Face-selective neurons maintain consistent visual responses across months

David B. T. McMahon[a,1], Adam P. Jones[a,b], Igor V. Bondar[c], and David A. Leopold[a,d]

[a]Section on Cognitive Neurophysiology and Imaging, Laboratory of Neuropsychology, National Institute of Mental Health, National Institutes of Health, Bethesda, MD 20892; [b]Department of Biology, University of Maryland, College Park, MD 20742; [c]Institute of Higher Nervous Activity and Neurophysiology, Moscow 117485, Russia; and [d]Neurophysiology Imaging Facility, National Institute of Mental Health, National Eye Institute, and National Institute of Neurological Disorders and Stroke, National Institutes of Health, Bethesda, MD 20892

Edited by Charles Gross, Princeton University, Princeton, NJ, and approved April 1, 2014 (received for review September 30, 2013)

Face perception in both humans and monkeys is thought to depend on neurons clustered in discrete, specialized brain regions. Because primates are frequently called upon to recognize and remember new individuals, the neuronal representation of faces in the brain might be expected to change over time. The functional properties of neurons in behaving animals are typically assessed over time periods ranging from minutes to hours, which amounts to a snapshot compared to a lifespan of a neuron. It therefore remains unclear how neuronal properties observed on a given day predict that same neuron's activity months or years later. Here we show that the macaque inferotemporal cortex contains face-selective cells that show virtually no change in their patterns of visual responses over time periods as long as one year. Using chronically implanted microwire electrodes guided by functional MRI targeting, we obtained distinct profiles of selectivity for face and nonface stimuli that served as fingerprints for individual neurons in the anterior fundus (AF) face patch within the superior temporal sulcus. Longitudinal tracking over a series of daily recording sessions revealed that face-selective neurons maintain consistent visual response profiles across months-long time spans despite the influence of ongoing daily experience. We propose that neurons in the AF face patch are specialized for aspects of face perception that demand stability as opposed to plasticity.

vision | fMRI | physiology

Many biological systems perform in a stable manner over time. The consistent behavior that is evident at the global level of the organism can persist despite continuous change in the system's component parts. A cup of coffee, for instance, tastes the same today as it did one year ago, despite the fact that the receptors in our taste buds are replaced every two weeks (1). In addition to the continuous replacement of individual cells, the protein constituents of cells likewise undergo nearly continuous turnover (2). What happens in the central nervous system? Because the lifespan of most neurons in the brain approaches that of the organism (3), it is conceivable that the population of face-selective cells that fire today when you see your mother, for instance, is identical to the population that fired under the same circumstances 10 years ago. However, in the absence of direct evidence, there is no reason to assume that this is the case, or more generally that stable network performance implies stable units. And indeed, both theoretical (4) and experimental results from motor cortex (5, 6) demonstrate that a network composed of unstable units may nonetheless exhibit stable performance. Recent recordings from the mouse hippocampus showed that reliable location signaling is driven largely by neurons that gradually enter and leave the population of functionally active place cells over the course of several days (7). In that study, only 15% of longitudinally monitored neurons were found to retain the same place fields in two sessions that were separated by one month. This finding cuts against the strong assumption that, in the absence of explicit learning pressure, the default is for neurons to remain the same.

Computational studies of neural network models have pointed to a problem known as the "stability–plasticity dilemma" that presumably confronts the human brain routinely in the course of object recognition (8). This issue arises from the opposing processes of storing new patterns in a memory network, which requires plasticity, and recalling previously stored patterns under the appropriate conditions, which requires stability (9–11). For highly social animals such as humans and other primates, face perception poses special challenges that likely affect the balance of stability and change in how faces are represented in the brain. Constraints favoring stability arise from the fact that we need to recognize the same individuals over long time periods. This task poses special challenges because, due to the pliability and expressiveness of the facial musculature, the same individual may present a drastically different appearance from one encounter to the next. The physical mutability of faces is an extreme case of the more general problem of invariant object recognition, which the brain must solve across a broad range of viewing conditions. Constraints favoring plasticity arise from the need to constantly update our repertoire of social knowledge. Consistent with this notion, recordings from human surgical patients reveal that single neurons can be so narrowly tuned that they only respond to the sight (or printed name) of specific individuals (12). Learning to recognize new individuals is likely to be mediated by synaptic changes within at least some components of the face-processing network.

Whereas theoretical and intuitive accounts suggest that the neuronal representation of faces is subject to opposing pressures for stability and flexibility, it remains unclear how this tradeoff is resolved at the level of individual neurons in the primate brain. Results from one early study suggested that the visual response

Significance

The brains of humans and other primates contain specialized regions dedicated to the perception of socially important objects such as faces. It is not known whether face representations remain stable over time, or alternatively whether they are subject to slow change in response to ongoing experience. By using implanted microwires to monitor the activity of single cells longitudinally across several weeks and up to one year, we demonstrate that face-selective neurons maintain the same distinctive selectivity patterns for as long as they were followed. The long-term consistency of these neurons might reflect a division of labor within the face-processing network for aspects of social perception that require stability as opposed to plasticity.

Author contributions: D.B.T.M., I.V.B., and D.A.L. designed research; D.B.T.M., A.P.J., and I.V.B. performed research; D.B.T.M., A.P.J., and D.A.L. analyzed data; and D.B.T.M. and D.A.L. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence should be addressed. E-mail: mcmahond@mail.nih.gov.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1318331111/-/DCSupplemental.

patterns in some face-selective neurons for a familiar set of images is altered when novel faces are introduced (13). However, the same study also reported that the majority of neurons maintained consistent visual responses, and the changes that did occur in a minority of neurons were typically small. The primate brain contains a number of discrete regions specialized for face processing (14–17). This functional organization raises the possibility that different aspects of face perception that make opposing demands for stability and flexibility could be processed in separate neuronal populations. Our ability to test this idea is limited by the fact that virtually all previous physiological studies of face-selective neurons were conducted with acute recording electrodes, which do not permit monitoring the same neuron for more than a few hours. Using chronically implanted microwire bundles, we recently showed that, in many cases, the visual properties of object-selective neurons in the inferotemporal cortex are consistent across four to five days (18). Here we exploited the stability of the chronic microwire technique to ask whether neurons in functionally identified face patches in the temporal lobe show consistent patterns of stimulus selectivity over months-long time scales.

Results

We localized face-selective regions in the temporal lobe of two macaque monkeys using a standard functional MRI (fMRI) block design in which alternating epochs of face and nonface stimuli were presented while the monkeys fixated (Fig. 1 *A* and *B*). In agreement with earlier studies (14, 19), face-selective clusters of voxels were detected in the superior temporal sulcus (STS) at ~6 mm and 16 mm anterior to the interaural canal, corresponding to the "middle fundus" (MF) and "anterior fundus" (AF) patches, respectively [following the nomenclature of Tsao et al. (19)]. Based on this functional mapping, chronic microwire electrodes were implanted in the anterior STS face patch. Again in keeping with previous reports (20–24), single neurons in the AF face patch showed strong preference for face stimuli when tested with a library of images contrasting face and nonface stimuli ($P << 0.01$ in both monkeys, two-way ANOVA followed by Tukey–Kramer test; Fig. 1 *C* and *D*).

