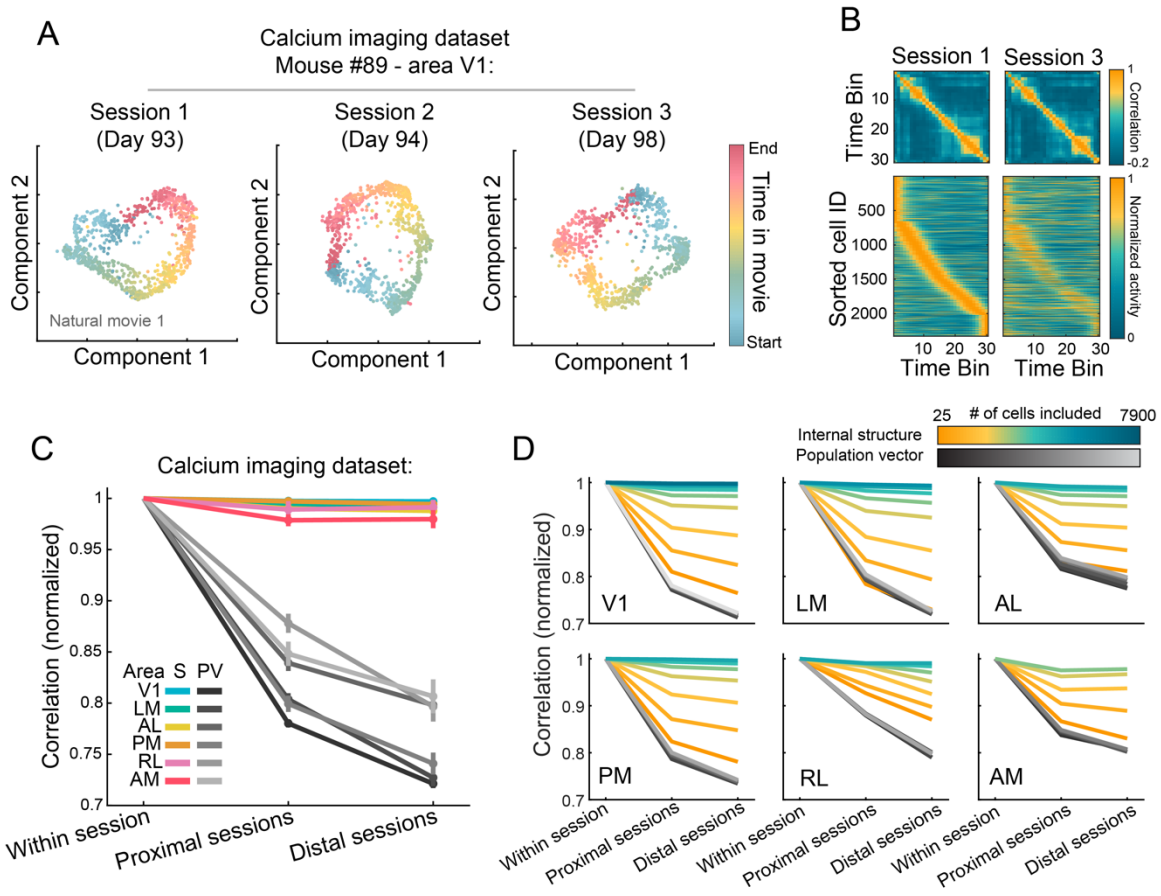

Figure 7. The internal structure of neuronal activity is stable over time. (A-D) Data from the Ca^{2+} imaging dataset during the presentation of 'Natural movie 1'. (A) A similar low-dimensional structure to the one shown in Figure 6A is seen in a different example mouse recorded from area V1. (B) While the internal structure in 'pseudo-area AL' is maintained across imaging sessions (top panels), the individual neurons whose activity patterns underlie the same internal structure drift across sessions (bottom panels). (C) Normalized correlation between the internal structures (colored lines) or the PVs (gray lines) between the two halves of the same session, between halves of two temporally proximal sessions and between halves of two temporally distal sessions. Data are mean ± SD across 1000 different pseudo-mice realizations. (D) Same as in panel C only colored according to the number of neurons included in the analysis. Data are mean across 1000 different pseudo-mice realizations . Correlations in panels C and D were normalized to the value of the 'within session' correlation. See also Figure S7.

STAR ★ METHODS

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources should be directed to the Lead Contact, Yaniv Ziv (yaniv.ziv@weizmann.ac.il).

Materials availability

This study did not generate new unique reagents.

Data and code availability

- This paper analyzes existing, publicly available data which is available via the AllenSDK at: https://allensdk.readthedocs.io
- Processed data and code supporting the current study is deposited in a GitHub repository and will be publicly available as of the date of publication at: https://github.com/zivlab/visual_drift.
- Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

We analyzed data from the publicly available Allen Brain Observatory: two-photon Ca^{2+} imaging[40] and electrophysiology (Neuropixels) datasets[41]. We used the default functions in Allen software development kit (AllenSDK) package to download the raw Neurodata Without Borders (NWB) files containing the neuronal and behavioral data from the relevant experiments. Their full data collection methodology can be found in the white paper (https://observatory.brain-map.org/visualcoding). Briefly, in the Ca^{2+} imaging dataset, 216 transgenic mice expressing GCaMP6f in laminar-specific subsets of cortical pyramidal neurons underwent intrinsic signal imaging to map their visual cortical regions before cranial windows were implanted above the desired visual region. Mice were habituated to head fixation before the three imaging sessions, in which they were shown a battery of natural scenes, natural movies, locally sparse noise, or gratings. In the Neuropixels dataset, 30 C57BL/6J wild-type mice and 28 mice from three transgenic lines (N = 8 Pvalb-IRES-Cre x Ai32, N = 12 Sst-IRES-Cre x Ai32, and N = 8 Vip-IRES-Cre x Ai32) were implanted with up to six Neuropixels silicone probes each. The dataset contains simultaneous recordings from up to 8 cortico-thalamic visual areas (as well as nearby regions, such as CA1). During each recording session, mice passively viewed a battery of natural and artificial stimuli, depending on their experimental group.

METHOD DETAILS

Data analysis and data exclusion.

Analysis was carried out using the AllenSDK package default functions (for data curation) and custom-written MATLAB scripts (for data analysis). In the Ca^{2+} imaging dataset, we used all available excitatory and inhibitory Cre lines, including all layers and brain areas. All

Ca^{2+} imaging dataset analyses were performed using the excitatory Cre lines unless stated otherwise. The dataset is structured into 'experiment containers' that group recordings from three different imaging sessions of the same field of view (FOV). We considered each such container as an individual mouse. We included only mice that passed a fixed criterion of at least 20 recorded cells in the compared time-points. Specifically, in the within-block and between-days analyses, we included only mice with at least 20 recorded cells in each of the three imaging sessions, and in the between-blocks analysis, mice with at least 20 recorded cells within the same session ('Session A'). In the analyses comparing different cortical layers, we employed the same inclusion criteria described above while grouping the mice based on the imaging depth of their FOV regardless of the identity of the recorded visual area (Layer 2/3: 150-250μm; Layer 4: 265-350μm; Layer 5: 365-500μm; Layer 6: 550-700μm). In the analyses comparing between inhibitory and excitatory Cre lines (Figure 4G-I), due to the relatively low number of cells in an average FOV of the inhibitory Cre lines mice, we included in the analyses all the mice from the inhibitory Cre lines while maintaining the same inclusion criteria described above for the excitatory Cre lines. For the analysis of the Neuropixels dataset, we used the AllenSDK package default functions to filter out units that were likely to be highly contaminated or missing many of spikes and to retrieve the relevant unit's identity according to its corresponding manually labeled brain areas. This resulted in excluding units with ISI violations larger than 0.5, an amplitude cutoff larger than 0.1 and a presence ratio smaller than 0.9 (Siegle et al., 2021). We then included in all analyses only data from areas with at least 15 recorded units.

Detection of Ca^{2+} events.

Neuropil-corrected fluorescence change ($\Delta F_{(t)}/F_0$) traces for each cell were extracted using automated, structural region of interest (ROI)-based methods. The full procedure appears in de Vries et al. (2020). We performed no further preprocessing on the $\Delta F_{(t)}/F_0$ traces after downloading them with the AllenSDK. We identified Ca^{2+} events by searching each trace for local maxima that had a peak amplitude higher than four times the trace absolute median while including only the frames that showed an increase in Ca^{2+} transients relative to their previous frame. All the $\Delta F_{(t)}/F_0$ values in the frames that passed the assigned filters were set to the value of 1, and the rest to a value of 0. It should be noted that significant representational drift across days was also evident when using the $\Delta F_{(t)}/F_0$ traces (Figure S5K,L).

Registration of cells across sessions.

We used each cell's match labels across sessions as provided by Allen Brain Institute in each experiment's NWB file. Briefly, an algorithm that combines the degree of spatial overlapping and closeness between the ROIs of different cells was used to create a unified label for each cell across all three sessions. The full registration procedure appears in de Vries et al., (2020). To verify that observed drift did not stem from errors in the specific cell registration algorithm used to create the cell's match labels across sessions, we used an additional, independent probabilistic cell registration algorithm[58]. This method models the distribution of centroid distances for neighboring cells from different recording sessions (candidates for being the same cell) as a weighted sum of the distributions of two subpopulations: same cells and different cells (Figure S6). Then, based on the model that best fits the data, the method estimates the probability of each candidate in the dataset to be the same cell (Psame). This allows estimating the overall rates of false-positive errors

(different cells falsely registered as the same cell) and false-negative errors (the same cell falsely registered as different cells), providing a Psame registration threshold that is optimized to the dataset of each mouse. The threshold used for registration controls the tradeoff between false-positive and false-negative errors. Therefore, we chose a registration threshold of Psame=0.5, which constitutes an appropriate balance between false-positive and false-negative registration errors[58]. In cases with multiple candidates that cross the registration threshold, only the pair with the highest Psame was registered as the same cell, thus avoiding some of the false positive errors, the result of which was lower percentages of false-positive errors than those estimated by the probabilistic model.

Visual stimuli.

Natural movies. For our main analyses, we used data during the presentation of 'Natural movie 1' (30-second clip) and 'Natural movie 3' (120-second clip) stimuli from the Allen Brain Observatory paradigm. In the Ca^{2+} imaging dataset, 'Natural movie 1' was presented across all three imaging sessions (ten repeats per session). 'Natural movie 3' was presented only in one of the sessions (session A), with ten repeats spanning two blocks (five repeats in each block). In the Neuropixels dataset, 'Natural movie 1' was presented with either 60 repeats spanning two blocks (30 repeats in each block) for the 'Functional Connectivity' group, or with 20 repeats spanning two blocks (10 repeats in each block) for the 'Brain Observatory' group. 'Natural movie 3' was presented with ten repeats spanning two blocks (five repeats in each block) only for the 'Brain Observatory' group.

Temporally shuffled natural movie. In addition to 'Natural movie 1', mice from 'Functional Connectivity' group in the Neuropixels dataset were also presented with a temporarily shuffled version of the 'Natural movie 1' stimulus (termed 'Shuffled natural movie 1'). The random shuffle of frames was performed only once resulting in the same sequence of frames presented across all mice and movie repeats (total of 20 movie repeats spanning two blocks). The relevant analyses are presented in Figure S3H-J and Figure S7E-F.

Full-field drifting gratings. In addition to natural movies, all mice in the calcium imaging dataset and mice from the 'Brain Observatory' group in the Neuropixels dataset were also presented with full-field drifting gratings. The drifting gratings were presented for 2 seconds, followed by a 1-second inter-stimulus interval (grey screen), with a spatial frequency of 0.04 cycles/deg, 80% contrast, 8 directions (0°, 45°, 90°, 135°, 180°, 225°, 270°, 315°) and 5 temporal frequencies (1, 2, 4, 8, and 15 Hz). Each combination of direction and temporal frequency (total of 40 combinations) was presented 15 times, resulting in a total of 600 drifting grating, divided and presented randomly for each mouse across three different blocks. We restricted our analysis only to the neuronal activity during the 2-second time window of gratings presentation. The relevant analyses are presented in Figure S4 and Figure S7G.

Population vector correlation.

