

Stimulus-dependent representational drift in primary visual cortex

Tyler D. Marks¹, Michael J. Goard^{1,2,3,*}

¹Neuroscience Research Institute,
University of California, Santa Barbara,
Santa Barbara, CA 93106, USA

²Department of Molecular, Cellular, and Developmental Biology,
University of California, Santa Barbara,
Santa Barbara, CA 93106, USA

³Department of Psychological & Brain Sciences,
University of California, Santa Barbara,
Santa Barbara, CA 93106, USA

*Correspondence to: michael.goard@lifesci.ucsb.edu

ABSTRACT

To produce consistent sensory perception, neurons must maintain stable representations of sensory input. However, neurons in many regions exhibit progressive drift across days.

Longitudinal studies have found stable responses to artificial stimuli across sessions in primary

5 sensory areas, but it is unclear whether this stability extends to naturalistic stimuli. We performed chronic 2-photon imaging of mouse V1 populations to directly compare the representational stability of artificial versus naturalistic visual stimuli over weeks. Responses to gratings were highly stable across sessions. However, neural responses to naturalistic movies exhibited progressive representational drift across sessions. Differential drift was present across
10 cortical layers, in inhibitory interneurons, and could not be explained by differential response magnitude or higher order stimulus statistics. However, representational drift was accompanied by similar differential changes in local population correlation structure. These results suggest representational stability in V1 is stimulus-dependent and related to differences in preexisting circuit architecture of co-tuned neurons.

15 **Keywords: Chronic imaging, Visual processing, Cortical plasticity, Representational drift**

INTRODUCTION

Ongoing experience-dependent and homeostatic synaptic plasticity suggests that neocortical connectivity is in a constant state of flux [1,2]. The potential for ongoing synaptic modification enables animals to rapidly adapt to a changing environment. However, in the face of ongoing plasticity, the cortex must somehow create stable representations of the external world and internal behavioral states in order to reliably represent the external environment and produce behaviors necessary for an animal's survival. To investigate how the brain handles the tradeoff between flexibility and long-term stability, researchers have sought to perform longitudinal measurements of the same neurons over long time periods to measure "representational drift" in neuronal response properties [3].

Many early chronic recording studies used extracellular recordings to track neurons in motor cortex and hippocampus (HPC). These regions serve as suitable targets due to their relevance in producing stable representations for stereotyped behavior, but the results have not been conclusive. Many studies in motor cortices reveal highly stable motor representations [4–6], but others report unstable individual M1 neuron tuning properties underlying stable ensemble-level representations of highly stereotyped motor actions [7–10]. It has often been difficult to draw strong conclusions from electrophysiology experiments due to low sample sizes, electrode drift, ambiguity in neuron identification across sessions, and potential sampling biases. In particular, high spike rate neurons are likely oversampled in blind electrophysiology recordings and may consequently give a biased impression of stability across the population [4,11,12]. Two-photon imaging has several advantages for chronic neural measurements. First, it has granted us insight into the dynamics of subcellular structures [13], such as the finding that sensory experience accelerates dendritic spine instability underlying synaptic turnover in sensory cortex [14] (though see [15]). Second, *in vivo* 2-photon calcium imaging enables the functional recording of large, structurally identified populations of neurons [16]. Although multi-photon imaging sacrifices temporal resolution afforded by precise spike measurement, its high spatial resolution mitigates sampling biases and reduces ambiguity in neuronal identification during chronic measurements, allowing for longitudinal studies of large neuronal populations [11,17,18].

These developments have produced the opportunity to expand the investigation of cortical stability to other brain areas over longer time periods. Recent studies in posterior parietal cortex revealed that stable learned associations can be achieved by neuronal populations in the presence of individual neurons whose coding properties continuously drift [19,20]. However, studies in sensory cortex have found more stable stimulus representation in single neurons [12,16,18,21–25]. While single-neuron representational drift has been theorized to play a role in learning in areas like motor cortex and HPC [3,8,26–28], it is less clear what purpose would be served by instability in early sensory areas, where circuitry is tasked with reliably representing external sensory information on a single exposure. Even so, there is

substantial trial-to-trial variability in single neurons [29–34], even in early sensory cortex. The evidence of stable stimulus representation in the face of this variability has led to the suggestion that the functional connectivity of sensory cortex may grant it robustness to noise while maintaining the ability to undergo experience-dependent plasticity [17].

Almost all chronic recording studies examining stability in visual cortex have used simple parameterized stimuli, such as oriented drifting gratings, as they present straightforward measurements for determining the stability of individual neuron response properties. These studies are largely in consensus, finding a high degree of stability with respect to orientation tuning, spatial frequency tuning, and size tuning [18,22,35]. Previous work has characterized the existence of small subnetworks of highly stimulus-responsive neurons coexisting against a backdrop of relatively unresponsive neurons [36,37], and a recent study reports these highly active neurons to be particularly stable amidst contextual modulation factors [38]. Responses to grating stimuli are also found to be stable following monocular deprivation [25] and even apical or basal dendritic ablation [39]. However, gratings are designed to optimally stimulate the receptive fields of visual cortical neurons, and have simple visual statistics compared to an animal's natural visual input. In addition, orientation tuning is widely believed to be 'hardwired' early in development [40–43], and iso-oriented neurons exhibit high connectivity [44,45] likely serving to stabilize orientation responses through ongoing Hebbian plasticity [45,46]. Ensembles of neurons driven by naturalistic stimuli are not necessarily iso-tuned, and almost certainly exhibit lower levels of intrinsic connectivity. As a result, there may be considerably more representational drift in response to naturalistic stimuli than to gratings, particularly over long time periods (weeks to months). A recent analysis of multi-day naturalistic movie responses from visual cortex found representational drift in several visual cortical areas, which was mainly attributed to shifts in ensemble firing rate patterns and, to a lesser extent, changes in ensemble stimulus tuning patterns [47]. These results are intriguing and unexpected, and prompt further exploration of the stability of naturalistic stimulus representations in visual cortex.

Here, we perform chronic 2-photon calcium imaging of thousands of neurons in primary visual cortex (V1) of awake, head-fixed mice viewing both oriented drifting grating stimuli as well as repeated presentations of a continuous naturalistic movie. We demonstrate stable orientation preference and high stability of grating responses across sessions, consistent with previous work. However, responses to repeated presentations of naturalistic movies exhibited progressive drift across sessions, involving the gain and loss of individual response peaks over the course of several weeks. The stimulus-dependent difference in response stability was true even for neurons exhibiting selectivity to both grating and natural movie stimuli. This representational drift was seen across cortical layers and in both excitatory and inhibitory cell types and could not be explained by response magnitude or eye movements. Finally, we found that representational drift in response to natural stimuli was accompanied by greater drift in the correlation structure of the local neuronal population during natural stimuli than for gratings.

Taken together, these results demonstrate that neurons can exhibit different levels of stability to distinct encoded features and suggests that local network connectivity may influence the stability of particular response properties.

RESULTS

Visual cortical neurons exhibit representational drift in response to natural stimuli

We performed chronic 2-photon imaging to measure single cell visual responses over many weeks in the primary visual cortex (V1) of awake, head-fixed transgenic mice expressing the calcium indicator GCaMP6s in excitatory neurons (13 fields in 12 mice, [48–50]). The mice passively viewed the same visual stimuli on every session, starting with a repeated sequence of periodic oriented drifting gratings (PDG), followed by repeated presentations of a 30 second naturalistic movie (MOV) (Figure 1A-B, Methods), including a single clip with no cuts from beginning to end. Recordings were performed at 7 ± 1 day intervals for a total of 5 to 7 weeks. To ensure our selected imaging fields solely contained V1 neurons, we used a widefield microscope to determine visual area boundaries through established retinotopic mapping procedures [51–53] (Figure 1B, Methods). Visual landmarks such as blood vessels were used to identify and align the 2-photon imaging field on a given recording session, and small differences in the alignment of the horizontal plane were corrected in post-processing using nonrigid image registration software [9]. Somatic regions of interest (ROIs) were defined based on the average fluorescence maps and average pixelwise activity maps across sessions (see Methods). To mitigate experimental artifacts stemming from misalignments of the field in the z-direction (depth) and to ensure unambiguous identification of individual neurons across sessions, we manually inspected each defined ROI across all sessions. We assigned each cell a quality rating describing the neuron's structural robustness and a presence indicator on a session-by-session basis (Figures 1C, S1A-E, [19]). ROIs were considered for analysis only if they met a sufficient quality threshold (quality index ≥ 3 ; 4,143 ROIs out of 5,689) and if the neuron was present on all analyzed sessions.

To characterize a neuron's responsiveness to a given visual stimulus, we calculated its 'reliability' on every session [54], defined by the Pearson correlation coefficient (CC) of the session-averaged activity of two random halves of trials, iterated many times and averaged (see Methods). Reliability for both stimuli followed a skewed distribution, with a higher number of neurons responding reliably to the MOV stimulus (Figure 1D). To determine which neurons were visually responsive we tested the actual reliability of the response to each stimulus against a null distribution of reliability calculated with circularly shifted data (Methods). We found that while the majority of neurons (68% of total) were visually responsive to the MOV stimulus and a smaller proportion (29% of total) were responsive to PDG, an even smaller subset (20% of total) was responsive to both stimuli, which we called 'dual-responsive' (Figure 1E). Additionally, a

neuron's response reliability for one of the stimuli was uncorrelated with that of the other
 130 (Pearson correlation between neurons' z-scored trial-averaged activity for both stimuli $r = 0.01$,
 $p > 0.05$).

Consistent with previous work characterizing the stability of visual responses to gratings
 [22,23,25,35,38], the PDG responsive neurons we recorded exhibited highly stable orientation-
 tuned responses (Figures 1G-F, S2). Orientation tuning curves were strongly aligned across
 135 sessions for both single neurons (Figure S2B) and for all tuned neurons across all mice (Figure
 1G). Changes in orientation selectivity (OSI) were minimal ($4.4 \pm 0.3\%$ change from reference
 session on average across all neurons present and tuned across all sessions) and did not
 increase as a function of elapsed time (Figure S2C). In the vast majority of cases (96% of
 neurons tuned on the first session), shifts in orientation preference over 4 - 6 weeks fell within \pm
 140 1 orientation in the sequence of presentations (Figure S2D). Neurons for which orientation
 preference shifted outside of this range made up only a small fraction of total cases (4% of
 tuned neurons).

Responses to the MOV stimulus, however, revealed striking differences when compared
 to the PDG responses (Figures S3, S4). Qualitatively, we found single neuron responses to
 145 MOV to be volatile across sessions, observing independent emergence and disappearance of
 individual response peaks between sessions (Figures 1F, S3). In some cases, this occurred
 suddenly or randomly between sessions, but in many cases such changes occurred gradually
 and continuously. Note that the unstable responses to MOV were not found only in neurons
 responding selectively to the MOV stimulus, but also in the subset of neurons responsive to
 150 both stimuli, such as the neuron in Figure 1F, in which representational drift in MOV responses
 could be directly contrasted with highly stable responses to the PDG stimulus.

In contrast to measures of reliability occurring across trials within a recording session,
 here we define 'stability' as the consistency of the average neuronal responses to the same
 visual stimulus over many sessions. To quantify changes in single-neuron response stability, we
 155 defined the 'Representational Drift Index' (RDI) between two sessions as the difference between
 the within-session between-trial correlation coefficient (CC_{WS}) and the between-session
 between-trial correlation coefficient (CC_{BS}), normalized by the sum of the two (Figure 1H).
 Pooling data from all dual-responsive neurons across all mice (824, 808, 793, 830, 761, 698
 neurons for sessions 2-7, from 13 populations in 12 mice), we found that RDI values for MOV
 160 responses were on average significantly greater than those for PDG on all sessions (paired-
 sample t-test, $p < 0.001$ for all sessions), and that the representational drift was not random, but
 progressively increased over sessions for most mice (Figures 1I, Figure S5; PDG vs. MOV
 average RDI curve slopes, $p < 0.05$, paired-sample t-test). The RDI values were not significantly
 different when including neurons responsive to single stimuli, when compared to only dual-
 165 responsive neurons (comparing single stimulus responsive to dual-responsive, $p > 0.05$ for all
 weeks for both stimuli, Wilcoxon ranksum tests). Additionally, the significant differences

between RDI values for each stimulus were also present when considering all neurons, regardless of stimulus responsiveness ($p < 0.001$ for all weeks, paired-sample t-test).

Despite a robust difference in RDI between PDG and MOV stimuli on average, we observed some variability between mice. While most clearly exhibited a difference in stability between the stimuli, this difference was either small or not significant in a subset of mice (Figure S5, mice 5, 11, 12). Additionally, while some mice exhibited a progressive divergence between the PDG and MOV RDI curves (mice 1, 3, 4, 6, 8, 9), others exhibited higher baseline RDI levels for MOV stimuli (mice 2, 7, 10). The reasons for these differences were unclear given that there were no identifiable correlated differences between groups of mice. Furthermore, the exclusion of neurons via thorough *a priori* manual inspection (Figure 1A-D) eliminated any artifacts that may have resulted from misalignment of neurons across sessions in some mice, and repeating our analysis using different ROI quality thresholds yielded robust results (Figure S1F). One potential source of variability between mice is that craniotomies for cranial window implants can cause inflammation and microglial activation [55], which may in turn result in increased synaptic turnover in superficial layers [56]. It is possible that exposure of the brain in the cranial window preparation may be partially responsible for the differences between mice we observe. Thin-skull preparations for transcranial 2-photon imaging have been shown to result in a relatively low degree of induced synaptic plasticity [15]. To test whether the surgical preparation contributed to the drift in MOV responses, we performed the same chronic imaging experiments in mice ($n = 4$) with a thin-skull preparation and found a similar level of differential stability between the two stimuli (Figure S6).

Responses to natural movies exhibit greater drift independent of magnitude

One factor that could explain differential stability on the level of single neurons is the wide diversity of stimulus responsiveness that exists in sensory cortices [36,37]. A recent study in mouse V1 found considerable differences in long-term stability between strongly and weakly visually responsive neurons [38]. In addition, it is possible that early accounts of stable cortical neurons from electrophysiology data may be influenced by biased sampling of particularly active neurons. We asked if a similar relationship between visual responsiveness and stability could be found in our data. On average, the neurons responsive to both stimuli with the strongest session-average z-scored responses were more stable than those with the weakest (Figures 2A, S7A). The difference was modest but statistically significant for both stimuli (0.04 difference in median RDI between top and bottom quartiles for PDG, $p < 0.01$; 0.07 difference for MOV, $p < 0.001$; Wilcoxon rank-sum test). Lack of striking, conclusive results from this analysis led us to ask if there may be further differences related to responsiveness that are not observable on the single-neuron level.

The sparse and episodic structure of V1 neuron responses to naturalistic stimuli has been well-characterized [57–59] and was qualitatively visible in our response data (Figure S3). Most, but not all, neurons responsive to MOV exhibited more than one time-distinct response ‘event’ across the course of the 30 s stimulus (Figures 1F, 2B, S3). We asked if a relationship between response magnitude and stability was present at the level of the individual visual response events within each neuron. For both stimuli, we used $\Delta F/F$ data to define response event periods based on the statistical significance of each frame’s single-trial responses compared to single-trial average baseline responses, performed separately for each session (see Methods). PDG and MOV responses in dual-responsive neurons yielded, on average, ~1.9 and ~2.4 events per neuron, respectively, with a sizable subset of neurons responding to MOV with 3 or more events (51.6% of dual-responsive neurons, 57.2% of MOV responsive neurons, Figure 2C). The event rate distribution for PDG was skewed further towards 1-2 events per neuron, reflecting selectivity towards only 1 or 2 grating orientations and occasional weak responses to neighboring orientations (Figure 2C). Events were categorized as either growing, decaying, or remaining static over time based on a statistical comparison of their single trial z-score values on the first two sessions versus the last two sessions. We found that on average across all imaging fields, PDG responses exhibited a much higher proportion of events that remained static between sessions than MOV responses did (~73% vs. ~54% for PDG and MOV respectively, $p < 0.001$, paired-sample t-test). Of the remaining 46% of MOV events, ~65% were events that decayed and ~35% were those that grew between sessions (Figure 2D). Examining the strongest event periods (events surrounding neurons’ peak response times) revealed that while a plurality of these events remained static over time, over half were classified as either growing or decaying (Figure 2E).

To compare an event’s magnitude and its stability more directly, we defined an event’s instability based on the normalized difference between its average ‘late’ z-score and its average ‘early’ z-score (last two versus first two sessions, respectively, see Methods). This revealed that for both stimuli, an event’s stability increased as a function of its overall response magnitude (Figure 2F, Figure S7A). This trend occurred to a similar extent for both stimuli, and event instability was greater on average for MOV events than they were for PDG events regardless of event magnitude (Figure 2G). As a result, differences in stability between the two stimuli cannot be explained by the presence of higher magnitude events in PDG responses compared to MOV responses. Additionally, we wondered if we could link an event’s stability to its redundancy across all neurons in its respective population, as previous findings show stability in sparse population activity [12]. Indeed, event stability seemed to decrease as a function of redundancy across all other neurons in the population (Figure S7B, significant only for PDG). In line with this, we found that event redundancy decreased significantly as a function of event response magnitude for both stimuli (Figure S7C), which contravenes the possibility that stronger events may be encoding highly salient visual features and thus may be represented by a greater fraction of neurons in the population.

Differential stability is not explained by behavioral variables

We performed a number of control experiments to account for external factors that could contribute to these results. A previous study reported that despite stable orientation tuning across days, pupil size was positively correlated with trial-to-trial variability [35]. We used an infrared camera to track the pupils of a subset of mice ($n = 4$) to determine whether the differential stability could be related to differences in a mouse's arousal state across sessions as measured by pupil size [60], average pupil movement, or pupil location. We observed that large eye movements were rare and mostly consisted of brief deviations from an otherwise stable average position that changed minimally between sessions (Figure S8A-B). On average, within-session pupil movement relative to a session's average pupil centroid location decreased over time (Spearman correlation $r = -0.37$, $p < 0.05$), and average pupil area decreased over time (Spearman correlation $r = -0.47$, $p < 0.01$), indicating a general decrease in arousal state over sessions (Figure S8C-D). Although we habituated our mice prior to imaging, this finding is consistent with further acclimation to the recording environment and the visual stimulus presentation. Within-session pupil movement was not correlated with RDI (Pearson correlation $r = -0.16$, $p > 0.05$), but relative pupil size showed a small anticorrelation with RDI (Pearson correlation $r = -0.38$, $p < 0.05$), suggesting that changes in arousal could be a factor contributing to the observed instability in the initial sessions. However, pupil size stabilized after the third session, and is unlikely to contribute to progressive drift observed in later sessions. Finally, in all individual mice, changes in average pupil centroid location between sessions were uncorrelated with RDI (Pearson correlation $r = -0.06$, $p > 0.05$, Figure S8E). In summary, though there were minor changes in pupil area across sessions, there were no changes in eye movements, and the minor fluctuations in mouse behavioral state are not sufficient to explain progressive increases in RDI, nor the difference in RDI between MOV and PDG stimuli.

Next, we considered the possibility that any subtle shifts in a neuron's spatial receptive field across sessions would manifest only in MOV responses and not PDG responses. While changes in a neuron's receptive field would likely capture different visual information for the MOV stimulus, this might not be the case for PDG due to its repeated spatial pattern. To account for this, we showed a subset of mice ($n = 4$) a spatial receptive field mapping stimulus (presented last on every session) to measure neurons' preferred altitude and azimuth on a given session (Methods). In all tested mice, changes in neurons' preferred altitude and azimuth were minimal and showed no correlation with higher RDI (Figure S9), indicating that receptive field shifts did not influence our findings.