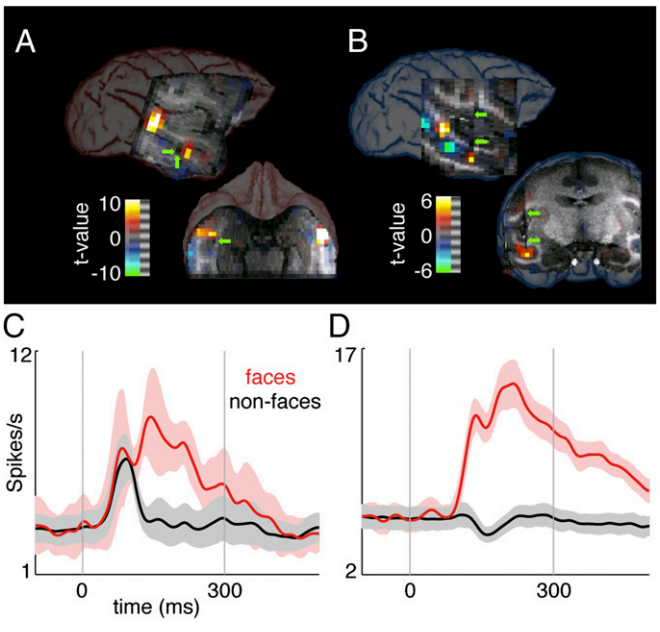

Fig. 1. Face-selective activity detected at the recording sites. (*A* and *B*) Location of face-selective voxels corresponding to face patches MF and AF within the STS in monkey 1 (*A*) and monkey 2 (*B*). Arrows indicate the position of the electrode track and tip within the anterior (AF) patch. (*C* and *D*) Face-selective responses from the population of neurons recorded on a single day in monkey 1 (*C*) and monkey 2 (*D*). Category selectivity was assessed using a library of images consisting of faces, objects, and body parts that differed from the stimulus sets used in daily longitudinal recordings. Firing rate histograms show the average evoked responses of 7 and 42 cells, respectively (error bars, SEM).

After the initial implantation procedure, the electrodes were advanced in small steps every few days. We previously showed that the microwires used in this study were capable of isolating spikes from single neurons from one day to the next (18). As in our initial report, spikes isolated on the same wire across days often showed stable waveform shapes and firing statistics (Fig. 2*A*). The strategy of the current study was specifically to exploit the stability of the microwire recordings (25, 26) to ask whether the functional properties of face-selective cells are consistent over long time scales. Accordingly, after briefly verifying that the basic visual properties of the targeted neurons conformed to expectations regarding face selectivity in recordings collected on a single day (Fig. 1 *C* and *D*), we commenced a series of longitudinal recording sessions in which fixed stimulus sets were presented repeatedly over multiple days. In each longitudinal experiment, the stimuli used were at first novel to the monkeys on the first day of recording, and gradually became more familiar over repeated exposures to the same images over the course of the subsequent days. This design allowed us to test how neural responses changed as novel stimuli became familiar to the animal, and then to further track any changes in the responses to the same stimuli over weeks and months.

Recordings from 144 neurons in the STS were obtained (85 from monkey 1, 59 from monkey 2). During these longitudinal experiments, the monkeys were rewarded for maintaining fixation while images from a library of face and nonface stimuli (from 54 to 206 images) were presented for 300 ms in random order on a computer screen. The monkeys were tested with these fixed stimulus sets for approximately one-hour sessions over the course of at least 13 days (mean, 37 days; max, 68 days). Because the aim of these recordings was to obtain dense longitudinal sampling of the neurons' selectivity patterns, the monkeys were tested daily whenever possible. In monkey 1, the electrode was typically advanced between finishing one longitudinal experiment and starting the next experiment. In monkey 2, the electrode was only advanced three times within the month following surgery and then allowed to remain in place. This permitted occasional sparse sampling of single unit responses over much longer intervals using two stimulus sets (93 days and 383 days). The aim of these intermittent sessions was to probe the upper limit of the time scales over which face-selective neurons show consistent visual responses.

Representative examples of face-selective visual responses from three neurons are shown in Fig. 2 *B*–*E*. Neurons showed various patterns of selectivity that included broad excitatory tuning with the greatest responses to monkey faces (Fig. 2 *B* and *C*), more narrowly tuned cells that also responded to nonface stimuli (Fig. 2*D*), and suppressive responses that were greatest for monkey faces (Fig. 2*E*). As the longitudinal trends shown in Fig. 2 illustrate, neurons in AF routinely gave very similar responses to the same stimuli from one day to the next and showed little if any change in their distinctive firing patterns for as long as we continued to follow them. More examples of individual neurons are shown in Fig. S1. Consistent functional properties were likewise maintained in neurons that were probed intermittently over much longer time periods, an extreme instance of which is shown in Fig. 3. The recordings from this cell were conducted at irregular intervals as opportunity permitted (i.e., when the monkey was not involved with other physiological or functional imaging experiments) and spanned more than one year (383 days). During this period, three other neurons were followed for at least 165 days and also maintained the same selectivity patterns throughout.

