Microcircuitry of Performance Monitoring

Amirsaman Sajad, David C. Godlove & Jeffrey D. Schall

SUPPLEMENTARY INFORMATION

Table S1

Within trial

Monkey	Trial type	RT (ms)	F	df	p	Reward	Count	P(Noncancel) Mean ± S.D.	RT (ms) Mean ± S.D.	F	df	p	SSRT Mean ± S.D.
Eu	No stop	313	424	8467	<0.001	High	2851		290 ± 115	6.21	5469	<0.001	118 ± 9
						Low	2620		339 ± 118				123 ± 6
	Noncanceled	259				High	1728	54 ± 0.7	244 ± 102	10.63	2996	<0.001	
						Low	1270	43 ± 1.8	281 ± 111				
X	No stop	263	439	17451	<0.001	High	5691		241 ± 102	16.89	11088	<0.001	106 ± 7
						Low	5399		286 ± 119				108 ± 7
	Noncanceled	229				High	3473	61 ± 0.7	213 ± 74	8.76	6361	<0.001	
						Low	2890	52 ± 0.6	249 ± 84				

Across trials

Monkey	Preceding trial	Mean ± S.D.	F	df	p
Eu	No stop	278 ± 114			
	Canceled	344 ± 135	230	3119	<0.001
	Noncanceled	320 ± 87	90	1979	<0.001
X	No stop	244 ± 109			
	Canceled	320 ± 132	430	4509	<0.001
	Noncanceled	254 ± 93	12	3850	<0.001

Table S1. Sensitivity of stop-signal task performance on reward magnitude and trial history for two monkeys. a) The number of trials recorded for the two saccade conditions, i.e., no-stop (correct, rewarded) and non-canceled (error, unrewarded), from each monkey and the mean ± SD of the reaction time (RT) across all 29 sessions (Eu: 12, X: 17) b) Mean ± SD of reaction time on no-stop trials that follow no-stop, canceled, and non-canceled trials for the two monkeys. Test statistics show results for slowing in RT following canceled and non-canceled saccades relative to the RT on no-stop trials.

a

All sites	Error	Loss		Gain		None	Total units (575)
		Suppressed	Facilitated	Suppressed	Facilitated		
Eu - P1	39	21	19	4	18	38	104
X - P2	3	57	44	8	22	46	133
X - P3	0	11	11	2	7	33	56
Eu - np	12	23	28	18	17	68	140
X - np	7	28	37	3	10	87	142
Chi-square test statistics χ^2 (9, N = 575)	86.86 p<0.001	32.8 p<0.001	9.97 p = 0.041	16.7 p = 0.002	7.8 p = 0.1		

b

Depth (channel units)	Error	Loss		Gain		none	Total units (293)
		Suppressed	Facilitated	Suppressed	Facilitated		
1 - 4	1	16	15	0	1	13	34
5 - 8	21	39	28	5	15	39	114
9 - 12	14	11	13	7	16	45	88
13 - 16	6	4	9	2	7	15	39
16 - 19	0	4	6	0	3	8	18
Chi-square test statistics χ^2 (9, N = 293)	6.33 p = 0.18	25.73 p <0.001	12.46 p = 0.014	0.2828 p = 0.87	5.28 p = 0.26		

Supplementary Table 2. Breakdown of signal type count across sites and layers. a) The count of various modulations reported in this study (see Supplementary Figure 2a, and Figure 4a) across the 575 recorded cells sampled from five sites in monkey Eu and monkey X. Three of the five penetrations were perpendicular (monkey Eu: P1, monkey X: P2 and P3) and two were not (Eu – np and Eu – X). Final row shows the test statistics of homogeneity chi-square test. For each type of activity, 5x2 contingency matrix was constructed based on the counts of units for each site (5 rows) with or without given response type. p < 0.05 indicates significant heterogeneity, suggesting that the distribution is significantly different from that predicted by chance. Error-related activity was largely recorded from one of the three perpendicular penetrations (site P1 in monkey Eu). The probability of encountering reinforcement-related neurons was notably higher in site P2 than other recording sites. Similarity in spatiotemporal profile of reinforcement-related neurons was highest between the two perpendicular sites P1 (in monkey Eu), and P2 (in monkey X) (Supplementary Figure 2d). While most reinforcement-related neurons (n=218) showed modulation in a single interval, some (n=47) showed significant modulation at two intervals: one starting after the tone but ending before reward delivery, and one starting after reward delivery. Thus, a total of 312 reinforcement-related modulations were detected. Most neurons (n=41) with two periods of modulation showed the same response specificity in both periods. The spiking activity during some modulation periods was sensitive to both positive and negative outcomes, but in opposite direction. b) similar conventions as (a), showing the number of significant modulations detected at five different depth intervals for units recorded from perpendicular penetrations (n = 293), and their respective test statistics for chi-square test of homogeneity similar to that performed in (b). Gain an loss neurons showed significantly different depth distributions (χ^2(9, N = 201) = 23.17, p < 0.001).