Representational drift is distributed across cortical layers and cell types

Because the distinct layers of visual cortex have stereotyped inputs, outputs, and connectivity within and between other layers [61], we wondered if they might show differences in their

capacities for representational drift. As L4 is the primary recipient of input from the lateral geniculate nucleus (LGN), we hypothesized that it would be comprised of neurons with relatively stable responses to the MOV stimulus compared to neurons in L2/3 and L5, which receive more processed input [62]. However, a recent study of natural movie responses in mouse visual cortex found higher stability of firing rate and stimulus tuning in L2/3 and L5 neurons compared to L4 neurons over multiple days [47]. To investigate differences between layers, we chronically implanted a subset of mice ($n = 4$) with a custom glass microprism that granted optical access to nearly the full cortical column [63] (Figure 3A-B), and allowed us to separate neurons by cortical layer based on ROI density (Figure 3C). Comparing average RDI values within each layer revealed that cortical layer did not have a significant effect on the instability of individual neurons (Figure 3D-E; two-way ANOVA, stimulus type $p = 3.32 \times 10^{-11}$, layer $p = 0.12$). This homogenous distribution of neurons exhibiting representational drift suggests that the transformation of presumably stable visual input to unstable neural responses occurs either at or before the first stage of information processing in cortex and propagates as the signal continues.

We next investigated if MOV representational drift varied by cell type. Inhibitory interneurons exhibit broader and less selective tuning than excitatory cells, believed to be due to promiscuous sampling from local excitatory populations [64]. This in turn may result in comparatively stable visual responses. Inhibitory interneurons neurons in mouse motor cortex [65] and zebra finch HVC [66] (but see [67]) have been shown to exhibit considerable long-term stability alongside less stable single excitatory neurons in the context of a motor task. To investigate this, we repeated our experiments in L2/3 of transgenic mice expressing GCaMP6s in GAD2+ inhibitory neurons ($n = 4$, Figure 4A-B). Inhibitory neurons exhibited clear visually-evoked responses, and as expected most neurons responsive to PDG were either very broadly orientation tuned or untuned (Figure 4C). As in our excitatory population results, of all 232 well-tracked neurons, only a subset ($14 \pm 5\%$) were visually responsive to both PDG and MOV (Figure 4D). We considered this subset of dual-responsive neurons as we did for the excitatory populations. Similar to the between-session changes observed in excitatory neurons, we observed the appearance and disappearance of individual visual response events within neurons across sessions (Figure 4C). As for excitatory neurons, we found that responses to MOV exhibited greater representational drift than responses to PDG (Figure 4E-F). Taken together, these results indicate that progressive drift of responses to naturalistic stimuli is present across the entire cortical network.

Representational drift is accompanied by changes in population correlation structure

To further investigate the underpinnings of the stimulus-dependent instability, we hypothesized that the greater representational drift in the MOV responses might be due to the higher order image statistics of the stimulus. We reasoned that the additional cortical processing required for

encoding more complex visual features might lead to more variable representation compared to that of oriented gratings. To test this, we performed parametric phase scrambling on the original MOV stimulus [68], which randomized the phase structure of the images in the movie while
 320 maintaining the amplitude spectrum and lower level image properties. To a subset of mice ($n = 2$), we presented the PDG stimulus in addition to a stimulus consisting of 3 versions of the MOV stimulus: 100% phase scrambled, 50% phase scrambled, and 0% phase scrambled (the original movie), presented for 20 repeats each in random trial order (Figure 5A). However, we observed no major differences between the RDI curves for the three MOV stimuli, indicating that the
 325 increased representational drift is not due purely to higher order image statistics (Figure 5B).

Our earlier results indicate that what separates stable and unstable visual responses in V1 is the identity of the visual input, and it would be inaccurate to conclude that a neuron itself is intrinsically stable or unstable. This raises the possibility that the stability in different conditions is a function of the particular functional inputs. We investigated whether or not the observed
 330 stimulus-dependent representational drift in single neurons could be linked to stimulus-dependent changes in the correlation structure of the local population, as even in primary sensory cortex neuronal responses are strongly influenced by local circuit activity [33,69–71]. Neurons co-responsive to naturalistic stimulus features may have less pre-existing connectivity compared to neurons co-tuned to specific orientations of the PDG stimulus, which have been
 335 found to exhibit high interconnectivity [45,46,72]. Does the MOV population correlation structure change across weeks while PDG correlation structure remains consistent? To answer this, on every session for both stimuli we calculated pairwise between-neuron signal correlations (SC), which capture shared responses to stimulus-driven input, and noise correlations (NC), which capture shared trial-to-trial variability (see Methods). Visualizing SC structure on the final
 340 recording session (D_{final}) versus the reference session (D_0) for individual mice showed larger differences for MOV than for PDG (Figure 6A), and absolute changes in mean SC across neurons were greater on average for MOV than for PDG in 12 out of 13 fields (Figure 6B-C). Consistent with another study investigating cross-session changes in grating stimulus correlation structure [35], both the MOV and PDG SC structures became gradually less similar
 345 to D_0 over time, as measured by the Pearson correlation between each matrix and the D_0 matrix (Figure 6D). However, the MOV correlation structure exhibited greater divergence as a function of time since D_0 (Figure 6D). Repeating these analyses for noise correlations revealed a weaker effect than that of SC, though the trend was in the same direction (Figure S10), possibly due to relatively low levels of noise correlations in our data. Previous studies indicate that NC
 350 measurements can vary significantly and may be affected by a variety of factors, such as spike rates and behavioral state [69,73]. For this reason, we cannot conclusively determine whether drift is present in NCs as well. In summary, the differential stability of responses to PDG and MOV stimuli does not appear to be related to higher order statistical features in the MOV stimulus, but rather to the decreased stability of the co-tuned neuronal ensembles driven by
 355 each stimulus.

DISCUSSION

Here we investigated the extent to which representational drift exists in populations of individual neurons through a comparison of neural responses to artificial (PDG) and naturalistic (MOV) visual stimuli in the primary visual cortex of awake mice over many weeks. We found that neurons exhibited highly consistent orientation tuned responses to the grating stimulus over all recording sessions (Figure 1, Figure S2), consistent with previous work [18,22,35,38]. However, many of the same neurons displayed less stable representations of the natural movie stimulus (Figures 1-2, Figures S3-5). Contrary to our hypothesis that geniculate inputs to L4 may grant it more consistent stimulus representation, translaminar imaging revealed that this differential stability existed not only in L2/3 and L5, but also in L4 (Figure 3). Repeating our experiments in populations of inhibitory interneurons yielded a similar difference in stability between the two stimuli (Figure 4). Importantly, the observed drift could not be explained by external factors such as eye movement, arousal, or surgical preparation (Figures S6,8,9). Additionally, the higher order image statistics of natural stimuli were not found to contribute to the differential drift between stimuli (Figure 5). Finally, we found that the drift observed in MOV responses was accompanied by changes in the population correlation structure, potentially due to lack of strongly connected iso-tuned subnetworks that respond to grating stimuli (Figure 6).

From work that utilizes artificial stimuli such as gratings, compelling evidence suggests that basic neuronal response properties such as orientation tuning, spatial frequency tuning, and size tuning are remarkably stable over time [18,22,35]. However, another recent study reports that naturalistic stimulus responses in mouse visual cortex exhibit significant representational drift [47]. By performing a direct comparison of artificial and naturalistic visual stimulus response in the same animals, our experiments allowed us to reconcile this discrepancy by showing that response stability is dependent on stimulus, and not simply an intrinsic property of a given neuron.

Further characterization of the representational drift observed in MOV responses revealed two important results. First, we showed that differential stability between stimuli existed not only at the single-neuron level but also at the scale of individual response events, wherein response events during the MOV stimulus were more likely to experience bidirectional amplitude changes across sessions than those that occurred during the PDG stimulus (Figure 2D, E). Second, consistent with previous work demonstrating a relationship between response magnitude and stability [38], we found such a correlation to be evident on the level of individual events (Figure 2F-G, Figure S7A). These results indicate representational drift may operate at the level of specific functional inputs, rather than at the level of single neurons. Deitch et al. found similar single neuron representational drift that appeared to be dominated by changes in firing rate rather than changes in stimulus tuning [47]. However, our results indicate that these two properties are interconnected in that changes in firing rate can manifest as changes in tuning over long time scales. We also noted an inverse relationship between event response

magnitude and event redundancy, wherein the strongest events are not only the most stable but also the least redundant across other neurons in the recorded population (Figure S7B-C). Neurons that primarily respond with strong, stable events that are relatively sparsely represented across the population may constitute a response class previously characterized as “soloists”, and those that primarily respond with weak, redundant events may represent “choristers” more strongly coupled to the local population [74,75]. This is consistent with the developing characterization of sensory cortices as being dominated by subnetworks of particularly stable and robustly responding neurons against a background of weakly responsive neurons [25,36,37,72].

One question still outstanding is the extent to which representational drift is present throughout the visual hierarchy. Presumably, repeated visual input is encoded in a consistent fashion within the retina and the earliest stages of the visual pathway. Based on the results that unstable responses arise as early as layer 4 of V1, one hypothesis might be that primary visual cortex as a whole represents a relatively intermediate processing stage in which single neurons are permitted greater coding flexibility, and that LGN, which receives input directly from the retina, exhibits greater stability. However, given recent results showing that retinogeniculate and retinocollicular boutons are modulated by arousal [76,77], representational drift might begin even earlier. Future studies using glass microprisms or GRIN lenses to allow chronic optical access to LGN [77] will allow elucidation of these possibilities. Finally, it remains to be seen whether stimulus representations stabilize as signals propagate to higher visual areas, or if the capacity for drift is preserved across the entire visual pathway. To answer this, another future direction would be to perform similar longitudinal imaging in higher visual areas concurrently with recordings from V1 in the same animals.

Representational drift appears in varying degrees across different brain structures, and most prevalently outside of sensory cortex. In motor and parietal cortices, evidence suggests that population-level activity patterns remain intact, even as the underlying neurons representing task or motor information experience significant drift [7,8,19,20,66], although see [4,6]. Several studies have identified strong representational drift in hippocampal place cells that exhibit dynamic cognitive maps of the environment [27,28,78,79], though results vary across hippocampal subregions [80]. One possible interpretation is that the opportunity for representational drift increases with distance from the sensory input, given that studies in sensory cortices have largely found relatively stable tuning properties. Additionally, in association areas, the capability of single neuron representations to drift could play a role in providing flexibility for coding learned associations [19]. However, our data show that this hypothesis may be overly simplistic, as stability appears to be not only a property of the individual neurons or brain region, but also a function of encoded information. As we show in V1, the same neuron can simultaneously exhibit stable tuning properties and more malleable

responses to naturalistic stimuli. It remains to be seen if this holds true for motor and association cortex as well.

How would stable tuning properties be maintained in the presence of representational drift? One possibility is that since subnetworks of neurons iso-tuned for orientation exhibit high connectivity (likely established early in development) [45,46,72], fundamental tuning properties would be stabilized throughout the lifetime of the neurons. However, neurons co-responsive to particular time points in a natural stimulus would not necessarily exhibit any higher connectivity than expected by chance. This could lead to the differential shifts in correlation structure we observed between the two stimuli, and the resulting differences in representational drift (working model illustrated in Figure 6E). Taken together, we propose that co-tuned subnetworks of neurons can preserve fundamental tuning properties while allowing for more flexible responses to complex naturalistic stimuli. Such a governing principle is potentially applicable to other cortical regions, where highly interconnected neuronal subnetworks may preserve stable encoding of fundamental input features while maintaining flexibility in their responses to more complex inputs.

METHODS

Animals

For cortex-wide calcium indicator expression, Emx1-IRES-Cre (Jax Stock #005628) x ROSA-LNL-tTA (Jax Stock #011008) x TITL-GCaMP6s (Jax Stock #024104) triple transgenic mice ($n = 10$) or Slc17a7-IRES2-Cre (Jax Stock #023527) x TITL2-GC6s-ICL-TTA2 (Jax Stock #031562) double transgenic mice ($n = 2$) were bred to express GCaMP6s in cortical excitatory neurons. For interneuron activity measurements, we used GAD2-IRES-Cre (Jax Stock #028867) x TITL2-GC6s-ICL-TTA2 (Jax Stock #031562) double transgenic mice ($n = 4$). For imaging experiments, 12 - 30 week old mice of both sexes (5 males and 7 females for transgenic excitatory populations, 2 males and 2 females for inhibitory populations) were implanted with a head plate and cranial window and imaged starting 2 weeks after recovery from surgical procedures and up to 10 months after window implantation. The animals were housed on a 12 hr light/dark cycle in cages of up to 5 animals before the implants, and individually after the implants. All animal procedures were approved by the Institutional Animal Care and Use Committee at University of California, Santa Barbara.

Surgical Procedures

All surgeries were conducted under isoflurane anesthesia (3.5% induction, 1.5 - 2.5% maintenance). Prior to incision, the scalp was infiltrated with lidocaine (5 mg kg⁻¹,

subcutaneous) for analgesia and meloxicam (2 mg kg⁻¹, subcutaneous) was administered preoperatively to reduce inflammation. Once anesthetized, the scalp overlying the dorsal skull was sanitized and removed. The periosteum was removed with a scalpel and the skull was abraded with a drill burr to improve adhesion of dental acrylic.

For standard cranial windows, a 4 mm craniotomy was made over the visual cortex (centered at 4.0 mm posterior, 2.5 mm lateral to Bregma), leaving the dura intact. A cranial window was implanted over the craniotomy and sealed first with silicon elastomer (Kwik-Sil, World Precision Instruments) then with dental acrylic (C&B-Metabond, Parkell) mixed with black ink to reduce light transmission. The cranial windows were made of two rounded pieces of coverglass (Warner Instruments) bonded with a UV-cured optical adhesive (Norland, NOA61). The bottom coverglass (4 mm) fit tightly inside the craniotomy while the top coverglass (5mm) was bonded to the skull using dental acrylic.

For columnar imaging, we used custom-designed microprisms (Tower Optical) that had a 200 μ m square base and a 700 μ m right angle prism (design available on our institutional lab website: <https://goard.mcdb.ucsb.edu/resources>). The hypotenuse of the right angle prism was coated with aluminum for internal reflectance. The microprism was attached to a 5 mm diameter coverglass (Warner Instruments) with a UV-cured optical adhesive (Norland, NOA61). Prior to implantation, a 3-4 mm craniotomy was made over primary visual cortex (centered at 4.0 mm posterior, 2.5 mm lateral to Bregma). A 1 mm length medial-to-lateral incision was then made through the dura and cortex to a depth of 1 mm with a sterilized diamond micro knife (Fine Science Tools, #10100-30) mounted on a manipulator, taking care to avoid blood vessels (approximately 4.0 mm posterior, 2.5 mm lateral to Bregma). Gelfoam (VWR) soaked in sterile saline was used to remove any blood from the incision site. Once the incision site had no bleeding, the craniotomy was submerged in cold sterile saline, and the microprism was lowered into the cortex using a manipulator, with the imaging face of the prism facing anterior. The microprism assembly was completely lowered until the coverglass was flush with the skull, then the edges of the window were sealed with silicon elastomer (Kwik-Sil, World Precision Instruments), then with dental acrylic (C&B-Metabond, Parkell) mixed with black ink. The microprism implant enabled imaging from 200 - 900 μ m below the coverglass surface, corresponding to approximately 100 - 800 μ m into the cortex due to approximately 100 μ m of dimpling near the top corner of the prism. Cortical layers 2-5 were visible in all recordings, with partial visibility of layer 1 and layer 6. The microprism implantations were stable for up to six months following the surgery, similar to previously published results [63].

For thinned-skull experiments, a 3 - 4 mm diameter patch of skull over V1 was carefully thinned using a drill burr and then a rubber polishing bit until fully translucent, being careful to keep the thinned skull wet with sterile saline. After the thinning was complete, the saline was wicked away, a drop of cyanoacrylate (Loctite 406) was placed on the thinned skull and a 3 mm coverglass (Warner) was lowered using a manipulator until flush with the thinned skull. The implant was then sealed with dental acrylic (C&B-Metabond, Parkell) mixed with black ink. The thinned-skull window implantations were stable for up to six months following the surgery, similar to previously published results [81].

After cranial window implantation, a custom-designed stainless steel head plate (eMachineShop.com) was affixed using dental acrylic (C&B-Metabond, Parkell) mixed with black ink. After surgery, mice were administered carprofen (5 - 10 mg kg⁻¹, oral) every 24 hr for 3 days to reduce inflammation. The full specifications and designs for head plate and head fixation hardware are available on our institutional lab web site (<https://goard.mcdb.ucsb.edu/resources>).

Visual Stimuli

All visual stimuli were generated with a Windows PC using MATLAB and the Psychophysics toolbox [82]. Stimuli used for widefield visual stimulation during retinotopic mapping were presented on an LCD monitor (43 x 24cm, 1600 x 900 pixels, 60Hz refresh rate) positioned 10 cm from the eye at a 30 degree angle to the right of the midline, spanning 130° (azimuth) by 100° (elevation) of visual space. For chronic two-photon imaging experiments, visual stimuli were presented on an LCD monitor (17.5 x 13 cm, 800 x 600 pixels, 60Hz refresh rate) positioned 6 cm from the eye at a 30 degree angle right of the midline, spanning 120° (azimuth) by 100° (elevation) of visual space. Physical bars affixed to the table and reference point distance measurements were used to ensure that the stimulus monitor was fixed in the exact same location for each experiment.

Retinotopic mapping stimuli consisted of a drifting bar with contrast reversing checkboard (0.05 cycles degree⁻¹ spatial frequency; 2 Hz temporal frequency) that was spherically corrected to account for visual field distortions due to the proximity of the monitor to the mouse's eye. The stimulus was swept in the four cardinal directions, repeated 20-60 times.

For passive drifting grating (PDG) stimulation, 12 full-contrast sine wave gratings (spatial frequency: 0.05 cycles/deg; temporal frequency: 2 Hz) were presented full-field, ranging from 0° to 330° from vertical in 30° increments. We presented 8 repeats of the drifting grating stimulus;

a single repeat of stimulus consisted of all 12 grating directions presented in order for 2 sec with a 4 sec inter-stimulus interval (gray screen) to allow calcium responses to return to baseline between presentations.

For natural movie visual stimulation, we displayed a grayscale 30 sec clip from *Touch of Evil* (Orson Wells, Universal Pictures, 1958) containing a continuous visual scene with no cuts (https://observatory.brain-map.org/visualcoding/stimulus/natural_movies). The clip was contrast-normalized and presented at 30 frames per second. We presented 30 repeats of the natural movie stimulus; each repeat started with 5 sec of gray screen, followed by the 30 sec of movie.

For phase scrambled natural movie visual stimulation, the original 30 sec natural movie clip from *Touch of Evil* (Orson Wells, Universal Pictures, 1958) was separated into phase and amplitude spectra using 2-D fast Fourier transform, and its phase structure was randomly scrambled before being reconstructed with the original amplitude data to form the new movie. For the 50% phase scrambled movie, only a random half of the image's phase elements were scrambled. We presented 20 repeats each of the original natural movie stimulus, the 50% phase scrambled movie, and the 100% phase scrambled movie in a random order on every session; each repeat started with 5 sec of gray screen, followed by 30 sec of one of the movies. Each version was contrast-normalized and presented at 30 frames per second.

The receptive field mapping stimuli consisted of full screen length vertical (azimuth) or horizontal (altitude) bars (20° width) with contrast reversing checkerboards (0.04 cycles degree⁻¹ spatial frequency; 5 Hz temporal frequency). These bars were presented for 1 sec at a number of locations spanning the height (altitude) and width (azimuth) of the screen in random order on every repeat (30 locations for altitude, 40 locations for azimuth) for a total of 10 repeats for each direction, with a 2 sec gray screen between repeats.