To determine the level of similarity between visual representations of the same stimulus on different presentations, we calculated for each mouse the population vector correlation between pairs of different movie repeats. First, we divided each movie repeat into 30 equal time bins (each bin spanning 1 second in 'Natural movie 1' and 4 seconds in 'Natural movie 3'). Then, for each temporal bin, we defined the population vector as the activity rate for each cell/unit. We calculated the Pearson's correlation between the population vector (PV

correlation) in one repeat with that of all temporal bins in another movie repeat and averaged the correlations over all pairs of corresponding time bins. For the between-blocks analysis (Figure 3B and Figure 4H), we created two sets of PVs for each of the two blocks (each block consists of 5 movie repeats): one set of PVs from the average activity of the first two 'Natural movie 3' repeats (repeats 1-2), and a second set of PVs from the average activity of the last three repeats (repeats 3-5). We than calculated the PV correlation across the four sets of vectors of both blocks (as described above) and measured the difference between the correlations within blocks and across blocks. The mean correlations between the two sets of PVs of the same blocks capture the 'within-block' stability, and the mean correlations between different blocks, the 'between-blocks' stability. The between-days analysis was done similarly to the between blocks analysis with minor changes: For each 'Natural movie 1' session, two sets of PVs were calculated, one set of vectors from the average activity of the first five 'Natural movie 1' repeats (repeats 1-5) and a second set of vectors from the average activity of the last five movie repeats (repeats 6-10). We then calculated the PV correlation between each pair of PVs sets, including only cells that were active in both compared time-points, and calculated the difference in PV correlations within sessions and across session. The mean correlations between the two sets of PVs of the same session capture the 'within-session' stability, and that between different sessions, the 'between-sessions' stability. For the analysis shown in Figure 3I, Figure 4C and Figure S5J, PV correlations were calculated after averaging the activity rates over all movie repeats in each session. For the drifting gratings stimulus (Figure S4), we calculated the average activity rate of each cell for each combination of direction and temporal frequency, resulting in a set of 40 population vectors for each of the three drifting grating blocks. We then calculated the Pearson's correlation between the vectors of different blocks and averaged the correlation values across all corresponding temporal frequencies for each orientation difference.

Ensemble rate correlation.

To quantify the similarities in activity patterns between different presentations of the same stimulus (regardless of the specific tuning of each neuron), we calculated for each mouse the ensemble rate correlation between pairs of different movie repeats. First, we calculated the overall activity rate for each neuron in each movie repeat. We then calculated for each pair of movie repeats the ensemble rate correlation as the Pearson's correlation between their vectors of activity rates. As in the PV correlation analysis, the differences in ensemble rate correlation for within and between blocks (or sessions) were calculated after averaging the activity rates of individual neurons over the first and second halves of movie repeats in each block (or session). For the analysis shown in Figure 3J and Figure S5A, ensemble rate correlations were calculated after averaging the activity rate over all movie repeats in each session. For the drifting gratings stimulus (Figure S4D,F), we calculated the overall activity rate for each neuron across all combinations of directions and temporal frequencies presented within the same block, resulting in a single vector for each of the three blocks. We then calculated for each pair of blocks the Pearson's correlation between their vectors of activity rates and tested whether the correlations of two proximal blocks were different than the correlation of two distal blocks.

Tuning curve correlation.

To quantify the similarities in the tuning preference of individual neurons across different

presentations of the same stimulus (regardless of changes in activity rates), we calculated for each neuron the tuning curve correlation between different movie repeats. As in the PV correlation analysis, we first divided each movie repeat into 30 equal time bins (each bin spanning 1 second in 'Natural movie 1' and 4 seconds in 'Natural movie 3'). Then, for each neuron, we defined the tuning curve as the mean activity rate in each temporal bin within the movie. We calculated the Pearson's correlation between the tuning curve of each individual neuron in one movie repeat and that of the same neuron in another movie repeat and used the median value across all neurons to capture the central tendency of the entire population. As in the PV correlation analysis, the differences in tuning curve correlation for within and between blocks (or sessions) were calculated after averaging the activity rates for the first and second halves of movie repeats in each block (or session). For the analysis shown in Figure 3K and Figure S5A, tuning curve correlations were calculated after averaging the activity rate of each individual neuron over all movie repeats in each session. Due to the sparseness of neuronal responses in the Ca^{2+} imaging dataset, we used the mean value across all cells (instead of the median) when computing the tuning curve correlation between individual movie repeats (Figure S1L and Figure 5D). For the drifting gratings stimulus (Figure S4D,F), we defined for each neuron the tuning curve as the mean activity rate for each of the 40 combinations of directions and temporal frequencies within the same block. Then, we calculated the Pearson's correlation between the tuning curve of each individual neuron in one block and that of the same neuron in another block and used the median value across all neurons to capture the central tendency of the entire population. Finally, we tested whether the tuning curve correlations for two proximal blocks were different than the correlation for two distal blocks.

Relationship between rate and tuning stability (Related to Figure S1G-J).
To assess how changes in the activity rates are related to changes in the cells' tuning curves, we first examined the linear relationship between the three chosen measurements of stability (PV correlation, ensemble rate correlation and tuning curve correlation) in terms of explained variance. In the short-timescale analysis (between movie repeats within a block) we used data of mice from the Neuropxiels 'Functional Connectivity' during the presentation of 'Natural movie 1'. For each mouse, we calculated the PV correlation, ensemble rate correlation and tuning curve correlation between pairs of movie repeats within the same block. This procedure resulted in three matrices, each is symmetric and 30-by-30 in size (capturing the similarities across pairs of movie repeats). Next, we employed a set multiple linear regression models to calculate for each mouse the coefficient of determination (R^2) as an estimate of the fraction of variation in the dependent variable that can be explained by the variation in the independent variables. The first model quantified the fraction of variance in the PV correlation values (dependent variable) explained by both the ensemble rate and tuning curve correlation values (independent variables) of the same mouse. The second and third models quantified the fraction of variance in the PV correlation values explained by the values of either the ensemble rate (second model) or tuning curve correlations (third model). The fourth model quantified the fraction of variance in the ensemble rate correlations values explained by the tuning curve correlations values. Since all measurements are affected by the interval of time between exposures, we ran each model using only the correlation values with the same interval between movie repeats (intervals of 1-20 repeats) and averaged the resultant R^2 values across all intervals and blocks. Similar results are obtained when no such procedure is employed (data not

shown). For the long-timescale analysis (between different blocks occurring on different days), we used mice from the calcium imaging dataset during the presentation of 'Natural movie 1'. For each mouse, we calculated the PV correlation, ensemble rate correlation or tuning curve correlation between all halves of all three recording sessions, resulting in three symmetric, 6-by-6 matrices (for the three measurements) capturing the similarities across pairs of session halves. Then, for each mouse, we used the vectored upper-half of the symmetric matrices and applied the same set multiple linear regression models as described above. Here, we could not control for the interval of time between comparisons due to the small number of comparisons between session halves. Therefore, the long-timescale analysis might overestimate the dependence between the different measurements.

To examine the linear relationship between different activity measures and tuning curve stability at the single-cell level, we calculated for each pair of movie repeats the fraction of variance in single cells' tuning curve correlations explained by the average activity rates of the same cells. We then averaged the R^2 values across all pairs of different movie repeats. This procedure was repeated using either the absolute difference in activity rate or the absolute difference in activity rate score (absolute difference in activity rate between the two movie repeats, divided by their sum) to predict the single cells' tuning curve correlations. For the long-timescale analysis, (between different blocks occurring on different days)-we calculated the relationship between the different measurements based on the neuronal activity in each session half.

Similarity index (Related to Figure 4C and Figure 5)**.**
Since different cortical layers and brain areas exhibit different ranges of correlation values, in order to compare levels of representational drift, we normalized the magnitude of change in correlations for each of the compared groups prior to the statistical analysis. To this end, we calculated the 'Similarity index', defined as the difference between the correlation coefficient value calculated for a given interval ($CC_{interval}$) and correlation coefficient value calculated for the smallest interval ($CC_{reference}$) divided by their sum:

$$Similarity\ index = \frac{CC_{interval} - CC_{\text{reference}}}{CC_{interval} + CC_{\text{reference}}}$$

Therefore, no change in correlation value relative to the correlation value of the smallest interval will result in a value of 0 and a complete decorrelation will result in a value of -1. Negative correlation coefficient values were rectified to zero prior the normalization. Similar results were obtained when we repeated the analyses reported in Figure 5 using either Fisher z-transformed Pearson's correlation coefficients or coefficient of determination (R^2) values (instead of Pearson's r correlations) prior to the normalization procedure (data not shown).

Time-lapse decoding analysis (related to Figure S1K and Figure S5H).
For the analysis presented in Figure S1K, we used a *k*-nearest neighbors classifier with K=1 to decode the time bin at a given movie repeat of 'Natural movie 1' based on the population vectors of a preceding movie repeat using the Euclidean distance between the response vectors. The performance of the decoder was defined as the percentage of correct classifications out of the 30 time-bins for each pair of movie repeats, and was compared to

that of the same decoder after shuffling the identities of the labels of the training data. For the analysis presented in Figure S5H, we applied the same decoder but used the average activity in the first half of the session (repeats 1-5) as training data and the average activity in the second half of the same session (repeats 6-10) as test data.

Internal structure of neuronal population activity.
Similar to the PV correlation analysis, we divided the movie into 30 equal time bins and calculated the average activity of each neuron in each time bin, yielding a matrix of 30 by the number of recorded neurons. Then, we calculated the Pearson's correlation across all population vectors, resulting in a symmetric, 30-by-30 matrix, representing the structure of similarities across the population activity patterns at all different time bins of the presented movie. We defined this matrix of similarities as the 'internal structure of neuronal population activity' (or 'internal structure'). Since this structure no longer holds the identities of individual neurons, it is possible to measure the resemblance between structures extracted from different datasets (e.g., movie repeats, natural movies, areas, mice, etc.) without relying on the ability to record from the same cells or requiring equal numbers of cells across measurements. Note that for visualization purpose, in Figure 6F, Figure 7A, and Figure S7 the population activity was divided into 90 equal time bins and underwent a non-linear dimensionality reduction. The quantitative analyses shown in Figure 6G,H, Figure 7C,D and Figure S7 were performed on the pairwise distances between the original (non-reduced) vectors of internal structures as described in the Methods.

Pseudo-mice and shuffled pseudo-mice (related to Figure 6-7 and Figure S7).
To reduce the effect of incidental differences in the coding properties of the sampled neurons on our ability to capture the true internal structure of each of the studied areas, we constructed 'pseudo-mice', which are a pooling of cells recorded from different mice of the same dataset. To create two independent pseudo-mice (i.e. pseudo-mice that have no overlap in their source of neuronal activity), we first randomly split the complete dataset (Neuropixels or Ca^{2+} imaging datasets depending on the analysis) into two equal non-overlapping groups of mice. Each mouse in each group contained the neuronal activity recorded from 1-6 brain areas. Pooling all the cells/units from each brain area across all mice of the same group yielded six distinct sets of neurons (one per area) for each of the two pseudo-mice (12 pseudo-areas in total). Since there is variability in the number of recorded areas and cells across mice, the pooling procedure resulted in a different number of cells in each pseudo-area. To ensure that differences between the internal structures of different areas did not stem from the size of the recorded neuronal population, we randomly subsampled an equal number of cells from the entire population of each area. The exact number of subsampled cells was determined based on the pseudo-area with the lowest number of cells among both pseudo-mice. To verify the uniqueness of the internal structure of each area, the analysis was compared to complementary 'shuffled pseudo-mice' that were created by the random redistribution of all cells across areas within each of the pseudo-mice.