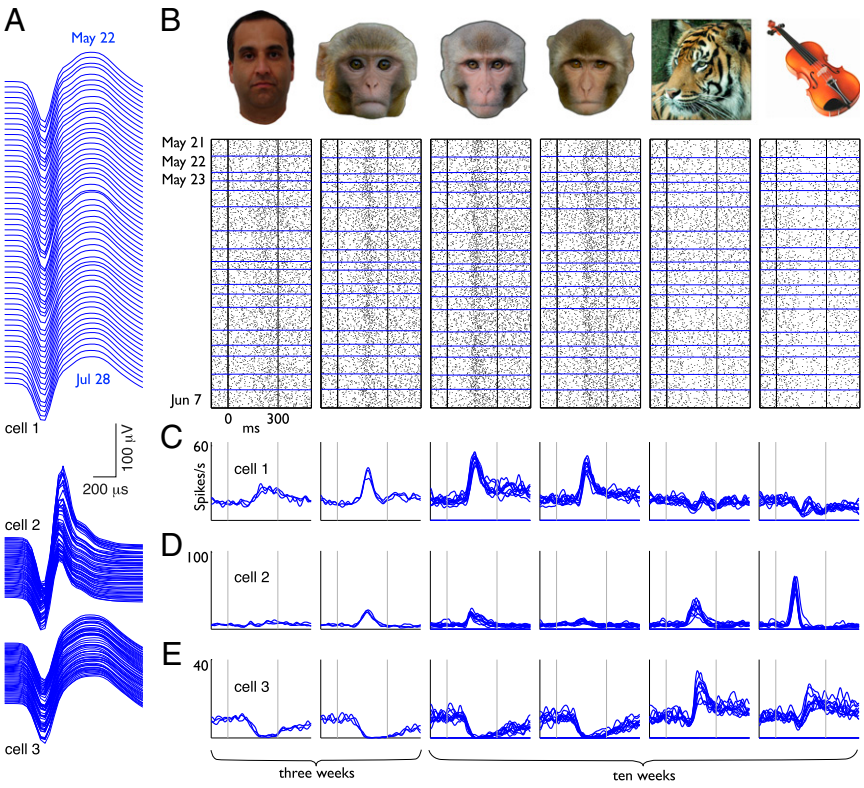

Fig. 2. Consistent visual responses from three face-selective neurons tracked in daily recording sessions over two months. (*A*) Spike waveforms of three neurons observed in a series of daily recording sessions over the course of 68 days. Principal components clusters are shown in Fig. S2. (*B*) Stimuli and corresponding visual responses from one neuron (cell 1, also shown in *A* and *C*) drawn from a series of daily recording sessions. Each raster plot consists of 436 trials collected over the first 18 consecutive days of recording. (*C*–*E*) Weekly averaged visual responses from the same neurons shown in *A* over three weeks (469 trials, left two stimuli) and 10 weeks (1,200 trials, right four stimuli).

These instances of very long-lasting stable neurons indicate that the upper limit on the duration of consistent visual responses extends far beyond the range that could feasibly be sampled in daily screening sessions.

In eight longitudinal recording sessions (five in monkey 1, three in monkey 2), the average lifespan of an isolated spike was 29 days (interquartile range, 9–37 days). In total, 59 neurons were held for at least 20 days (19 neurons in monkey 1, 40 neurons in monkey 2). To assess the day-to-day consistency of visual responses, we divided the trials into two subsets and computed a split-halves correlation coefficient for each day. This value (r_{within}) was based on the averaged responses to all stimuli presented on odd versus even trials (Fig. 4*A*). An identical procedure was used to compute two additional correlation coefficients based on split-halves drawn from two consecutive days (r_{across}) or after shuffling the identity of stimuli ($r_{shuffle}$). On average, selectivity patterns assessed across days were only marginally (albeit significantly) less correlated with each other than selectivity patterns measured within the same day (mean, 0.64 and 0.59, respectively; Fig. 4*B*), whereas both correlations were significantly greater than the stimulus-shuffled control (mean $r = 0.006$; $P << 0.01$, two-sampled t test; Fig. 4*B*, *Inset*). Equivalent results were obtained when r_{across} was computed based on the first and last days of recording, both in the cases of the four example cells (Fig. 4*A*, thick red lines) and in the population (Fig. S3). This pattern of results indicates that the subjective impression of consistent visual responses conveyed by the example neurons shown in Figs. 2 and 3 is representative of the population of longitudinally monitored neurons in AF.

Because our conclusion that face-selective cells are consistent across days rests in large part on the fact that we observed similar patterns of stimulus selectivity, we asked whether the columnar organization of the inferotemporal cortex might provide an alternative explanation for our findings (27–29). Specifically, we considered the possibility that if face patches showed the same tendency for neurons in close proximity to have closely matching patterns of face selectivity, then two different nearby neurons encountered on successive days might mistakenly be considered the same neuron. To assess the likelihood of this error, we first quantified the within-cell consistency of visual responses by computing a correlation coefficient between responses to all stimuli on odd versus even trials (Fig. 4*C*; median r, 0.94; interquartile range, 0.74–0.98; $n = 144$). We then repeated the same calculation for all pair-wise combinations (also for odd vs. even trials) of different neurons that were screened with the same stimulus set (median $r =$ 0.10; interquartile range, –0.09–0.37; $n = 1{,}561$ pairs). The correlations between different neurons, although weak, were significantly greater than zero ($P << 0.01$, one-sampled t test), confirming that the AF face patch does indeed follow the trend that nearby neurons have similar selectivity patterns. However, the split-halves correlations were much weaker between different neurons than within the same neurons ($P << 0.01$, two-sampled t test). Based on these results, a receiver operating characteristic (ROC) analysis showed that two different neurons would only be incorrectly identified as the same 7% of the time (d-prime, 0.75; Fig. 4*C*, *Inset*). We conclude that the consistent patterns of visual responses that we observed in face-selective cells cannot be explained by the columnar structure of face patch AF.

Discussion

Biological systems require the capacity for both stability and flexibility. Whether consistency or change predominates is likely

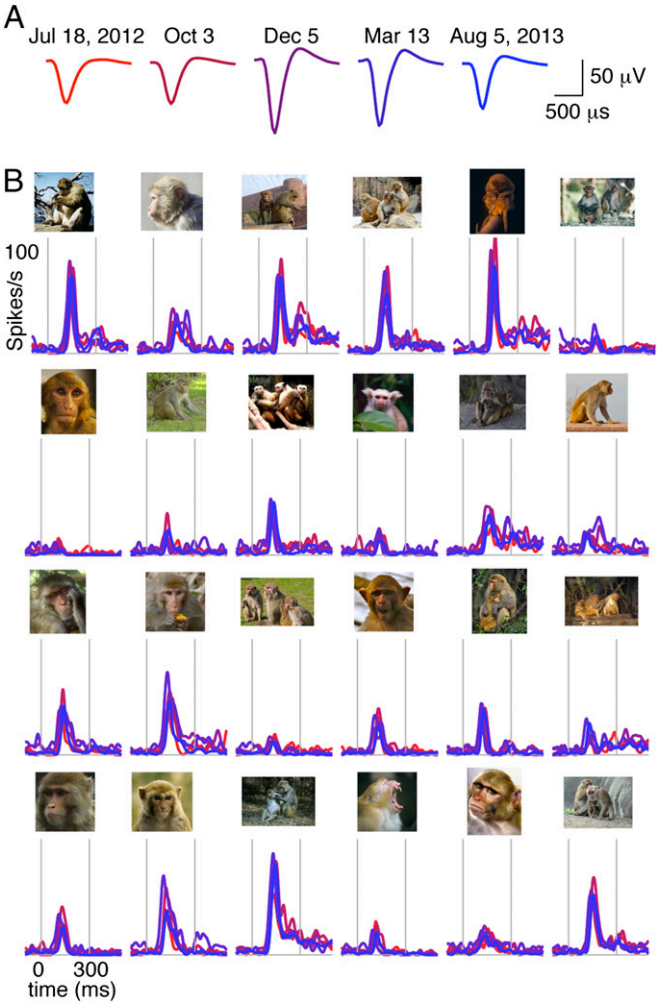

Fig. 3. Consistent visual responses from one face-selective neuron recorded intermittently over one year. (*A*) Spike waveforms from a single neuron recorded on five experimental sessions spaced 63–145 days apart from July 18, 2012 to August 5, 2013. Principal components clusters are shown in Fig. S2. (*B*) Visual responses from the same neuron to 24 stimuli observed over the course of one year. The full stimulus set included 96 images that were presented in each session. Images included both close-up pictures of faces and distal views of whole bodies for both macaque monkeys and new world monkeys.