Eye Tracking

To confirm that representational drift was not due to eye movements, we performed eye tracking experiments on 4 mice. These mice were head-fixed identically to imaging experiments, but an IR camera (Thorlabs DCC1645C with IR filter removed; Computar T10Z0513CS 5-50mm f/1.3 lens) was placed such that the image sensor was located immediately lateral to the stimulus monitor. Video was acquired at 15 fps and images were analyzed offline in MATLAB (Mathworks).

The pupil was identified for each frame using an automated procedure. In brief, raw images were binarized based on pixel brightness, and the resulting images were

morphologically cleaned by removing isolated pixels. For the initial frame, the pupil was manually chosen. For subsequent frames, the pupil was chosen from potential low intensity regions based on a linear combination of size, location, and eccentricity of the pupil in the previous frame.

Widefield Imaging

After >2 weeks of recovery from surgery, GCaMP6s fluorescence was imaged using a custom widefield epifluorescence microscope [51]. In brief, broad spectrum (400-700 nm) LED illumination (Thorlabs, MNWHL4) was band-passed at 469nm (Thorlabs, MF469-35), and reflected through a dichroic (Thorlabs, MD498) to the microscope objective (Olympus, MVPLAPO 2XC). Green fluorescence from the imaging window passed through the dichroic and a bandpass filter (Thorlabs, MF525-39) to a scientific CMOS (PCO-Tech, pco.edge 4.2). Images were acquired at 400 x 400 pixels with a field of view of 4.0 x 4.0 mm, leading to a pixel size of 0.01 mm pixel⁻¹. A custom light blocker affixed to the head plate was used to prevent light from the visual stimulus monitor from entering the imaging path.

Widefield Post-processing

Images were acquired with pco.edge camera control software and saved into multi-page TIF files. All subsequent image processing was performed in MATLAB (Mathworks). Visual field sign maps were derived and segmented as previously described [51]. After processing, borders were drawn around each patch, and resulting patches were compared against published sign maps for both size and sign to label each patch as a visual area. Visual areas V1, LM, AL, PM, LI, RL, and AM were present in all mice. Area V1 was targeted for all further recordings.

2-Photon imaging

After >2 weeks' recovery from surgery, GCaMP6s fluorescence was imaged using a Prairie Investigator 2-photon microscopy system with a resonant galvo scanning module (Bruker). For fluorescence excitation, we used a Ti:Sapphire laser (Mai-Tai eHP, Newport) with dispersion compensation (Deep See, Newport) tuned to $\lambda = 920$ nm. Laser power ranged from 40 - 75 mW at the sample depending on GCaMP6s expression levels. Photobleaching was minimal (<1% min⁻¹) for all laser powers used. For collection, we used GaAsP photomultiplier tubes (Hamamatsu). A custom stainless-steel light blocker (eMachineShop.com) was mounted to the head plate and interlocked with a tube around the objective to prevent light from the visual stimulus monitor from reaching the photomultiplier tubes. For imaging, we used a 16x/0.8 NA microscope objective (Nikon) at 1.2 (690 x 690 μ m) to 2x (425 x 425 μ m) magnification. During

imaging experiments, the polypropylene tube supporting the mouse was suspended from the
behavior platform with high tension springs (Small Parts) to reduce movement artifacts.

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

2-Photon Post-processing

Images were acquired using PrairieView acquisition software and converted into TIF files. All subsequent analyses were performed in MATLAB (Mathworks) using custom code (<https://goard.mcdub.ucsb.edu/resources>). First, images were corrected for X-Y movement within each session by registration to a reference image (the pixel-wise mean of all frames) using 2-dimensional cross correlation. Next, to align recordings to the reference session, we used semi-automated rigid registration, similar to prior work [9,83]. First, anchor points were automatically generated from matching image features between average projections detected by the 'Speeded-Up Robust Features' (SURF) algorithm (Computer Vision Toolbox, Mathworks). The anchor points were manually corrected through visual inspection, and additional anchor points were added when necessary. These anchor points defined a predicted displacement vector field that would be used to map coordinates from one session to the other. For each coordinate, the predicted vector was defined by the average (weighted inversely by distance) of the vectors for all defined anchor points. This vector field was then applied to every frame of the recording to warp the coordinates of each image to the reference coordinate plane.

To identify responsive neural somata, a pixel-wise activity map was calculated using a modified kurtosis measure. Neuron cell bodies were identified using local adaptive threshold and iterative segmentation, using average activity maps across sessions. Automatically defined regions of interest (ROIs) were then manually checked for proper segmentation in a graphical user interface.

To ensure that the response of individual neurons was not due to local neuropil

contamination of somatic signals, a corrected fluorescence measure was estimated according to:

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

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

$$\Delta F/F = \frac{F_n - F_0}{F_0}$$

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

Analysis of 2-Photon Imaging Data

To minimize potential artifacts introduced by misalignments of the imaging fields across sessions, we manually inspected the average projection and pixel-wise activity maps underlying every defined ROI across all sessions. We assigned each ROI a quality rating based on its appearance and included only ROIs of sufficient quality in our analyses (threshold quality of 3 unless indicated otherwise). Briefly, we defined ROI quality as follows: ROIs rated a quality of 4 or 5 were cells that were clearly present across sessions, and the cell structure could be clearly resolved in both the average projection and activity map. ROIs rated a quality of 3 were also cells unambiguously tracked across sessions but had average maps that were often noisier than cells rated 4 or 5 (for example, they may be identifiable solely by their appearance on the activity map). ROIs rated a quality of 2 were either cells that were not well-tracked or were not unequivocally neuronal somata. ROIs rated a quality of 1 were cells that were not present on the reference session. Each ROI was also marked as either present or not present on each session.

Reliability on a given session was calculated as the Pearson correlation coefficient (CC) between trial-averaged activity taken from two random halves of trials. For MOV, the 30 trials were randomly subsampled to 8 to match PDG, a CC value was found using this subsampled data, and then this was repeated 10 times and averaged for a final reliability value.

A neuron's responsiveness to a stimulus was determined based on a collective measure of the reliability of the neurons in a given field using time-shuffled data. First, a neuron's activity

on each trial was circularly shuffled by a random amount. Next, a reliability value was calculated using this shuffled data. This was repeated 1,000 times to yield a distribution of reliability values, and the 99th percentile of this distribution was stored. This 99th percentile threshold was found for every neuron. If a neuron's actual average reliability across sessions was statistically greater than the average of these 99th percentile values (one-sample t-test), it was classified as responsive to the stimulus. For PDG, a neuron's orientation selectivity index (OSI) on a given session was calculated as:

$$OSI = \frac{R_{pref} - R_{orth}}{R_{pref} + R_{orth}}$$

Where R_{pref} is the neuron's average response to its preferred orientation, and R_{orth} is the average response to the orthogonal orientations. To map preferred orientation for Figure S2D, a neuron's average orientation response vector was wrapped (averaged between opposite orientations), linearly interpolated onto a 180-degree scale, and then fit with a gaussian curve to determine peak response.

The Representational Drift Index (RDI) with respect to a given session and the reference session was calculated as:

$$RDI = \frac{CC_{ws} - CC_{bs}}{CC_{ws} + CC_{bs}}$$

Where CC_{ws} is the Pearson correlation of the trial-averaged activity of two random halves of trials within a session, and CC_{bs} is the Pearson correlation of the trial-averaged activity of two random halves of trials across the compared sessions. For these calculations, negative CC values were rectified to zero. Control RDI values were calculated by treating half of the trials on the first session as the 'test session' and the other half as the 'reference session'.

Visual response events were identified using $\Delta F/F$ data and refined using inferred spike rate data. For each neuron, each frame spanning the length of a stimulus was evaluated for visual response significance by comparing the distribution of its activity values across trials with the distribution of frame-averaged gray-screen period baseline fluorescence values across trials (Wilcoxon sign rank test, right-tailed). Frames were evaluated on a session-by-session basis: if a frame was determined to be significantly responsive in this way on at least two sessions, it was considered to be an event period frame. After each frame was evaluated, the resulting events (periods of contiguous significant frames) were cleaned up by discarding any events consisting of fewer than 5 frames (500 ms) and combining any events 2 or fewer frames (200

ms) apart. Next, a second set of event periods were defined independently using deconvolved spike rate data. A peristimulus time histogram was calculated using spike rate data from all trials across sessions, and frames above a threshold of 2 spikes/sec were treated as event periods. To avoid merging distinct response events due to fluorescence tails, the original event periods
690 determined using $\Delta F/F$ data were further refined using these spike rate event periods such that detected $\Delta F/F$ events lasted no longer than 10 frames (1 s) past the end of the leading spike rate event period.

Event types were determined by z-scoring an event's responses on every trial against baseline activity and comparing the distributions of these z-score values between the first two
695 sessions (60 trials) and the last two sessions (60 trials) (two-sided Wilcoxon rank-sum test). Event instability, or normalized delta z-score, was calculated as:

$$Norm. |\Delta zScore| = \frac{|Z_{late} - Z_{early}|}{Z_{late} + Z_{early}}$$

Where Z_{early} represents the average z-score of an event's trial-averaged activity on the first two sessions, and Z_{late} represents the average z-score of an event's trial-averaged activity on the
700 last two sessions. For boxplot comparisons in event analyses, the first quartile of data was compared against the last quartile (Wilcoxon rank-sum test).

To draw boundaries between cortical layers, layer 4 was identified by finding peak ROI density along a user-defined translaminar axis of the window. A 140 μm window centered at this point was defined as the L2/3 / L4 boundary and the L4 / L5 boundary, and 150 μm deeper than
705 the L4 / L5 boundary was defined as the L5 / L6 boundary [84].

Pairwise signal correlations were calculated as the Pearson correlation coefficient between two neurons' trial-averaged activity on a given session. For pairwise noise correlations, a neuron's trial-average activity was first subtracted from its activity on every trial. All frames were then concatenated into a continuous signal, and the Pearson correlation coefficient was
710 found between these full traces (for MOV, only the first 5760 frames were used to match the length of the PDG recording). Calculation of noise correlations using the average correlation of response vectors across stimulus presentations produced similar results [35].

Statistical Information

To test statistical significance of single groups compared to a group of zero mean, one-
715 sample t-tests (normally distributed data) or Wilcoxon signed-rank (non-parametric) tests were

performed. For comparing experimental groups, two-sample paired t-tests were performed for paired groups and either unpaired t-tests (normally distributed data) or Wilcoxon ranksum (non-parametric) tests were performed for unpaired groups. One-way and two-way ANOVA were performed for comparing more than two groups and testing the effects of multiple factors, respectively. All tests were performed two-tailed unless indicated otherwise.

REFERENCES

1. Chklovskii, D. B.; Mel, B. W.; Svoboda, K. (2007). Cortical rewiring and information storage. *Nanosci. Nanotechnologies Nanophysics* 431, 583–618.
2. Holtmaat, A., and Svoboda, K. (2009). Experience-dependent structural synaptic plasticity in the mammalian brain. *Nat. Rev. Neurosci.* 10, 647–658.
3. Lee, C.R., Najafizadeh, L., and Margolis, D.J. (2020). Investigating learning-related neural circuitry with chronic in vivo optical imaging. *Brain Struct. Funct.* 225, 467–480.
4. 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.
5. Stevenson, I.H., Cherian, A., London, B.M., Sachs, N.A., Lindberg, E., Reimer, J., Slutzky, M.W., Hatsopoulos, N.G., Miller, L.E., and Kording, K.P. (2011). Statistical assessment of the stability of neural movement representations. *J. Neurophysiol.* 106, 764–774.
6. Chestek, C.A., Batista, A.P., Santhanam, G., Yu, B.M., Afshar, A., Cunningham, J.P., Gilja, V., Ryu, S.I., Churchland, M.M., and Shenoy, K. V. (2007). Single-neuron stability during repeated reaching in macaque premotor cortex. *J. Neurosci.* 27, 10742–10750.
7. Carmena, J.M., Lebedev, M.A., Henriquez, C.S., and Nicolelis, M.A.L. (2005). Stable ensemble performance with single-neuron variability during reaching movements in primates. *J. Neurosci.* 25, 10712–10716.
8. Rokni, U., Richardson, A.G., Bizzi, E., and Seung, H.S. (2007). Motor Learning with Unstable Neural Representations. *Neuron* 54, 653–666.
9. 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.
10. Masamizu, Y., Tanaka, Y.R., Tanaka, Y.H., Hira, R., Ohkubo, F., Kitamura, K., Isomura, Y., Okada, T., and Matsuzaki, M. (2014). Two distinct layer-specific dynamics of cortical ensembles during learning of a motor task. *Nat. Neurosci.* 17, 987–994.
11. Lütcke, H., Margolis, D.J., and Helmchen, F. (2013). Steady or changing? Long-term monitoring of neuronal population activity. *Trends Neurosci.* 36, 375–384.
12. Margolis, D.J., Lütcke, H., Schulz, K., Haiss, F., Weber, B., Kügler, S., Hasan, M.T., and Helmchen, F. (2012). Reorganization of cortical population activity imaged throughout long-term sensory deprivation. *Nat. Neurosci.* 15, 1539–1546.

13. Holtmaat, A., Bonhoeffer, T., Chow, D.K., Chuckowree, J., Paola, D., Hofer, S.B., Hübener, M., Keck, T., Knott, G., Lee, a, *et al.* (2009). Through a Chronic Cranial Window. 755
14. Trachtenberg, J.T., Chen, B.E., Knott, G.W., Feng, G., Sanes, J.R., Welker, E., and Svoboda, K. (2002). Long-term in vivo imaging of experience-dependent synaptic plasticity in adult cortex. *Nature* 420, 788–794.
15. Grutzendler, J., Kasthuri, N., and Gan, W.B. (2002). Long-term dendritic spine stability in the adult cortex. *Nature* 420, 812–816. 760
16. 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.
17. 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.
- 765 18. Mank, M., Santos, A.F., Drenth, S., Mrsic-Flogel, T.D., Hofer, S.B., Stein, V., Hendel, T., Reiff, D.F., Levelt, C., Borst, A., *et al.* (2008). A genetically encoded calcium indicator for chronic in vivo two-photon imaging. *Nat. Methods* 5, 805–811.
19. 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.
- 770 20. 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, 1–16.
21. 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.
- 775 22. 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.
23. Tolia, 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. 780
24. Mayrhofer, J.M., Haiss, F., Helmchen, F., and Weber, B. (2015). Sparse, reliable, and long-term stable representation of periodic whisker deflections in the mouse barrel cortex. *Neuroimage* 115, 52–63.
25. 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, 1319–1322. 785
26. Mankin, E.A., Sparks, F.T., Slayyeh, B., Sutherland, R.J., Leutgeb, S., and Leutgeb, J.K. (2012). Neuronal code for extended time in the hippocampus. *Proc. Natl. Acad. Sci. U. S. A.* 109, 19462–19467.
- 790 27. Mankin, E.A., Diehl, G.W., Sparks, F.T., Leutgeb, S., Jill, K., Diego, S., Jolla, L., Diego, S., and Jolla, L. (2016). Hippocampal CA2 activity patterns change over time to a larger extent than between spatial contexts. 85, 190–201.
28. Manns, J.R., Howard, M.W., and Eichenbaum, H. (2008). Gradual changes in

hippocampal activity support remembering the order of events. *Memory* 56, 530–540.

- 795 29. 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.
30. Steveninck, R.R.D.R. Van, Lewen, G.D., Strong, S.P., Koberle, R., and Bialek, W. (1997). Reproducibility and Variability in coded almost exponentially (Fig. 1D), so that Neural Spike Trains. 275.
- 800 31. DeWeese, M.R., Hromádka, T., and Zador, A.M. (2005). Reliability and representational bandwidth in the auditory cortex. *Neuron* 48, 479–488.
32. Mainen, Z.F., and Sejnowski, T.J. (1995). Reliability of spike timing in neocortical neurons. *Science* (80-.). 268, 1503–1506.
- 805 33. 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.
34. Faisal, A.A., Selen, L.P.J., and Wolpert, D.M. (2008). Noise in the nervous system. *Nat. Rev. Neurosci.* 9, 292–303.
- 810 35. 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.
36. Shoham, S., O'Connor, D.H., and Segev, R. (2006). How silent is the brain: Is there a “dark matter” problem in neuroscience? *J. Comp. Physiol. A Neuroethol. Sensory, Neural, Behav. Physiol.* 192, 777–784.
- 815 37. Yassin, L.; Benedetti, B.L.; Jouhanneau, J.-S.; Wen, J.; Poulet, J.F.A.; Barth, A.L. (2008). An embedded subnetwork of highly active neurons in the neocortex. *Bone* 23, 1–7.
38. Ranson, A. (2017). Stability and Plasticity of Contextual Modulation in the Mouse Visual Cortex. *Cell Rep.* 18, 840–848.
- 820 39. Park, J., Papoutsis, A., Ash, R.T., Marin, M.A., Poirazi, P., and Smirnakis, S.M. (2019). Contribution of apical and basal dendrites to orientation encoding in mouse V1 L2/3 pyramidal neurons. *Nat. Commun.* 10, 1–11.
40. Chapman, B., and Stryker, M.P. (1993). Development of orientation selectivity in ferret visual cortex and effects of deprivation. *J. Neurosci.* 13, 5251–5262.
- 825 41. Chapman, B., Stryker, M.P., and Bonhoeffer, T. (1996). Development of orientation preference maps in ferret primary visual cortex. *J. Neurosci.* 16, 6443–6453.
42. Kuhlman, S.J., Tring, E., and Trachtenberg, J.T. (2011). Fast-spiking interneurons have an initial orientation bias that is lost with vision. *Nat. Neurosci.* 14, 1121–1133.
43. White, L.E., Coppola, D.M., and Fitzpatrick, D. (2001). The contribution of sensory experience to the maturation of orientation selectivity in ferret visual cortex. *Nature* 411, 1049–1052.
- 830 44. Ringach, D.L., Mineault, P.J., Tring, E., Olivas, N.D., Garcia-Junco-Clemente, P., and Trachtenberg, J.T. (2016). Spatial clustering of tuning in mouse primary visual cortex. *Nat. Commun.* 7, 1–9.

