Vectorized instructive signals in cortical dendrites during a brain-computer interface task

Valerio Francioni[1,2], Vincent D Tang[1,2], Norma J. Brown[1,2], Enrique H.S. Toloza[1,2], Mark Harnett[1,2]*

[1]McGovern Institute for Brain Research, MIT, Cambridge, MA, USA

[2]Department of Brain and Cognitive Sciences, MIT, Cambridge, MA, USA

*Corresponding author. E-mail: harnett@mit.edu

Abstract

Backpropagation of error is the most widely used learning algorithm in artificial neural networks, forming the backbone of modern machine learning and artificial intelligence[1,2]. Backpropagation provides a solution to the credit assignment problem by vectorizing an error signal tailored to individual neurons. Recent theoretical models have suggested that neural circuits could implement backpropagation-like learning by semi-independently processing feedforward and feedback information streams in separate dendritic compartments[3–7]. This presents a compelling, but untested, hypothesis for how cortical circuits could solve credit assignment in the brain. We designed a neurofeedback brain-computer interface (BCI) task with an experimenter-defined reward function to evaluate the key requirements for dendrites to implement backpropagation-like learning. We trained mice to modulate the activity of two spatially intermingled populations (4 or 5 neurons each) of layer 5 pyramidal neurons in the retrosplenial cortex to rotate a visual grating towards a target orientation while we recorded GCaMP activity from somas and corresponding distal apical dendrites. We observed that the relative magnitudes of somatic versus dendritic signals could be predicted using the activity of the surrounding network and contained information about task-related variables that could serve as instructive signals, including reward and error.

The signs of these putative teaching signals both depended on the causal role of individual neurons in the task and predicted changes in overall activity over the course of learning. These results provide the first biological evidence of a backpropagation-like solution to the credit assignment problem in the brain.

Main

Dendrites as a biological substrate for credit assignment in the brain

Learning is attributed to changes in the strength of synaptic connections between neurons[8–15]. In complex hierarchical networks, such as those found in the brain, these changes can have difficult-to-predict effects on downstream readouts of network output. When presented with a specific task, the challenge of determining how individual synapses should be altered to increase performance is known as the credit assignment problem[16–20].

One way to solve credit assignment in artificial systems is through the backpropagation-of-error algorithm, which converts a global performance measure into a cell-specific vectorized teaching signal[1,2,16]. Backpropagation-like learning has long been thought to be biologically-implausible, in part due to requirements such as temporal segregation of feedforward and feedback[16–18,21] information streams. Recent theoretical papers have suggested that the biophysical properties of pyramidal neurons - specifically, active dendritic processing - could resolve this issue by segregating feedforward and feedback information streams into distinct functional compartments, rather than distinct temporal phases as is currently done in machine learning[3–7]. This idea is consistent with anatomical and circuit evidence that feedforward inputs are received at perisomatic compartments of pyramidal neurons while feedback inputs are received in distal dendritic compartments[22–30]. However, experimental evidence for dendritic involvement in backpropagation-like signaling is lacking.

Dendritic implementation of backpropagation-like learning should fulfill three experimentally testable requirements. First, dendritic activity should contain information not present in somatic activity alone. While theoretically somas could transmit gradients using qualitatively different spiking patterns (e.g., single spikes and bursts[4,6,31]), the cable properties of dendrites predict some level of independence between somatic and dendritic activity. Second, dendritic activity should encode information about task performance that could serve as instructive

signals, such as reward and error-representations. Third, the activity of a specific dendrite should reflect the contribution of that neuron's activity to task performance (i.e., the reward function).

Specifying a reward function using a brain-computer interface task

Testing biological neural networks for evidence of backpropagation (or any supervised learning process) has proven challenging for a number of reasons. In particular, a teaching signal can only be defined relative to a reward function that maps neural activity to task performance. Yet it is unclear if such functions are explicitly represented in the brain; even if they are, experimenters are blind to their specific formulation[17].