a
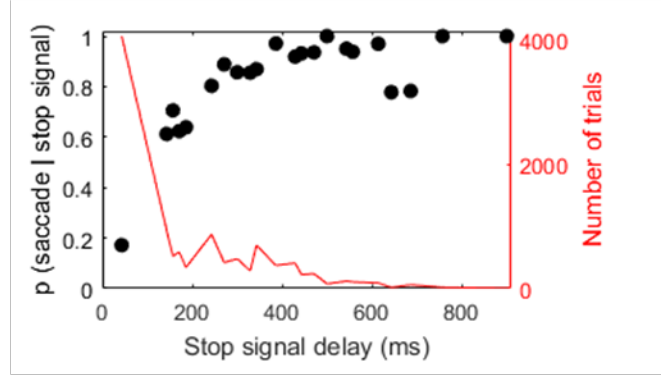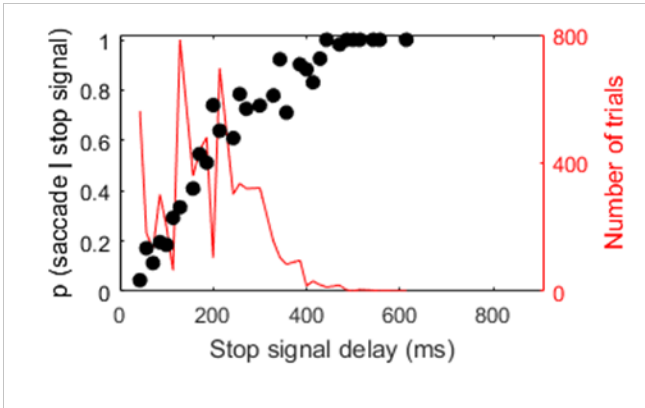
b
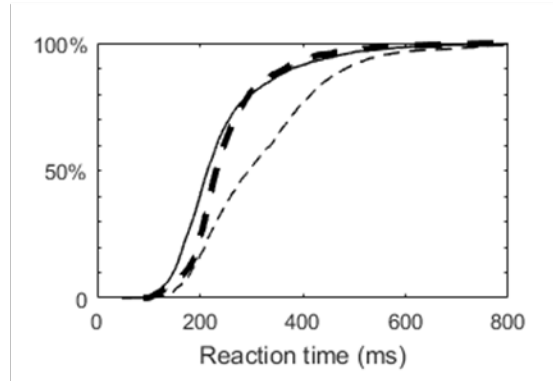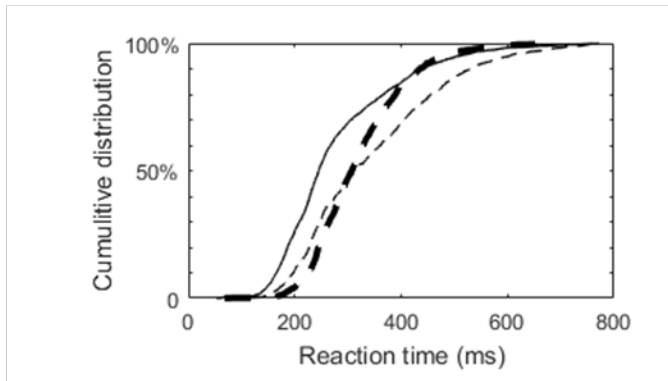
Supplementary Figure 1. Countermanding performance in the two monkeys a) Probability of saccade given stop signal (black dots, left axis) and number of trials (red trace, right axis) for all tested stop signal delays across the 29 sessions in monkey Eu (left) and monkey X (right). Both monkeys performed the task in a manner that could be fit with the Logan race model. **b)** influence of trial history on RT. Cumulative distributions of RT for correct no-stop trials following no-stop trials (thin solid), correct trials (thin dashed), and error non-canceled trials (thick dashed). Monkey Eu (left) and monkey X (right) show RT slowing following correct canceled trials and error non-canceled trials. Quantitative details in Table S1.

a

b Modulation patterns,

c

d
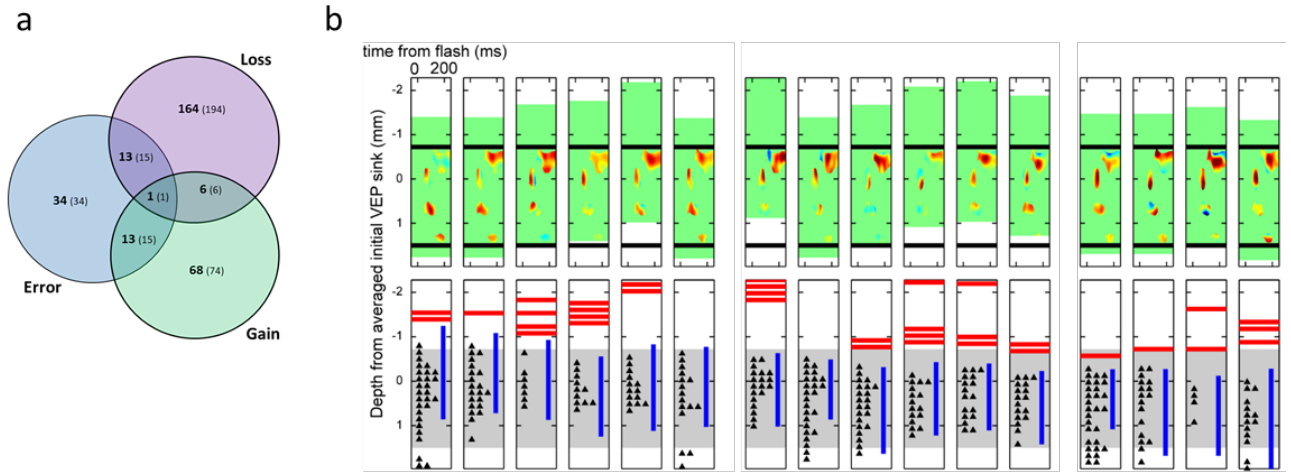
**Supplementary Figure 2. depth alignment, and sample consistency. a)** Across all recording sessions (perpendicular and non-perpendicular penetrations) 575 units were recorded. In 299 units a total of 373 significant modulation intervals for error or reinforcement were detected. Values in the Venn diagram show the number of units (in bold) and the number of detected modulation intervals (in brackets) within each segment of the Venn diagram. Overlapping segments distinguish the 69 units multiplexed across two or more signal types. The distribution of different response types between the different recording sessions is shown in Table 2a. Of the 575 neurons recorded, 293 were sampled from perpendicular penetrations. **b-c)** show alignment of electrode depths across sessions with perpendicular penetration. **b)** Alignment based on CSD sinks following passive visual stimulation. Top: heat-map represents the magnitude of the current sinks across 16 perpendicular penetrations. Sessions were aligned based on CSD (see supplementary methods for details). Bottom: schematic representing the laminar position of isolated units (black triangles), the heart-beat signal (red bars), and the normalized gamma power (blue vertical line) obtained from this alignment. The gray shade represents the estimated gray matter (modified from Godlove et al., 2014). **c)** Alignment based on the LFP gamma power. Similar conventions as **(b)** (modified from Ninomiya et al., 2015). The alignment scheme in **(a)** was used in our study to match those reported previously (Godlove et al., 2014). Units that fall outside of the designated range of gray matter estimate were also reported, allowing for measurement error. **d)** Spiking activity for units stacked in cortical depth recorded in different sessions. The color map represents the difference spike density function between error (unrewarded) and correct (rewarded) trials, with purple representing higher activity on error (unrewarded) trials. Left panels show the difference function for post-saccadic period of depth-aligned neurons sampled in two different sessions recording from the same site (site P1 in monkey Eu). Right panels show the difference function centered around juice delivery time for depth-aligned neurons sampled from in two sessions recording from different monkeys (site P1 in monkey Eu and site P2 in monkey X). Thus, similarities in spatiotemporal pattern of spiking activity were observed across sessions recording from the same and different sites.