to depend on both the level of analysis (from molecules to whole organism) and the time scale (from seconds to years) at which the system is observed. Most previous research at the level of single neuron activity in behaving mammals considered time scales of a few hours or less, largely due to the technical limitations of conventional recording methods. In previous work using chronically implanted microwires, we found that object selectivity patterns in the inferotemporal cortex (primarily sampling TEa and TEav) commonly remain consistent for as long as they were followed in that study—a few days on average, and in a handful of cases past two weeks (18). In the current study, we focused specifically on neurons in the AF face patch and found that functional properties are consistent over time scales an order of magnitude longer than those tested previously: nearly one month on average, and in the longest cases over one year. Moreover, in contrast to Bondar et al. (2009) (18), in which the stimuli were highly familiar from the beginning, the current study was able to track the course of responses over the course of familiarization, from the first presentation of each stimulus. Our observation that response profiles were stable throughout the entire testing period, including the initial familiarization, is surprising in light of earlier work showing that inferotemporal neurons are commonly sensitive to whether stimuli are novel or familiar (30–33).

We speculate that the discrepancy between our current findings and earlier studies pointing to more labile selectivity patterns might be due to the functional heterogeneity of the inferotemporal cortex (34, 35). A recent study showed that the prevalence of identity tuning was greater in the anterior-most face patch (AM), compared with the more posterior face patches (AL, ML, and MF, meaning anterior lateral, middle lateral, and middle fundus, respectively) in which viewpoint tuning prevailed (22). This division of labor among face patches lends credence to the idea that some face patches might maintain consistent visual responses whereas others are more subject to change. For instance, it is reasonable to suspect that regions dedicated to remembering new acquaintances might be more susceptible to plasticity than regions dedicated to the spatial aspects of face perception. If this idea is correct, the opposing constraints of maintaining stable representations of known faces and acquiring new social knowledge without disrupting previously acquired response patterns could be resolved by the anatomical segregation of different aspects of face processing. The results of the current study suggest that the AF face patch might be well suited for the aspects of face processing that demand temporal consistency.

Because longitudinal studies of individual neurons across days are uncommon in the literature, it remains unclear how the long-term consistency of face selectivity found in AF compares with the functional properties of neurons in other brain areas. Development of perceptual systems in the brain is generally understood to involve a plastic phase that tunes the sensitivities of a naive network to match the statistics of sensory experience, followed by a stable phase in which functional properties are held fixed by local abatement of plasticity at synapses or through the global influence of attractor dynamics (36, 37). Some evidence from longitudinal recordings from primary visual and somatosensory cortices supports this idea (26, 38). Results from the motor cortex are more variable, with some groups reporting consistent relations between behavior and the preferred direction of single neurons (39–41) and others finding some degree of change across days (5, 42) or even within the course of a single session (6). Because activity outside of the primary sensory and motor cortical areas is less tightly coupled to action in the periphery, we might expect to see more change over time in higher level association areas. Neurons in the prefrontal cortex show a high degree of adaptability in acquiring new functional properties in response to learning pressures (43–45), but they also maintain broadly consistent response patterns during tasks involving well-established rules (46). A remarkable early longitudinal study in the rodent hippocampus showed one instance of a place cell that maintained a consistent place field over 152 days (47). Recent work confirmed that some hippocampal neurons do maintain stable place fields over the course of one month (7). However, the same study also showed that the majority of place cells in area CA1 are more transient and gain or lose their location selectivity over the course of a few days. Whereas the population of place-selective cells appears to undergo continuous turnover in CA1, the neighboring hippocampal area CA3 contains neurons that show more consistent place fields across days (48). It thus appears that different subregions within the same brain structure can employ distinct strategies to resolve the opposing pressures favoring stability versus plasticity. The current finding that neurons commonly give similar responses upon seeing the same faces months apart raises the possibility that some neurons might respond the same way to the same individual faces over most of the animal's lifespan.

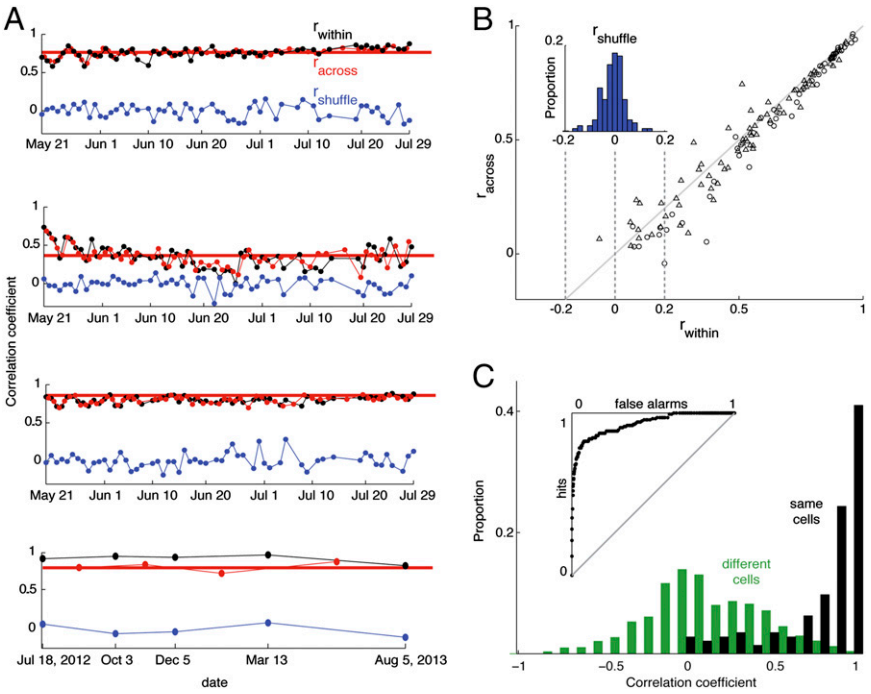

Fig. 4. Correlation-based assessment of visual response consistency. (*A*) Correlation coefficients derived from split-halves analysis (odd vs. even trials) comparing visual responses recorded within a single day (black), between two days (red), and after shuffling the stimulus order (blue). Each point is based on 206 stimuli presented over 68 days (top three, same cells as in Fig. 2) or 96 stimuli shown over 383 days (bottom plot, same cell as in Fig. 3). The position of the thick horizontal red line indicates the r_{across} correlation between the first and last days of recording. (*B*) Population summary of 144 neurons (85 from monkey 1, triangles; 59 from monkey 2, circles) showing correlations computed between odd and even trials recorded on the same day (r_{within}, *x* axis) and on different days (r_{across}, *y* axis). (*Inset*) Distribution of $r_{shuffle}$ values for the same 144 neurons. (*C*) Population summary comparing split-halves correlations computed for trials drawn from the same neurons versus different neurons. Black, odd versus even trials from the same neuron (144 cells); green, odd versus even trials from all pair-wise combinations of different neurons for which responses to the same stimuli were available (1,561 pairs, eight stimulus sets). (*Inset*) ROC curve plotting proportion of false alarms (*x* axis) against hits (*y* axis).