We developed a neurofeedback brain-computer interface (BCI) task to explicitly define the relationship between neural activity and task performance, thereby specifying the reward function to be optimized (Fig. 1a, b and c, and Extended Data Fig. 1 and Extended data Fig. 2). We trained head-fixed mice under a 2-photon microscope to control the activity of two spatially intermingled sets of GCaMP7f-labelled layer 5 pyramidal neurons, selected for having a high frequency of $\Delta F/F_0$ transients with a high signal-to-noise ratio, in the retrosplenial cortex (RSC), designated P+ and P- (Extended data Fig. 3). The difference in mean somatic GCaMP activity of P+ versus P- neurons rotated a visual grating towards a rewarded target angle[32–37] (Fig, 1d, e and f and Extended data Fig. 1). Previous studies have shown that rodents can learn BCI neurofeedback tasks across a range of brain areas[32–37], and we selected RSC due to the optical accessibility of layer 5 and previous demonstration of independent dendritic events in this area[38]. We recorded GCaMP activity at 15 Hz in the proximal trunk dendrite as a proxy for somatic activity; this allowed imaging of many neurons while simultaneously reducing signal contamination due to the more precise spatial footprint and faster signal kinetics of the apical trunk[39–41]. We measured task performance with two metrics: accuracy, which represented the fraction of rewarded trials; and

speed, which represented the numbers of rewards obtained per minute. All mice (n = 6) learned the task by both metrics (Fig. 1g, Extended Data Fig. 4 and Extended Data Fig. 5).

We compared activity levels of P+ and P- populations, as well as the population of surrounding neurons not directly involved in the rotation of the stimulus (termed P_0), across days of task performance. We imaged the same neurons longitudinally throughout all experiments. We found that learning was accompanied by the differential regulation in the activity of P+ and P- neurons over days (Fig. 1h, i). While on average changes in activity in P_0 neurons resembled changes in P+ neurons (Fig. 1i), selecting the subpopulation of P_0 neuron matching activity levels of P+ and P- neurons on day 1 revealed that changes in activity in P_0 neurons fell in between P+ and P- neurons (Extended data Fig. 6). As the most active neurons on day 1 were also those most strongly downregulated (Extended data Fig. 6b), our results are consistent with a model of learning by sparsification, an energy-efficient solution to the task[42]. Although it is not possible, by definition, to rule out that these changes in neuronal activity are due to untracked changes in behavior, the increase in task performance was not correlated with changes in locomotion across days (Extended Data Fig. 7). Moreover, the P+ and P- populations were spatially intermingled, and had the same GCaMP transient frequency on day 1 (Extended Data Fig. 3 and Extended Data Fig. 6a), ruling out the possibility of learning the task by simply engaging a non-specific gain modulation mechanism.

Dendrites contain information not found in the soma alone

To determine whether apical dendritic activity contained information not encoded in somatic activity alone, we used an electrically tunable lens (ETL) to semi-simultaneously (15 Hz per plane) record activity in proximal and distal trunk dendrites across learning (Fig. 2a). We paired proximal and distal dendrites based on the Pearson correlation of their GCaMP signals, thresholded at r = 0.6 as in previous studies[39–41]. Previous work in brain slices demonstrated that dendritic GCaMP

signals are larger when current is injected in the distal trunk and smaller when current is injected at the soma[39] (controlling for the same number of triggered corresponding action potentials). This indicates that differences in somatic versus dendritic magnitude for coincident GCaMP events reflect the spatial bias of the different inputs that target these two compartments. To estimate the magnitude of somatic and dendritic events, we first deconvolved the GCaMP traces of somas and dendrites using CASCADE[43]. Deconvolution allowed us to compensate for the well-described problem of different signal kinetics in different dendritic compartments[44]. Next, we utilized an area-under-the-curve approach to quantify the magnitude of individual transients (additionally validated for all main results using a $\Delta F/F_0$-based approach to estimation of transient's magnitude, see Methods, Extended Data Fig. 8) and defined events as coincident whenever they occurred within 500 ms of each other. Since these coincident events represent the vast majority of GCaMP transients[38–41,44–50], we focused all subsequent analysis on events for which a transient was detected in both compartments.

Empirically, we observed that the relative magnitude of coincident events in somas and dendrites varies dramatically, despite overall event correlation being very high (Fig. 2b; consistent with prior studies[38,39,41,44,45,47]). Since event magnitudes in soma and dendrites are best described by a linear relationship (Extended Data Fig. 9 and Extended Data Fig. 10b) we assessed the relative degree of dendritic amplification versus attenuation with a best-fit line through all events, and then calculated the somato-dendritic (SD) residual associated with individual transients (Fig. 2b, c)[41]. This captured the variance of dendritic responses for a given somatic event magnitude. We then defined positive and negative residuals as dendritically amplified and attenuated events, respectively.