4
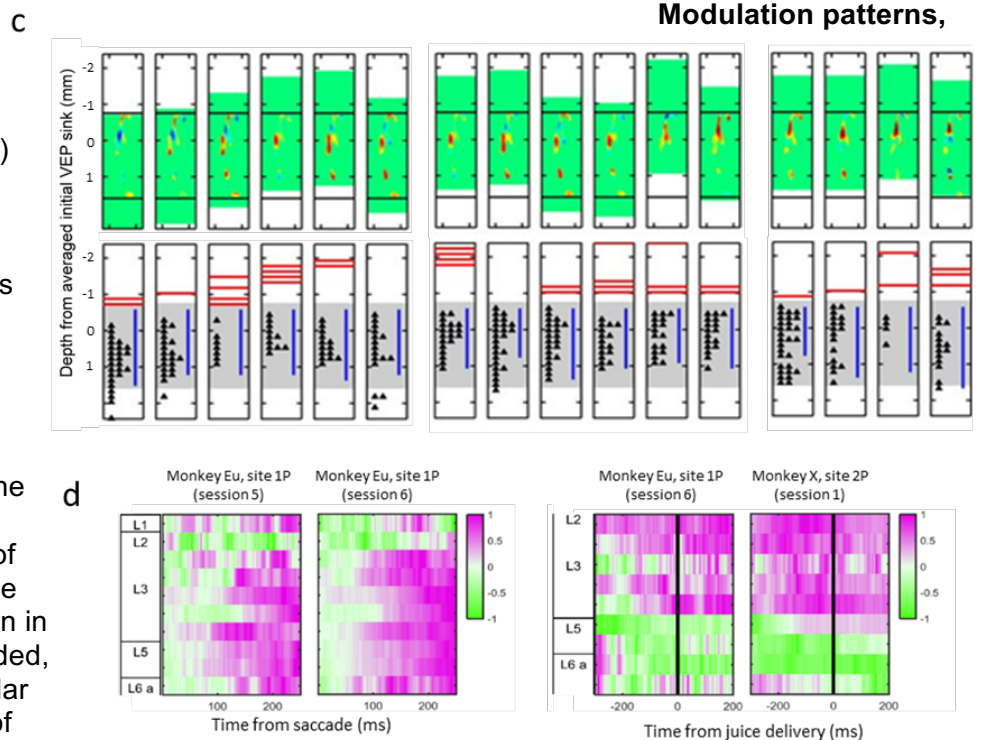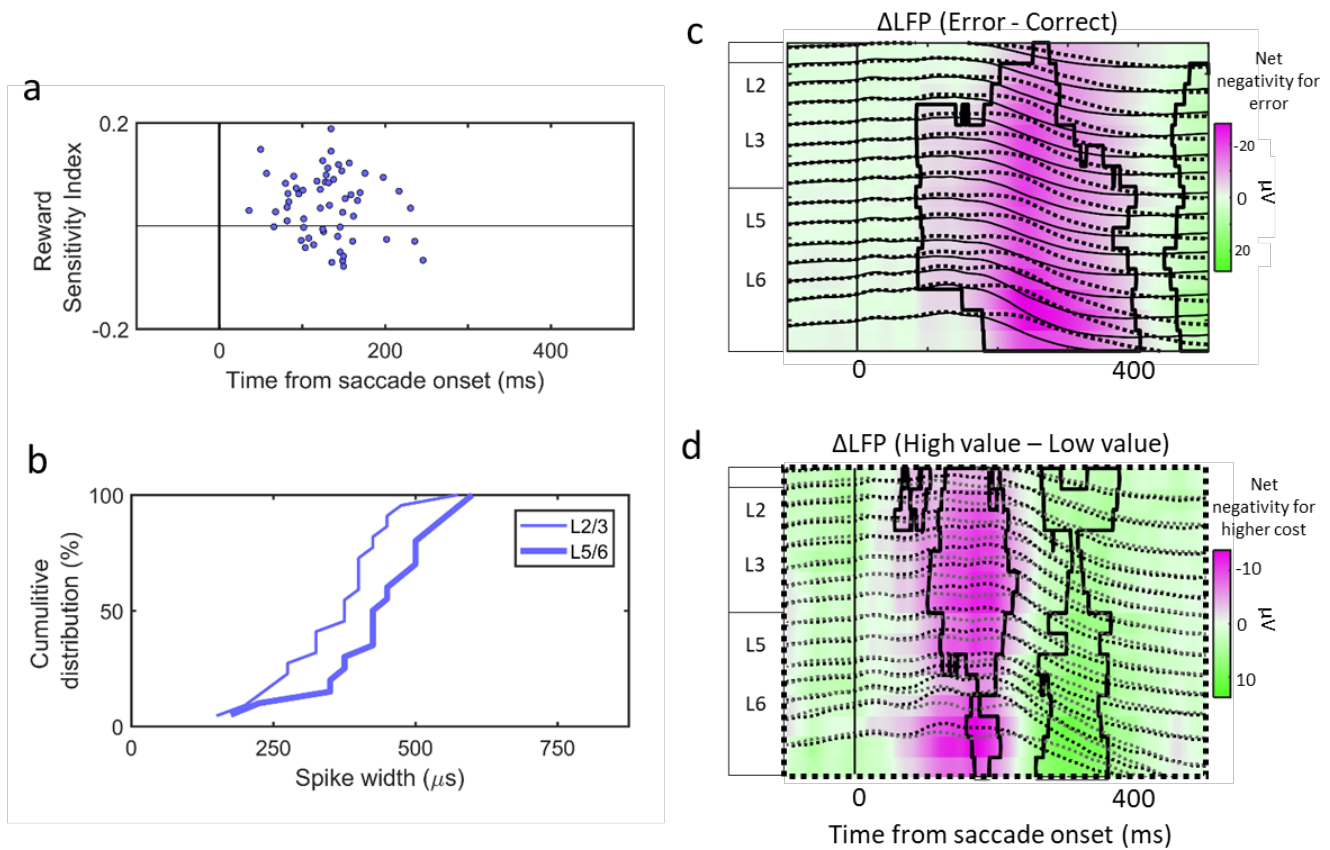
Supplementary Figure 3. Extended analysis of error-related activity. a) Reward sensitivity index of spiking activity as a function of latency of error-related modulation. Positive values indicate higher discharge rate on high-reward than low-reward error trials, i.e., higher activity upon larger error cost. **b)** Cumulative distribution spike widths (trough-to-peak duration) for units in L2/3 (thin) and L5/6 (thick) layers. **c)** Grand average polarization of local field potentials stacked in recorded cortical depth for error (dotted) and correct (solid) trials and their difference (ΔLFP) shown in color map. Upward field polarizations correspond to negative polarizations, following conventions in ERP studies. The purple color represents net negativity in field polarization for error trials, homologous to the ERN. **d)** Grand average polarization of local field potentials for error trials with high error cost (black, dotted) and low error cost (gray, dotted) trials and the difference (ΔLFP). Purple represents net negativity in field polarization for higher cost. The black contours represent patches with significant grand average ΔLFP.