Between 'pseudo-mice' decoding (related to Figure 6G,H and Figure S7C,D,F,G).
To decode the identity of the recorded brain areas between pseudo-mice based on the similarities across their internal structures, we calculated the internal structure of each area in each of the two pseudo-mice based on the averaged activity over all movie repeats. This

yielded 12 symmetric, 30-by-30 matrices (6 areas x 2 pseudo-mice) capturing the similarities across activity patterns of the different time-bins in the movie. We then considered all 720 possible classifications (6! permutations) of brain areas across pseudo-mice. Each permutation defined a one-to-one mapping between the six internal structures of one pseudo-mouse to the six internal structures of the other pseudo-mouse. We then chose the permutation that maximized the similarities between internal structures across pseudo-mice (Pearson's correlation, sum over all six pairs of internal structures):

$$\hat{O} = \underset{<a_1,a_2,a_3 \dots a_6>}{\operatorname{argmax}} \left\{ \sum_{i=1}^{6} corr\left(V_i^A, V_{a_i}^B\right) \right\},$$

where $\hat{O}$ is the inferred brain area labels for six internal structure vectors in pseudo-mouse B, $< a_1, a_2, a_3 \dots a_6 >$ is a possible permutation of the six patterns' internal structure vectors of pseudo-mouse B, V_i^A is the internal structure vector of the i^{th} visual area in pseudo-mouse A, and $V_{a_i}^B$ is the internal structure vector of the ${a_i}^{th}$ visual area in pseudo-mouse B. This procedure was repeated 1000 times to obtain representative results across different realizations of pseudo-mice (different realizations of dividing the mice population into two random subsets) and was compared to the results obtained when using shuffled pseudo-mice. For the analyses presented in Figure 6G, H and Figure S7D, we included all mice from both 'Brain Observatory' and 'Functional Connectivity' groups of the Neuropixels dataset in response to 'Natural movie 1'. Since the two groups of mice were presented with a different number of 'Natural movie 1' repeats (20 and 60, respectively), the analyses were performed using the average activity across the first 20 movie repeats. For the analysis presented in Figure S7C, we included all mice from the calcium imaging dataset in response to 'Natural movie 1' using the average activity across all 30 movie repeats. For the analysis presented in Figure S7F, we included all the mice from the 'Functional Connectivity' group of the Neuropixels dataset during the presentation of 'Shuffled natural movie 1' (SNM1) and 'Natural movie 1' (NM1). Since the two stimuli were presented with a different number of movie repeats (10 repeats per block for SNM1 and 30 repeats per block for NM1), the analysis was performed using the average activity across 20 movie repeats (10 subsequent repeats from each block). Finally, for the analysis presented in Figure S7G, we used all mice from the calcium imaging dataset during the presentation of 'Natural movie 3' and 'Drifting gratings'. The 'Drifting gratings' stimuli consists of 40 unique combinations of direction and temporal frequencies, resulting in an internal structure matrix in the size of 40-by-40. To control for the effects of matrix size on the results, we divided each 'Natural movie 3' repeat into 40 equal time bins (instead of 30 bins as done in previous analysis). In addition, since 'Natural movie 3' was presented only across two blocks (10 minutes each, separated by approximately 20 minutes) and the 'Drifting grating' stimulus was presented across three blocks (10 minutes each, separated by approximately 15 minutes), we performed the analysis using the average activity across the first two blocks for both stimuli. Note that similar results were obtained when using all three blocks of 'Drifting gratings' or without controlling for the size of the matrices (data not shown).

Internal structures across recording technologies (related to Figure 6E). In this analysis, we used all the Neuropixels and two-photon Ca^{+2} imaging 'Natural movie 1' data to create two pseudo-mice, one for each recording technique. Since different mice in the Neuropixels dataset were presented with a different number of movie repeats, 20 repeats in the 'Brain

Observatory' group and 60 repeats in the 'Functional Connectivity' group, we used only the first 20 repeats for 'Functional Connectivity' group. First, we calculated for each brain area of each pseudo-mouse the internal structure per movie repeat. We then calculated the median internal structures over all movie repeats to create 12 internal structures (6 areas x 2 pseudo-mice). Finally, we normalized (z-score) the internal structures within each pseudo-mouse and calculated the Pearson's correlation distance matrix across areas of the two pseudo-mice (Neuropixels pseudo-mouse and Ca^{2+} imaging pseudo-mouse).

Temporally shuffled internal structure (related to Figure S7D). To verify that the internal structures contain information beyond the similarities between adjacent time bins, we repeated the analysis presented in Figure 6H after performing a different random cyclic temporal shuffle (prior the vectorization procedure) to each of the internal structures of each pseudo-mouse. This procedure was repeated 1000 times to obtain representative results across different cyclic shuffles and realizations of pseudo-mice.

Internal structure stability (related to Figure 7C-D and Figure S7H,I).
We created different realizations of pseudo-mice by randomly sampling 70% of the total pool of cells that were active in all three Ca^{2+} imaging sessions in response to 'Natural movie 1'. Then, for each area, we calculated the population vectors for all time bins and internal structures based on the average activity rates over the first half of the session (repeats 1-5) or the second half of the session (repeats 6-10). Lastly, we calculated for each of the two measurements the Pearson's correlations between the two halves of the same session, between halves of two temporally proximal sessions, and between halves of two temporally distal sessions. This procedure was repeated 1000 times to obtain representative results across different realizations of pseudo-mice.

Structure of pairwise similarities (related to Figure S7J-L).
Similar to the PV correlation analysis, we divided each movie repeat into 30 equal time bins and calculated the population vector for each time bin, yielding a matrix of 30 by the number of recorded neurons (30-by-n). Then, we calculated the Pearson's correlation across vectors of all neurons, resulting in a symmetric n-by-n matrix. This matrix represented the structure of pairwise similarities between individual neurons tuning curves. We defined this matrix of similarities as the 'structure of pairwise similarities'.

t-distributed stochastic neighbor embedding (tSNE; related to Figure 6C-F, Figure 7A and Figure S7A,E). For visualizing the relationships between internal structures, the vectors of pairwise correlations across activity patterns were embedded in three dimensions using tSNE[83,84]. We used the exact tSNE algorithm with similar embedding settings for all visualizations (cosine distance metric, using 10 PCA components, exaggeration 4 (default), and learning rate 500 (default)). The perplexity (effective number of local neighbors of each point) was chosen for each visualization based on multiplication of the minimal number of movie repeats used in the analysis (30, 60 and 20 for Figure 6C-E, respectively). For visualizing the low-dimensional manifold of neuronal population activity, the PVs of all the movie repeats of the same block (or session) were embedded in either two (Figure 6A, Figure 7A and Figure S7A) or three dimensions (Figure 6F and Figure S7E). We used the exact t-SNE algorithm with similar embedding settings for all visualizations (cosine distance metric, using 20 PCA components, perplexity 200, exaggeration 4 (default), and learning rate 500 (default)). Note that embedding in the reduced space is used for visualization

purposes only and all quantifications were performed based on the pairwise distances between the original (non-reduced) vectors of internal structures as described in other sections of the Methods.

QUANTIFICATION AND STATISTICAL ANALYSIS

All statistical details, including the specific statistical tests, are specified in the corresponding figure legends. In general, two-tailed non-parametric Wilcoxon rank sum tests (unpaired data), Wilcoxon signed rank tests (paired data), and Friedman's tests (repeated measures) were performed and corrected for multiple comparisons (using Holm–Bonferroni method) when required. A one-sided Pearson's correlation coefficient was used to estimate the effect of elapsed time on ensemble rate and tuning curve stability (Figure 3J,K). In all tests, significance was defined at $\alpha=0.05$. Aside from mice with a low number of recorded cells (see 'Data analysis' section in the Methods), no neuronal data were excluded from the analysis. All statistical analyses were conducted using MATLAB 2017b (Mathworks).

References

1. Alvarez-Castelao, B., and Schuman, E.M. (2015). The regulation of synaptic protein turnover. J. Biol. Chem. *290*, 28623–28630.

2. Yasumatsu, N., Matsuzaki, M., Miyazaki, T., Noguchi, J., and Kasai, H. (2008). Principles of long-term dynamics of dendritic spines. J. Neurosci. *28*, 13592–13608.

3. Holtmaat, A., and Svoboda, K. (2009). Experience-dependent structural synaptic plasticity in the mammalian brain. Nat. Rev. Neurosci. *10*, 647–658.

4. Minerbi, A., Kahana, R., Goldfeld, L., Kaufman, M., Marom, S., and Ziv, N.E. (2009). Long-term relationships between synaptic tenacity, synaptic remodeling, and network activity. PLoS Biol. *7*.

5. Loewenstein, Y., Kuras, A., and Rumpel, S. (2011). Multiplicative dynamics underlie the emergence of the log- normal distribution of spine sizes in the neocortex in vivo. J. Neurosci. *31*, 9481–9488.

6. Ziv, Y., Burns, L.D., Cocker, E.D., Hamel, E.O., Ghosh, K.K., Kitch, L.J., Gamal, A. El, and Schnitzer, M.J. (2013). Long-term dynamics of CA1 hippocampal place codes. Nat. Neurosci. *16*, 264–266.

7. Okun, M., Lak, A., Carandini, M., and Harris, K.D. (2016). Long term recordings with immobile silicon probes in the mouse cortex. PLoS One *11*, 1–17.

8. Dhawale, A.K., Poddar, R., Wolff, S.B.E., Normand, V.A., Kopelowitz, E., and Ölveczky, B.P. (2017). Automated long-Term recording and analysis of neural activity in behaving animals. Elife *6*, 1–40.

9. Andermann, M.L., Kerlin, A.M., and Reid, R.C. (2010). Chronic cellular imaging of mouse visual cortex during operant behavior and passive viewing. Front. Cell. Neurosci. *4*, 1–16.

10. Poort, J., Khan, A.G., Pachitariu, M., Nemri, A., Orsolic, I., Krupic, J., Bauza, M., Sahani, M., Keller, G.B., Mrsic-Flogel, T.D., et al. (2015). Learning Enhances Sensory and Multiple Non-sensory Representations in Primary Visual Cortex. Neuron *86*, 1478–1490.

11. Peron, S.P., Freeman, J., Iyer, V., Guo, C., and Svoboda, K. (2015). A Cellular Resolution Map of Barrel Cortex Activity during Tactile Behavior. Neuron *86*, 783–799.

12. Huber, D., Gutnisky, D.A., Peron, S., O'Connor, D.H., Wiegert, J.S., Tian, L., Oertner, T.G., Looger, L.L., and Svoboda, K. (2012). Multiple dynamic representations in the motor cortex during sensorimotor learning. Nature *484*, 473–478.

13. Bondar, I. V., Leopold, D.A., Richmond, B.J., Victor, J.D., and Logothetis, N.K. (2009). Long-term stability of visual pattern selective responses of monkey temporal lobe neurons. PLoS One *4*.

14. Rokni, U., Richardson, A.G., Bizzi, E., and Seung, H.S. (2007). Motor Learning with Unstable Neural Representations. Neuron *54*, 653–666.

15. Tolias, A.S., Ecker, A.S., Siapas, A.G., Hoenselaar, A., Keliris, G.A., and Logothetis, N.K. (2007). Recording chronically from the same neurons in awake, behaving primates. J.

Neurophysiol. *98*, 3780–3790.

16. Rose, T., Jaepel, J., Hübener, M., and Bonhoeffer, T. (2016). Cell-specific restoration of stimulus preference after monocular deprivation in the visual cortex. Science (80-.). *352*, 1319 LP – 1322.

17. Tolhurst, D.J., Movshon, J.A., and Dean, A.F. (1983). The statistical reliability of signals in single neurons in cat and monkey visual cortex. Vision Res. *23*, 775–785.

18. Arieli, A., Sterkin, A., Grinvald, A., and Aertsen, A. (1996). Dynamics of ongoing activity: Explanation of the large variability in evoked cortical responses. Science (80-.). *273*, 1868–1871.

19. Faisal, A.A., Selen, L.P.J., and Wolpert, D.M. (2008). Noise in the nervous system. Nat. Rev. Neurosci. *9*, 292–303.

20. Engel, T.A., and Steinmetz, N.A. (2019). New perspectives on dimensionality and variability from large-scale cortical dynamics. Curr. Opin. Neurobiol. *58*, 181–190.

21. Clopath, C., Bonhoeffer, T., Hübener, M., and Rose, T. (2017). Variance and invariance of neuronal long-term representations. Philos. Trans. R. Soc. B Biol. Sci. *372*.

22. Lütcke, H., Margolis, D.J., and Helmchen, F. (2013). Steady or changing? Long-term monitoring of neuronal population activity. Trends Neurosci. *36*, 375–384.

23. Schölvinck, M.L., Saleem, A.B., Benucci, A., Harris, K.D., and Carandini, M. (2015). Cortical state determines global variability and correlations in visual cortex. J. Neurosci. *35*, 170–178.

24. Cohen, M.R., and Maunsell, J.H.R. (2010). A neuronal population measure of attention predicts behavioral performance on individual trials. J. Neurosci. *30*, 15241–15253.

25. Montijn, J.S., Goltstein, P.M., and Pennartz, C.M.A. (2015). Mouse V1 population correlates of visual detection rely on heterogeneity within neuronal response patterns. Elife *4*, 1–31.

26. Rubin, A., Geva, N., Sheintuch, L., and Ziv, Y. (2015). Hippocampal ensemble dynamics timestamp events in long-term memory. Elife *4*, 1–16.

27. Sheintuch, L., Geva, N., Baumer, H., Rechavi, Y., Rubin, A., and Ziv, Y. (2020). Multiple Maps of the Same Spatial Context Can Stably Coexist in the Mouse Hippocampus. Curr. Biol. *30*, 1467-1476.e6.

28. Driscoll, L.N., Pettit, N.L., Minderer, M., Chettih, S.N., and Harvey, C.D. (2017). Dynamic Reorganization of Neuronal Activity Patterns in Parietal Cortex. Cell *170*, 986-999.e16.

29. Rule, M.E., Loback, A.R., Raman, D. V, Driscoll, L.N., Harvey, C.D., and O'Leary, T. (2020). Stable task information from an unstable neural population. Elife *9*.