To validate that SD residuals are biologically meaningful, we used activity from all the somas in our field of view in the two seconds preceding individual GCaMP events in a neuron of interest (P+ and P- neurons on days 1 to 14) to predict whether these events were dendritically amplified or attenuated (Fig. 2d). To do so, we used a linear Support Vector Machine (SVM), a

common algorithm to both classify and regress using high-dimensional data. We found that the performance of our binary classifier on individual neurons strongly correlated with the decoder's ability to capture the magnitude of dendritic amplification/attenuation in the classification confidence (Fig. 2e, g, h, Extended Data Fig. 10c, d and Extended Data Fig. 11a, b). This was an emergent property since the decoder was trained for binary classification only and had no information about the magnitude of dendritic amplification/attenuation. Among 466 neurons, approximately 20% showed a significant correlation between classification confidence and the magnitude of SD residual (Fig. 2h, Extended Data Fig. 10c, d and Extended Data Fig. 11a, b). We found that in these neurons, we could accurately decode 61% of the events as being either amplified or attenuated, significantly above the 50% chance level (Fig. 2j, Extended Data Fig. 10e and Extended Data Fig. 11c). Additionally, at the single-cell level we found a statistically-significant positive Pearson correlation between classification confidence and SD residual, demonstrating that the surrounding network of neurons can be used to predict the amplitude of the residual for coincident somato-dendritic transients (Fig. 2k, Extended Data Fig. 10f and Extended Data Fig. 11d). Importantly, our analysis approach completely decorrelates somatic event magnitude from SD residuals (Fig. 2f, I and Extended Data Fig. 10a), indicating that mismatches in somato-dendritic coupling are predicted independently from somatic activity and represent information encoded *de novo* in the dendrites. Additionally, our results demonstrate that P_0 neurons could be decoded at the same level as P+ and P- neurons (Extended Data Fig. 12), and that decoding does not depend on somatic responses to visual stimuli across the three subpopulations (Extended Data Fig. 13).

Finally, we found that dendritically amplified events consistently peaked earlier than dendritically-attenuated events compared to the soma (Fig. 2l, m, Extended Data Fig. 10g and Extended Data Fig. 12e). This finding is congruent with the previously mentioned results in brain slices[39].

Somato-dendritic residuals decode reward and trial outcome at the population level

Next, we evaluated if changes in SD residual during coincident events contained information about task-related variables that could serve as putative teaching signals. We first tested whether changes in SD residual at the population level contained reward-related information. For each imaging session, we decoded rewarded versus unrewarded trials by comparing the 2 s following neural activity reaching target activation on rewarded trials with the analogous 2 s timeout period during unrewarded trials (Fig, 3a, b, c). Using a linear SVM trained on SD residuals (see Methods), we were able to decode at 63% accuracy on average, above both chance and shuffle performance (Fig. 3d, e, Extended Data Fig. 14 and Extended Data Fig. 15a, b).

Next, we tested if inputs onto the apical tuft dendrites represent instructive signals during learning. We used SD residuals to decode successful versus unsuccessful trials in the 2 s periods preceding successful target activation versus timeout, respectively. Once again, we found that our decoder performed significantly above chance at 57% accuracy on average (Fig. 3c, f, g, Extended Data Fig. 14 and Extended Data Fig. 15c, d), demonstrating that individual neurons encode information about the network states that correspond to successful versus unsuccessful outcomes in their SD residuals both before and after reward delivery. Since the trial time we analyzed is pre-outcome, our results indicate that the SD residuals encode instructive signals based on the task-associated reward function.

Somato-dendritic residuals reflect neuron-specific task error signals

We exploited the explicit definition of error and of functionally-opposite classes of neurons in our experimental design to test whether error signals are received at apical dendrites and if so, whether they differ between neurons according to their causal role in the task (Fig. 4a, b). We reasoned that a scalar error signal would manifest as amplified dendritic activity during periods of error reduction for both P+ and P- neurons and as attenuated during times of error increase.