- 835 45. Ko, H., Mrsic-Flogel, T.D., and Hofer, S.B. (2014). Emergence of Feature-Specific Connectivity in Cortical Microcircuits in the Absence of Visual Experience. *J. Neurosci.* 34, 9812–9816.
46. Ko, H., Cossell, L., Baragli, C., Antolik, J., Clopath, C., Hofer, S.B., and Mrsic-Flogel, T.D. (2013). The emergence of functional microcircuits in visual cortex. *Nature* 496, 96–100.
- 840 47. Deitch, D., Rubin, A., and Ziv, Y. (2020). Representational drift in the mouse visual cortex. *bioRxiv*, 2020.10.05.327049.
48. Chen, T.W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V., *et al.* (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* 499, 295–300.
- 845 49. Harris, J.A., Hirokawa, K.E., Sorensen, S.A., Gu, H., Mills, M., Ng, L.L., Bohn, P., Mortrud, M., Ouellette, B., Kidney, J., *et al.* (2014). Anatomical characterization of Cre driver mice for neural circuit mapping and manipulation. *Front. Neural Circuits* 8.
50. Madisen, L., Garner, A.R., Shimaoka, D., Chuong, A.S., Klapoetke, N.C., Li, L., van der Bourg, A., Niino, Y., Egolf, L., Monetti, C., *et al.* (2015). Transgenic mice for intersectional targeting of neural sensors and effectors with high specificity and performance. *Neuron* 85, 942–958.
- 850 51. Sit, K.K., and Goard, M.J. (2020). Distributed and retinotopically asymmetric processing of coherent motion in mouse visual cortex. *Nat. Commun.* 11, 1–14.
52. Wechselblatt, J.B., Flister, E.D., Piscopo, D.M., and Niell, C.M. (2016). Large-scale imaging of cortical dynamics during sensory perception and behavior. *J. Neurophysiol.* 115, 2852–2866.
- 855 53. Zhuang, J., Ng, L., Williams, D., Valley, M., Li, Y., Garrett, M., and Waters, J. (2017). An extended retinotopic map of mouse cortex. *Elife* 6, 1–29.
54. Goard, M., and Dan, Y. (2009). Basal forebrain activation enhances cortical coding of natural scenes. *Nat. Neurosci.* 12, 1444–1449.
- 860 55. Bacskai, B.J., Kajdasz, S.T., Christie, R.H., Carter, C., Games, D., Seubert, P., Schenk, D., and Hyman, B.T. (2001). Imaging of amyloid- β deposits in brains of living mice permits direct observation of clearance of plaques with immunotherapy. *Nat. Med.* 7, 369–372.
- 865 56. Zuo, Y., Lin, A., Chang, P., and Gan, W.B. (2005). Development of long-term dendritic spine stability in diverse regions of cerebral cortex. *Neuron* 46, 181–189.
57. Tolhurst, D.J., Smyth, D., and Thompson, I.D. (2009). The sparseness of neuronal responses in ferret primary visual cortex. *J. Neurosci.* 29, 2355–2370.
58. Vinje, W.E., and Gallant, J.L. (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* (80-.). 287, 1273–1276.
- 870 59. Weliky, M., Fiser, J.O., Hunt, R.H., and Wagner, D.N. (2003). Coding of natural scenes in primary visual cortex. *Neuron* 37, 703–718.
60. 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.

- 875 61. Douglas, R.J., and Martin, K.A. (1991). A functional microcircuit for cat visual cortex. *J. Physiol.* **440**, 735–769.
62. Johnson, R.R., and Burkhalter, A. (1997). A polysynaptic feedback circuit in rat visual cortex. *J. Neurosci.* **17**, 7129–7140.
- 880 63. Andermann, M.L., Gilfoy, N.B., Goldey, G.J., Sachdev, R.N.S., Wölfel, M., McCormick, D.A., Reid, R.C., and Levene, M.J. (2013). Chronic Cellular Imaging of Entire Cortical Columns in Awake Mice Using Microprisms. *Neuron* **80**, 900–913.
64. Kerlin, A.M., Andermann, M.L., Berezovskii, V.K., and Reid, R.C. (2010). Broadly Tuned Response Properties of Diverse Inhibitory Neuron Subtypes in Mouse Visual Cortex. *Neuron* **67**, 858–871.
- 885 65. Peters, A.J., Chen, S.X., and Komiyama, T. (2014). Emergence of reproducible spatiotemporal activity during motor learning. *Nature* **510**, 263–267.
66. Liberti, W.A., Markowitz, J.E., Perkins, L.N., Liberti, D.C., Leman, D.P., Guitchounts, G., Velho, T., Kotton, D.N., Lois, C., and Gardner, T.J. (2016). Unstable neurons underlie a stable learned behavior. *Nat. Neurosci.* **19**, 1665–1671.
- 890 67. Katlowitz, K.A., Picardo, M.A., and Long, M.A. (2018). Stable Sequential Activity Underlying the Maintenance of a Precisely Executed Skilled Behavior. *Neuron* **98**, 1133–1140.e3.
68. Stojanoski, B., and Cusack, R. (2014). Time to wave good-bye to phase scrambling: Creating controlled scrambled images using diffeomorphic transformations. *J. Vis.* **14**, 6–6.
- 895 69. Cohen, M.R., and Kohn, A. (2011). Measuring and interpreting neuronal correlations. *Nat. Neurosci.* **14**, 811–819.
70. Seriès, P., Latham, P.E., and Pouget, A. (2004). Tuning curve sharpening for orientation selectivity: Coding efficiency and the impact of correlations. *Nat. Neurosci.* **7**, 1129–1135.
- 900 71. Hofer, S.B., Ko, H., Pichler, B., Vogelstein, J., Ros, H., Zeng, H., Lein, E., Lesica, N.A., and Mrsic-Flogel, T.D. (2011). Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. *Nat. Neurosci.* **14**, 1045–1052.
72. Cossell, L., Iacaruso, M.F., Muir, D.R., Houlton, R., Sader, E.N., Ko, H., Hofer, S.B., and Mrsic-Flogel, T.D. (2015). Functional organization of excitatory synaptic strength in primary visual cortex. *Nature* **518**, 399–403.
- 905 73. Ecker, A.S., Berens, P., Keliris, G.A., Bethge, M., Logothetis, N.K., and Tolias, A.S. (2010). Decorrelated neuronal firing in cortical microcircuits. *Science* (80-.). **327**, 584–587.
74. Okun, M., Steinmetz, N.A., Cossell, L., Iacaruso, M.F., Ko, H., Barthó, P., Moore, T., Hofer, S.B., Mrsic-Flogel, T.D., Carandini, M., *et al.* (2015). Diverse coupling of neurons to populations in sensory cortex. *Nature* **521**, 511–515.
- 910 75. Sweeney, Y., and Clopath, C. (2020). Population coupling predicts the plasticity of stimulus responses in cortical circuits. *Elife* **9**.
- 915 76. Schröder, S., Steinmetz, N.A., Krumin, M., Pachitariu, M., Rizzi, M., Lagnado, L., Harris, K.D., and Carandini, M. (2020). Arousal Modulates Retinal Output. *Neuron* **107**, 487–

495.e9.

77. Liang, L., Fratzl, A., Reggiani, J.D.S., El Mansour, O., Chen, C., and Andermann, M.L. (2020). Retinal Inputs to the Thalamus Are Selectively Gated by Arousal. *Curr. Biol.* 30, 3923-3934.e9.

920 78. Kentros, C.G., Agnihotri, N.T., Streater, S., Hawkins, R.D., and Kandel, E.R. (2004). Increased attention to spatial context increases both place field stability and spatial memory. *Neuron* 42, 283–295.

925 79. 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.

80. Hainmueller, T., and Bartos, M. (2018). Parallel emergence of stable and dynamic memory engrams in the hippocampus. *Nature* 558, 292–296.

930 81. Drew, P.J., Shih, A.Y., Driscoll, J.D., Knutsen, P.M., Blinder, P., Davalos, D., Akassoglou, K., Tsai, P.S., and Kleinfeld, D. (2010). Chronic optical access through a polished and reinforced thinned skull. *Nat. Methods* 7, 981–984.

82. Brainard, D.H. (1997). The Psychophysics Toolbox. *Spat. Vis.* 10, 433–436.

83. Pho, G.N., Goard, M.J., Woodson, J., Crawford, B., and Sur, M. (2018). Task-dependent representations of stimulus and choice in mouse parietal cortex. *Nat. Commun.* 9, 1–16.

935 84. Ji, W., Gămănuț, R., Bista, P., D’Souza, R.D., Wang, Q., and Burkhalter, A. (2015). Modularity in the Organization of Mouse Primary Visual Cortex. *Neuron* 87, 632–643.

ACKNOWLEDGEMENTS

We would like to thank Ji Xia and Ralf Wessel for comments on manuscript. This work was supported by grants to M.J.G. from the Whitehall Foundation, Larry L. Hillblom Foundation, NSF (1707287), and NIH (R00 MH104259).

940

AUTHOR CONTRIBUTIONS

T.D.M. and M.J.G. designed the experiments; T.D.M. conducted the experiments and analyzed the data, with guidance from M.J.G.; T.D.M. and M.J.G. wrote the manuscript.

945 COMPETING INTERESTS

The authors declare no competing interests.

DATA AVAILABILITY

Most of the hardware designs can be found on our institutional lab website

950 (<https://goard.mcdb.ucsb.edu/resources>). Raw data from Figures 1-6 are available on FigShare (doi: TBA, final version will be uploaded upon manuscript acceptance). All other raw data are available upon request to corresponding author.

CODE AVAILABILITY

955 All of the code used for analysis is available on the following GitHub page: (TBA, final version to be uploaded upon manuscript acceptance)

FIGURES

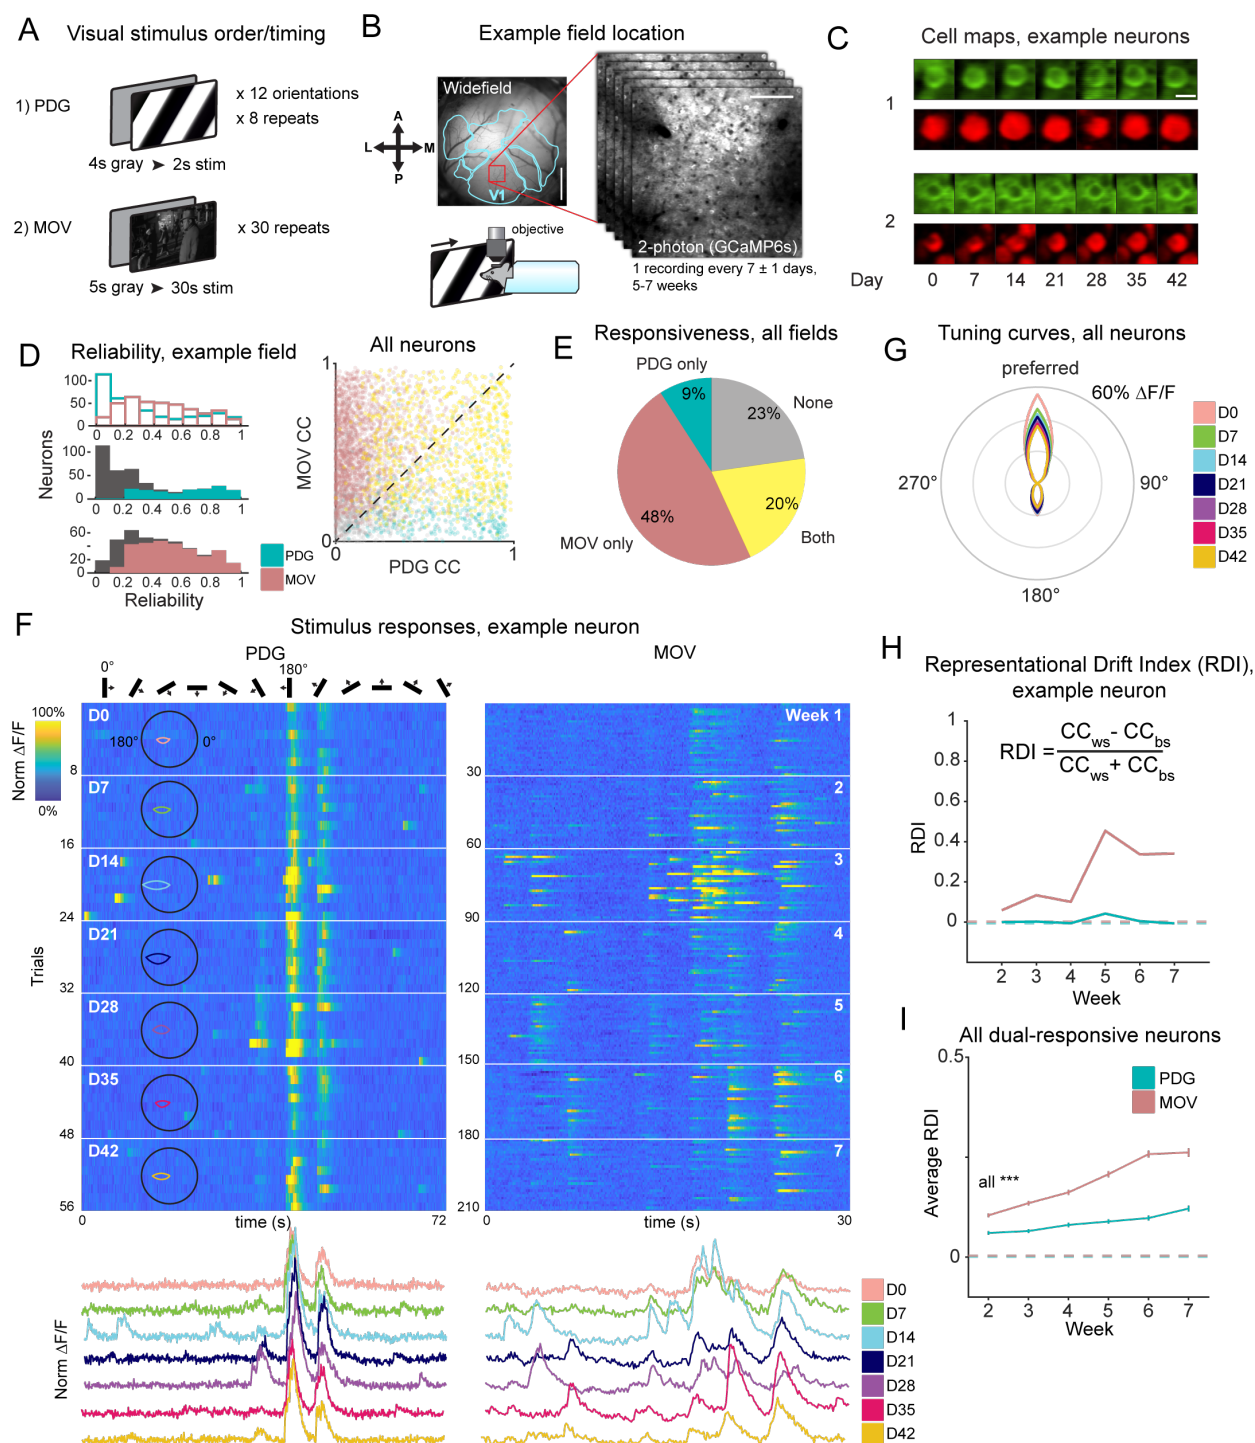


Figure 1: Chronic 2-photon imaging reveals differential stability of visual responses in single cells.

(A) Visual stimuli: top screen depicts a drifting grating of one orientation as presented in the passive drifting grating (PDG) stimulus, bottom screen depicts a single frame from the natural movie (MOV) stimulus. PDG is presented first as 8 repeats of a 12 orientation sequence, followed by 30 repeats of MOV.

(B) Location of one example recording field in primary visual cortex (V1). Left: widefield fluorescence image of a 4 mm cortical window for one example Emx1-cre x ROSA-tTA x TITL-GCaMP6s mouse, with overlay (light blue) of visual area boundaries determined by retinotopic mapping (see methods); red box indicates the approximate location of 2-photon recordings in V1; scale bar = 1mm. Right: example average projection of a 2-photon imaging field; scale bar = 100 μ m. Bottom: schematic of head-fixed mouse watching the PDG stimulus.

(C) Example images of registered cells (see methods) from imaging field in (B) on all recording days. Top row, green: average projection of GCaMP fluorescence channel. Bottom row, red: pixel-wise activity map (see Methods). Scale bar = 15 μ m.

(D) Top left: single cell reliability distributions for PDG and MOV stimuli on the first recording session for one example mouse. Reliability is defined as the Pearson correlation coefficient (CC) of trial-averaged activity from two halves of the trials. Middle left: PDG reliability distribution; subset of PDG responsive neurons is colored. Bottom left: MOV reliability distribution; subset of MOV responsive neurons is colored. Right: each neuron's between-trial CC for PDG vs. MOV, for neurons present on the reference session across all mice (4,142 neurons). Dots are colored by significant responsiveness to stimuli, as in (E).

(E) Average percentage of neurons significantly responsive to each stimulus (MOV only: $47.7 \pm 2.1\%$ s.e.m., PDG only: $9.1 \pm 0.9\%$, both: $20.4 \pm 2.3\%$, none: $22.7 \pm 2.2\%$).

(F) Fluorescence traces ($\Delta F/F$) for one example neuron. Trials are concatenated across sessions. Left: responses to the PDG stimulus; overlay: orientation tuning curves for each recording day. Right: responses to the MOV stimulus. White horizontal lines in each heatmap indicate divisions between recording sessions (8 trials per session for PDG, 30 trials per session for MOV). Heatmaps for each stimulus are co-normalized. Below each heatmap are trial-averaged responses colored by session.

(G) Orientation tuning curves colored by session, averaged across all orientation-tuned neurons in all imaging fields and aligned to 0 degrees based on preferred orientation. Neurons are only included if they are present on a given session and orientation tuned (740, 710, 694, 670, 701, 659, 596 neurons per session).

(H) Top: Representational Drift Index (RDI) curves for each stimulus for example neuron shown in (F); values closer to 0 indicates a more stable response (similar to the first recording session), and closer to 1 indicates greater response drift (see methods); inset: RDI formula: CC_{WS} = within-session correlation coefficient, CC_{BS} = between-session correlation coefficient; dotted line indicates control RDI for this cell, determined using half the trials of the session 1 as the reference and the other half as a test data set (see methods).

(I) Average RDI curves from all neurons across all imaging fields. Values are calculated using neurons that are present on any given session and visually responsive to both stimuli ('dual-responsive'; $n = 824, 808, 793, 830, 761, 698$ neurons for sessions 2-7, respectively); dotted line indicates control RDI, as in (H); error bars are \pm s.e.m. Significance marker indicates the comparison of average RDI distributions within each session ($***p < 0.001$, paired-sample t-test).