30. Chambers, A.R., and Rumpel, S. (2017). A stable brain from unstable components: Emerging concepts and implications for neural computation. Neuroscience *357*, 172–184.

31. Mankin, E.A., Diehl, G.W., Sparks, F.T., Leutgeb, S., and Leutgeb, J.K. (2015). Hippocampal CA2 Activity Patterns Change over Time to a Larger Extent than

between Spatial Contexts. Neuron *85*, 190–201.

32. Rule, M.E., O'Leary, T., and Harvey, C.D. (2019). Causes and consequences of representational drift. Curr. Opin. Neurobiol. *58*, 141–147.

33. Josselyn, S.A., Köhler, S., and Frankland, P.W. (2015). Finding the engram. Nat. Rev. Neurosci. *16*, 521–534.

34. Tonegawa, S., Morrissey, M.D., and Kitamura, T. (2018). The role of engram cells in the systems consolidation of memory. Nat. Rev. Neurosci. *19*, 485–498.

35. Haak, K. V., Morland, A.B., and Engel, S.A. (2015). Plasticity, and its limits, in adult human primary visual cortex. Multisens. Res. *28*, 297–307.

36. Jeon, B.B., Swain, A.D., Good, J.T., Chase, S.M., and Kuhlman, S.J. (2018). Feature selectivity is stable in primary visual cortex across a range of spatial frequencies. Sci. Rep. *8*, 1–14.

37. Montijn, J.S., Meijer, G.T., Lansink, C.S., and Pennartz, C.M.A. (2016). Population-Level Neural Codes Are Robust to Single-Neuron Variability from a Multidimensional Coding Perspective. Cell Rep. *16*, 2486–2498.

38. Ranson, A. (2017). Stability and Plasticity of Contextual Modulation in the Mouse Visual Cortex. Cell Rep. *18*, 840–848.

39. Betzel, R.F., Wood, K.C., Angeloni, C., Geffen, M.N., and Bassett, D.S. (2019). Stability of spontaneous, correlated activity in mouse auditory cortex. PLoS Comput. Biol. *15*, 1–25.

40. de Vries, S.E.J., Lecoq, J.A., Buice, M.A., Groblewski, P.A., Ocker, G.K., Oliver, M., Feng, D., Cain, N., Ledochowitsch, P., Millman, D., et al. (2020). A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex. Nat. Neurosci. *23*, 138–151.

41. Siegle, J.H., Jia, X., Durand, S., Gale, S., and Bennett, C. (2019). A survey of spiking activity reveals a functional hierarchy of mouse corticothalamic visual areas. 1–53.

42. Jun, J.J., Steinmetz, N.A., Siegle, J.H., Denman, D.J., Bauza, M., Barbarits, B., Lee, A.K., Anastassiou, C.A., Andrei, A., Aydin, Ç., et al. (2017). Fully integrated silicon probes for high-density recording of neural activity. Nature *551*, 232–236.

43. David, S. V., Vinje, W.E., and Gallant, J.L. (2004). Natural stimulus statistics alter the receptive field structure of V1 neurons. J. Neurosci. *24*, 6991–7006.

44. Kampa, B.M., Roth, M.M., Göbel, W., and Helmchen, F. (2011). Representation of visual scenes by local neuronal populations in layer 2/3 of mouse visual cortex. Front. Neural Circuits *5*, 1–12.

45. Talebi, V., and Baker, C.L. (2012). Natural versus synthetic stimuli for estimating receptive field models: A comparison of predictive robustness. J. Neurosci. *32*, 1560–1576.

46. Musall, S., Kaufman, M.T., Juavinett, A.L., Gluf, S., and Churchland, A.K. (2019). Single-trial neural dynamics are dominated by richly varied movements. Nat. Neurosci. *22*, 1677–1686.

47. Vinck, M., Batista-Brito, R., Knoblich, U., and Cardin, J.A. (2015). Arousal and Locomotion Make Distinct Contributions to Cortical Activity Patterns and Visual Encoding. Neuron *86*, 740–754.

48. Niell, C.M., and Stryker, M.P. (2010). Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex. Neuron *65*, 472–479.

49. Keller, G.B., Bonhoeffer, T., and Hübener, M. (2012). Sensorimotor Mismatch Signals in Primary Visual Cortex of the Behaving Mouse. Neuron *74*, 809–815.

50. Dipoppa, M., Ranson, A., Krumin, M., Pachitariu, M., Carandini, M., and Harris, K.D. (2018). Vision and Locomotion Shape the Interactions between Neuron Types in Mouse Visual Cortex. Neuron *98*, 602-615.e8.

51. Ayaz, A., Saleem, A.B., Schölvinck, M.L., and Carandini, M. (2013). Locomotion controls spatial integration in mouse visual cortex. Curr. Biol. *23*, 890–894.

52. Polack, P.O., Friedman, J., and Golshani, P. (2013). Cellular mechanisms of brain state-dependent gain modulation in visual cortex. Nat. Neurosci. *16*, 1331–1339.

53. Erisken, S., Vaiceliunaite, A., Jurjut, O., Fiorini, M., Katzner, S., and Busse, L. (2014). Effects of locomotion extend throughout the mouse early visual system. Curr. Biol. *24*, 2899–2907.

54. Ruff, D.A., and Cohen, M.R. (2014). Attention can either increase or decrease spike count correlations in visual cortex. Nat. Neurosci. *17*, 1591–1597.

55. Mineault, P.J., Tring, E., Trachtenberg, J.T., and Ringach, D.L. (2016). Enhanced spatial resolution during locomotion and heightened attention in mouse primary visual cortex. J. Neurosci. *36*, 6382–6392.

56. Kohn, A. (2007). Visual adaptation: Physiology, mechanisms, and functional benefits. J. Neurophysiol. *97*, 3155–3164.

57. Grill-Spector, K., Henson, R., and Martin, A. (2006). Repetition and the brain: Neural models of stimulus-specific effects. Trends Cogn. Sci. *10*, 14–23.

58. Sheintuch, L., Rubin, A., Brande-Eilat, N., Geva, N., Sadeh, N., Pinchasof, O., and Ziv, Y. (2017). Tracking the Same Neurons across Multiple Days in Ca2+ Imaging Data. Cell Rep. *21*, 1102–1115.

59. Harris, J.A., Mihalas, S., Hirokawa, K.E., Whitesell, J.D., Choi, H., Bernard, A., Bohn, P., Caldejon, S., Casal, L., Cho, A., et al. (2019). Hierarchical organization of cortical and thalamic connectivity. Nature *575*, 195–202.

60. Rubin, A., Sheintuch, L., Brande-Eilat, N., Pinchasof, O., Rechavi, Y., Geva, N., and Ziv, Y. (2019). Revealing neural correlates of behavior without behavioral measurements. Nat. Commun. *10*.

61. Juavinett, A.L., and Callaway, E.M. (2015). Pattern and Component Motion Responses in Mouse Visual Cortical Areas. Curr. Biol. *25*, 1759–1764.

62. Andermann, M.L., Kerlin, A.M., Roumis, D.K., Glickfeld, L.L., and Reid, R.C. (2011). Functional specialization of mouse higher visual cortical areas. Neuron *72*, 1025–1039.

63. Glickfeld, L.L., Andermann, M.L., Bonin, V., and Reid, R.C. (2013). Cortico-cortical projections in mouse visual cortex are functionally target specific. Nat. Neurosci. *16*, 219–226.

64. Murakami, T., Matsui, T., and Ohki, K. (2017). Functional segregation and development of mouse higher visual areas. J. Neurosci. *37*, 9424–9437.

65. Marks, T.D., and Goard, M.J. (2020). Stimulus-dependent representational drift in primary visual cortex. bioRxiv.

66. Ledochowitsch, P., Huang, L., Knoblich, U., Oliver, M., Lecoq, J., Reid, C., Li, L., Zeng, H., Koch, C., Waters, J., et al. (2019). On the correspondence of electrical and optical physiology in in vivo population-scale two-photon calcium imaging. bioRxiv, 1–32.

67. Siegle, J.H., Ledochowitsch, P., Jia, X., Millman, D., Ocker, G.K., Caldejon, S., Casal, L., Cho, A., Denman, D.J., Durand, S., et al. (2020). Reconciling functional differences in populations of neurons recorded with two-photon imaging and electrophysiology. bioRxiv, 1–42.

68. Wei, Z., Lin, B.J., Chen, T.W., Daie, K., Svoboda, K., and Druckmann, S. (2020). A comparison of neuronal population dynamics measured with calcium imaging and electrophysiology. PLoS Comput. Biol. *16*, 1–29.

69. Huang, L., Ledochowitsch, P., Knoblich, U., Lecoq, J., Murphy, G.J., Reid, R.C., de Vries, S.E.J., Koch, C., Zeng, H., Buice, M.A., et al. (2021). Relationship between simultaneously recorded spiking activity and fluorescence signal in gcamp6 transgenic mice. Elife *10*, 1–19.

70. Maffei, A., and Turrigiano, G.G. (2008). Multiple modes of network homeostasis in visual cortical layer 2/3. J. Neurosci. *28*, 4377–4384.

71. Hengen, K.B., Lambo, M.E., VanHooser, S.D., Katz, D.B., and Turrigiano, G.G. (2013). Firing rate homeostasis in visual cortex of freely behaving rodents. Neuron *80*, 335–342.

72. Keck, T., Keller, G.B., Jacobsen, R.I., Eysel, U.T., Bonhoeffer, T., and Hübener, M. (2013). Synaptic scaling and homeostatic plasticity in the mouse visual cortex in vivo. Neuron *80*, 327–334.

73. Slomowitz, E., Styr, B., Vertkin, I., Milshtein-Parush, H., Nelken, I., Slutsky, M., and Slutsky, I. (2015). Interplay between population firing stability and single neuron dynamics in hippocampal networks. Elife *2015*, 1–21.

74. El-Boustani, S., Ip, J.P.K., Breton-Provencher, V., Knott, G.W., Okuno, H., Bito, H., and Sur, M. (2018). Locally coordinated synaptic plasticity of visual cortex neurons in vivo. Science (80-.). *360*, 1349–1354.

75. Felipe, Y., Kossio, K., Goedeke, S., and Klos, C. (2020). Drifting Assemblies for Persistent Memory neurons . For faithful storage these assemblies are assumed to consist of the same neurons. 1–21.

76. Rule, M.E., and O'leary, T. (2021). Self-Healing Neural Codes. bioRxiv, 2021.03.08.433413.

77. Gallego, J.A., Perich, M.G., Chowdhury, R.H., Solla, S.A., and Miller, L.E. (2020). Long-

term stability of cortical population dynamics underlying consistent behavior. Nat. Neurosci. *23*, 260–270.

78. Pashkovski, S.L., Iurilli, G., Brann, D., Chicharro, D., Drummey, K., Franks, K., Panzeri, S., and Datta, S.R. (2020). Structure and flexibility in cortical representations of odour space. Nature *583*, 253–258.

79. Bolding, K.A., Nagappan, S., Han, B.X., Wang, F., and Franks, K.M. (2020). Recurrent circuitry is required to stabilize piriform cortex odor representations across brain states. Elife *9*, 1–23.

80. Xia, J., Marks, T.D., Goard, M.J., and Wessel, R. (2021). Stable representation of a naturalistic movie emerges from episodic activity with 2 gain variability.

81. Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C.B., Carandini, M., and Harris, K.D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. Science (80-.). *364*.

82. Schoonover, C.E., Ohashi, S.N., Axel, R., and Fink, A.J.P. (2021). Representational drift in primary olfactory cortex. Nature *594*, 541–546.