However, a vectorized error signal consistent with backpropagation-like learning would exhibit selective P+ versus P- dendritic activation, since the activity of each group is causally related to error in opposite ways. To disambiguate between these scenarios, we averaged the error in 2 s windows throughout the task and defined each window as an error increase or decrease epoch, as the angle of the visual stimuli presented to the animals represented the instantaneous task-associated error (Fig. 4a). Next, we calculated the SD residuals for P+ and P- neurons for coincident soma-dendrite events in each window during error decrease and error increase epochs. We found that the dendrites of P+ neurons were relatively amplified during error-reduction compared to error-increase epochs (Fig. 4c). Dendrites in P- neurons exhibited the converse relationship: relative dendritic attenuation and amplification occurred during error-reduction and error increase, respectively (Fig. 4d, e, f, Extended Data Fig. 16 and Extended Data Fig.17). This relationship could be observed in 6 out of 6 mice trained in the task (Extended Data Fig.18) and remained intact when we restricted our analysis to neurons whose somatic activity was the same during epochs of error increase and reduction (Extended Data Fig. 19). Additionally, the same inverted relationship between dendritic signals and task-associated errors was found in the dendrites of P_0 neurons which were functionally correlated to P+ and P- neurons (Extended Data Fig. 20). In contrast to instructive signals found in classical implementations of backpropagation, SD residuals represented error derivatives but not errors (Extended Data Fig. 21). These results show that during learning, dendrites receive and independently process task-related error-representing signals which are tailored in a neuron-specific manner according to their functional role in the task, consistent with vectorized teaching signals, a fundamental principle of error backpropagation in ANNs.

Discussion

Here, we provide the first biological evidence of a vectorized solution to the credit assignment problem in the brain via cortical dendrites. Our data is consistent with a model of credit assignment in which learning is instructed by an instantaneous, vectorized instructive signal received onto the distal dendrites of pyramidal neurons[3–7]. This spatial-segregation mechanism allows cortical circuits to overcome the biologically implausible temporal separation of feedforward and feedback streams in the canonical backpropagation algorithm.

The data presented here reveal magnitude differences in coincident somato-dendritic events that can be predicted using activity in the surrounding network of neurons, demonstrating that individual neurons track circuit dynamics through computation onto their dendrites. At the population level, differences in somato-dendritic coupling encode *de novo* information along orthogonal dimensions to the soma. This information could be used by individual neurons as instructive signals, such as reward and task error, providing novel evidence that individual neurons can explicitly access the reward function of a learning task through independent dendritic computation. We further demonstrate that cell-specific changes in SD residuals correlate with the functional role of individual neurons as well as with subsequent changes in activity levels during learning.

While we cannot rule out that dendritic activity reflects moment-to-moment computational signals - as opposed to teaching signals for synaptic weight changes - our results demonstrate the existence of a signed, vectorial dendritic input which is tailored in a condition-specific manner to individual neurons: the hallmark of the backpropagation algorithm in ANNs[1,2,16]. Future work is needed to assess whether these dendritic signals result from glutamatergic inputs from higher-order cortical areas, from neuromodulation, or as a product of recurrent excitatory and inhibitory local computation. Additionally, future experiments should explore whether errors signals are calculated locally at each hierarchical layer, or are transmitted across layers as in the classical formulation of backpropagation[2].

The error signals we observed have appealing connections to the gradient calculations found via the backpropagation algorithm. Establishing a direct linkage between SD residuals and a gradient derived weight change requires explicit knowledge of the exact loss function represented in the brain (i.e., how subcellular computation assigns synapse specific weight adjustments). While the BCI approach described in this paper allows us to specify the reward function of the task, it does not allow us to control the exact specification of the loss function that is used for learning. In particular, the asymmetry of our task (due to the target angle only being approachable from the negative direction) could lead to an internal representation of the loss function that is also asymmetric (e.g. $L = \text{target angle} - \text{control angle}$ as opposed to the more commonly implemented $L = (\text{target angle} - \text{control angle})^2$). Such a loss function could explain why the dendrites in our task received signals that unexpectedly bore signatures of error derivative rather than error itself, a mechanism that could lead to unstable learning in a symmetric version of our task. Future studies will be needed to disentangle the behavioral and computational conditions leading to the selection of one specific cost function over another and to make further comparisons to the backpropagation algorithm, as well as to other theories involving the computation of error signals in somatic and dendritic activity mismatches[51].

Together, our results reconcile early findings and theories of dendritic function, which focused on single dendritic branches as the building blocks for independent computation, with in vivo findings highlighting the co-occurrence of dendritic and somatic events[17,24,49,52]. By demonstrating that apical dendrites locally compute reward and error-related signals, our results present a framework for dendritic computation which does not require fully independent dendrites to perform credit assignment for adaptive behavior and highlight new directions for the development of biologically-inspired ANNs.

Methods

Animals

All experiments were compliant with guidance and regulation from the NIH and the Massachusetts Institute of Technology Committee on Animal Care. Male and female Rbp4-Cre heterozygous mice were maintained on a 12-hour light/dark cycle in a temperature- and humidity-controlled room with ad libitum food access and were used for experiments at 8-15 weeks of age. Animals were water-deprived by decreasing water intake from 3 ml to 1.2 ml over the course of 10-14 days and maintained at 1.2 ml thereafter, for 5-7 days before experiments and throughout training.