5
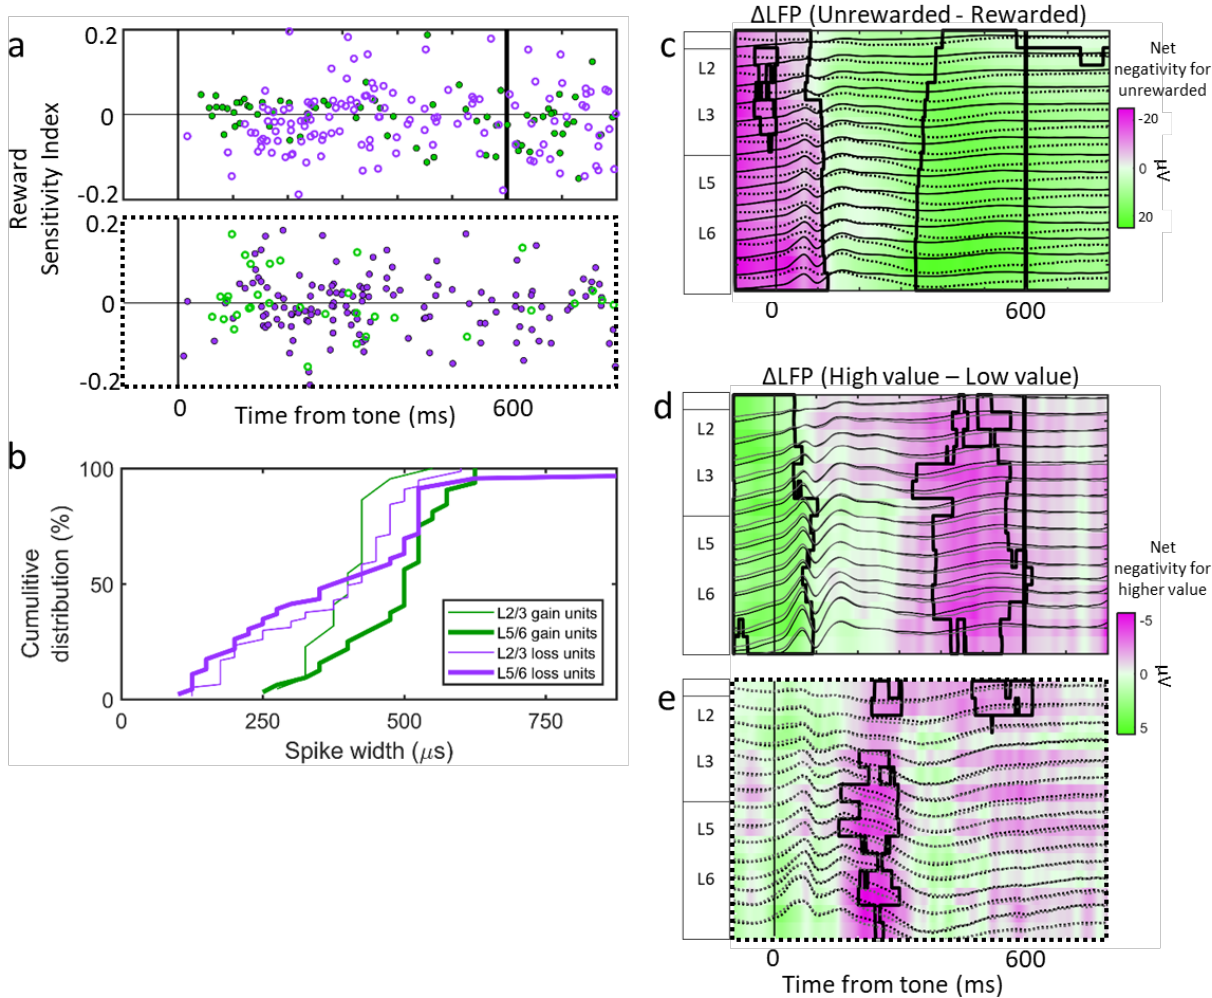
Supplementary Figure 4. Extended analysis of reinforcement-related activity. a) Reward sensitivity index (RSI) of gain (green) and loss (purple) units on rewarded (upper panel) and unrewarded (lower panel) trials as a function of modulation onset time relative to feedback tone. Positive RSI indicates higher discharge rate on high-reward than low-reward trials. **b)** Cumulative distribution spike widths (trough-to-peak duration) for units in L2/3 (thin) and L5/6 (thick) layers for gain (green lines) and loss (purple lines) units. **c)** Grand average polarization of local field potentials stacked in recorded cortical depth for unrewarded (dotted) and rewarded (solid) trials and their difference (ΔLFP) represented by the color map. Upward field polarization corresponds to negative polarizations, following conventions in ERP studies. The purple color represents net negativity in field polarization for unrewarded trials. **d-e)** Grand average polarization of local field potentials for rewarded trials with high (black) and low (gray) gain value **(d)** and for unrewarded trials with high (black) and low (gray) loss value **(e)** and their respective difference LFP (ΔLFP) represented by the color map. Purple represents net negativity in field polarization for higher cost. The black contours represent patches with significantly higher than chance grand average ΔLFP.