Methods

Subjects. Two rhesus macaque monkeys (laboratory designations SI and TO, both females, and 5.0 kg and 5.6 kg, respectively) were used in the experiments. All procedures were approved by the Animal Care and Use Committee and complied with the regulations of the National Institute of Mental Health and National Institutes of Health. The subjects were maintained in colony rooms that housed 16 (monkey 1) and 24 (monkey 2) rhesus macaques in total. Over the course of the experiments, the colony membership changed little or not at all. The monkeys likewise encountered mostly the same humans (primarily scientists and members of the animal care staff) from day to day.

Functional Imaging. Functional scans were collected on a 4.7 Tesla vertical scanner (Bruker BioSpin) using T2*-weighted echo planar image (EPI) acquisition sequences using a two-channel surface coil that was centered over the anterior temporal lobes. Voxels resolution was 1.5 mm isotropic, with 20 coronal slices providing coverage from +2 to +32 mm anterior to the interaural canal. Higher resolution (0.5 mm isotropic) anatomical scans were collected using modified-driven equilibrium Fourier transform (MDEFT) scans in separate sessions and registered to the EPI. Functional imaging was based on detection of monocrystalline iron oxide nanocompound (MION) (49) signal in m1 and blood oxygenation level-dependent (BOLD) signal in m2. Functional mapping was achieved through presentation of 24-second blocks consisting of baseline fixation, faces, and nonface objects. Twelve blocks were presented during each scan, with testing sessions consisting of 10–12 such scans. The functional data were run through a standard preprocessing sequence, including magnetic field distortion correction using the PLACE (phase labeling for additional coordinate encoding) algorithm (50), motion correction implemented in AFNI (analysis of functional neuroimages) (51), and high-pass filtering with a cutoff of 0.01 Hz. The effects of the implant on stimulus- and category-related activity were evaluated using custom-written software in Matlab (Mathworks).

Physiological Recordings. Two chronic electrodes consisting of bundles of 32 or 64 microwires were implanted in the AF face patch of the left hemisphere of both monkeys. The apparatus, surgical implantation, and electrophysiological recording procedures will be described in detail in a forthcoming publication. Briefly, microwire electrodes were implanted in both animals together with a custom-made sealed plastic recording chamber and MRI-compatible microdrive in a surgical procedure while anesthesia was maintained under isoflurane. An MRI-compatible head post was also implanted in either the same surgery or in an earlier procedure. The implant was embedded in an acrylic cap and anchored to the skull by ceramic bone screws. Daily recording sessions began two to four weeks after the initial surgery. The electrode was advanced in small steps (350–700 microns) every few days until visually responsive spiking activity was encountered on multiple channels, and then allowed to settle for a period of several weeks. Physiological recordings were carried out in a radio frequency-shielded room. Data were collected using either a Multichannel Acquisition Processor (Plexon Inc.) with 32-channel capacity or a RS4 Bioamp Processor (Tucker-Davis Technologies) with 128-channel capacity. Broadband electrophysiological responses (0.5 Hz–5 kHz) were collected, from which individual spikes were extracted and analyzed. All aspects of the task related to stimulus presentation, eye position monitoring, and reward delivery were controlled by custom software courtesy of David Sheinberg (Brown University, Providence, RI) running on a QNX computer. The monitor (either a ViewSonic 18″ CRT display or a Eizo 17″ LCD Monitor) was placed 91 cm in front of the monkey. Visual stimuli were presented by a graphics slave computer running the psychophysics toolbox (52) in Matlab. Eye position was monitored via an infrared camera and computed into analog voltage signals corresponding to *x* and *y* coordinates by the software EyeLink (SR Research).

Experimental Design. Two stimulus sets of 54 (monkey 1) or 96 (monkey 2) unique face, object, and body part images were used for the initial assessment of category selectivity. For longitudinal recording sessions, eight stimulus sets comprising 54–206 unique images were used for the duration of each longitudinal dataset. On the first day of each longitudinal recording experiment, a new stimulus set was introduced that was at first wholly novel to the animal and then gradually became highly familiar through repeated exposure over subsequent days. The images used in the longitudinal experiments were a varied set that typically included pictures of human faces, monkey

faces and whole bodies, other animals, and nonface objects, but were not designed for the purpose of assessing face selectivity. In all cases, the animals depicted on the screen were unknown to the subjects. The monkeys initiated each trial of the task by fixating within a 3-degree window around a 0.3-degree white fixation spot, and were rewarded with a drop of fruit juice or water for maintaining fixation throughout four full stimulus presentation cycles. Visual stimuli were presented one at a time in random order behind the fixation spot at a duty cycle of 300 ms on and 300 ms off. The stimuli fell within a 3.5-degree square. Typically 16–25 trials per stimulus were collected in each recording session (range, 12–100). Data analysis procedures were as described in *SI Methods*.

ACKNOWLEDGMENTS. We thank Katy Smith, Yemi Afuwape, and Heba Elnaiem for assistance with physiological data collection; Frank Ye, Charles Zhu, and David Yu for assistance with imaging experiments; and Rebecca Berman for comments on the manuscript. This work was supported by the National Institute of Mental Health, National Institute of Neurological Disorders and Stroke, and National Eye Institute Intramural Research Programs, and by National Institutes of Health Grant EY018028 (to D.B.T.M.).