Surgery

Mice were initially anaesthetized using 4% isoflurane and subsequently maintained at 1-2% isoflurane through the rest of the surgery. Body temperature was maintained at physiological levels using a closed-loop heating pad. Additional heating was provided for post-surgical recovery. To protect eyes from dryness, eye cream (Bepanthen, Bayer) was applied. Animals were injected with Dexamethasone (4mg/kg), Carprophen (5mg/kg) and Buprenorphine (slow release, 0.5mg/kg) subcutaneously. The scalp was shaved using hair-removal cream and cleaned afterwards using iodine solution and ethanol. Next, the skull was exposed. For in vivo imaging, a 3mm-wide craniotomy was performed. At 3-4 different sites, we injected 100 nl of AVV1-syn-FLEX-jGCaMP7f-WPRE (Addgene, catalog # 104492-AAV1, 2-5x10^{12} vg/ml concentration after a 1:10 dilution from the original concentration) at 400 µm from the surface of the brain in the left hemisphere of the Retrosplenial cortex (2.5 mm caudal of bregma). The dura was left intact. Cranial windows consisted of two stacked 3mm coverslips (inserted within the craniotomy) attached to a larger 5 mm coverslip which was subsequently fixed to the skull using cyanoacrylate glue and dental cement. A custom metal headplate was implanted in order to perform imaging under head-fixed conditions. At the end of the procedure, a single dose of 25mL/kg of Ringer's solution was injected subcutaneously to rehydrate the animal. Recordings started 4-6 weeks post-surgery.

Two photon imaging

A Neurolabware 2-photon microscope equipped with GaAsP photomultiplier tubes was used for data acquisition. Imaging was performed at 980 nm using an ultrafast pulsed laser (Spectra-Physics, Insight DeepSee) coupled to a 4x pulse splitter to reduce photodamage and bleaching. For excitation and photon collection we used a 16x Nikon objective with 0.8 numerical aperture. Bidirectional scanning was performed (512x796 pixels) semi-simultaneously in two separate planes using an electrically-tunable lens at 30.92 Hz (15.46 Hz for each plane). Laser intensity was independently optimized at each imaging plane using an electro-optical modulator. A custom light shield was attached to the headplate in order to avoid light contamination. Animals were habituated to human handling for 5-10 minutes every day and to head-fixation for 15 minutes a day for at least 3 days directly preceding imaging. Small 10% sucrose water rewards were randomly dispensed during habituation. Daily water intake of at least 1.2 mL was maintained throughout the behavioral experiments. The animal's locomotion was recorded using an optical encoder (E6, US Digital, 2500 cycles per revolution) tracking the rotation of a cylindrical treadmill 19 cm wide in radius and acquired using the Scanbox software interfaced to a custom-built Arduino system. To maximize the number of units recorded while simultaneously reducing signal contamination, we imaged the trunk of layer 5 pyramidal neurons at two different planes: proximal to the soma and right below the nexus (tuft bifurcation point).

Brain computer interface task

Similar to previous implementations of brain-computer interface learning paradigms[32,34,36], mice were trained so that they obtained rewards by modulating the activity of 8 or 10 layer 5 pyramidal neurons in the retrosplenial cortex to control the rotation of a grating Gabor patch. The 8 or 10 neurons were equally divided into 2 subpopulations, P+ neurons whose activity rotated the stimulus towards a target angle of 90-degrees (horizontal) and P- neurons whose activity rotated the stimulus away from the target angle, towards a 0-degree (vertical) orientation. Neural activity