Current-Source Density

Current source density (CSD) derived from the local field potential (LFP) reveals the laminar sequence of neural activation across cortical layers (Freeman and Nicholson, 1975; Mitzdorf and Singer, 1979; Buzsaki et al., 1986). Using linear electrode arrays to compare LFP at adjacent sites, one can remove far-field potentials and referencing artifacts and observe local current flow (Kajikawa and Schroeder, 2012). This approach reveals the temporal structure of activity across cortical layers as ensembles of dendrites depolarize in sequence (Di et al., 1990; Schroeder et al., 1998; Lakatos et al., 2007; Lipton et al., 2010). CSD thus provides a functional readout of cortical microcircuitry, encompassing a wider field of view than single-unit and optogenetics approaches. Combined with measures of spiking activity, CSD can also reveal the laminar origin of single- and multi-unit activity. However, the very properties that make agranular cortex useful for testing the generality of the CCM may render CSD recordings impractical. The CSD approach has been effective in sensory areas precisely because of their anatomical structure; dendrites from neural ensembles arborize together in well-defined layers and depolarize in unison, allowing the summation of current flow to be observed at the mesoscopic scale (Freeman and Nicholson, 1975; Riera et al., 2012).

SUPPLEMENTARY DETAILED METHODS
Monkey care and surgical procedures

Data were collected from one male (monkey Eu, *M. radiata* 8.8 kg) and one female (monkey X *M. mulatta* 6 kg) macaque. Animal care exceeded policies of the USDA and Public Health Service Policy on Humane Care and Use of Laboratory Animals. All procedures were approved by the Vanderbilt Institutional Animal Care and Use Committee. MRIs were acquired to aid in placement of recording chambers (Godlove et al., 2014), with a Philips Intera Achieva 3 tesla scanner using SENSE Flex-S surface coils placed above and below the head. T1-weighted gradient-echo structural images were obtained with a 3D turbo field echo anatomical sequence (TR = 8.729 ms; 130 slices, 0.70 mm thickness). Cilux recording chambers (Crist Instruments, Hagerstown, MD) were implanted normal to the cortex (17° monkey Eu, 9° monkey X relative to stereotaxic vertical) centered on midline 30 mm (monkey Eu) and 28 mm (monkey X) anterior to the interaural line. Surgical details have been described in detail (Godlove et al., 2011, 2014).

Cortical mapping and electrode placement

The SEF is located in the dorsal medial convexity in macaques, making it readily accessible for laminar electrode array recordings perpendicular to the cortical layers. The data reported here are from a total of five recording sites, two in Monkey Eu and three in monkey X. Three out of five were perpendicularly penetrated into the cortex. In monkey Eu, the perpendicular penetrations sampled data at site P1 located 31 mm anterior to the interaural line and 5 mm lateral to the midline. In monkey X the perpendicular recordings were obtained from site P2 and P3 located 5 mm lateral and 29 and 30 mm anterior, respectively. During mapping of the bank of the medial wall of cortex, we noted that both monkeys had chambers placed ~1 mm to the right with respect to midline of the brain. This was confirmed in the co-registered CT/MRI data. Thus, our stereotaxic estimates of 5 mm lateral actually placed our electrodes ~4 mm lateral with respect to the cortical (as opposed to the skull-based stereotaxic) midline.

In addition to the data reported here, during every recording session we collected data while monkeys were presented with passive visual stimulation (Godlove et al., 2014). A detailed method for chamber placement, identification of the SEF, and electrode placement has been described here (Godlove et al., 2014). In short, chambers implanted over medial frontal cortex were mapped using tungsten microelectrodes (2-4 MΩ, FHC, Bowdoin, ME) to apply 200 ms trains of biphasic microstimulation (333 Hz, 200 µs pulse width). SEF was identified as the area from which saccades could be elicited using < 50 µA of current. (Schlag and Schlag-Rey, 1987; Schall, 1991; Martinez-Trujillo et al., 2004).

We found the lateral positions granting access to SEF perpendicular to the cortical layers by consulting MRI scans. These positions were further refined through mapping the three-dimensional orientation of gray matter within the chamber by monitoring spontaneous neural activity as a function of depth, using a Grass Technologies (Warwick, RI) audio monitor.

We discriminated the gray-to-white matter transition (GWT) by the sudden paucity of units and the overall decrease in audible hash. When entering cortex obliquely, the GWT was encountered >2 mm after contacting the pial surface. We found the position allowing a penetration perpendicular to the cortical surface by advancing electrodes at gradually more lateral positions until we discerned the GWT 2 mm after contacting

the pial surface. To confirm that these coordinates placed electrodes perpendicular to gray matter, we conducted CT scans with guide tubes in place and co-registered these data with structural MRIs using a point-based method implemented in OsiriX (Geneva Switzerland). Images were reconstructed at 512 x 512 x 512 with a voxel size of 0.252 x 0.252 x 0.122 mm^3.