1. Beidler LM, Smallman RL (1965) Renewal of cells within taste buds. *J Cell Biol* 27(2): 263–272.
2. Price JC, Guan S, Burlingame A, Prusiner SB, Ghaemmaghami S (2010) Analysis of proteome dynamics in the mouse brain. *Proc Natl Acad Sci USA* 107(32):14508–14513.
3. Spalding KL, Bhardwaj RD, Buchholz BA, Druid H, Frisén J (2005) Retrospective birth dating of cells in humans. *Cell* 122(1):133–143.
4. Maass W, Natschläger T, Markram H (2002) Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Comput* 14(11):2531–2560.
5. Carmena JM, Lebedev MA, Henriquez CS, Nicolelis MAL (2005) Stable ensemble performance with single-neuron variability during reaching movements in primates. *J Neurosci* 25(46):10712–10716.
6. Rokni U, Richardson AG, Bizzi E, Seung HS (2007) Motor learning with unstable neural representations. *Neuron* 54(4):653–666.
7. Ziv Y, et al. (2013) Long-term dynamics of CA1 hippocampal place codes. *Nat Neurosci* 16(3):264–266.
8. Carpenter GA, Grossberg S (1993) Normal and amnesic learning, recognition and memory by a neural model of cortico-hippocampal interactions. *Trends Neurosci* 16(4):131–137.
9. Treves A, Rolls ET (1992) Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus* 2(2):189–199.
10. O'Reilly RC, McClelland JL (1994) Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off. *Hippocampus* 4(6):661–682.
11. Hasselmo ME, Anderson BP, Bower JM (1992) Cholinergic modulation of cortical associative memory function. *J Neurophysiol* 67(5):1230–1246.
12. Quiroga RQ, Reddy L, Kreiman G, Koch C, Fried I (2005) Invariant visual representation by single neurons in the human brain. *Nature* 435(7045):1102–1107.
13. Rolls ET, Baylis GC, Hasselmo ME, Nalwa V (1989) The effect of learning on the face selective responses of neurons in the cortex in the superior temporal sulcus of the monkey. *Exp Brain Res* 76(1):153–164.
14. Bell AH, Hadj-Bouziane F, Frihauf JB, Tootell RBH, Ungerleider LG (2009) Object representations in the temporal cortex of monkeys and humans as revealed by functional magnetic resonance imaging. *J Neurophysiol* 101(2):688–700.
15. Tsao DY, Freiwald WA, Knutsen TA, Mandeville JB, Tootell RBH (2003) Faces and objects in macaque cerebral cortex. *Nat Neurosci* 6(9):989–995.
16. Pinsk MA, DeSimone K, Moore T, Gross CG, Kastner S (2005) Representations of faces and body parts in macaque temporal cortex: A functional MRI study. *Proc Natl Acad Sci USA* 102(19):6996–7001.
17. Kanwisher N, McDermott J, Chun MM (1997) The fusiform face area: A module in human extrastriate cortex specialized for face perception. *J Neurosci* 17(11): 4302–4311.
18. Bondar IV, Leopold DA, Richmond BJ, Victor JD, Logothetis NK (2009) Long-term stability of visual pattern selective responses of monkey temporal lobe neurons. *PLoS ONE* 4(12):e8222.
19. Tsao DY, Moeller S, Freiwald WA (2008) Comparing face patch systems in macaques and humans. *Proc Natl Acad Sci USA* 105(49):19514–19519.
20. Bell AH, et al. (2011) Relationship between functional magnetic resonance imaging-identified regions and neuronal category selectivity. *J Neurosci* 31(34):12229–12240.
21. Tsao DY, Freiwald WA, Tootell RBH, Livingstone MS (2006) A cortical region consisting entirely of face-selective cells. *Science* 311(5761):670–674.
22. Freiwald WA, Tsao DY (2010) Functional compartmentalization and viewpoint generalization within the macaque face-processing system. *Science* 330(6005):845–851.
23. Desimone R, Albright TD, Gross CG, Bruce C (1984) Stimulus-selective properties of inferior temporal neurons in the macaque. *J Neurosci* 4(8):2051–2062.
24. Perrett DI, Rolls ET, Caan W (1982) Visual neurons responsive to faces in the monkey temporal cortex. *Exp Brain Res* 47(3):329–342.
25. Krüger J, Caruana F, Volta RD, Rizzolatti G (2010) Seven years of recording from monkey cortex with a chronically implanted multiple microelectrode. *Front Neuroeng* 3:6.
26. Nicolelis MA, Ghazanfar AA, Faggin BM, Votaw S, Oliveira LM (1997) Reconstructing the engram: Simultaneous, multisite, many single neuron recordings. *Neuron* 18(4): 529–537.
27. Fujita I, Tanaka K, Ito M, Cheng K (1992) Columns for visual features of objects in monkey inferotemporal cortex. *Nature* 360(6402):343–346.
28. Sato T, Uchida G, Tanifuji M (2009) Cortical columnar organization is reconsidered in inferior temporal cortex. *Cereb Cortex* 19(8):1870–1888.
29. Kreiman G, et al. (2006) Object selectivity of local field potentials and spikes in the macaque inferior temporal cortex. *Neuron* 49(3):433–445.
30. Li L, Miller EK, Desimone R (1993) The representation of stimulus familiarity in anterior inferior temporal cortex. *J Neurophysiol* 69(6):1918–1929.
31. Freedman DJ, Riesenhuber M, Poggio T, Miller EK (2006) Experience-dependent sharpening of visual shape selectivity in inferior temporal cortex. *Cereb Cortex* 16(11): 1631–1644.
32. Woloszyn L, Sheinberg DL (2012) Effects of long-term visual experience on responses of distinct classes of single units in inferior temporal cortex. *Neuron* 74(1):193–205.
33. Riches IP, Wilson FA, Brown MW (1991) The effects of visual stimulation and memory on neurons of the hippocampal formation and the neighboring parahippocampal gyrus and inferior temporal cortex of the primate. *J Neurosci* 11(6):1763–1779.
34. Seltzer B, Pandya DN (1978) Afferent cortical connections and architectonics of the superior temporal sulcus and surrounding cortex in the rhesus monkey. *Brain Res* 149(1):1–24.
35. Baylis GC, Rolls ET, Leonard CM (1987) Functional subdivisions of the temporal lobe neocortex. *J Neurosci* 7(2):330–342.
36. Vallabha GK, McClelland JL (2007) Success and failure of new speech category learning in adulthood: Consequences of learned Hebbian attractors in topographic maps. *Cogn Affect Behav Neurosci* 7(1):53–73.
37. Hensch TK (2004) Critical period regulation. *Annu Rev Neurosci* 27:549–579.
38. Mioche L, Singer W (1989) Chronic recordings from single sites of kitten striate cortex during experience-dependent modifications of receptive-field properties. *J Neurophysiol* 62(1):185–197.
39. Chestek CA, et al. (2007) Single-neuron stability during repeated reaching in macaque premotor cortex. *J Neurosci* 27(40):10742–10750.
40. Richardson AG, Borghi T, Bizzi E (2012) Activity of the same motor cortex neurons during repeated experience with perturbed movement dynamics. *J Neurophysiol* 107(11):3144–3154.
41. Ganguly K, Carmena JM (2009) Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol* 7(7):e1000153.
42. Fraser GW, Schwartz AB (2012) Recording from the same neurons chronically in motor cortex. *J Neurophysiol* 107(7):1970–1978.
43. Freedman DJ, Riesenhuber M, Poggio T, Miller EK (2001) Categorical representation of visual stimuli in the primate prefrontal cortex. *Science* 291(5502):312–316.
44. Wallis JD, Anderson KC, Miller EK (2001) Single neurons in prefrontal cortex encode abstract rules. *Nature* 411(6840):953–956.
45. Hussar CR, Pasternak T (2009) Flexibility of sensory representations in prefrontal cortex depends on cell type. *Neuron* 64(5):730–743.
46. Greenberg PA, Wilson FAW (2004) Functional stability of dorsolateral prefrontal neurons. *J Neurophysiol* 92(2):1042–1055.
47. Thompson LT, Best PJ (1990) Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. *Brain Res* 509(2): 299–308.
48. Mankin EA, et al. (2012) Neuronal code for extended time in the hippocampus. *Proc Natl Acad Sci USA* 109(47):19462–19467.
49. Leite FP, et al. (2002) Repeated fMRI using iron oxide contrast agent in awake, behaving macaques at 3 Tesla. *Neuroimage* 16(2):283–294.
50. Xiang Q-S, Ye FQ (2007) Correction for geometric distortion and N/2 ghosting in EPI by phase labeling for additional coordinate encoding (PLACE). *Magn Reson Med* 57(4):731–741.
51. Cox RW (1996) AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Comput Biomed Res* 29(3):162–173.
52. Brainard DH (1997) The psychophysics toolbox. *Spat Vis* 10(4):433–436.