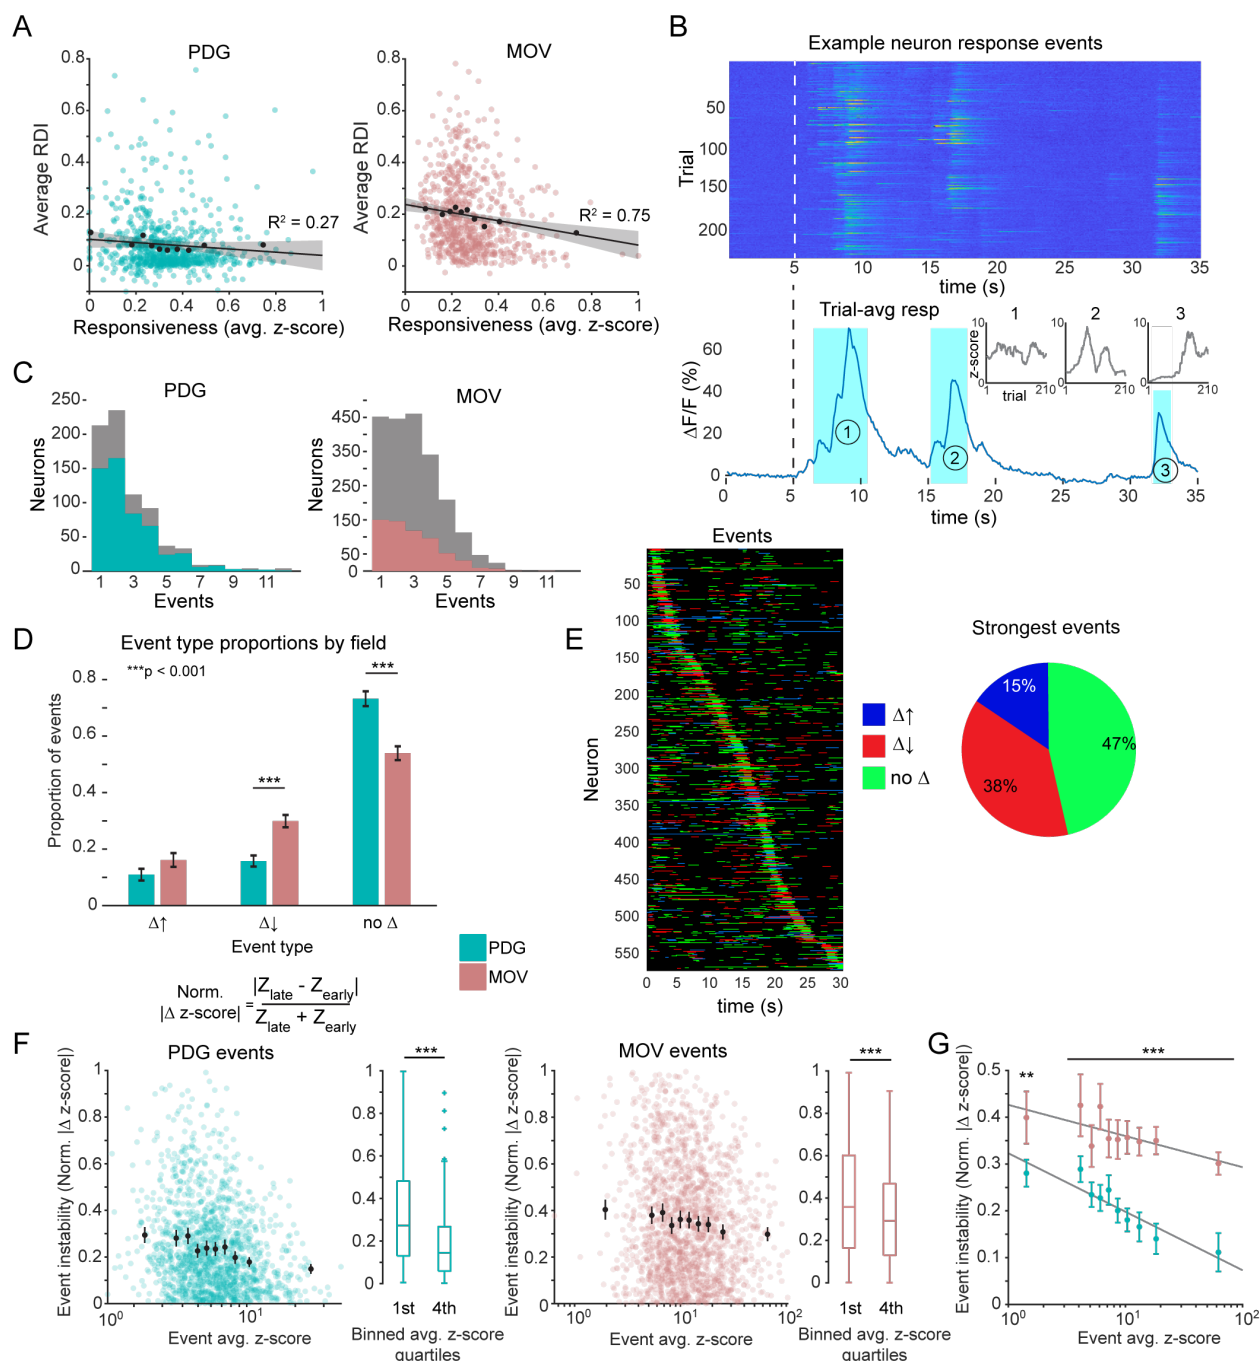


Figure 2: Characterization of dynamic response events underlying single cell responses.

(A) Single neuron RDI as a function of responsiveness (session-average z-score of $\Delta F/F$ activity). Each colored dot is one neuron; black dots are 10th percentile binned means; black line is a linear fit of the binned data; shaded area indicates 95th percent confidence interval. Neurons are z-scored using the entire recording on a given session. Data shown for all neurons responsive to both PDG and MOV ($n = 736$).

(B) Response events from one example neuron. Top: $\Delta F/F$ responses, all trials across all recording sessions. Bottom: Trial-averaged response. Shaded areas indicate identified events. Insets: z-score trajectories (smoothed using 30-point moving average) across all trials for the three events in the example neuron.

(C) Number of response events per neuron. Left: all PDG responsive neurons shown in grey, dual-responsive neurons shown in color. Right: all MOV responsive neurons shown in grey, dual-responsive neurons shown in color.

(D) Proportions of event types (growing, decaying, and static) in responses to both stimuli. Event type is determined by z-scoring an event waveform's single trial responses and comparing the distributions of these values between the first two sessions (60 trials) and the last two sessions (60 trials; Wilcoxon rank-sum test). Data shown are mean proportions across all imaging fields \pm s.e.m. ($n = 13$ imaging fields in 12 mice). Proportion comparisons made with independent-samples t-test (growing events: $t = 1.6$, $p = 0.12$; decaying events: $t = 4.8$, $p = 6.88 \times 10^{-5}$; static events: $t = 5.4$, $p = 1.61 \times 10^{-5}$; $n = 13$ fields).

(E) Visualization of MOV response event magnitude changes. Left: bands indicate event periods for each neuron, colored by event type. Neurons are ordered by time of maximum trial-averaged response. Right: proportions of event types for neurons' peak responses (diagonal of left plot). Data shown for all dual-responsive neurons.

(F) Event instability (normalized delta z-score) as a function of event magnitude (session-average event z-score). Each colored dot is one event; black dots are 10th percentile binned means \pm 95th percent confidence interval. Box plots are first quartile of data tested against 4th quartile ($***p < 0.001$, Wilcoxon ranksum test, $n = 282$ events per quartile for PDG, 402 events per quartile for MOV). Data shown for all dual-responsive neurons.

(G) MOV events are less stable than PDG events independent of event magnitude. Binned means (10th percentiles) using data from (F) for both stimuli shown together. Data from all events across both stimuli were pooled to determine bin edges, events from each stimulus were then binned separately. Error bars represent 95th percent confidence interval, grey lines are linear fits of the data. Significance markers indicate comparison of PDG and MOV values in each bin ($**p < 0.01$ for first bin, $***p < 0.001$ for all other bins; Wilcoxon ranksum test).

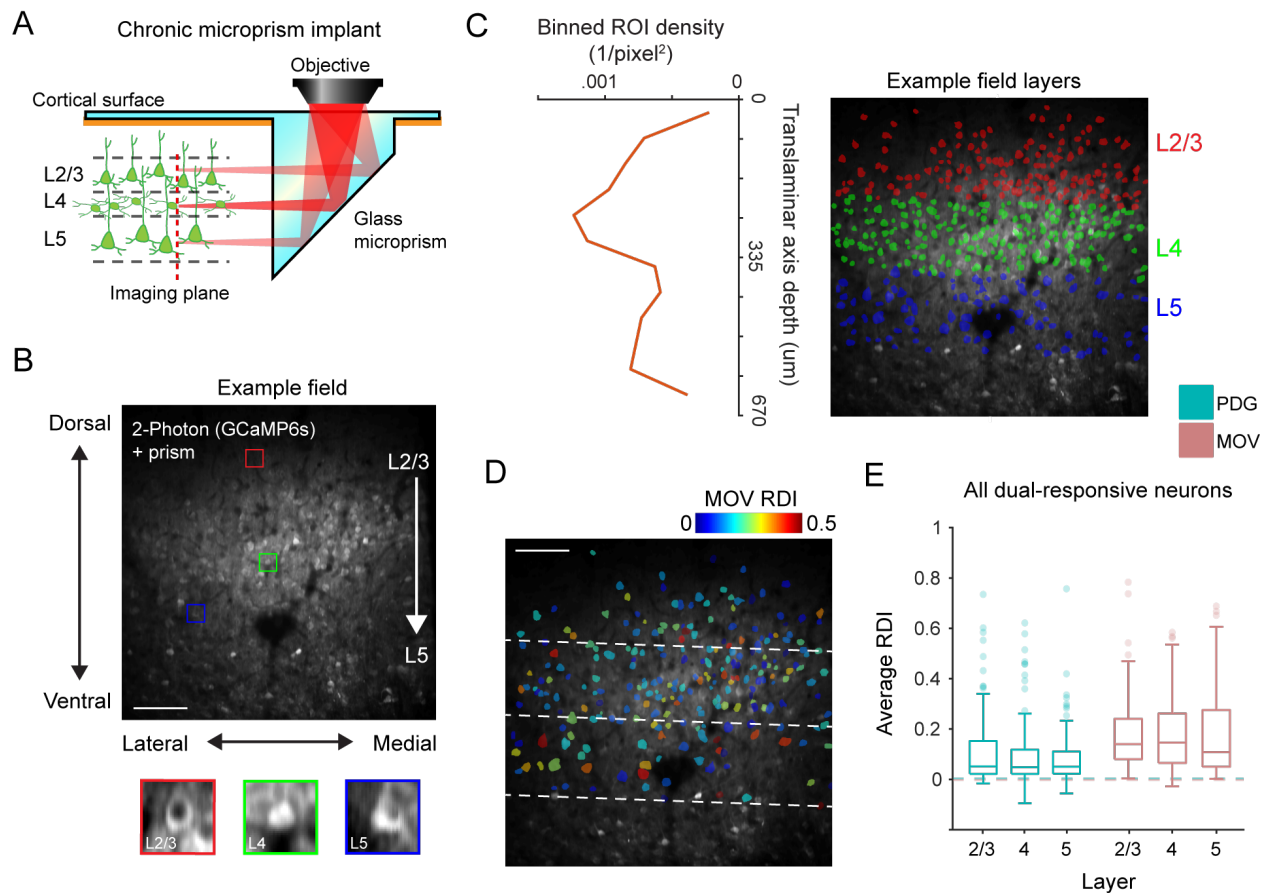


Figure 3: Translaminar imaging shows equal RDI distributions across layers.

(A) Schematic of glass microprism placement in V1. Red dotted line depicts the translaminar imaging plane, which is rotated 90 degrees from the original horizontal plane and spans ~700 µm of the cortical column, capturing neurons in L2-5 for a typical recording.

(B) Example average fluorescence image from a chronic imaging session of a prism field. Colored boxes are zoomed-in example cells (red: L2/3 pyramidal neuron, green: L4 stellate neuron, blue: L5 pyramidal neuron). Scale bar = 100µm.

(C) Delineation of cortical layers. Left: ROI density of binned pixel windows perpendicular to the cortical column axis (corresponds to right subfigure); layers are determined by finding peak density and assigning a 140 µm window around it as L4, and then a further 150 µm from the L4 deep boundary as L5. Right: example field shown in (B) with overlay of all ROIs colored by layer; dotted line is the translaminar axis.

(D) Example field shown in (B, C) with overlay of ROIs of all well-tracked neurons responsive to MOV, colored according to MOV RDI. Dotted lines are layer boundaries.

(E) Session-averaged RDI distributions by stimulus and layer, using all dual-responsive neurons recorded in prism fields (n = 4 mice, 64 L2/3 neurons, 111 L4 neurons, 121 L5 neurons). No significant difference found between layers, significant difference found between stimuli (layer $F = 2.06$, $p = 0.12$; stimulus $F = 45.7$, $p = 3.32 \times 10^{-11}$; two-way ANOVA). Dotted lines indicate control RDI (as in Figure 1H, I) using all dual-responsive neurons.

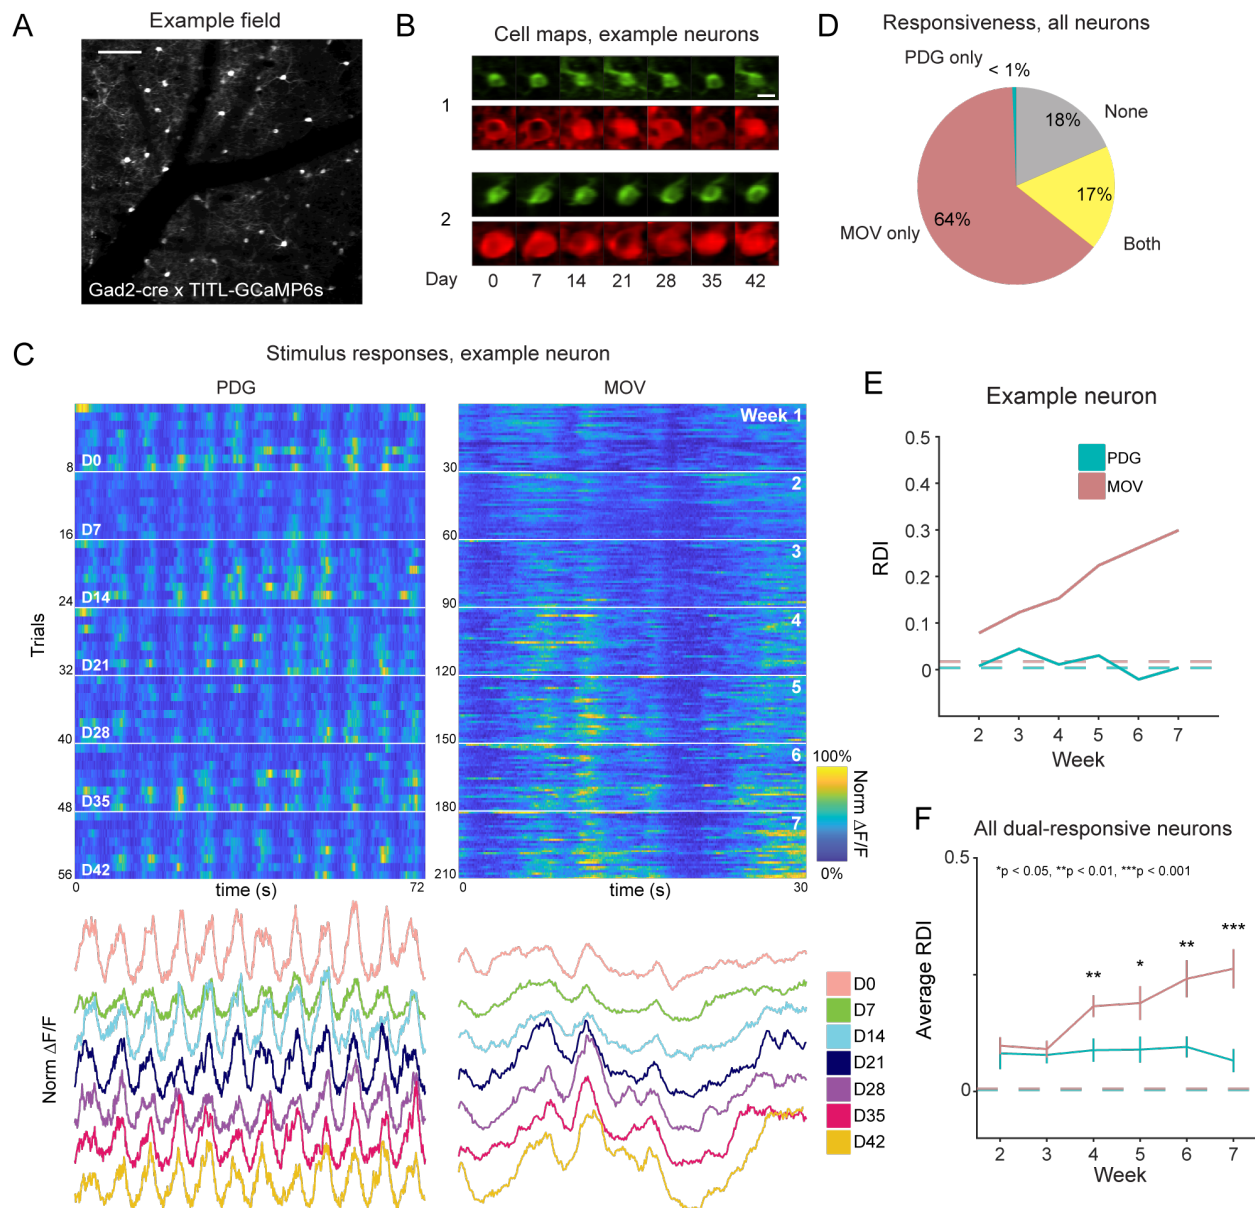


Figure 4: Inhibitory neuron populations also exhibit representational drift to MOV stimuli.

(A) Average fluorescence image from an example field of inhibitory neurons in L2/3 of V1 of GAD2-Cre x TITL2-G6s mice. Scale bar = 100 μ m.

(B) Example images of well-tracked neurons (see Methods) from the field in (A) on all recording days, green and red colors are the same as in Figure 1C. Scale bar = 15 μ m.

(C) Fluorescence traces ($\Delta F/F$) for one example neuron. Trials are concatenated across sessions. Left: responses to PDG. Right: responses to MOV. White horizontal lines indicate divisions between recording sessions. Heatmaps for each stimulus are co-normalized. Below each heatmap are trial-averaged responses colored by session.

(D) Average percentage of neurons responsive to each stimulus (MOV only: $59.4 \pm 1.9\%$ s.e.m., PDG only: $0.6 \pm 0.6\%$, both: $14.4 \pm 5.7\%$, none: $25.6 \pm 6.3\%$).

(E) RDI curves for the example neuron shown in (C); dotted line indicates control RDI for this cell (see methods).

(F) Average RDI curves from neurons across all imaging fields; error bars are \pm s.e.m.; data shown for dual-responsive neurons present on any given session (neurons per session: 33, 34, 34, 33, 33, 31); significance markers indicate the comparison of each session's PDG RDI values and MOV RDI values (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, paired-sample t-test); dotted lines indicate control RDI as in previous figures.

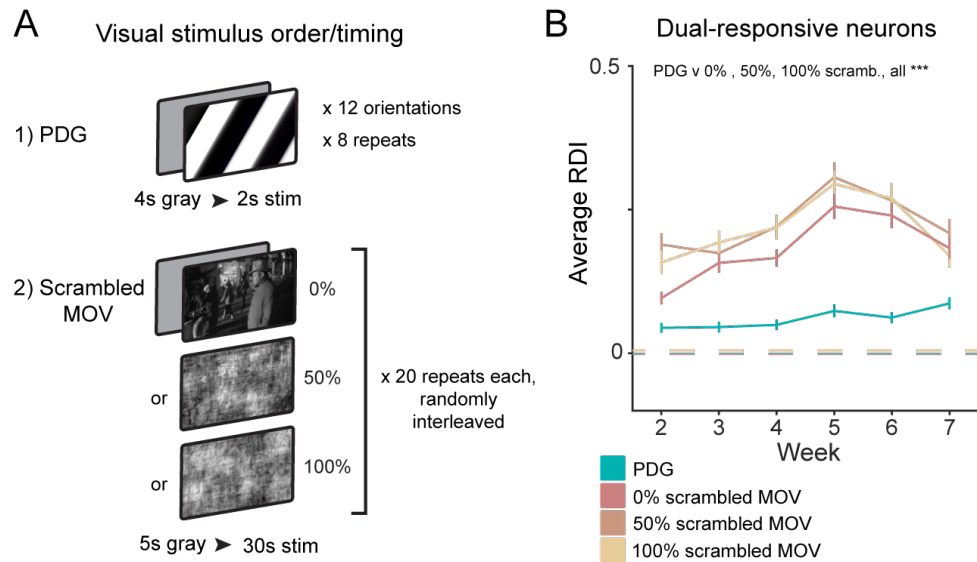


Figure 5: Stability is not dependent on the higher order statistics of the visual stimulus.

(A) Visual stimuli, where the original 30 repeats of MOV is replaced with 20 repeats each of 0%, 50% and 100% phase-scrambled versions of the original movie, randomly interleaved. Bottom screens depict the same freeze-frame from each of the movie versions.

(B) Average RDI curves from all neurons; error bars are \pm s.e.m.; values are calculated using only neurons that are present on any given session (120, 118, 118, 119, 113, 111 neurons per session after first session). Average PDG RDI is significantly different from average MOV RDI for all three movie versions ($***p < 0.001$, paired-sample t-tests); dotted lines indicate control RDI, as in previous figures.