We used 4 points on the skull that could be easily seen in both CT and MR images to carry out the co-registration; 1) the crest of the bone surface on the brow ridge midway between the supraorbital processes, 2) the point where the interior of the skull protrudes between the base of the occipital lobe and the cerebellum, and 3-4) the most lateral positions of the interior aspects of the left and right zygomatic arches. In both the MRI and the CT data, points 1 and 2 could be easily identified in a midline sagittal section. Points 3 and 4 could be identified in both imaging modalities by gradually advancing more lateral through sagittal slices and marking the location in the first slice where the anterior and posterior aspect of the zygomatic arch merged into one. These points were advantageous for several reasons. First, because each of these points represents an area of bone surrounded by soft tissue (as opposed to air filled sinuses) they were readily apparent in both imaging modalities. Second, these points are widely separated encompassing the majority of the skull in all 3 dimensions. Because these points bound the outer limits of the skull, and because our guide tube was positioned in between and close to the middle of all of these points, deviations in their placement resulted in comparatively small deviations in the guide tube position relative to cortex. By inspecting the data in all 3 dimensions we found that the visual landmarks on the skull yielded excellent co-registrations (Godlove et al., 2014).

Estimation of electrode track angles

We segmented the pial surface and the GWT in coronal and sagittal slices directly beneath the guide tube for each monkey without reference to the co-registered CT images, and transferred these boundaries to the co-registered data. We then implemented a custom algorithm in MATLAB to estimate angles perpendicular to gray matter. For every pixel representing the pial surface in the 2D image, the algorithm found and recorded the closest pixel in Euclidean space representing the GWT. This resulted in a network made up of triangular webs since a single GWT pixel was often found to be closest to several pial surface pixels. The algorithm then worked in reverse, matching every GWT pixel to its closest pial surface counterpart. Finally, the algorithm recorded the average angle of all connections between the pial surface and the GWT in a sliding window. Smoothing across 25 angles provided a balance between angle accuracy and spatial resolution.

Data collection protocol

During recordings, monkeys sat in enclosed primate chairs with heads restrained 45 cm from CRT monitor (Dell, P1130 background luminance of 0.10 cd/m2) running at 70 Hz subtending 46° x 36° of visual angle. We carried out an identical daily recording protocol across monkeys and sessions. After advancing the electrode array to the desired depth, we waited 3-4 hours until recordings stabilized across contacts. This waiting period resulted in consistently stable recordings; single units could usually be held indefinitely. After achieving recording stability, we recorded 1 hour of resting-state activity in near-total darkness with the CRT monitor unplugged. These data will be the subject of a future report. This was followed by a task involving visual flash and period of total darkness in alternating blocks (Godlove et al., 2014). Finally, the monkeys performed ~2000-3000 trials of a saccade stop-signal task, reported here.

Data acquisition

Intracranial data were recorded using a 24-channel Plexon U-probe (Dallas, TX) with 150 μm inter-electrode spacing. The U-probes had 100 mm probe length with 30 mm reinforced tubing, 210 μm probe diameter, 30° tip angle, 500 μm to first contact. Contacts were referenced to the probe shaft, and grounded to the headpost. We used custom built guide tubes consisting of 26 gauge polyether ether ketone (PEEK) tubing (Plastics One, Roanoke, VA) cut to length and glued into 19 gauge stainless steel hypodermic tubing (Small Parts Inc., Logansport, IN) that had been cut to length, deburred, and polished to support the U-probes as they penetrated dura and entered cortex. The stainless steel guide tube provided mechanical support, while the PEEK tubing electrically insulated the shaft of the U-probe, and provided an inert, low-friction interface that aided in loading and penetration. We used microdrive adapters that were fit to our recording chambers with < 400 μm of tolerance and locked in place at a single radial orientation (Crist Instruments, Hagerstown, MD). After setting up hydraulic microdrives (FHC, Bowdoin, ME) on these adapters, pivot points were locked in place

by means of a custom mechanical clamp and neither guide tubes nor U-probes were removed from the microdrives once recording commenced within a single monkey. These methods ensured that we were able to sample neural activity from precisely the same location relative to the chamber on repeated sessions.

All data were streamed to a single data acquisition system (MAP, Plexon, Dallas, TX). Time stamps of trial events were recorded at 500 Hz. Eye position data were streamed to the Plexon computer at 1 kHz using an EyeLink 1000 infrared eye-tracking system (SR Research, Kanata, Ontario, Canada). LFP and spiking data were processed with unity-gain high-input impedance head stages (HST/32o25-36P-TR, Plexon). LFP data were bandpass filtered at 0.2-300 Hz and amplified 1000 times with a Plexon preamplifier, and digitized at 1 kHz. Spiking data were bandpass filtered between 100 Hz and 8 kHz and amplified 1000 times with a Plexon preamplifier, filtered in software with a 250 Hz high-pass filter and amplified an additional 32,000 times. Waveforms were digitized from -200 to 1200 μs relative threshold crossings at 40 kHz. Thresholds were typically set at 3.5 standard deviations from the mean. Single units were sorted online using a software window discriminator and refined offline using principal components analysis implemented in Plexon offline sorter.

LFP and CSD analysis
Data analyses were carried out in MATLAB (Mathworks, Natick, MA). LFPs were time locked to task event (saccade onset for error-related, and auditory tone for reinforcement-related field potentials). Field potentials on each electrode contact was baseline corrected to the 200 ms interval preceding stimulus onset.

After constructing event-related LFPs, we estimated CSD by approximating the 2nd spatial derivative at each point in time using the equation: $CSD=-S\,(\Phi(x-h)-2\Phi(x)+\Phi(x+h))/h^2$, where Φ = the observed voltage, h = the inter-electrode spacing (150um in our case) and S = the average conductance of primate gray matter (0.4 S/m) (Logothetis et al., 2007). The CSD reveals local dendritic current flow in gray matter where neural ensembles arborize together and depolarize in unison, allowing the summation of current flow to be observed at the mesoscopic scale (Freeman and Nicholson, 1975; Riera et al., 2012). CSD values were scaled to nA/mm^3 to afford comparison to values obtained in primary visual cortex and passive visually-evoked CSD in the SEF (Maier et al., 2010; Maier et al., 2011; Godlove et al., 2014). CSD values for each session were then aligned in depth to their closest channel units derived based on alignment of passive visual stimulation CSDs (Godlove et al., 2014; Ninomiya et al. 2015). Notably, successive CSD recordings could be offset by increments smaller than 150 μm (inter-electrode spacing) but given that we perform alignment with a resolution equal to the inter-electrode spacing the alignment method used in this study can result in a maximum offset of 75 μm (0.5 channel units).