Supporting Information

McMahon et al. 10.1073/pnas.1318331111

SI Methods

Spikes from single units were sorted using the OfflineSorter software package (Plexon Inc.) by projecting waveforms into principal component space and identifying isolated clusters. Spikes were detected either online (in recordings using the Plexon acquisition system) or offline (in recordings using the Tucker-Davis acquisition system) by setting a trigger threshold to capture a snippet of the spike-band–filtered voltage trace. Waveforms detected on the same channel on different days were tentatively classified as belonging to the same neuron if the waveform and spiking statistics matched between consecutive days, following the approach described in detail previously (1). Visual responses were displayed by aligning neuronal activity on stimulus onset, averaging across trials, and convolving the mean firing rate histogram with a 10-ms Gaussian kernel. Face selectivity was assessed by computing a two-way ANOVA with category (face or nonface stimuli) and neuron as factors. Activity at the recording site was considered face selective if there was, first, a significant main effect of category and, second, a significant difference between face and nonface stimuli as determined by a Tukey–Kramer post hoc test. Consistency of stimulus selectivity was assessed by splitting datasets into two halves (odd vs. even runs) and computing correlation coefficients between the conditions of interest (within-day vs. across days, same neuron vs. different neurons, etc.) as described in the main text.

1. Bondar IV, Leopold DA, Richmond BJ, Victor JD, Logothetis NK (2009) Long-term stability of visual pattern selective responses of monkey temporal lobe neurons. *PLoS ONE* 4(12):e8222.