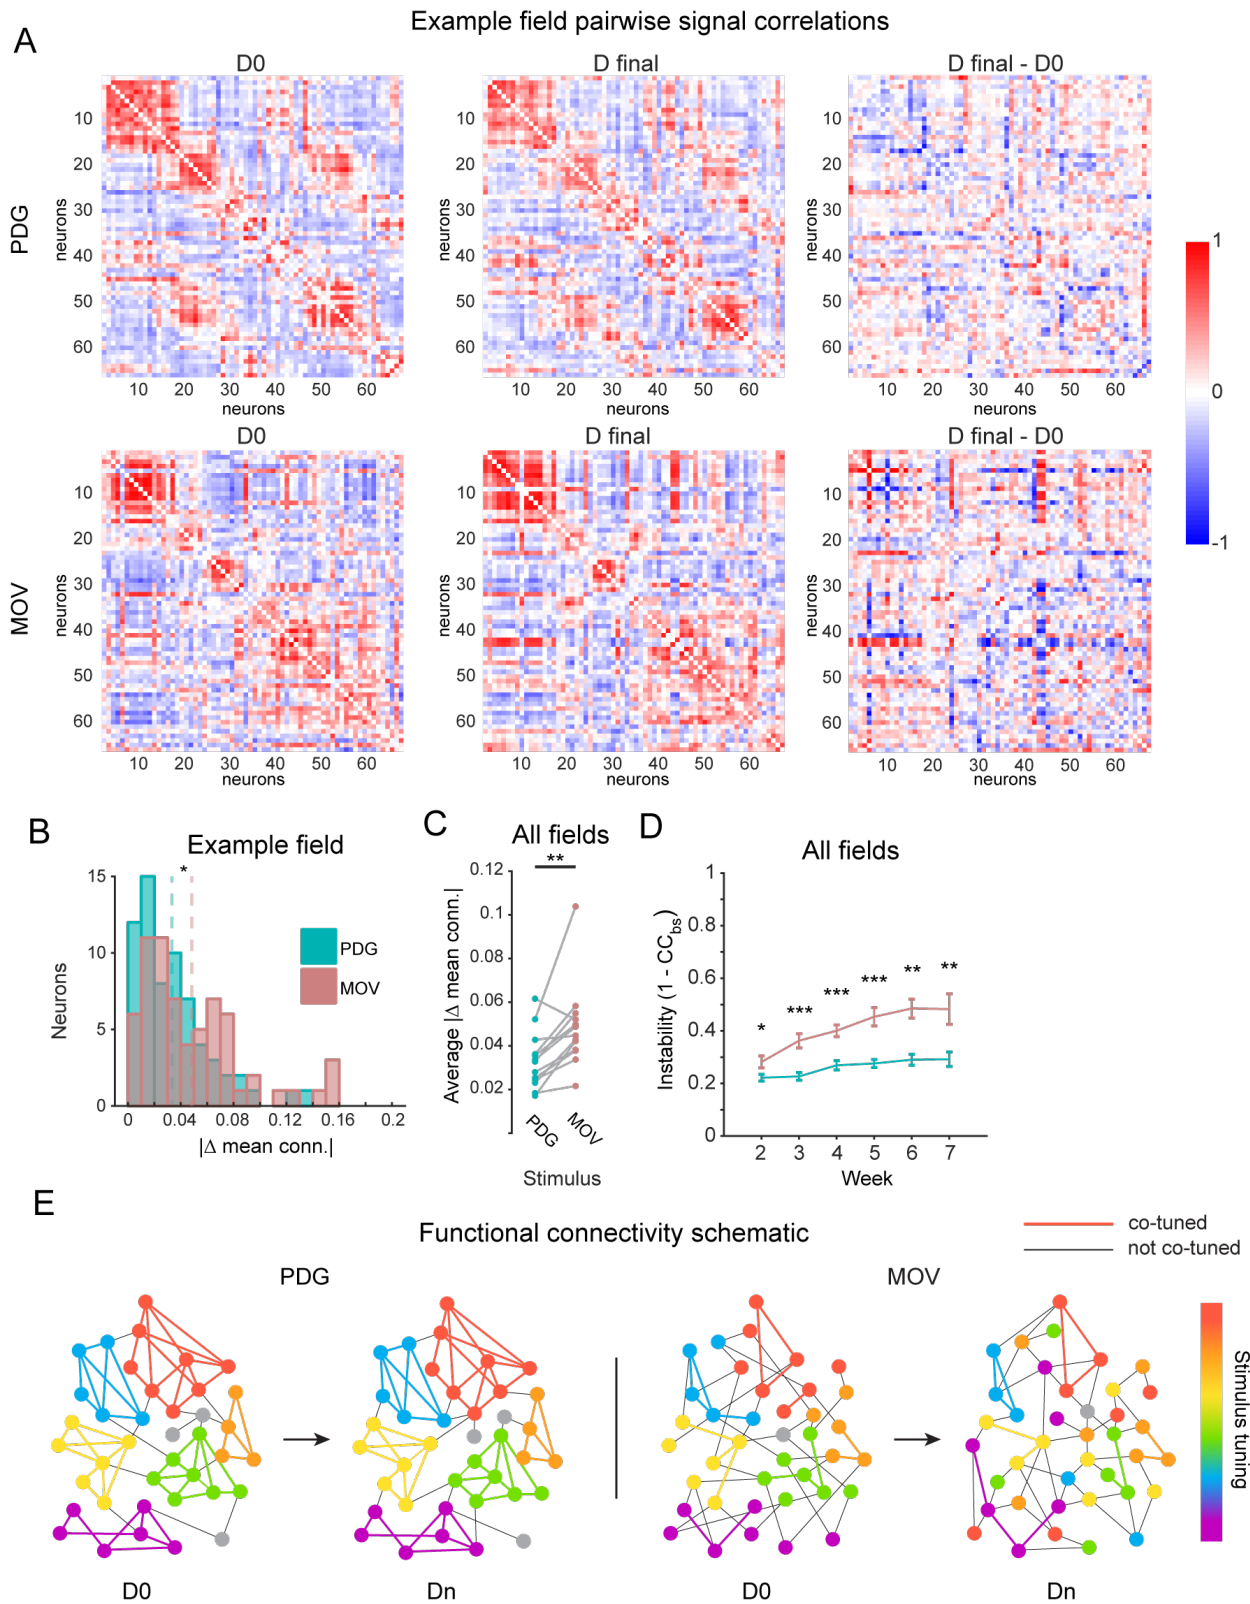


Figure 6: Between-neuron signal correlation stability is stimulus-dependent

(A) Pairwise signal correlations on session 1 (left), final session (middle), and their difference (right) for one example field. Neurons are sorted by time of peak response on D0 for each stimulus. Data shown for all neurons responsive to both stimuli.

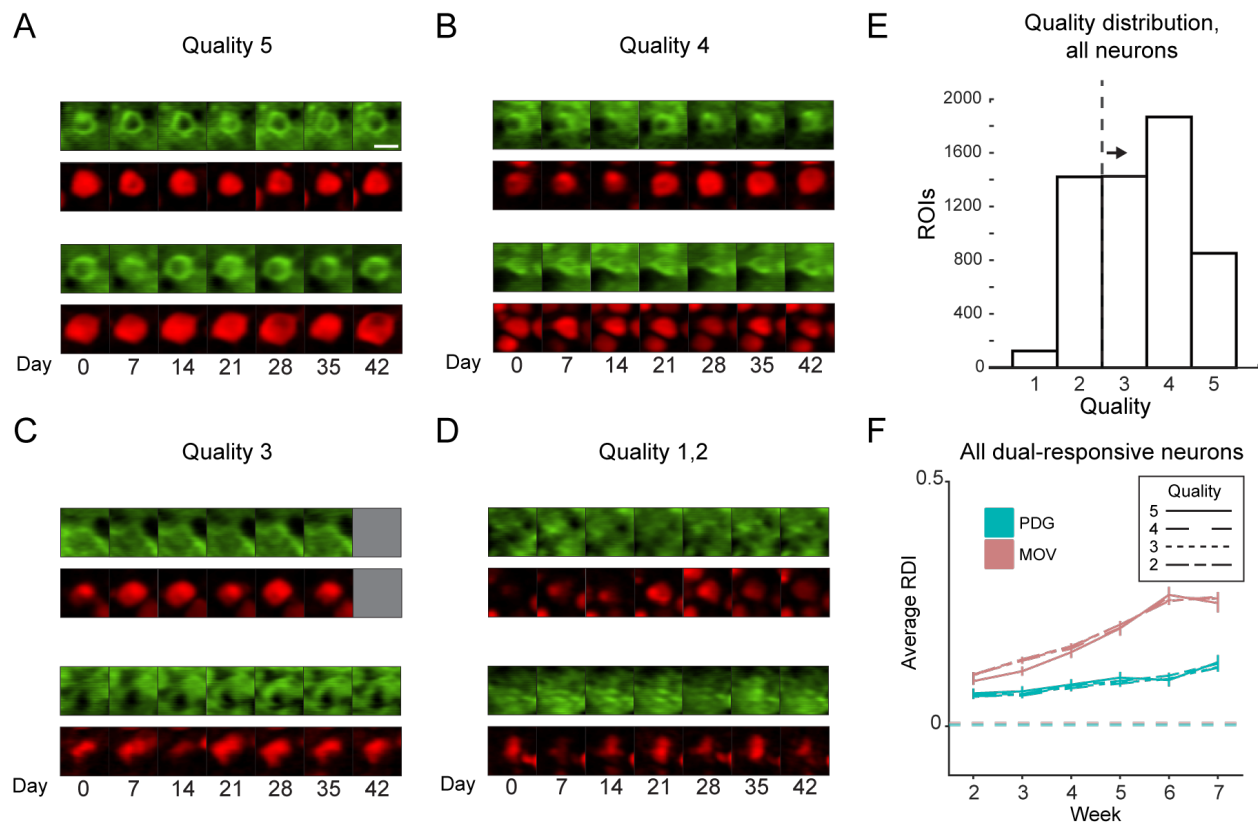
(B) Distributions of single-neuron average change in signal correlations between first and final sessions for example field in A. Dotted lines are means for each stimulus ($*p < 0.05$, Wilcoxon rank-sum test).

(C) Field-average changes in signal correlation between first and final sessions. Data shown for all fields. ($**p < 0.01$, paired-sample t-test, $n = 13$ fields).

(D) Average instability of signal correlation matrices with respect to the first session over time ($1 - CC_{bs}$, where CC_{bs} is the 2D cross correlation between signal correlation matrices). Data shown for all fields. Error bars are \pm s.e.m.; significance markers indicate comparison of PDG and MOV values on given session ($*p < 0.05$, $**p < 0.01$, $***p < 0.001$, paired-sample t-test, $n = 13$ fields).

(E) Schematic depicting the relationship between stimulus tuning stability and shifts in functional connectivity over time.

SUPPLEMENTAL INFORMATION

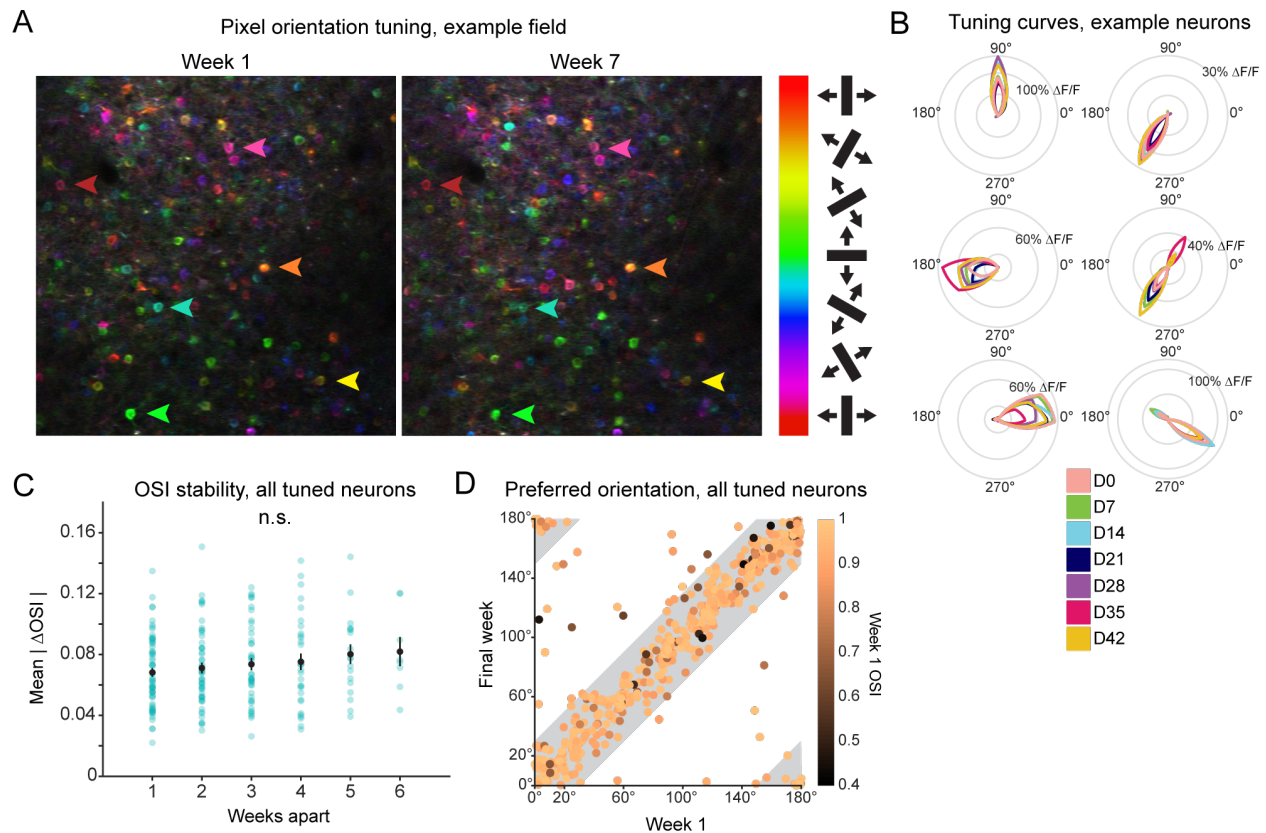


Supplementary Figure 1: Example neuron images and ROI quality distribution.

(A) - (D) Example ROIs from each quality rating. See methods for descriptions of each rating. Top row, green: average projection of GCaMP6s fluorescence channel. Bottom row, red: activity map of active pixels (see Methods). Scale bar = 15 μ m.

(E) Distribution of quality ratings. Data shown for all ROIs across all mice. For all analyses, ROIs of quality 2 or below are discarded.

(F) Average RDI curves for both stimuli as shown in Figure 1I, for different ROI quality rating inclusion thresholds.



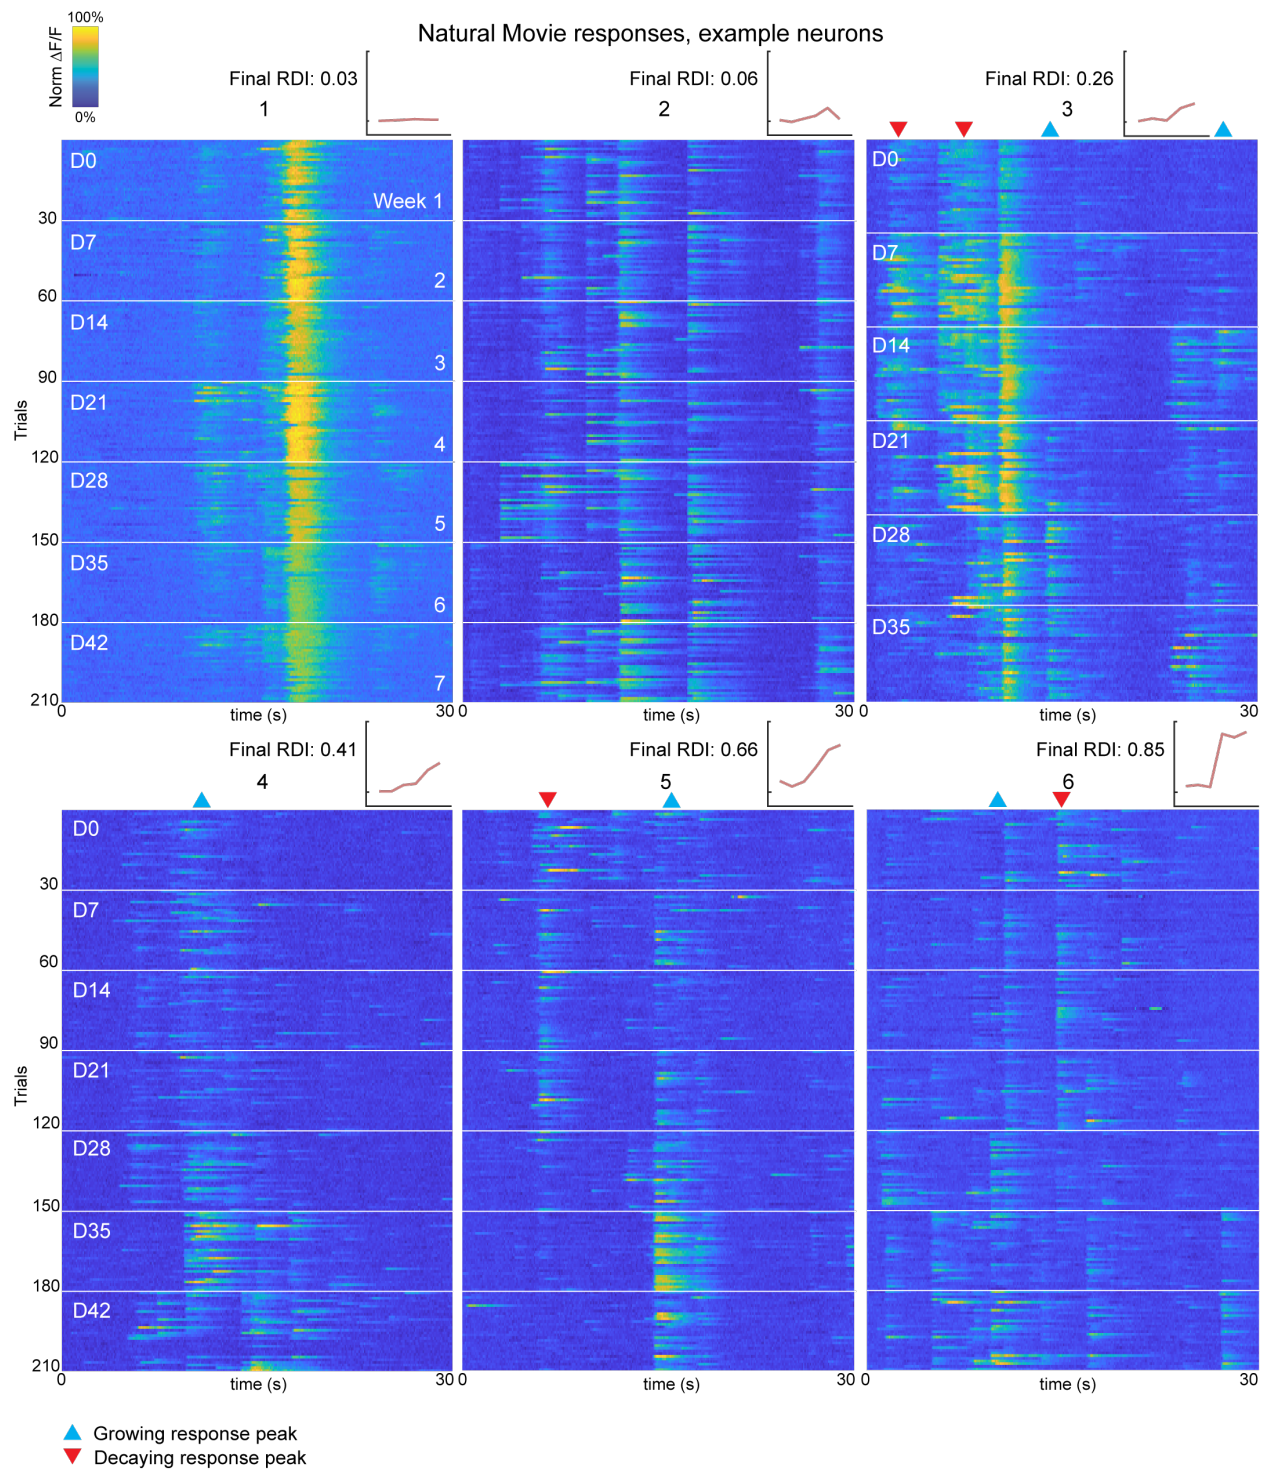
Supplementary Figure 2: Orientation tuning and selectivity are highly stable.