83. Maaten, L. Van Der, and Hinton, G. (2008). Visualizing Data using t-SNE. *9*, 2579–2605.

84. Maaten, L. Van Der (2014). Accelerating t-SNE using Tree-Based Algorithms. *15*, 3221–3245.

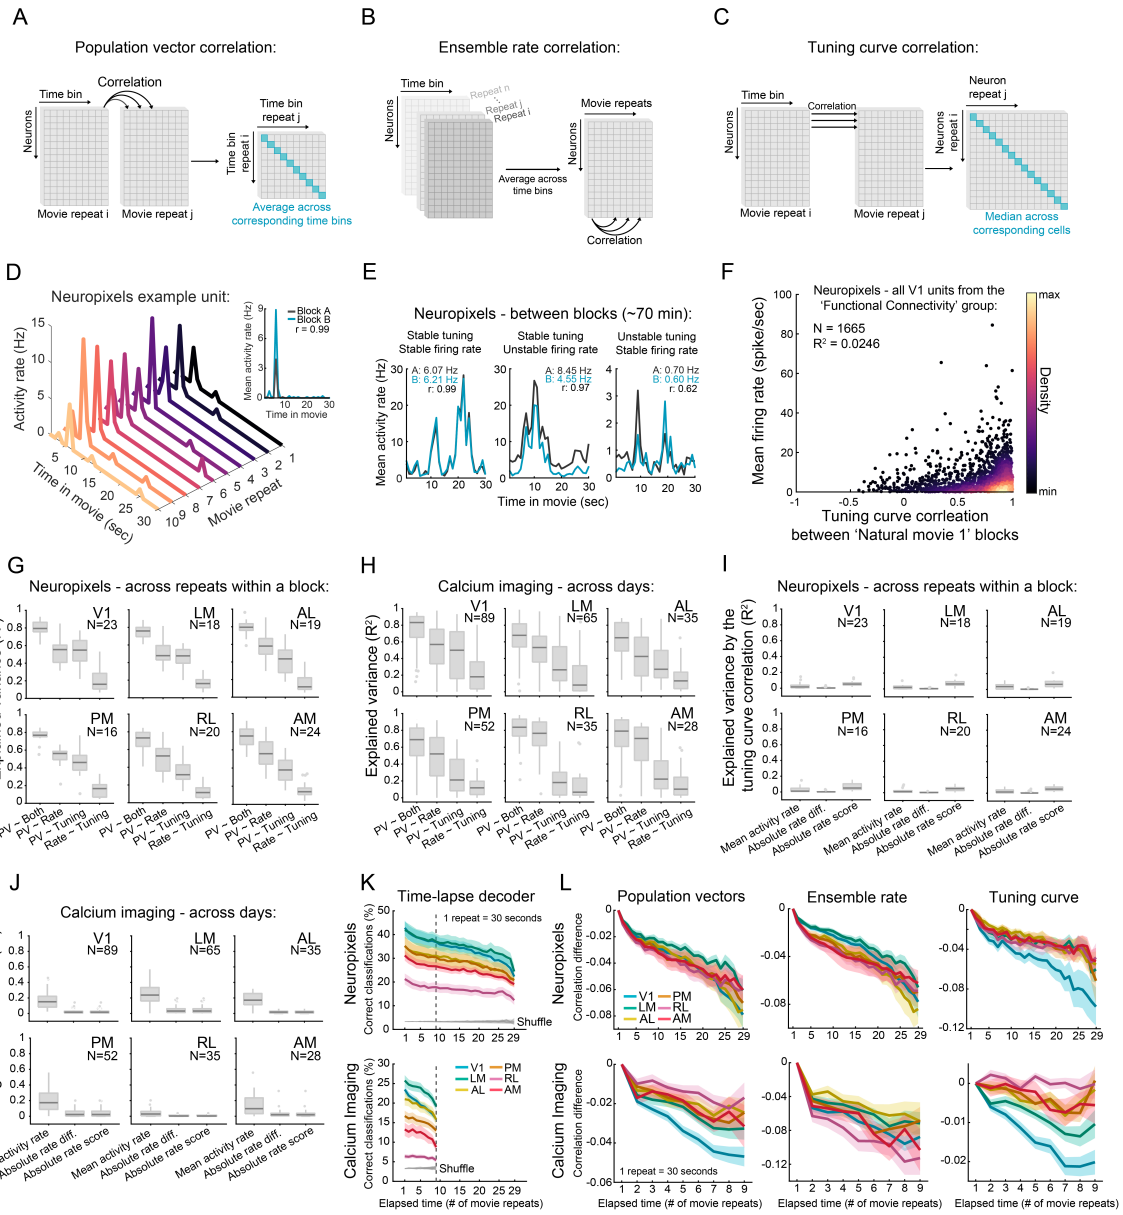

Figure S1. Changes in the tuning and activity rates of individual neurons differentially contribute to drift in visual representations. Related to Figure 2. (A-C) Workflow for calculating the population vector correlation (A), ensemble rate correlation (B), and tuning curve correlation (C) measurements (see Methods). (D) A single representative unit recorded from area V1 using Neuropixels probes showing high degree of tuning curve stability across different movie repeats and blocks (inset) of Natural movie 1. Note that the activity rate of the unit can fluctuate both within and across blocks irrespective of its tuning curve stability. (E) Responses of three V1 example units across different blocks of Natural movie 1 separated by ~70 minutes within the same recording session. Each unit exhibits a different degree of tuning curve and activity rate stability across the two blocks. (F) Distribution of tuning curve correlations between blocks as a function of the average activity rate across blocks of Natural movie 1 for all V1 units recorded using Neuropixels probes. Each data point represents a single unit. We found only a weak linear relationship between the average activity rate of each unit and

its tuning curve stability across blocks. (G-H) Quantifying the relationship between the different stability measurements (PV correlation, ensemble rate correlation and tuning curve correlation) in terms of explained variance for both short (G) and long timescales (H). For the short timescale analysis (between movie repeats within a block) we used data of mice from the Neuropxiels Functional Connectivity group during the presentation of Natural movie 1 and for the long timescale analysis (between different blocks occurring of different days) we used data from the calcium imaging dataset during the presentation of Natural movie 1. Using a multiple linear regression model, we found that values of each of the two measurements (ensemble rate correlation and tuning curve correlation) contributed differentially to the variance explained in the PV correlation values. Only a small fraction (<15%) of the variance in the values of the ensemble rate correlation measurement can be explained using the tuning curve correlation values. (I-J) Quantifying the relationship between different single cell activity properties (average activity rate, absolute activity rate difference and absolute activity modulation index) and single cell tuning curve stability in terms of explained variance for both short (I) and long timescales (J). Only a small fraction (< 20%) of the variance in the values of the tuning curve correlation measurement can be explained by the measurements used to assess the stability of single cells' excitability. (K) Percentage of correct classifications as a function of the elapsed time between the train and test movie repeats for both mice from the Neuropixels Functional Connectivity group (top) and mice from the calcium imaging dataset (bottom) during the presentation of Natural movie 1. The difference in correct classifications between the minimal and maximal interval of movie repeats was significant for all areas ($p \leq 0.011$, two-tailed Wilcoxon signed-rank test with Holm–Bonferroni correction). (L) Difference in PV correlations (left), ensemble rate correlations (middle) and tuning curve correlations (right) as a function of elapsed time for all six visual cortical areas of both Neuropixels recorded mice (top) and Ca^{2+} imaged mice (bottom) during the presentation of Natural movie 1. The box plots in panels G-J show the data range (whiskers), 25th and 75th percentiles (box), and median (dark line). Outliers are marked by gray dots.

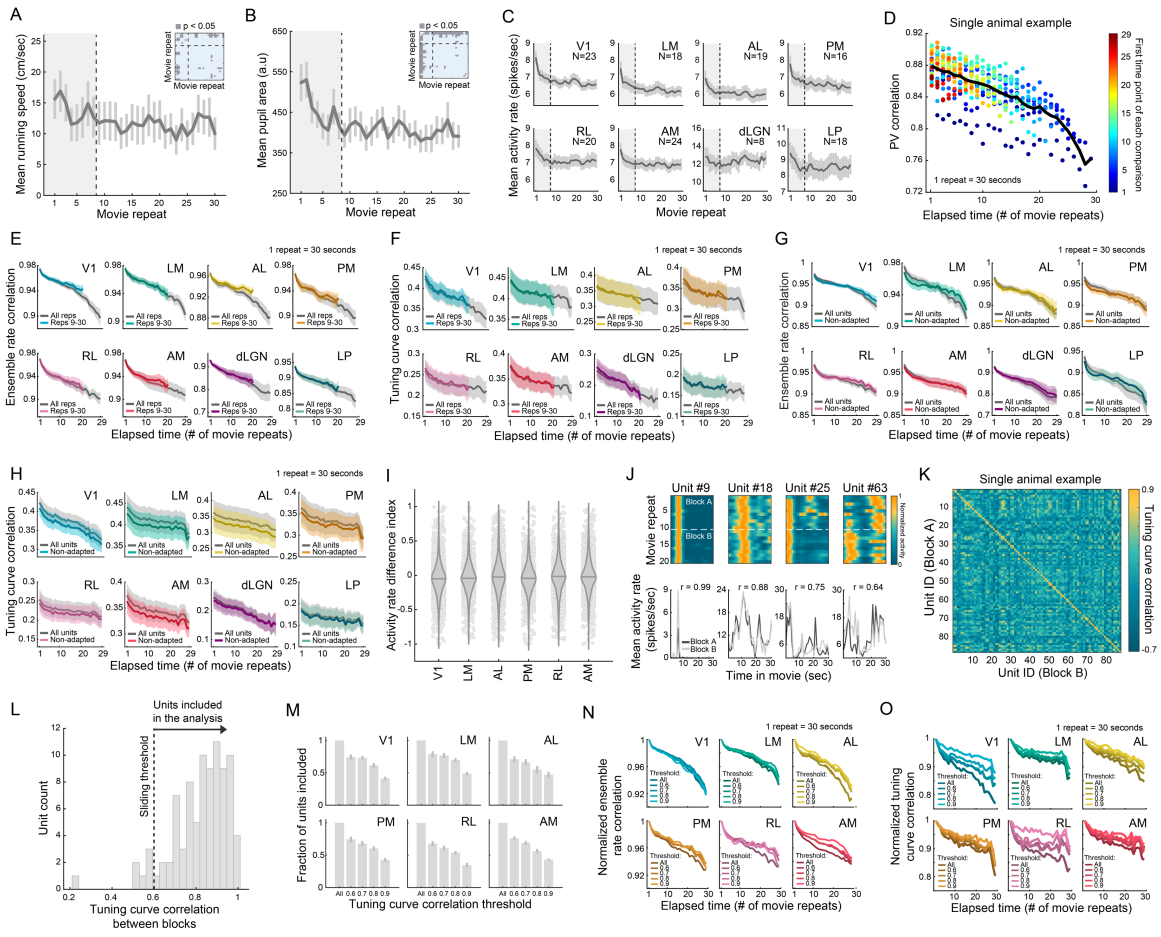

Figure S2. Representational drift is not driven by changes in behavioral state, reduction in global activity rates or recording instability. Related to Figure 2. (A-O) Analyses using data from the Neuropixels dataset during the presentation of Natural movie 1. (A) Mean running speed for each movie repeat across animals. (B) Mean pupil area for each movie repeat across animals. Insets in A and B indicate a significant difference between movie repeats (paired t-test, $p<0.05$, two-tailed, without correction for multiple comparisons). (C) Mean activity rates for each movie repeat across animals for each brain area. (D) Mean PV correlation as a function of the elapsed time. Each data point represents the mean PV correlation value for a single pair of movie repeats from Figure 2B and is colored according to the first time point of each comparison. Note that the comparisons containing the first few repeats (dark blue) generally have lower correlation values but still follow the trend of decorrelation over increasing intervals. The low similarity between the first movie repeats and the rest of the movie repeats might stem from the differences in arousal and activity rates presented in A-C. These early repeats are overrepresented in comparisons of longer intervals and contribute to the relative increase in the slope found in these longer intervals. (E) Ensemble rate correlation as a function of elapsed time, performed on a subset of movie repeats (repeats 9-30; colored lines). Ensemble rate correlations of this subset of the data gradually declined with the interval between movie repeats, similarly to the ensemble rate correlations of the full dataset (gray lines) from all movie repeats. All areas showed a significant decrease in ensemble rate correlations as function of elapsed time ($\chi^2_{(20)} \geq 142$, $p < 10^{-3}$, Friedman's tests with Holm–Bonferroni correction). (F) Tuning curve correlation as a function of time, performed on a subset of movie repeats (repeats 9-30; colored lines). Tuning curve correlations of this subset of the data gradually declined with

the interval between movie repeats, similarly to the tuning curve correlations of the full dataset (gray lines) from all movie repeats. All areas showed a significant decrease in tuning curve correlations as function of elapsed time ($\chi^2_{(20)} \geq 51.52$, $p < 10^{-3}$, Friedman's tests with Holm–Bonferroni correction). (G) Ensemble rate correlation as a function of time, performed using the subset of non-adapted units (colored lines). Ensemble rate correlations of this subset of the data gradually declined with the interval between movie repeats, similarly to the ensemble rate correlations when using all units (gray lines). All areas showed a significant decrease in ensemble rate correlations as function of time ($\chi^2_{(28)} \geq 217.57$, $p < 10^{-3}$, Friedman's tests with Holm–Bonferroni correction). (H) Tuning curve correlation as a function of time, performed using the subset of non-adapted units (colored lines). Tuning curve correlations of this subset of the data gradually declined with the interval between movie repeats, similarly to the tuning curve correlations when using all units (gray lines). All areas showed a significant decrease in tuning curve correlations as function of time ($\chi^2_{(28)} \geq 86.04$, $p < 10^{-3}$, Friedman's tests with Holm–Bonferroni correction). (I) Distribution of normalized activity rate difference for each of the six visual areas. For each unit, the mean activity rates (spikes/sec) of repeats 1-5 was subtracted from the mean activity of repeats 26-30, and divided by their sum. This procedure was done separately for each of the two Natural movie 1 blocks resulting in two points in the graph for each unit. (J) Responses of four V1 example cells from the same representative mouse across different repeats of Natural movie 1, spanning two blocks within the same recording session. Each unit exhibits a different degree of tuning curve stability across the two blocks (indicated by the Pearson's correlation values in the bottom panels). (K) Tuning curve correlation between blocks for all the units of the same representative mouse shown in J. (L) Distribution of the tuning curve correlation values of the main diagonal in K. Units that show high tuning curve correlation across blocks are unlikely to represent cells whose identity is unstable within blocks. A sliding threshold was used to include different subsets of units with high tuning stability between blocks. (M) Fraction of units included in the analysis as a function of their tuning curve correlation between blocks. (N) Repeating the within-block stability analysis (shown in Figure 2H) while subsampling units based on their tuning curve correlation between blocks. (O) Repeating the within-block stability analysis (shown in Figure 2I) while subsampling units based on their tuning curve correlation between blocks. In all relevant panels, data are mean ± SEM across mice from the Neuropixels 'Functional Connectivity' group.