Grand average LFP and CSD calculation
For each session the baseline-corrected LFP or CSD was z-transformed by their respective standard deviation of all recorded values across all channels in that session that spanned the gray matter (which exhibited a normal distribution). Thus, the same normalization factor was used for all channels which conserved the relationship in magnitude between channels but corrected for the session-by-session differences in signal gain. To obtain the grand average CSD and LFP the normalized CSD and LFP values were averaged and the resultant was multiplied by the mean of the session-specific normalization factors. Qualitatively, virtually identical results were obtained for the grand average LFPs and CSDs even without normalization but this approach ensured that the contribution of no one session to this average value was inflated.

An additional step for the grand average CSD was to approximate CSD continuously across space. This step involved the interpolation of CSD between electrode contacts using nearest neighbors to a density of 10 μm and convolving the results with a Gaussian filter (σ = 100 μm) (Godlove et al., 2014; Pettersen et al., 2006). A pixel-by-pixel Wilcoxon test was performed on grand average LFPs and CSDs to determine pixels with significant different from baseline. To correct for multiple comparisons only continuous patches that contained significant value for at least 50 ms and spanned a minimum of 50 microns were deemed significant.

Automated depth alignment technique
Our recording depths were jittered from session to session, both intentionally by advancing the electrode array to different levels, and unintentionally by unavoidable day-to-day deviations in cortical dimpling caused by viscoelasticity of the neural tissue. This made it necessary to develop methods to align the CSD

recordings across sessions in an unbiased fashion. Microdrive depth measures are not sufficiently reliable because they do not account for variable cortical dimpling. We adopted methods similar to those used previously (Maier et al., 2010), who aligned and averaged consecutive recording sessions relative to the peak of the initial visually evoked sink that is readily apparent in primary visual cortex following presentation of a flashed visual stimulus. Although we consistently observed a sink in consecutive recording sessions around 50ms after presentation of a flashed stimulus, this sink was smaller in magnitude resulting in a lower signal-to-noise ratio relative to findings from visual cortex. This lower signal-to-noise ratio may have biased our results if we had adopted a manual alignment procedure. Thus, we devised an automated depth alignment procedure to minimize differences between recording sessions using all available source and sink information in a given time window.

Our goal was to find the optimal vector of N recording contact depths, $\Delta = (\delta_1, \delta_2, \ldots \delta_N)$, to align CSD values over time across N sessions. This was accomplished by obtaining the least squared error of the CSD across time (t) and depth (d). For each possible vector Δ of depth shifts we computed an error term (ϵ) by summing all pairwise CSD differences across *nd* depths, *nt* time points, and *N* pairwise comparison between recording sessions:

$$\epsilon(\Delta) = \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} \frac{1}{\alpha_{\Delta(i,j)}} \sum_{d}^{nd} \sum_{t}^{nt} (CSD_{i,d+\Delta(i),t} - CSD_{j,d+\Delta(j),t})^2$$

where $\alpha_{\Delta(i,j)}$ is the maximal integer number of overlapping measures of current flow between a given pair of sessions, determined by the shift for each session specified in Δ. This method differs from that used by Self and colleagues (2013) in which all data were matched to a single representative session. Because ϵ sums across all possible pairwise comparisons, the optimal depth solution is not dependent on a choice of representative session. This code is available upon request.

We implemented a genetic algorithm for parameter optimization to minimize ϵ (using the `ga()` function in the MATLAB global optimization toolbox). We fit the period of time 50-100 ms after stimulus onset, since we systematically observed large visually evoked activity during this time across monkeys and sessions. To ensure that adequate data were used to calculate ϵ, we also constrained the depth parameters to guarantee at least 50% overlap between every recording session. Given our microdrive depth records, this was a conservative estimate. We carried out this procedure separately for data from each monkey and aligned the subject data simply by maximizing the cross correlation of the CSDs averaged across sessions within animals from 50 to 100 ms after stimulus onset.

These methods were useful for aligning recording sessions with respect to one another, but the relationship between the CSD and the underlying anatomy must still be inferred. In our data, sinks could be observed across only 2 mm of tissue in the grand-averaged CSD. Sinks were not observed outside of this limited extent. To our knowledge, only one study has quantified the thickness of SEF and its associated layers. Matelli and colleagues (1991) reported that SEF is almost exactly 2 mm (1,992 +/-31 µm) in thickness, and also detailed the precise location and thickness of each layer. Comparing the 2 mm of measured active sinks to the 2 mm of SEF defined histologically (Matelli et al., 1991), we observed excellent correspondence between current sinks in the grand-averaged CSD, and the locations and thickness of histological layers in SEF. Accordingly, we estimated the laminar origin of current sinks in the grand-averaged CSD based on this correspondence. Because each individual recording session contributes to the grand-averaged CSD with a specific alignment, we were able to work backwards from the grand-average CSD and assign laminar origins to the spiking data. Single- and multi-unit data were binned based on their proximity to the center of mass of the nearest visually-evoked current sink in the grand-averaged CSD. Specifically, to delineate functionally defined borders between layers, we split the distance between the minimum of each active sink. Using the minimum of the initial visually-evoked sink in layer 3 as the zero depth measure, this method identified the following depths as laminar boundaries: layer 1 to 2/3 at 0.21 mm, layer 3 to 5 at 0.36 mm, and layer 5 to 6 at 1.02 mm. These boundaries form the basis for identification of spiking activity with specific laminae. It should be noted that this method is only as precise as the automated alignment algorithm detailed above. To the extent that individual recording sessions are aligned with the grand average, it is safe to consider spiking activity originating from the location of a specific sink and therefore from a specific layer. Blurring of these boundaries may occur when the

alignment of individual recording sessions deviates from that of the grand-averaged CSD. Inspection of the alignment of each individual session suggests that this type of blurring was minimal (Godlove et al., 2014). However, the laminar assignments of spiking activity should be regarded as averages rather than strictly delineated bins.