(A) Example PDG recording pseudocolored by pixel-wise preferred orientation on the first recording week and the seventh recording week. Arrowheads indicate a subset of example neurons that were well tracked across recordings.

(B) Tuning curves of example orientation-tuned neurons from one mouse across all recording sessions.

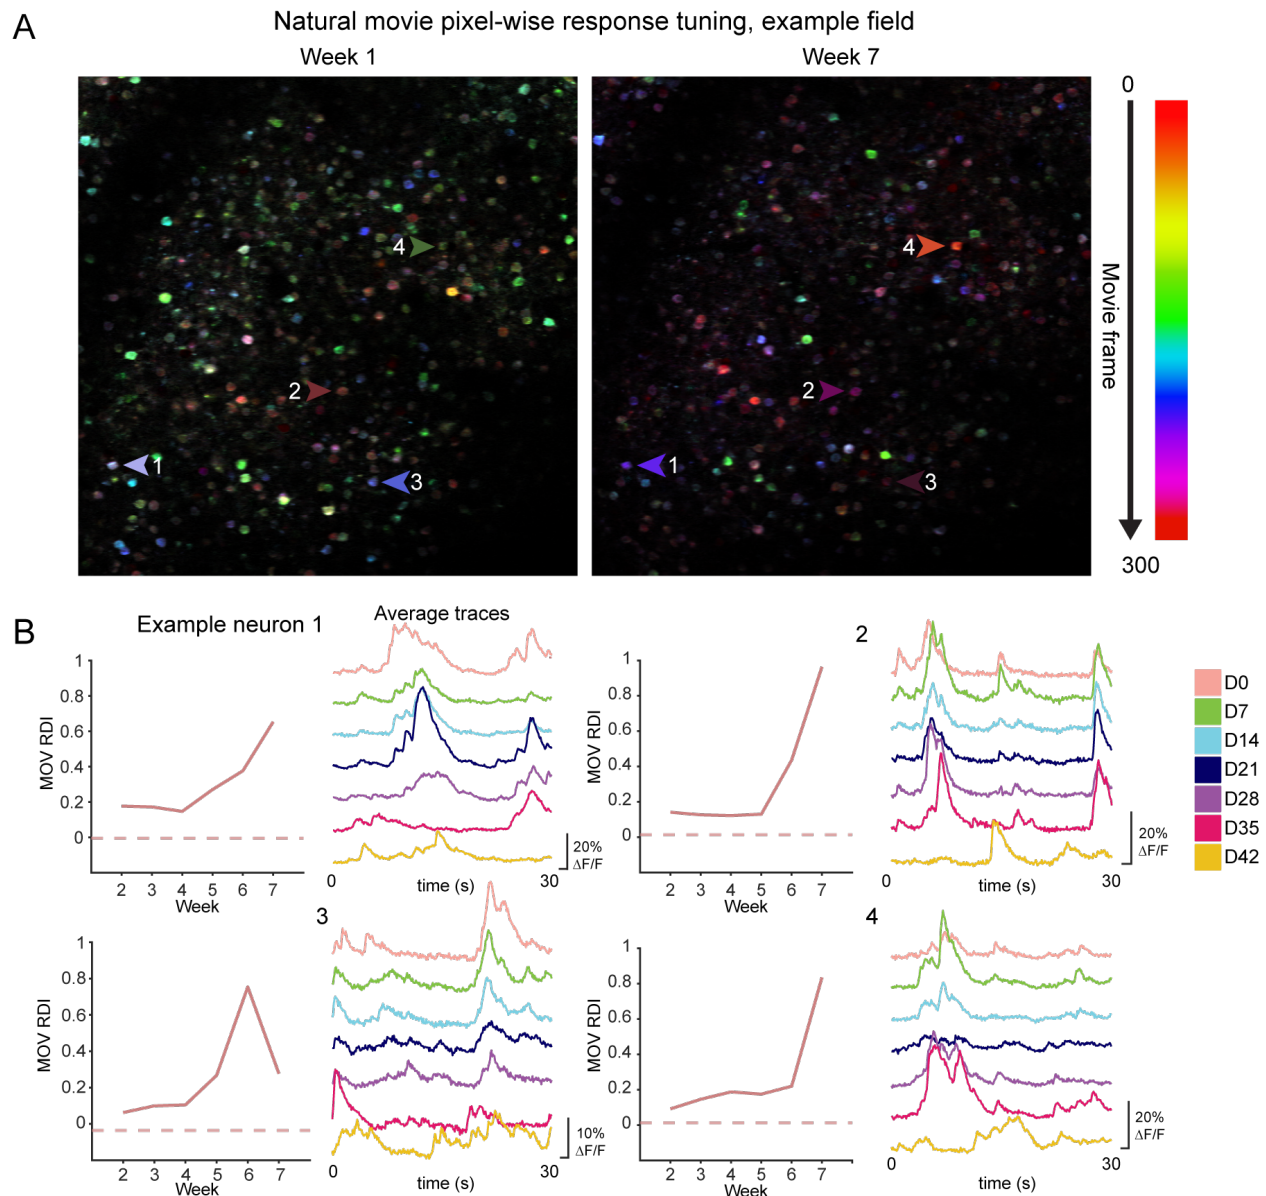
(C) Average change in orientation selectivity index (OSI) as a function of time between recordings. Data shown for all mice, neurons must be present on both sessions and orientation tuned on the first recording session to be included. Blue dots are individual pairwise comparisons between recordings, black dots indicate mean \pm s.e.m. of each distribution. No weeks are significantly different from average mean $|\Delta\text{OSI}|$ ($p > 0.05$, one-sample t-test, $n = 66, 53, 43, 31, 18, 8$).

(D) Preferred orientation on the first recording session vs. final recording session (see Methods). Data shown for all mice using neurons present on all sessions and orientation tuned on the first session ($n = 603$ neurons). Each dot is one neuron, colored by its OSI on week 1. Shaded areas indicate $\pm 30^\circ$ change between recordings (± 1 orientation step, 96% of all neurons tuned on the first session).



Supplementary Figure 3: Example natural movie responses.

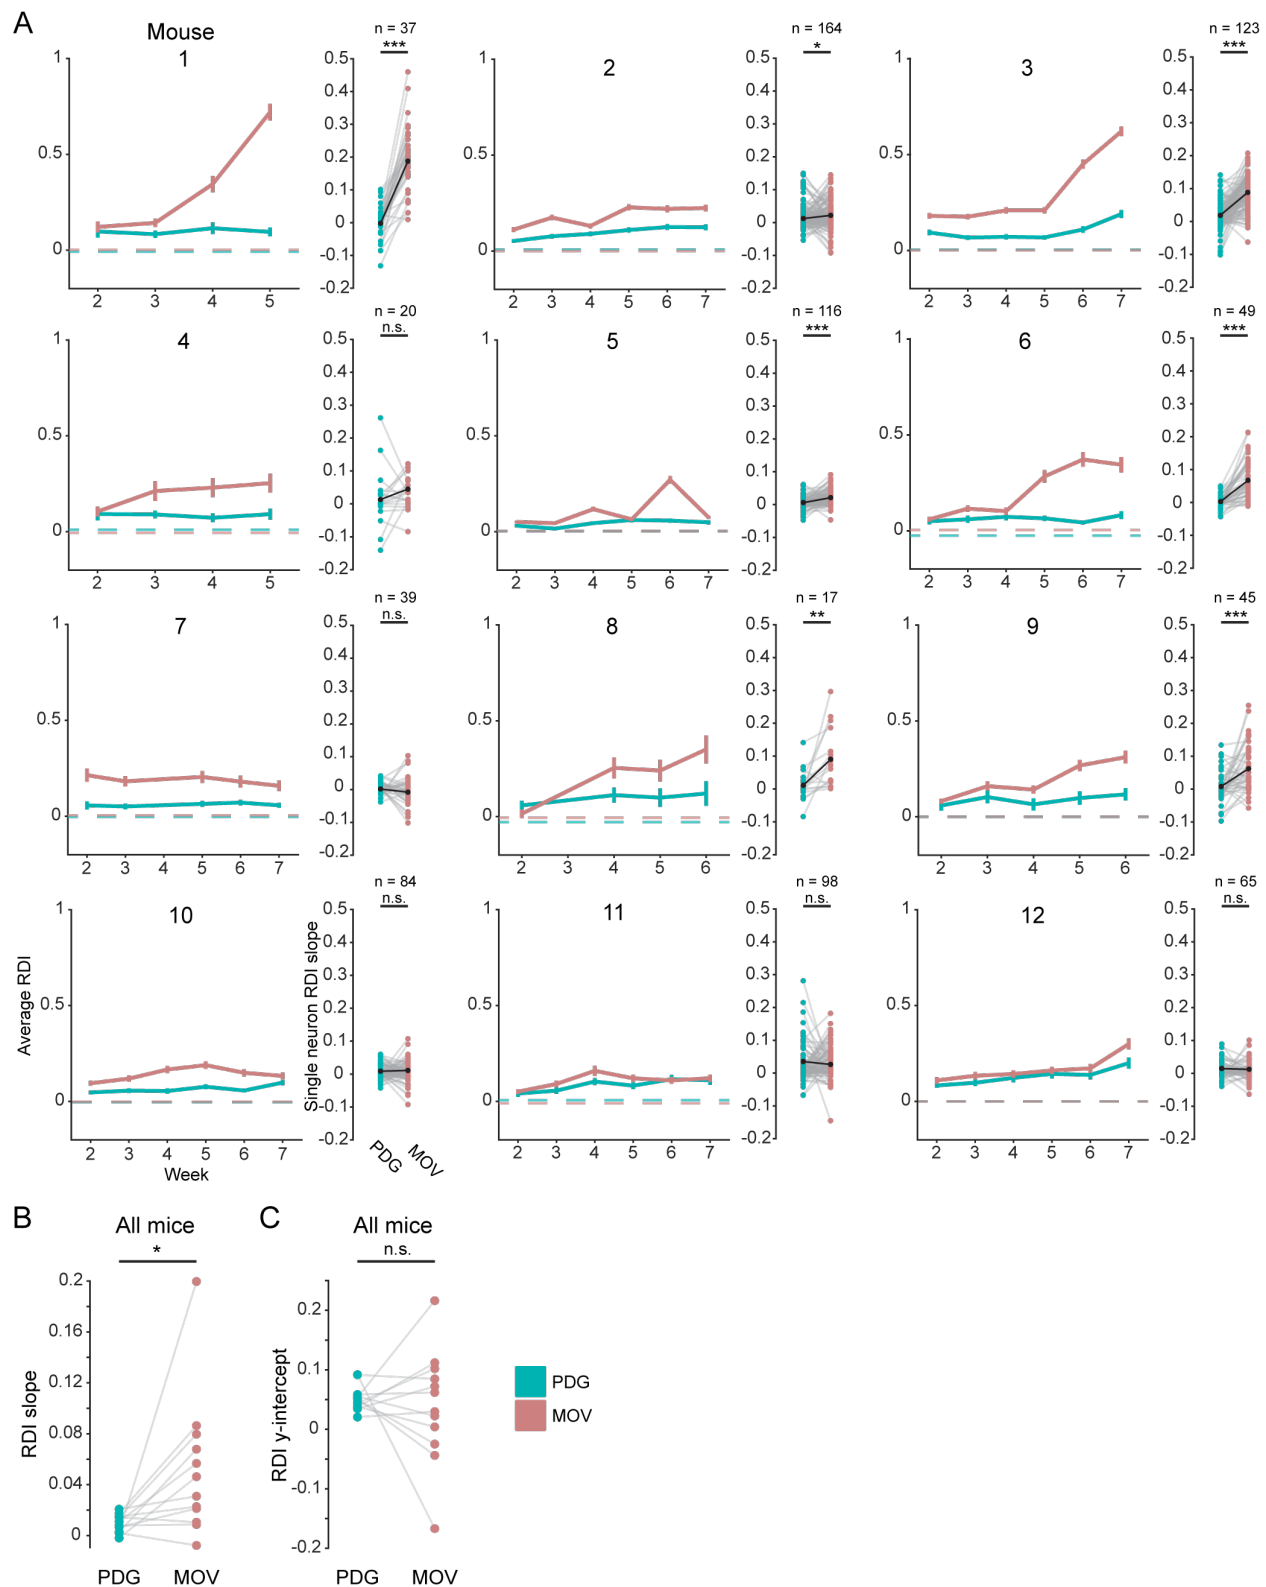
Additional examples of single neuron responses to the MOV stimulus. White borders indicate the separation of recording sessions. Each neuron's MOV RDI curve is shown on the top right of each heatmap. Data is normalized across sessions.



Supplementary Figure 4: Population responses to MOV stimulus.

(A) Pixels of an example MOV recording pseudo-colored by response timing on the first recording week and the seventh recording week. Arrowheads indicate a subset of example neurons that were well tracked across recordings.

(B) Example RDI curves and trial-averaged $\Delta F/F$ traces across sessions for neurons shown in (A).



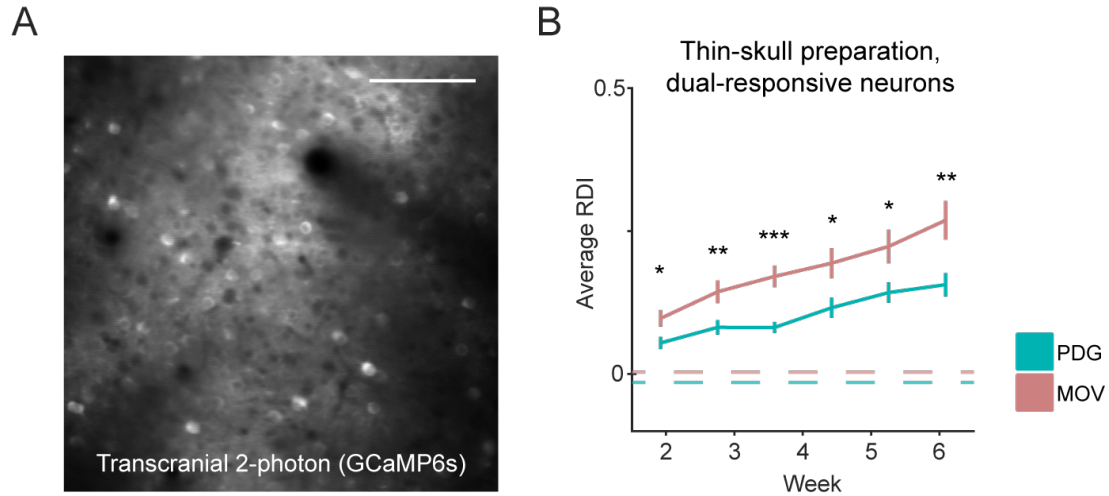
Supplementary Figure 5: Individual mice exhibit variability in amount of representational drift.

(A) Line plots: Average RDI curves for every imaged mouse. Values are calculated using only neurons that are present on any given session and responsive to both stimuli. Error bars are \pm s.e.m.; dotted line

indicates control RDI, as in previous figures. Scatter plots: pairwise comparison of RDI curve slopes for single neurons, each linked pair is one neuron (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, paired-sample t-test).

(B) Comparison of average RDI curve slopes for each stimulus across all mice. Each linked pair is one mouse ($t = 2.60$, $p = 0.02$, paired-sample t-test).

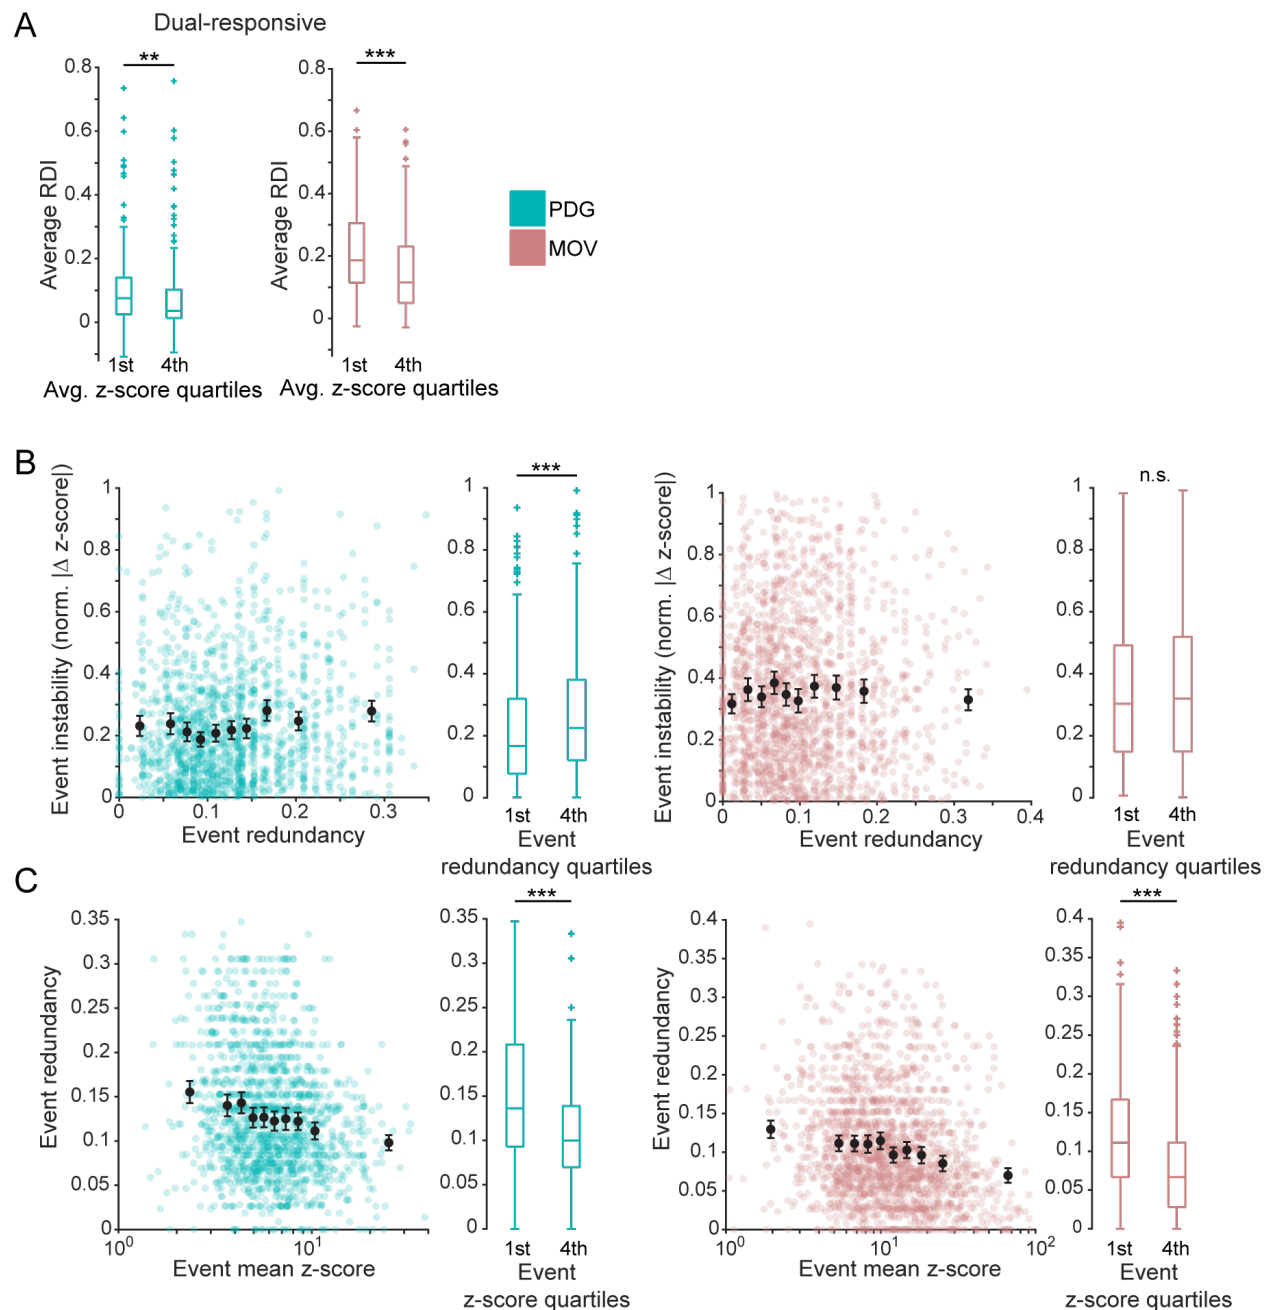
(C) Comparison of average RDI curve y-intercepts for each stimulus across all mice. Each linked pair is one mouse ($t = 0.44$, $p = 0.67$, paired-sample t-test).