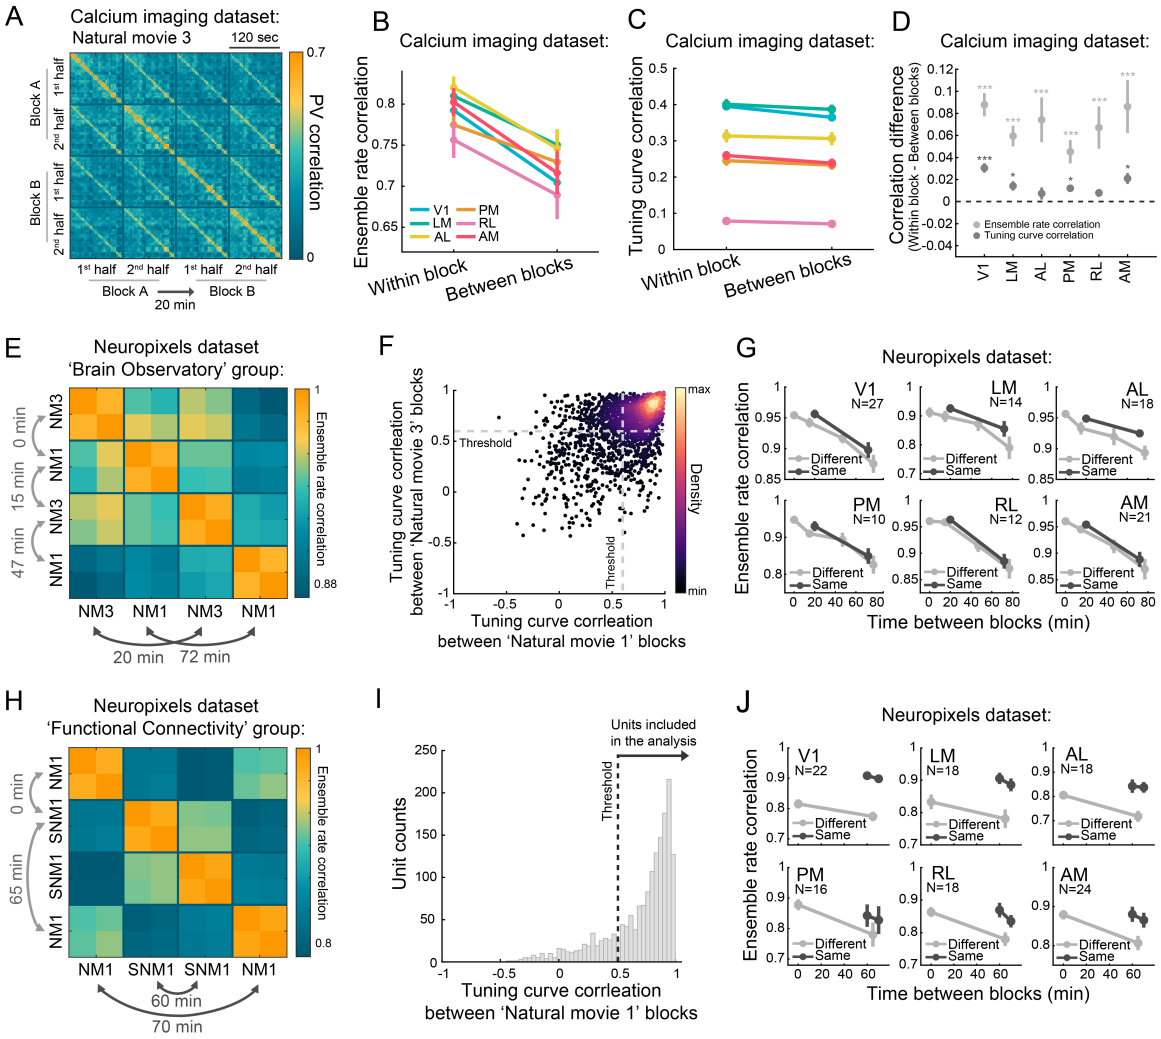

Figure S3. Visual representations change over timescales of tens of minutes. Related to Figure 3. (A-D) Analyses using data from the calcium imaging dataset during the presentation of Natural movie 3. (A) PV correlation between the 1st (repeats 1-2) and 2nd (repeats 3-5) halves of two different blocks of Natural movie 3 in a single visual area. The presented example is the average correlation matrix across all mice recorded in area AM using two-photon Ca^{2+} imaging. (B) Ensemble rate correlation between the two halves of the same block ('within block') and between halves of different blocks ('between blocks') of Natural movie 3 using the Ca^{2+} imaging dataset ($p < 10^{-3}$ for all areas, two-tailed Wilcoxon signed-rank test with Holm–Bonferroni correction). (C) Tuning curve correlation between the two halves of the same block ('within block') and between halves of different blocks ('between blocks') of Natural movie 3 using the Ca^{2+} imaging dataset ($p \leq 0.03$ for all areas, except areas AL and RL in which $p > 0.05$, two-tailed Wilcoxon signed-rank test with Holm–Bonferroni correction). (D) The difference in ensemble rate and tuning curve correlations within a block and between blocks of the Natural movie 3; Two-tailed Wilcoxon signed-rank test with Holm–Bonferroni correction; * p<0.05, ** p<0.01, *** p<0.001. (E-G) Analyses were done using the mice from the Neuropixels Brain Observatory group during the presentation of Natural movie 1 and Natural movie 3. (E) Ensemble rate correlations between halves of Natural movie 1 and Natural movie 3 blocks within the same session (NM1, Natural movie 1; NM3, Natural movie 3). The presented example is the average

correlation matrix across all mice from the Brain Observatory group, recorded with Neuropixels probes in area V1. (F) Tuning curve correlations between blocks of Natural movie 1 and Natural movie 3 for all V1 units across mice. Each data point represents a single unit. The units included in the analysis are those with tuning curve correlation r≥0.6 for both movies. (G) Ensemble rate correlations between blocks of the same and different natural movies decay with elapsed time. Note that ensemble rate correlations continuously decline with time, both between blocks of the same movie (dark gray) and between blocks of different movies (light gray). (H-J) Analyses using data from the Neuropixels Functional Connectivity group during the presentation of Natural movie 1 and Shuffled natural movie 1. (H) Ensemble rate correlations between halves of Natural movie 1 and Shuffled natural movie 1 blocks within the same session (NM1, Natural movie 1; SNM1, Shuffled natural movie 1). The presented example is the average correlation matrix across all mice from the Functional Connectivity group recorded with Neuropixels probes in area V1. (I) The V1 units included in this analysis showed tuning curve correlation r≥0.5 across the two blocks of Natural movie 1. (J) Similarly to the results presented in panel G, the ensemble rate correlations across different blocks of Natural movie 1 and different blocks of Shuffled natural movie 1 declined with time. Data in panels B-D, G and J are mean ± SEM across mice.

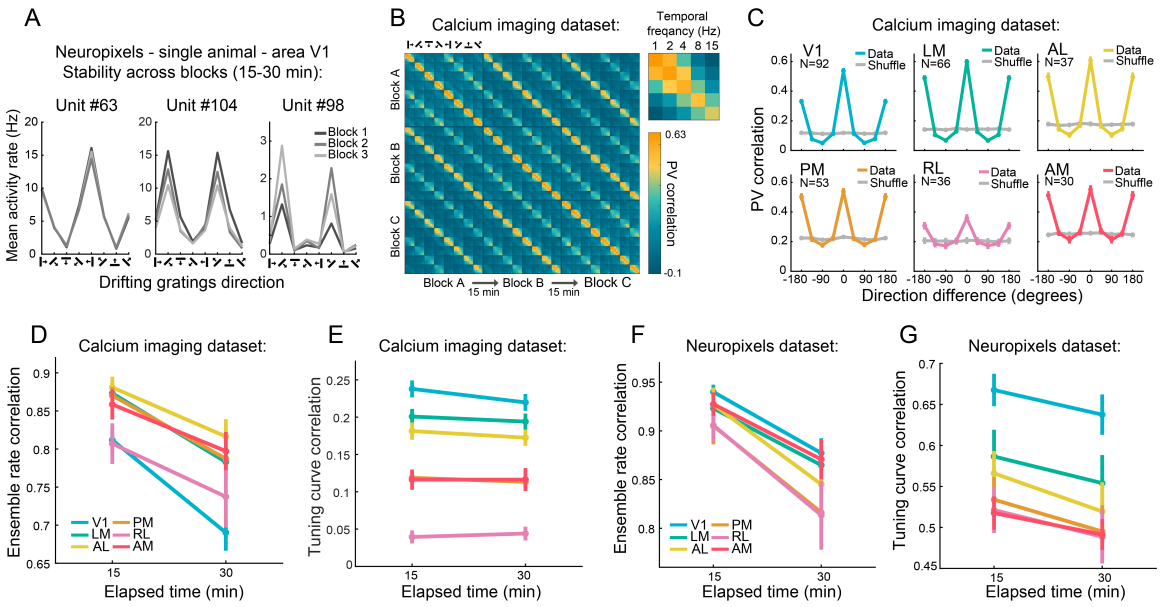

Figure S4. Characterizing the stability of visual representations of a synthetic stimulus. Related to Figure 3. (A) Responses of three V1 example units from the same representative mouse recorded across three different blocks of drifting gratings. Each unit exhibits a different degree of tuning curve and activity rate stability across the three blocks separated by ~15 minutes. (B-E) Analyses using data from the calcium imaging dataset during the presentation of drifting gratings. (B) PV correlation between the three blocks of drifting gratings in a single visual area (see Methods). The presented example is the average correlation matrix across all mice recorded in area V1 using two-photon Ca^{2+} imaging. Inset: the average PV correlation over all pairs of matching directions across different blocks, reveals selectivity to temporal frequency in addition to direction. (C) Mean PV correlation as a function of orientation difference across blocks for all six visual areas using two-photon Ca^{2+} imaging. All visual areas exhibit higher PV correlation values between matching orientations relative to opposite (±180°) orientations, orthogonal (±90°) orientations, and shuffled data. (D) Ensemble rate correlation between proximal blocks (separated by 15 minutes) and between distal blocks (separated by 30 minutes) of drifting gratings using the Ca^{2+} imaging dataset. The difference in ensemble rate correlations of two proximal blocks and that of two distal blocks was significant in all six visual areas ($p<10^{-3}$ for all areas, two-tailed Wilcoxon signed-rank test with Holm–Bonferroni correction). (E) Tuning curve correlation between proximal blocks and between distal blocks of drifting gratings using the Ca^{2+} imaging dataset. There was no significant difference in tuning correlations of two proximal blocks and that of two distal blocks in all visual areas except of area V1 (V1 ($Z = 3.37$, $p = 0.004$), LM ($Z = 0.85$, $p = 0.78$), AL ($Z = 2.14$, $p = 0.157$), PM ($Z = 1.49$, $p = 0.543$), RL ($Z = -1.2$, $p=0.679$), AM ($Z = -0.13$, $p=0.893$), two-tailed Wilcoxon signed-rank test with Holm–Bonferroni correction). (F-G) Analyses using data from the Neuropixels Brain Observatory group during the presentation of drifting gratings. (F) Ensemble rate correlation between proximal blocks and between distal blocks of drifting gratings using the Neuropixels dataset. The difference in ensemble rate correlations of two proximal blocks and that of two distal blocks was significant in all six visual areas ($p<10^{-3}$ for all areas, two-tailed Wilcoxon signed-rank test with Holm–Bonferroni correction). (G) Tuning curve correlation between proximal blocks and between distal blocks of drifting gratings using the Neuropixels dataset. There was no significant difference in tuning correlations of two proximal blocks and that of two distal blocks in all visual areas (V1 ($Z = 2.54$, $p = 0.065$), LM ($Z = 1.85$, $p = 0.098$), AL ($Z = 2.49$, $p = 0.065$), PM ($Z = 2.05$,

p = 0.12), RL (Z = 2.29, p=0.086), AM (Z = 1.96, p=0.12), two-tailed Wilcoxon signed-rank test with Holm–Bonferroni correction). Data in panels C-G are mean ± SEM across mice.