Assigning Depths of Contacts

To estimate the depth of the electrode array relative to gray matter several physiological measures were considered. First, an artifact associated with the cardiac rhythm (hereafter referred to as the pulse artifact) was observed on a superficial channel. This signal indicated where the electrode was in contact with either the dura mater or the epidural saline in the recording chamber, which pulsated visibly with the monkeys' heartbeat. Second, across all sessions, a marked increase of power in the gamma frequency range (40-80 Hz) was observed at several electrode contacts, which diminished gradually at deeper locations. Several recent studies have shown elevated gamma power in superficial and middle layers relative to deep layers (Maier et al., 2010; Xing et al., 2012; Smith and Sommer, 2013), suggesting that this measure provides another useful marker for estimating depth. Finally, using an automated depth alignment procedure which maximized the similarity profile of CSDs evoked by passive visual stimulation (which were consistent across sites and monkeys) (Godlove et al., 2014).
The latter (CSD alignment) was the main basis for alignment used in this study, however the other physiological signals provided converging evidence to evaluate the electrode position with only minor misalignments relative to that estimated by the CSDs (Godlove et al. 2014). (Supplementary Fig. 2)

Due to the placement of the electrode array, deeper channels were sampled less often than superficial channels (N = 16 channels 1-12, N = 14 channel 13, N = 12 channel 14-16). Note that in this study the depths were estimated to their closest channel-units.

Analysis of spiking activity

Spike were convolved with a kernel resembling a postsynaptic potential with an area = 1 (Thompson et al., 1996) to construct spike density functions for individual trials. Included trials for the calculation of spike density function had at least one spike during the trial (from 600ms prior to target onset until 900ms after feedback tone; see behavioral task). All analyses of spiking activity are based on the mean spike density function.

Error-related activity

Error-related activity were identified by a comparison in the discharge rate of neurons between non-canceled trials with presented stop-signal and correct (no-stop) saccade trials. Error trials in which the stop-signal appeared after saccade onset were not included in this analysis. Significant difference was defined as the interval when the difference function exceeded 2 S.D. above the pre-target baseline (300ms period before target onset) provided that this difference lasted for at least 100ms, or 50ms with the contingency that it reached at least 6 S.D. above the baseline. Through a saccade-matching procedure only saccades from the two trial types with similar RT (within 10ms) and direction were used for comparison.

On some sessions, after the error saccade the monkey shifted gaze back to the fixation point. These trials were excluded unless this resulted in too few trials for meaningful interpretation of spiking activity. In this case additional tests were performed to exclude the possibility that these movements influence our results. First, we tested whether the onset time of differential activity shifts with the timing of the second saccade or remains locked to task-related error saccade. For this, error trials were divided into two pools based on the latency of the second saccade (measured from the error saccade) and the slope between the onset of differential activity and median latency in each pool was calculated. If the slope of the line was smaller than 0.5 the putative error activity was considered saccade-related and hence was rejected as error-related. Second, it was confirmed that the removal of the trials with the second saccade still maintained the polarity of the difference function.

To test whether error-related activity depended on the magnitude of the lost reward, the magnitude of error-related activity (area of the difference function during significant deviation from baseline) on high reward trials was compared to that on low reward trials. Once again saccade-matching was performed. Then a reward sensitivity index was calculated (RSI = (Reward$_{High}$ − Reward$_{Low}$)/ (Reward$_{High}$ + Reward$_{Low}$)). RSI of +1 (or -1) occurs when the neuron is only active for higher (or lower) reward amount and is zero when there is no change

in discharge rate on different reward amount. For cases where error saccades were followed by a second saccade, an additional saccade-matching was performed between the second saccades of high-reward and low reward error trials.

Reinforcement-related activity

Reinforcement-related unit activity was determined by examining the differential discharge rate between rewarded (no-stop) and unrewarded (non-cancelled) saccade trials (i.e., *difference function*) in the period intervening the onset of the tone and 800ms after the tone (i.e., 200ms after juice delivery on rewarded trials). Significant difference was defined using similar criteria as described above (see error-related activity). The rewarded No-stop trials were used for calculating this differential function because the temporal structure of task events in these trials more closely matched that of non-canceled trials. Differential activity was only considered reinforcement-related if the difference between rewarded cancelled (C) and unrewarded NC trials in this period also had the same polarity and was significant (spike count comparison, unpaired Wilcoxon test, $p < 0.05$).

To control for activity related to saccades during feedback period we took advantage of the fact that there was no perfect one-to-one correspondence between the number of saccades in the post-tone period and trial outcome. This allowed us to reject putative reinforcement-related modulations if their strength did not correlate with the proportion of rewarded trials. First, we determined the time interval of significant differential activity. Then, with a bootstrapping approach (n=1000) we randomly selected a subset of trials from the total pool of no-stop (rewarded) and non-canceled (unrewarded) trials. The area under the spike density function and the total number of saccades in this interval, and the percentage of rewarded trials was calculated. This was repeated for all 1000 iterations of the bootstrap resulting in 1000 values that described the relationship between spiking activity (z), number of saccades (x), and proportion of rewarded trials (y). If the partial correlation between z and y (i.e., r(Activity difference, % rewarded trials | number of saccades)) was significant we took that activity as significantly modulated for reward amount.

Saccade matching

We observed a dependence of RT on both laterality of the saccade and the expected reward amount. Therefore, for any analysis that involved a comparison in neural activity between two conditions saccade matching was performed. This entailed matching the same-direction saccades in RT with a 10ms temporal resolution. In cases where several matching saccades were possible, the trial with the closest timestamp was selected.