Supplementary Figure 6: Method of surgical preparation does not influence RDI curves

(A) Average fluorescence image from example transcranial imaging field. Scale bar = 100 μ m.

(B) Average RDI curves for 4 mice (81, 84, 84, 79, 75, 50 dual-responsive neurons per session shown) imaged through a thin skull preparation instead of a cranial window. Error bars are \pm s.e.m. Significance markers indicate the comparison of each session's PDG RDI values and MOV RDI values (* p < 0.05, ** p < 0.01, *** p < 0.001, paired-sample t-test); dotted line indicates control RDI, as in previous figures.

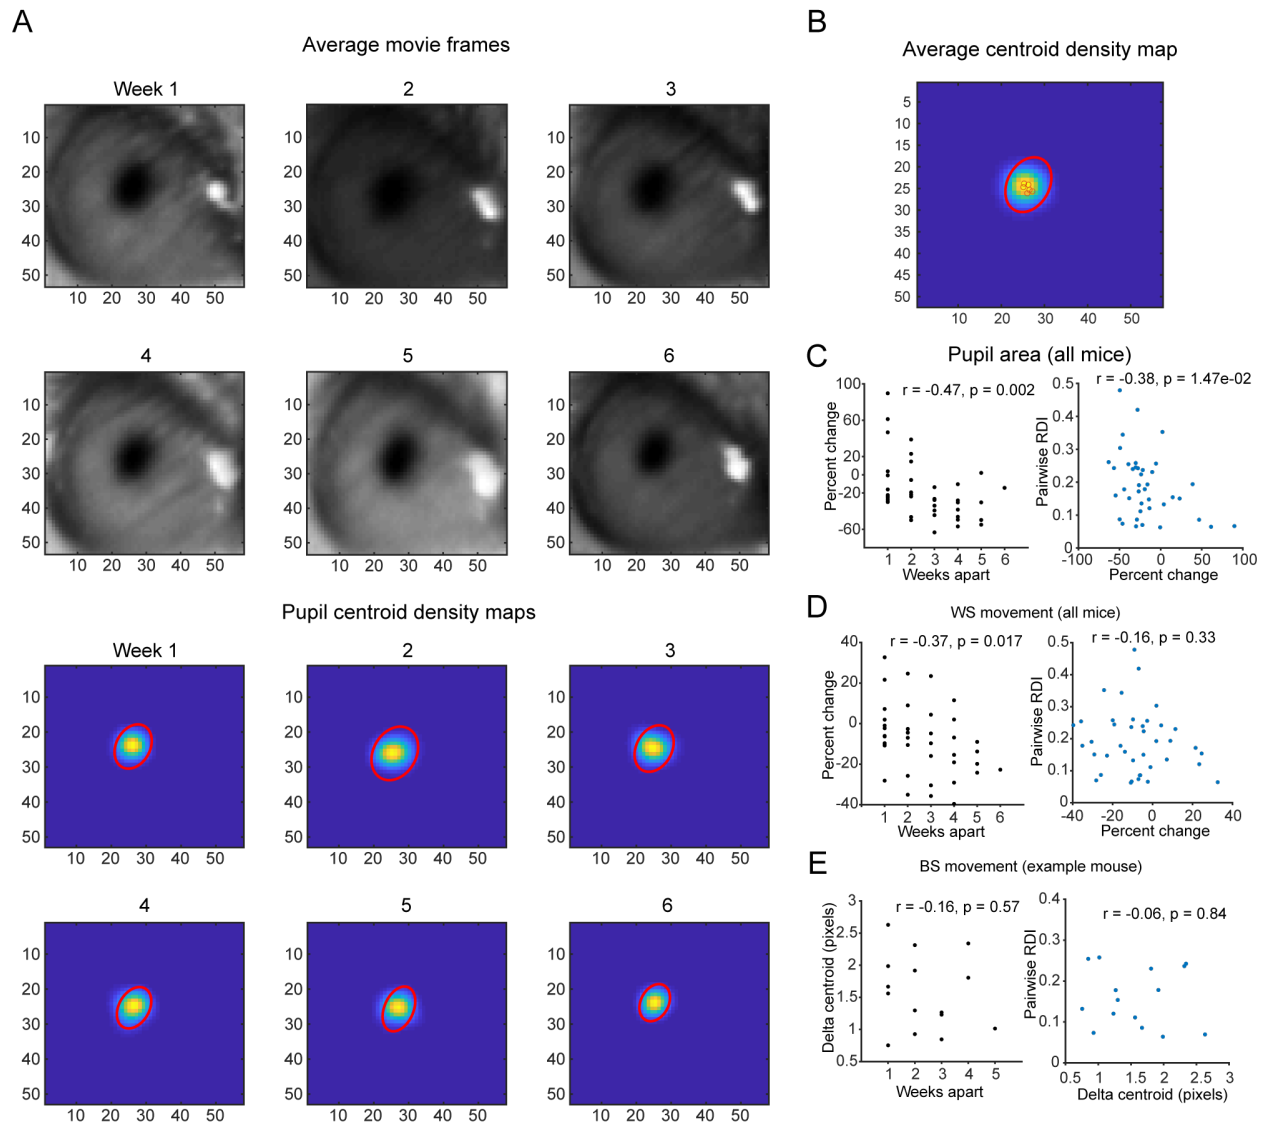


Supplementary Figure 7: Characterization of individual response event stability.

(A) Comparison of first and last quartiles of data from Figure 2A. Highly active neurons are more stable on average than weakly active neurons. (** $p < 0.01$, *** $p < 0.001$, Wilcoxon rank-sum test).

(B) Event instability as a function of event redundancy. Redundancy is defined as the fraction of total neurons with at least one event that overlaps by at least 75% of a given event's duration. Each colored dot is one event; black dots are 10th percentile binned means \pm 95th percent confidence interval. Boxplots compare first and last quartiles of data (Wilcoxon rank-sum test, *** $p < 0.001$). Data shown for all dual-responsive neurons.

(C) Event redundancy as a function of session-average event z-score. Each colored dot is one event; black dots are 10th percentile binned means \pm 95th percent confidence interval. Boxplots compare first and last quartiles of data (Wilcoxon rank-sum test, *** $p < 0.001$). Data shown for all dual-responsive neurons.



Supplementary Figure 8: Pupil size, but not eye movement, correlates with response stability.

(A) Top: Average frames from each of the example recording session's eye tracking videos. Bottom: Density heatmaps of the pupil centroid location on each recording session. Red ellipse depicts the average size, orientation and eccentricity of the pupil, centered on the average centroid location.

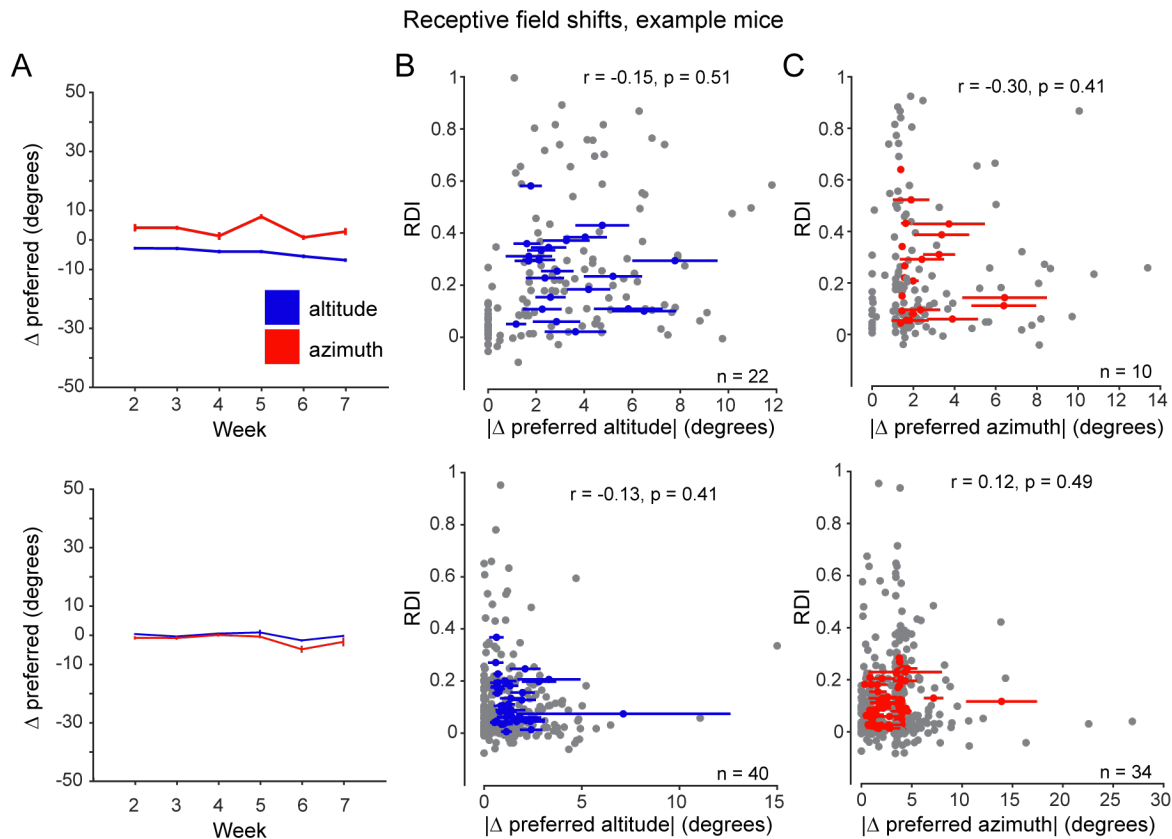
(B) Session-average pupil centroid density map for example recording in (A). Red ellipse depicts the size, orientation and eccentricity of the pupil on week 1. Small red circles depict average pupil centroid location on each recording session.

(C) Relationship between pupil area and RDI. Left: percent change of pupil area as a function of weeks between pairs of recording sessions, where the reference session is always the earlier of the two. Pupil area tends to decrease over the first few sessions (Spearman $r = -0.47, p < 0.01$). Right: percent change of pupil area between two sessions versus their pairwise RDI. Each dot is one pair of recording sessions. RDI is anticorrelated with pupil area (Pearson $r = -0.38, p < 0.05$). Data shown for all eye-tracked mice.

(D) Relationship between within-session pupil movement and RDI. Left: percent change of average distance from mean pupil centroid location within each session as a function of weeks between pairs of sessions. Average amount of movement within each session tends to decrease over time (Spearman $r = -$

0.37, $p < 0.05$). Right: percent change of within-session pupil movement between sessions versus pairwise RDI. Each dot is one pair of recording sessions. There is no association between relative degree of within-session movement and RDI (Pearson $r = -0.16$, $p > 0.05$). Data shown for all eye-tracked mice.

(E) Relationship between between-session pupil movement and RDI, for one example mouse. Left: Distance between mean pupil centroid location as a function of weeks between recording sessions. Average pupil location does not drift over time (Spearman $r = -0.16$, $p > 0.05$). Right: between-session pupil movement versus pairwise RDI. Each dot is one pair of sessions. Change in pupil centroid location is not correlated with RDI (Pearson $r = -0.06$, $p > 0.05$).

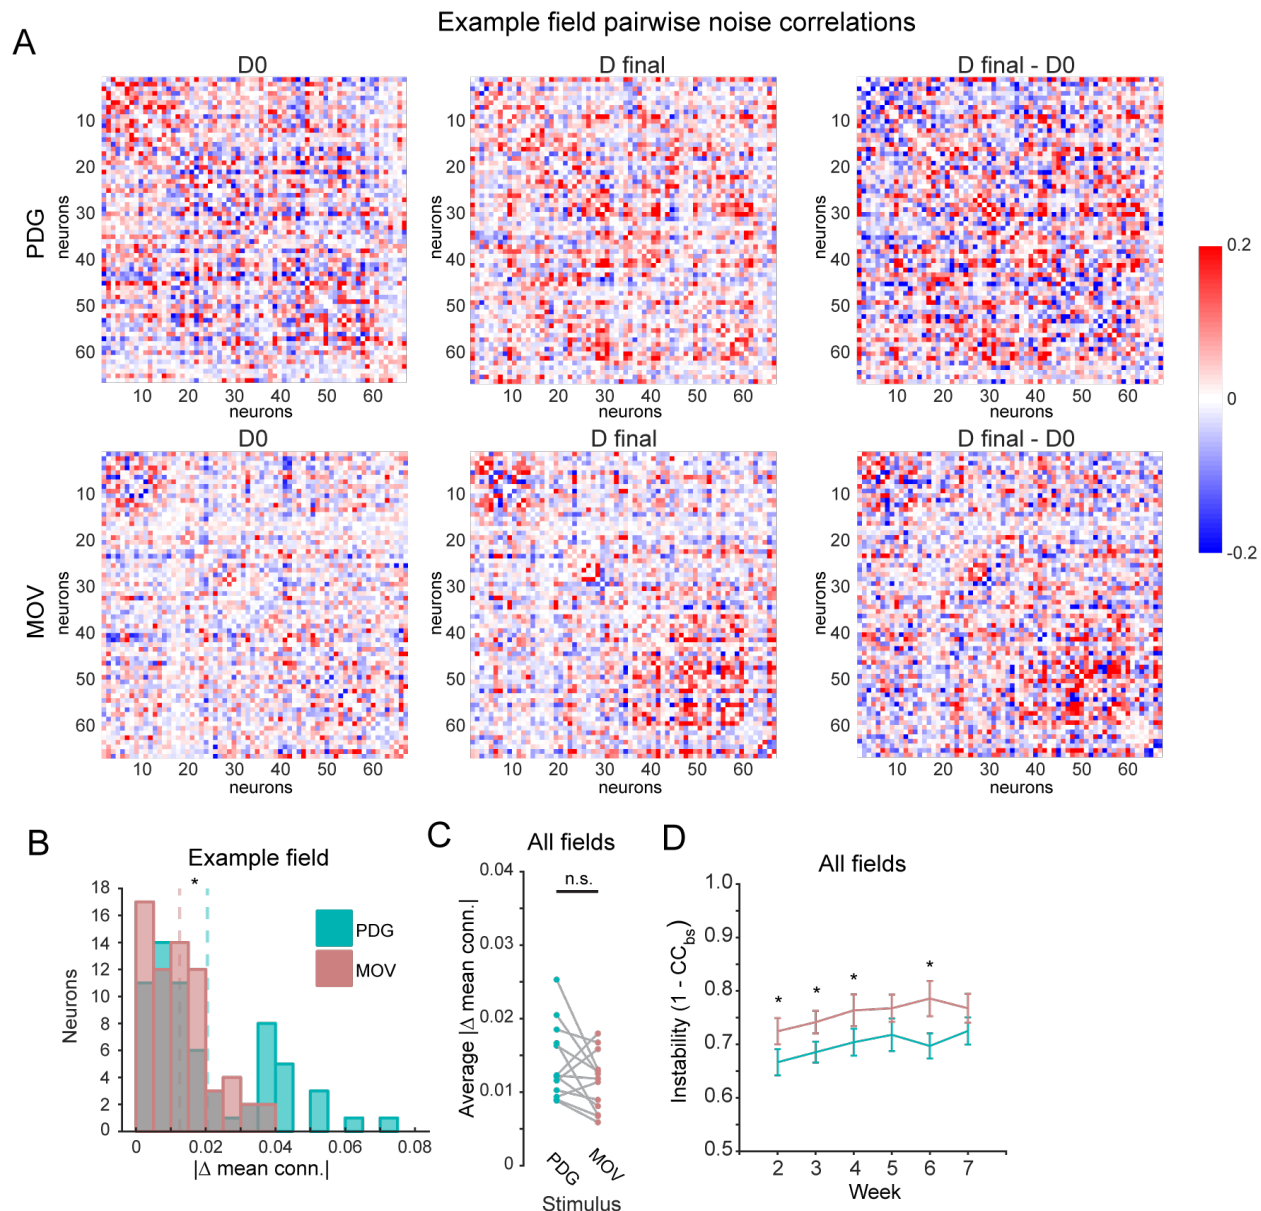


Supplementary Figure 9: RDI is not associated with changes in spatial receptive field location.

(A) Average change in preferred altitude (blue) and azimuth (red) across weeks. Error bars are \pm s.e.m. The visual stimulus screen spans 100° in the altitude direction and 130° in the azimuth direction.

(B) Effect of a neuron's change in spatial receptive field on its MOV RDI value. Each blue dot is one neuron, depicting its session-average absolute change in preferred altitude versus its session-average MOV RDI; horizontal error bars are \pm s.e.m. Each grey dot is a neuron's absolute change in preferred altitude on a given session versus its MOV RDI value on that session (there are multiple grey dots for every blue dot). Pearson correlation displayed on each plot uses colored data points. Data shown for all well-tracked neurons spatially tuned to altitude and visually responsive to MOV.

(C) Same as in (B), but for preferred azimuth. Data shown for all well-tracked neurons spatially tuned to azimuth and visually responsive to MOV.



Supplementary Figure 10: Population noise correlation stability is not clearly stimulus-dependent

(A) Pairwise noise correlations on session 1 (left), final session (middle), and their difference (right) for one example field. Neurons are sorted by time of peak response on D0 for each stimulus. Data shown for all neurons responsive to both stimuli.

(B) Distributions of single-neuron average change in noise correlations between first and final sessions for example field in A. Dotted lines are means for each stimulus (* $p < 0.05$, Wilcoxon rank-sum test).

(C) Field-average changes in noise correlation between first and final sessions. Data shown for all fields. ($p > 0.05$, paired-sample t-test, $n = 13$ fields). Although the example field showed a differential noise correlation stability between stimuli, there is no consistent difference across fields.

(D) Average instability of noise correlation matrices with respect to the first session over time ($1 - CC_{bs}$, where CC_{bs} is the 2D cross correlation between noise correlation matrices). Data shown for all fields. Error bars are \pm s.e.m.; significance markers indicate comparison of PDG and MOV values on given session (* $p > 0.05$, paired-sample t-test, $n = 13$ fields).