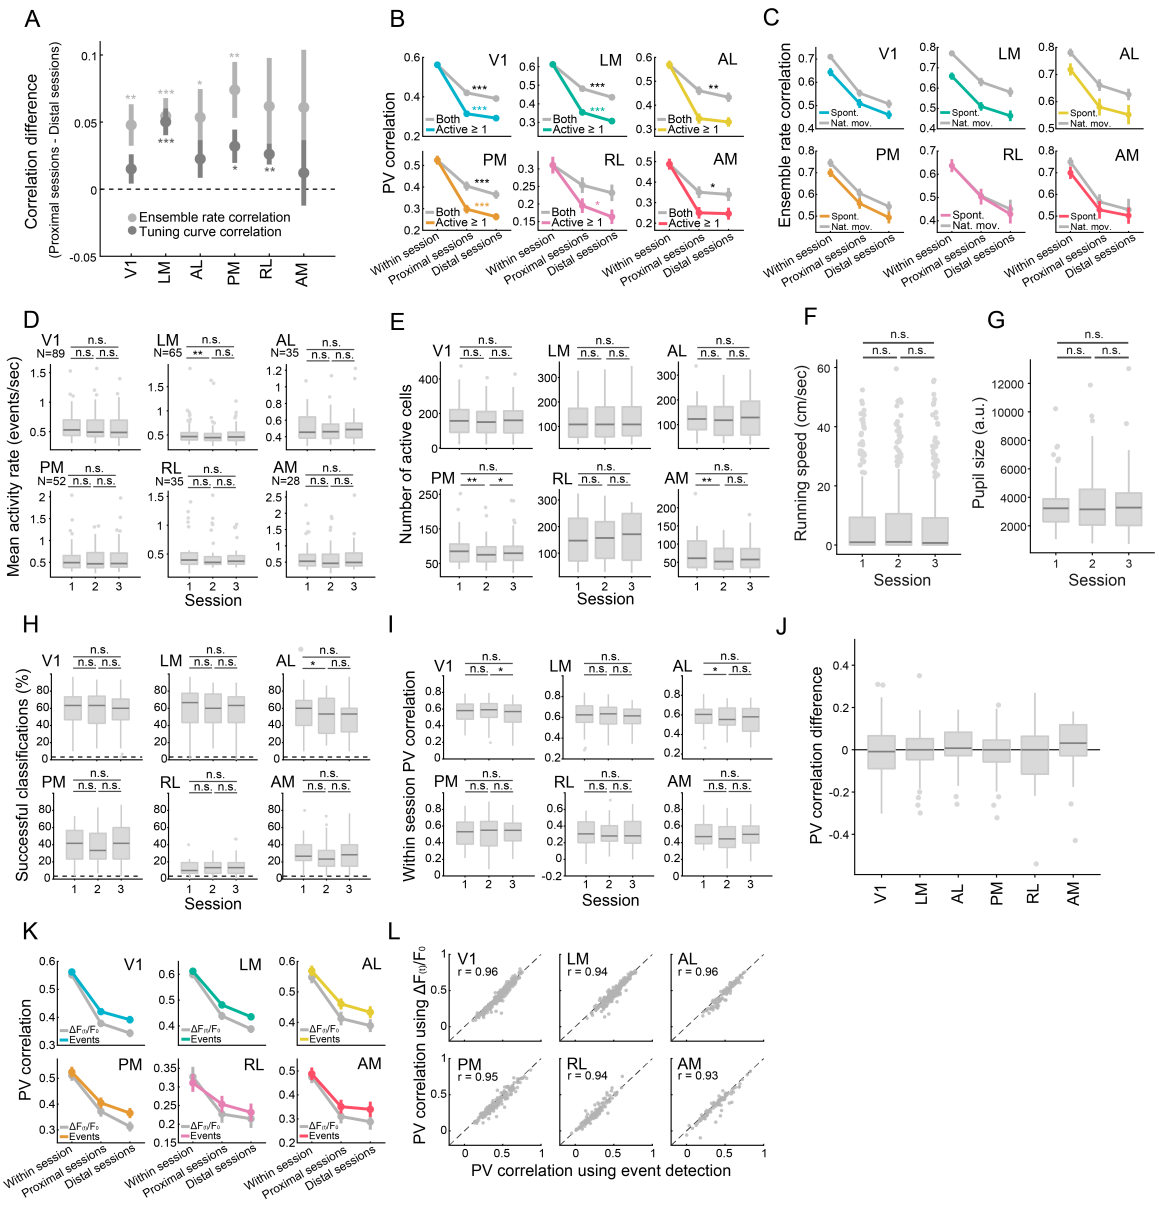

Figure S5. Stability of visual representations over days. Related to Figure 3. (A-L) Analyses using data from the calcium imaging dataset during the presentation of Natural movie 1. (A) The difference between the similarity in the representation of two temporally proximal sessions and that of two distal sessions for both ensemble rate and tuning curve correlations; V1 (Z_{Rate} = 3.31, p = 0.002; Z_{Tuning} = 2.09, p = 0.053), LM (Z_{Rate} = 4.27, $p < 10^{-4}$; Z_{Tuning} = 4.39, $p < 10^{-4}$), AL (Z_{Rate} = 2.57, p = 0.014; Z_{Tuning} = 1.77, p = 0.075), PM(Z_{rate} = 3.34, p = 0.002; Z_{tuning} = 2.65, p = 0.0159), RL (Z_{rate} = 1.53, p = 0.068; Z_{tuning} = 3.03, p = 0.005), AM (Z_{rate} = 1.83, p = 0.068; Z_{tuning} = 0.87, p = 0.19), one-tailed Wilcoxon signed-rank test with Holm–Bonferroni correction; *p < 0.5, ** p < 0.01; ***p < 0.001. (B) Repeating the analysis presented in Figure 3I for cells active in both compared time points ('active both'), and for cells that were active in at least one of the compared time points ('active ≥ 1'); one-tailed Wilcoxon signed-rank test for the difference between the correlation values between halves of two temporally proximal sessions ('proximal sessions') and between halves of two temporally distal sessions ('distal sessions'); *p < 0.5, ** p < 0.01; ***p < 0.001. (C) Ensemble rate correlation between the two halves of the same session ('within

session'), between halves of two temporally proximal sessions ('proximal sessions') and between halves of two temporally distal sessions ('distal sessions') during blocks of spontaneous activity (colored lines) and Natural movie 1 (gray lines); The difference in ensemble rate correlations between proximal sessions and distal sessions was significant for all areas during blocks of spontaneous activity (V1 ($Z = 3.41$, $p=0.001$), LM ($Z = 2.58$, $p= 0.014$), AL ($Z = 1.76$, $p = 0.046$), PM ($Z = 3.15$, $p=0.004$), RL ($Z = 2.85$, $p=0.008$), AM ($Z = 1.99$, $p=0.046$), one-tailed Wilcoxon signed-rank test with Holm–Bonferroni correction). (D) Distribution of the mean activity rates across mice across sessions. (E) Distribution of the number of active cells (at least one calcium event) for each session. (F) Distribution of the average running speed for each session. (G) Distribution of the average pupil area for each session. (H) Distribution of the within-day decoder performance for each session. Dashed lined indicate chance level. (I) Distribution of the within-day PV correlation values (correlation between the first half and second half of trials) for each session. (J) Distribution of the differences in the PV correlation values between pairs of subsequent sessions (i.e., the similarity between sessions 1 and 2 compared to that of sessions 2 and 3). V1 ($Z = -0.46$, $p = 0.64$), LM ($Z = 0.11$, $p = 0.90$), AL ($Z = 1.21$, $p = 0.22$), PM ($Z = -0.57$, $p = 0.56$), RL ($Z = -0.65$, $p=0.51$), AM ($Z = 1.45$, $p=0.14$), one-tailed Wilcoxon signed-rank without correction for multiple comparisons. (K) PV correlation between the two halves of the same session ('within session'), between halves of two temporally proximal sessions ('proximal sessions') and between halves of two temporally distal sessions ('distal sessions') using either Ca^{2+} events detection (colored lines, see Methods) or using neuropil-corrected fluorescence change ($\Delta F_{(t)}/F_0$) traces (gray lines); The difference in PV correlations between proximal sessions and distal sessions was significant for most areas when using the neuropil-corrected fluorescence change ($\Delta F_{(t)}/F_0$) traces (V1 ($Z = 3.78$, $p<10^{-3}$), LM ($Z = 5.06$, $p<10^{-3}$), AL ($Z = 1.69$, $p = 0.045$), PM ($Z = 4.41$, $p<10^{-3}$), RL ($Z = 0.64$, $p=0.258$), AM ($Z = 2.17$, $p=0.014$), one-tailed Wilcoxon signed-rank test). (L) Pearson's correlation between the PV correlation values calculated using Ca^{2+} events detection and the PV correlation values using the neuropil-corrected fluorescence change ($\Delta F_{(t)}/F_0$) traces. Data in panels A-C and K are mean ± SEM across mice. The box plots in panels D-J show the data range (whiskers), 25th and 75th percentiles (box), and median (dark line). Each data point represents an individual mouse. Outliers are marked by gray dots. In panels D-J the difference between imaging sessions was assessed by performing two-tailed Wilcoxon signed-rank tests without correction for multiple comparisons; * $p<0.05$, ** $p<0.01$.