was transformed into a visual stimulus angle according to the following method: At the beginning of each session, we measured the baseline responses of P+ and P- neurons to 7 randomly presented oriented gratings (0-, 15-, 30-, 45-, 60-, 75-, 90-degree) for approximately 13 minutes (12000 frames). $\Delta F/F_0$ was calculated for individual P+ and P- neurons and averaged across each population. The mean P- population signal was subtracted from the mean P+ population signal. Next, we randomly resampled 200 trials (435 frames each) from the aforementioned 12000-frame baseline recording and iteratively searched (in 0.005 $\Delta F/F_0$ incremental steps) for the subtracted $\Delta F/F_0$ value producing a 50% success rate. That value was set as the threshold value for target activity during the closed-loop phase of the BCI task. Next, we calculated the mean and standard deviation of the subtracted $\Delta F/F_0$ signal distribution and created a new distribution by mirroring the left side to the right. On day 1, we estimated the z-score corresponding to the $\Delta F/F_0$ threshold value on the mirrored distribution. On the following days, we estimated the subtracted $\Delta F/F_0$ signal distribution and its corresponding left-mirrored distribution in the same way as described above in the same way as described above, and utilized the $\Delta F/F_0$ value corresponding to the z-score used on day 1 as the task target activity during the closed-loop phase of the task. In this way mice could learn the task by either decreasing activity of P- neurons or increasing activity in P+ neurons (or both). The mapping between neuronal activity and visual feedback angle was defined as follows: 0-degree angle corresponded to the minimum value in the subtracted $\Delta F/F_0$ signal distribution while target, or 90-degree angle was reached at subtracted $\Delta F/F_0$ value corresponding to threshold defined as described above. Activity in between was split into 7 equally spaced bins each corresponding to a 15-degree interval between 0 and 90 degrees. The screen refreshed every time a 2-photon frame at the soma was acquired (at 15Hz). At each screen refresh, the angle presented reflected the mapping between angle bins and the subtracted $\Delta F/F_0$ signal averaged over the last 3 frames. In each trial, mice had 28 seconds to reach target activity. If they did, a reward, consisting of 4 μL of 10% sucrose water was delivered 1 second after. Additionally,

after reaching target activity, the stimulus froze to a 90-degree angle for 2 seconds. After that, mice saw a black screen for 1 additional second and a new trial was initiated. All new trials were initiated by a 0.5 s iso-luminant grey stimulus. If a mouse did not reach target activity within the 28 seconds of the trial, a 3 seconds timeout was given to them consisting of a black screen. To avoid the problem of drifting baselines, $\Delta F/F_0$ for each neuron was calculated as $(F_i - F_i0)/F_i0$ where F_i0 was the 10th percentile of fluorescence in the previous 30 seconds.

Online motion-correction

In order to avoid drifts in x and y out of our selected regions of interest, we used a Fast-Fourier transformation approach to live motion-correct our movies. To do so, at the beginning of each recording session we acquired a reference image by averaging 20-40 seconds (300-600 frames) collected onto our field of view. To motion-correct each subsequent frame, we selected 4 smaller central areas to register independently from one another (2-D rigid translation) against the corresponding 4 areas in the reference image[53]. We finally rigid-translate the entire 2-D image by taking the average translation in x and y for these 4 subregions.

Visual stimuli

Visual stimuli were generated using the Psychophysics Toolbox package for MATLAB (MathWorks, MA)[54] and displayed on a monitor 20 cm away from the contralateral eye. Visual stimuli consisted of a rotating Gabor patch at 7 angles spaced 15 degrees apart from 0 to 90.

Offline image analysis and signal extraction

To correct for brain motion after image acquisition, as well as to automatically detect ROIs, we used the Suite2P pipeline[55]. For each field of view (FOV), we removed duplicates by excluding ROIs whose signal correlation was above 0.6 and whose center was within 20 μm of distance. In order to separate trunk signals from potential neuropil contamination, fluorescence signals of our ROIs were processed using FISSA[56] with the following parameters: 4 neuropil subregions and

alpha = 0.1. To estimate $\Delta F/F_0$ after neuropil subtraction, we calculated $\Delta F/F_0$ at time point i as (Fi – F0)/F0. F0 is defined as the 10th percentile of a 120 seconds long sliding window to remove fluorescence drifts over the course of imaging. Next, we performed spike inference using the CASCADE model Global_EXC_15Hz_smoothing200ms[43].

Field of view matching and ROI registration across days

Registration of neurons across days for BCI training was performed manually at the beginning of each session with the help of a custom-designed software. On day 1, a mean intensity reference image of our field of view of interest was acquired. Using a custom-design software, we manually drew 10-20 reference ROIs which included any recognizable brain structure including dendrites, cell somas and sharp-contrast blood vessels. On the following days, after manually finding the same approximate area for the field of view imaged on day 1, a more accurate manual registration was performed by aligning our reference ROIs drawn on day 1 with their corresponding structures on following days. As the relative x and y distance between structures varies along the z-plane, our approach allowed us to consistently match our field of view on day 1 across x, y and z dimensions on any given day. Offline registration of ROIs across days on the other hand, was initially performed using the ROIMatchPub implementation for Suite2P followed by an exhaustive manual curation.