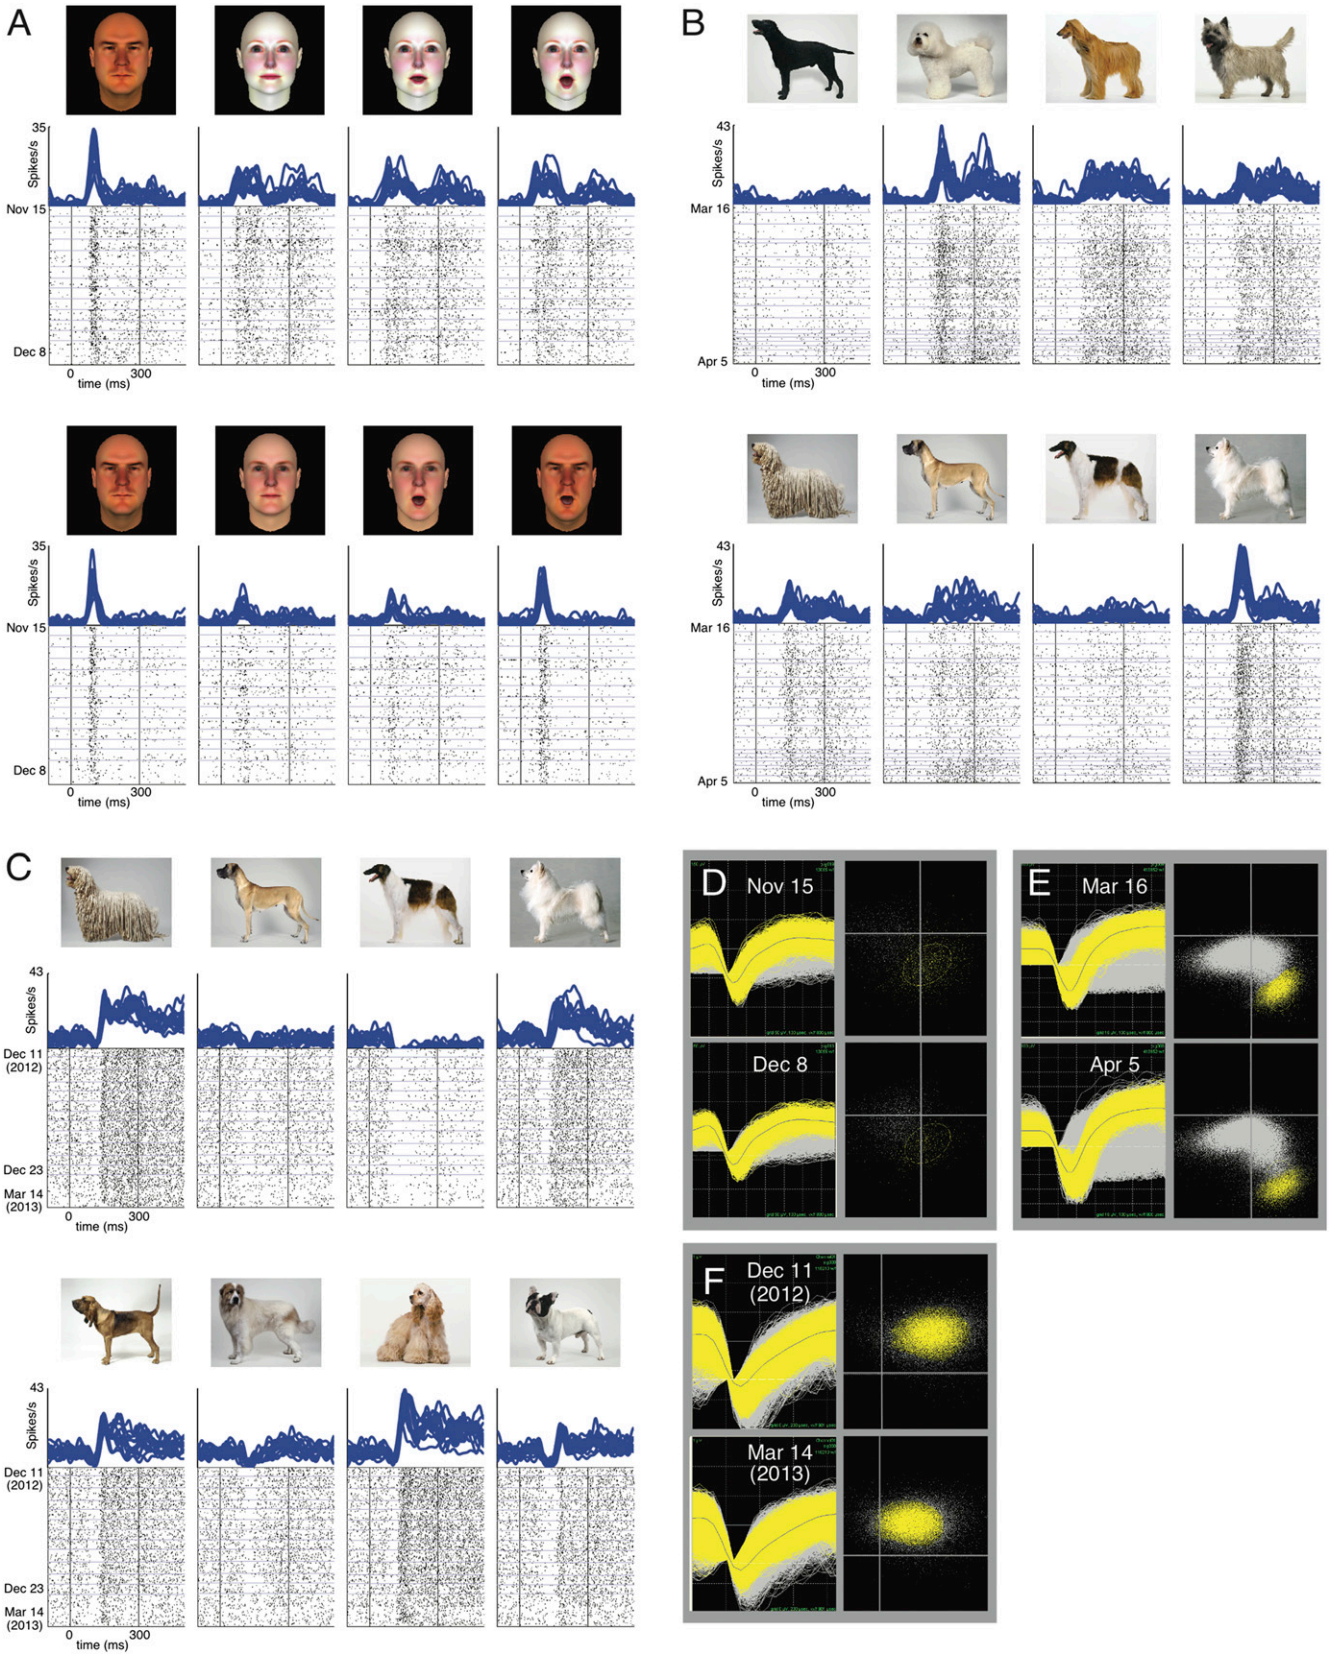

Fig. S1. Visual responses recorded from three longitudinally recorded neurons from monkey 1 (*A* and *B*) and monkey 2 (*C*). Neruons were followed for 23, 21, and 93 days, as indicated by the dates next to the raster plots. The corresponding spike waveforms and clusters in principal components space are shown in *D*–*F*.

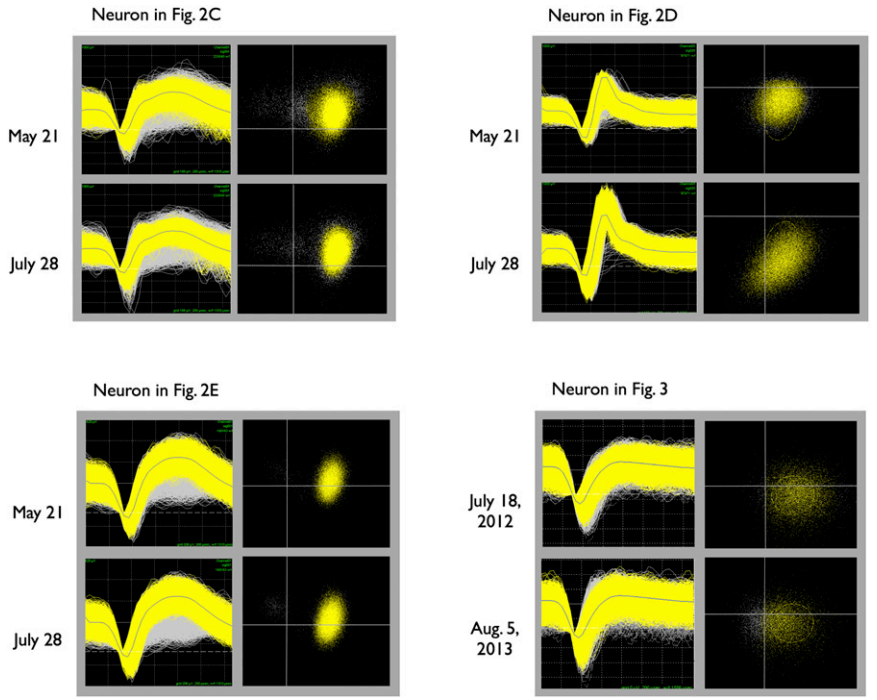

Fig. S2. Stability of spike waveforms over the duration of longitudinal recordings. Each panel shows the voltage traces (*Left*) of isolated spikes and corresponding clusters in principal component space (*Right*; PC1 and PC2 on *x* and *y* axes, respectively) for the four example neurons shown in the main text. The first day (*Upper*) and last day (*Lower*) of recording are shown in each panel.

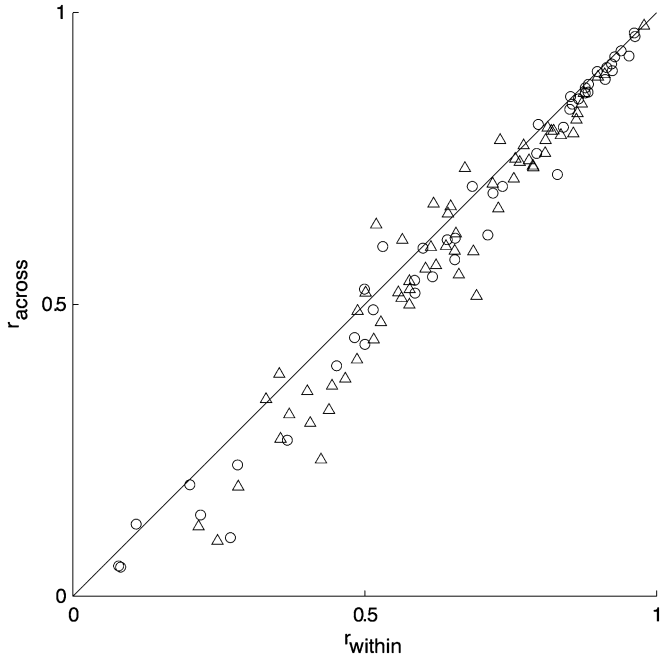

Fig. S3. Longitudinal consistency of stimulus selectivity patterns for population of recorded neurons (triangles, monkey 1; circles, monkey 2). The method of analysis is identical to Fig. 4*C*, except that only the first and last days of recording are considered, rather than recordings on all sequential days.