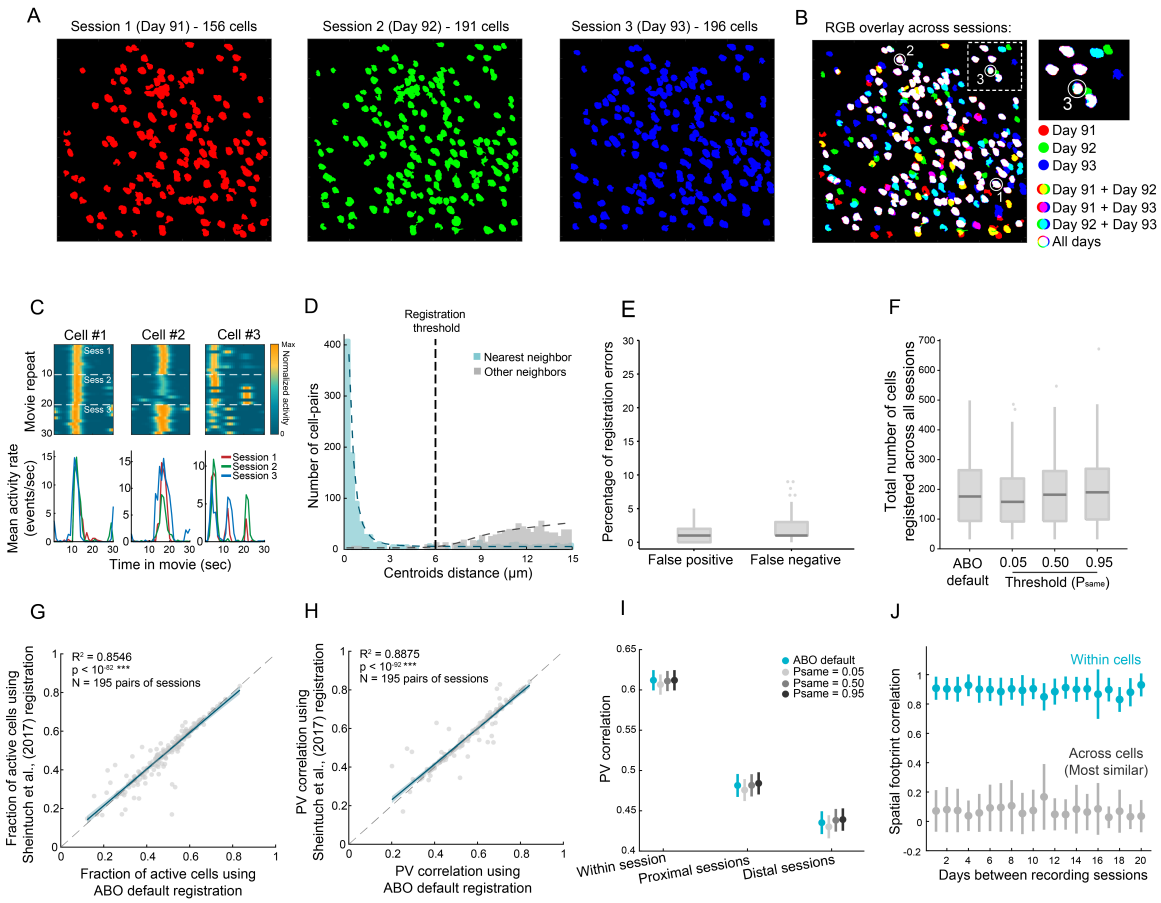

Figure S6. Verification of cell registration across sessions. Related to Figure 3. (A-J) Analyses were done using mice from in the calcium imaging dataset recorded in area LM. (A) The projection of all detected cells of a single representative mouse recorded in LM using two-photon Ca^{2+} imaging across three different recording sessions (session 1 (left), session 2 (middle), session 3 (right)). (B) Red-Green-Blue overlay of the three sessions shown in panel A after they were aligned. Inset: magnification of a selected region in the field of view. Note the clear separation between individual cells and the uniform color within cells, indicate consistent position and shape across sessions which underlie their reliable registration. (C) Responses of three V1 example cells shown in panel B across different repeats of 'Natural movie 1' spanning three recording sessions occurring on different days. Each neurons exhibits a different degree of tuning curve and activity rate stability. Note that these neurons are well isolated from the rest of the population of cells in the field of view suggesting that the observed dynamics are not due to inability to detect and register them across days. (D) Distribution of centroid distances between pairs of nearest neighbor cells (blue) and other neighboring cells (gray) from different sessions for the same representative mouse shown in panel A. Note that, while some of the neighboring cell pairs have intermediate centroid distances values, the vast majority of cell pairs exhibits either very low centroid distances (suggesting they are the same cells) or very high centroid distances (suggesting they are different cells). Black-dashed line shows the value of the centroid distance at the intersection between the models of same cells and different cells (where the probability to be the same cell P_{same} = 0.5), providing a registration threshold that is optimized to the specific dataset. (E) Estimated percentage of false-positive and false-negative registration errors per mouse (registration threshold of P_{same} = 0.5, N=65 mice imaged from area LM). (F) Total number of cells registered using the Allen Brain Observatory (ABO) default registration and the Sheintuch et al., (2017) registration with different registration thresholds (Psame values

of 0.05,0.5 and 0.95; N=65 imaged from area LM). (G) The fraction of active cells in both compared sessions using the Sheintuch et al., (2017) registration (with threshold of Psame = 0.5) relative to those obtained using ABO default registration. (H) PV correlation between pairs of sessions using the Sheintuch et al., (2017) registration (with threshold of Psame = 0.5) relative to those obtained using the ABO default registration. (I) PV correlation between the two halves of the same session ('within session'), between halves of two temporally proximal sessions ('proximal sessions') and between halves of two temporally distal sessions ('distal sessions') using both the ABO default registration and the Sheintuch et al., (2017) registration with different Psame thresholds. PV correlations decreased between sessions using all registration methods and thresholds. Data shown are mean ± SEM across mice (N=65). (J) Pearson's correlation between the imaged spatial footprints of a given single cell with its own spatial footprint on a subsequent session (blue), and between the spatial footprint of a given single cell and the most similar single cell to it on a subsequent session (gray). Data shown are mean ± SEM across pair of cells. In panel A, the age of the mouse (in days) is indicated in parenthesis. The box plots in panels E and F show the data range (whiskers), 25th and 75th percentiles (box), and median (dark line). Outliers are marked by gray dots. In panels G and H, each mouse is represented by three data points, corresponding to the three different comparisons between pairs of sessions, with a regression line (blue) ± CI of 95% (two-tailed Pearson's correlation).

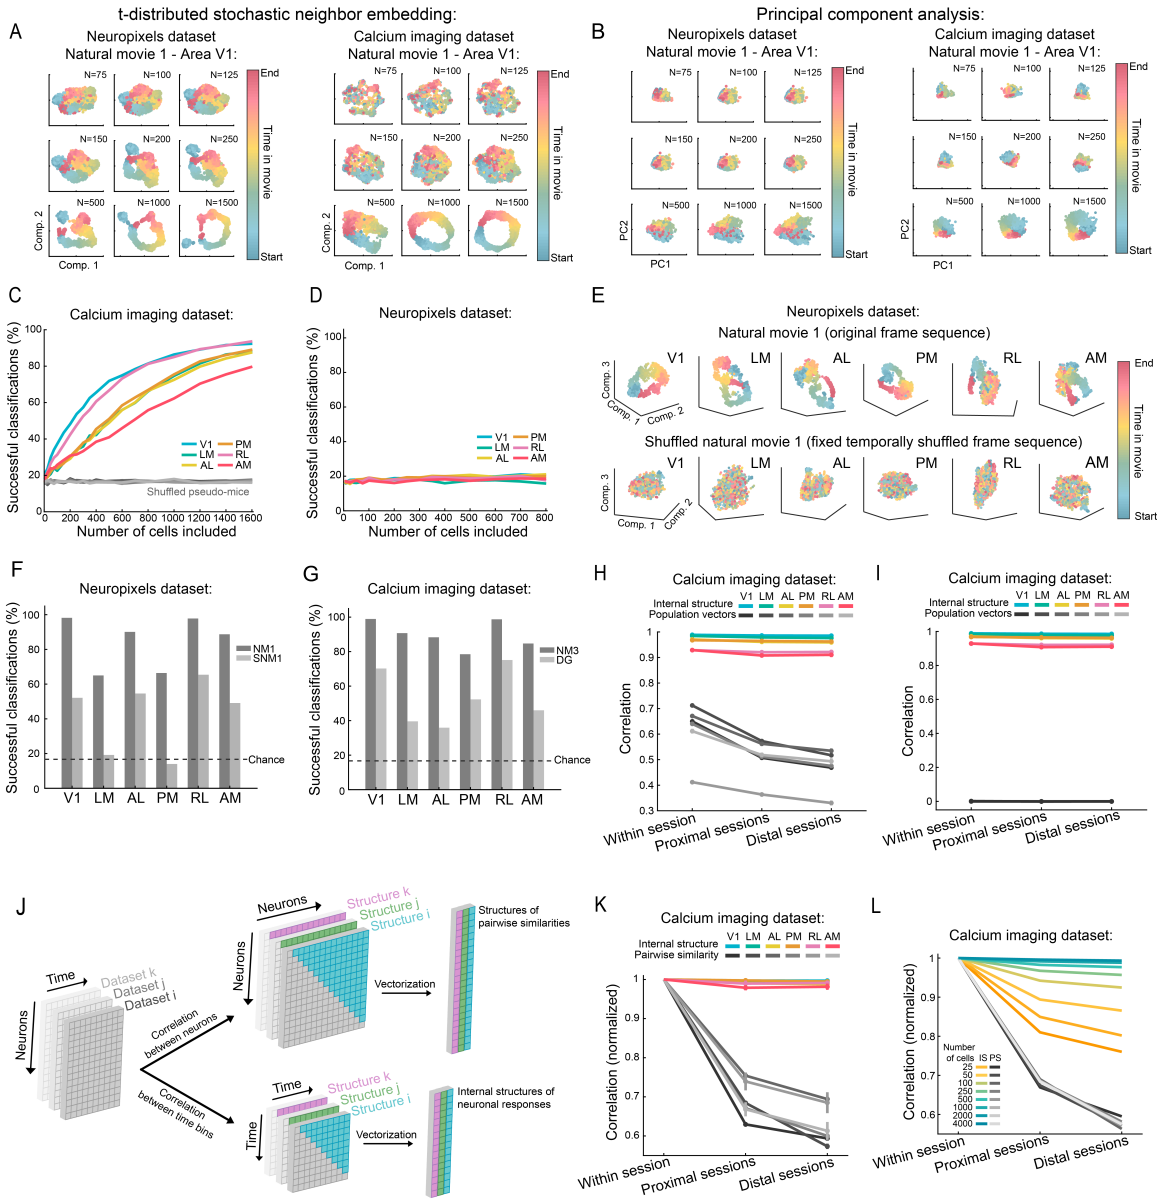

Figure S7. The unique coding properties of each visual area underlie a stereotypic and stable internal structure of neuronal population activity. Related to Figures 6 and 7. (A) Non-linear dimensionality reduction (t-distributed stochastic neighbor embedding; tSNE) applied on the population activity of a single example pseudo-mouse recorded in area V1 with either Neuropixels probes (left) or Ca^{2+} imaging (right) recovers a low-dimensional structure. The geometry of the recovered structure depends on the number of cells included in the analysis. Each point represents a single time-point of population activity during the presentation of a single Natural movie 1 repeat, and is colored according to time in the presented movie. (B) Linear dimensionality reduction (principal component analysis; PCA) applied on the population activity of a single example pseudo-mouse recorded in area V1 with either Neuropixels probes (left) or Ca^{2+} imaging (right). (C) Percentage of successful classifications of the internal structures of Natural movie 1 to their corresponding visual areas across pairs of Ca^{2+} imaging pseudo-mice as a function of the number of cells included in the analysis (data are mean across n=2000 iterations). (D) Percentage of successful classifications of the internal structures to their

corresponding visual areas across pairs of Neuropixels pseudo-mice after performing a cyclic temporal shuffle on the internal structures of each pseudo-mouse. The performance of the decoder did not exceed chance level for any number of cells included in the analysis (data are mean across n=2000 iterations). (E) Top: Dimensionality reduction (tSNE) on the population activity of a single pseudo-mouse recorded using Neuropixels probes during the presentation of Natural movie 1 recovers a distinct low-dimensional structure for each visual area. Each point represents a single time-point of population activity of a single Natural movie 1 repeat, and is colored according to time in the presented movie. Bottom: Running the same algorithm with the same parameters on the neuronal activity of the same pseudo-mouse shown in the top panels in response to the Shuffled natural movie 1, failed to recover distinct low-dimensional structure for the different visual areas. (F) Percentage of successful classifications of the internal structures of Natural movie 1 (NM1) and those of Shuffled natural movie 1 (SNM1) to their corresponding visual areas across pairs of Neuropixels pseudo-mice (data are mean across n=1000 different realizations of pseudo-mice). While the decoder performed above chance in most visual areas in both presented stimuli, its performance was better when using the internal structure of Natural movie 1 compared to that of Shuffled natural movie 1. (G) Percentage of successful classifications of the internal structures of Natural movie 3 (NM3) and those of Drifting grating (DG) to their corresponding visual areas across pairs of Ca^{+2} imaging pseudo-mice (data are mean across n=1000 different realizations of pseudo-mice). While the decoder performed above chance in most visual areas in both presented stimuli, its performance was better when using the internal structure of Natural movie 3 compared to that of Drifting grating. (H) The same analysis as in Figure 7C but without normalizing the correlation values. Data are mean ± SD across N=1000 different realizations of Ca^{+2} imaging pseudo-mice. (I) The same analysis as in H but after shuffling the identities of recorded cells in each time point before calculating the correlation between the internal structures (colored lines) or the PVs (gray lines). Data are mean ± SD across N=1000 different realizations of pseudo-mice of Ca^{+2} imaging pseudo-mice. (J) Workflow for the extraction of the structure of pairwise similarities (top) from the population neuronal responses. Starting with a matrix (n x t) containing the mean neuronal activity in each temporal bin for each dataset (e.g., movie repeat, session, mouse, stimuli etc.). Correlating each neuron with the rest of the neurons within a given dataset produces equally sized (n x n) matrices across datasets. Vectorizing the upper half of these matrices produces vectors representing the structure of pairwise similarities (vector size = $(n^2-n)/2$)). To extract the internal structure of the population activity (bottom), a similar procedure was performed, but with correlating each temporal bin with the rest of the temporal bins within a given dataset. (K) Normalized correlation between the internal structures (colored lines) or the structure of pairwise similarities (gray lines) between the two halves of the same session ('within session'), between halves of two temporally proximal sessions ('proximal sessions) and between halves of two temporally distal sessions ('distal sessions'). Data are mean ± SD across N=1000 different realizations of Ca^{+2} imaging pseudo-mice. (L) Normalized correlation between the internal structures (colored lines) or the structure of pairwise similarities (gray lines) between the two halves of the same session ('within session'), between halves of two temporally proximal sessions ('proximal sessions) and between halves of two temporally distal sessions ('distal sessions'), colored according to the number of neurons included in the analysis. Data are mean across N=1000 different realizations of Ca^{+2} imaging pseudo-mice recorded from area LM. Correlations in panels K and L were normalized to the value of the 'within session' correlation.