Quantification of event frequency, magnitude, and timing

Events were detected for each ROI using the MATLAB function *findpeaks* on the spike-inferred signal. For analysis of the spike inferred signal, we estimated the integral of individual peaks by multiplying the height and width of individual transients. Event occurrence was defined as the time at which spike probability peaked. For $\Delta F/F_0$ analysis, once we found an event, we utilized a 2-seconds backward sliding window to identify the frame at which the derivative of the $\Delta F/F_0$ signal became consecutively positive for 300 ms. This was considered the transient onset frame while

the peak of the transient was considered the maximum $\Delta F/F_0$ value in the 2 seconds following peak detection. We therefore estimated the integral of the $\Delta F/F_0$ signal by multiplying the height (maximum $\Delta F/F_0$ value – $\Delta F/F_0$ value at transient onset) and the width (frame at maximum $\Delta F/F_0$ value – frame at transient onset) of the $\Delta F/F_0$ signal. The backward and forward detection windows were limited in time by the presence of a precedent or subsequent event detected using the spike-inferred signal. Proximal trunks were paired to their correspondent distal trunk whenever their $\Delta F/F_0$ signal correlation was equal above 0.6. Whenever we found more than 1 distal dendrite correlated with the same proximal trunk, we selected the one with the best signal-to-noise ratio, so to always have a single distal dendrite associated with a proximal trunk. Coincident events were defined as two events occurring (independently detected) within a 500 ms window in the two compartments. To quantify the somato-dendritic magnitude mismatches of coincident events, we first fit a best-fit line against the somatic and dendritic magnitudes of all events. For each event, we calculated the residual from the best-fit line, and defined residuals larger than 0 as dendritically-amplified and residuals smaller than 0 as dendritically-attenuated.

Decoders

To decode whether individual transients would be amplified or attenuated, we trained a support vector machine binary classifier (SVM, linear kernel) using stochastic gradient descent[57] (as implemented by MATLAB *fitclinear*). For each coincident event in the soma and dendrites, we averaged the spike-inferred activity of each neuron in our field of view (excluding the neuron of interest) in the 2 preceding seconds, and we used this average activity to create an n-dimensional population activity vector where n corresponds to the number of isolated units in our field of view. The binary classifier was trained to separate dendritic amplification from dendritic attenuation (see above) using a leave-one-out approach. Accuracy was determined as the fraction of correctly classified events. For imbalanced datasets, we used a Synthetic Minority Oversampling Technique (SMOTE, k neighbors = 5) to train (not test) using a balanced dataset. SMOTE was

applied after separating our train and test dataset. To control for any potential data leakage, our shuffle control went through the exact same procedure as our test dataset, including SMOTE oversampling with the only difference that labels were randomly shuffled before separating the train and test data. We calculated the confidence of a prediction as the Euclidean distance from the hyperplane. Reward-associated and reward-instructive epochs were defined as 2 seconds before and 2 seconds after the reach of target activity, respectively for successful trials, and 2 seconds before and 2 seconds after the end of a trial for unsuccessful trials. To decode successful from unsuccessful trials, we generated a n-dimensional somato-dendritic residual vector by taking the residual for each neuron for which we identified a somato-dendritic pair (see above) in these two seconds epochs. Neurons inactive in the two seconds epochs were assigned a value of 0. The binary classifier was trained in the same manner as described above.

Statistics

All analysis was performed using MATLAB 2020a using custom-written scripts and functions. All error bars in figures represent standard error of the mean (SEM). Statistical tests and independent samples are described in figure legends.

Data and code availability

All analysis, BCI code, and data are available upon request.

References

1. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
2. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).
3. Sacramento, J., Bengio, Y., Costa, R. P. & Senn, W. Dendritic cortical microcircuits approximate the backpropagation algorithm. in *Advances in Neural Information Processing Systems* (2018).
4. Payeur, A., Guerguiev, J., Zenke, F., Richards, B. A. & Naud, R. Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nat. Neurosci.* **24**, 1010–1019 (2021).
5. Guerguiev, J., Lillicrap, T. P. & Richards, B. A. Towards deep learning with segregated dendrites. *eLife* **6**, (2017).
6. Greedy, W., Zhu, H. W., Pemberton, J., Mellor, J. & Costa, R. P. Single-phase deep learning in cortico-cortical networks. Preprint at https://doi.org/10.48550/arXiv.2206.11769 (2022).
7. Körding, K. P. & König, P. Supervised and Unsupervised Learning with Two Sites of Synaptic Integration. *J. Comput. Neurosci.* **11**, 207–215 (2001).
8. Magee, J. C. & Grienberger, C. Synaptic Plasticity Forms and Functions. *Annu. Rev. Neurosci.* (2020) doi:10.1146/annurev-neuro-090919-022842.
9. Bittner, K. C., Milstein, A. D., Grienberger, C., Romani, S. & Magee, J. C. Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science* **357**, 1033–1036 (2017).
10. Caporale, N. & Dan, Y. Spike Timing–Dependent Plasticity: A Hebbian Learning Rule. *Annu. Rev. Neurosci.* (2008) doi:10.1146/annurev.neuro.31.060407.125639.

11. Artola, A., Bröcher, S. & Singer, W. Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature* **347**, 69–72 (1990).

12. Kirkwood, A., Rioult, M. G. & Bear, M. F. Experience-dependent modification of synaptic plasticity in visual cortex. *Nature* **381**, 526–528 (1996).

13. Bliss, T. V. P. & Lømo, T. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* **232**, 331–356 (1973).

14. Bienenstock, E. L., Cooper, L. N. & Munro, P. W. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J. Neurosci.* **2**, 32–48 (1982).

15. Abbott, L. F. & Blum, K. I. Functional significance of long-term potentiation for sequence learning and prediction. *Cereb. Cortex* (1996) doi:10.1093/cercor/6.3.406.

16. Lillicrap, T. P., Santoro, A., Marris, L., Akerman, C. J. & Hinton, G. Backpropagation and the brain. *Nat. Rev. Neurosci.* (2020) doi:10.1038/s41583-020-0277-3.

17. Richards, B. A. & Lillicrap, T. P. Dendritic solutions to the credit assignment problem. *Curr. Opin. Neurobiol.* (2019) doi:10.1016/j.conb.2018.08.003.

18. Richards, B. A. *et al.* A deep learning framework for neuroscience. *Nat. Neurosci.* (2019) doi:10.1038/s41593-019-0520-2.

19. Minsky, M. Steps toward Artificial Intelligence. *Proc. IRE* **49**, 8–30 (1961).

20. Lansdell, B. J., Prakash, P. R. & Kording, K. P. Learning to solve the credit assignment problem. (2019).

21. Crick, F. The recent excitement about neural networks. *Nature* **337**, 129–132 (1989).

22. Makino, H. & Komiyama, T. Learning enhances the relative impact of top-down processing in the visual cortex. *Nat. Neurosci.* **18**, 1116–1122 (2015).

23. Larkum, M. E., Zhu, J. J. & Sakmann, B. A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature* (1999) doi:10.1038/18686.

24. Larkum, M. A cellular mechanism for cortical associations: An organizing principle for the cerebral cortex. *Trends Neurosci.* **36**, 141–151 (2013).

25. Lafourcade, M. *et al.* Differential dendritic integration of long-range inputs in association cortex via subcellular changes in synaptic AMPA-to-NMDA receptor ratio. *Neuron* (2022) doi:10.1016/j.neuron.2022.01.025.

26. Marques, T., Nguyen, J., Fioreze, G. & Petreanu, L. The functional organization of cortical feedback inputs to primary visual cortex. *Nat. Neurosci.* (2018) doi:10.1038/s41593-018-0135-z.

27. Manita, S. *et al.* A Top-Down Cortical Circuit for Accurate Sensory Perception. *Neuron* **86**, 1304–1316 (2015).

28. Rockland, K. S. & Pandya, D. N. Laminar origins and terminations of cortical connections of the occipital lobe in the rhesus monkey. *Brain Res.* **179**, 3–20 (1979).

29. Coogan, T. A. & Burkhalter, A. Conserved patterns of cortico-cortical connections define areal hierarchy in rat visual cortex. *Exp. Brain Res.* **80**, 49–53 (1990).

30. Cauller, L. J., Clancy, B. & Connors, B. W. Backward cortical projections to primary somatosensory cortex in rats extend long horizontal axons in layer I. *J. Comp. Neurol.* **390**, 297–310 (1998).

31. Naud, R. & Sprekeler, H. Sparse bursts optimize information transmission in a multiplexed neural code. *Proc. Natl. Acad. Sci. U. S. A.* (2018) doi:10.1073/pnas.1720995115.

32. Clancy, K. B., Koralek, A. C., Costa, R. M., Feldman, D. E. & Carmena, J. M. Volitional modulation of optically recorded calcium signals during neuroprosthetic learning. *Nat. Neurosci.* (2014) doi:10.1038/nn.3712.

33. Neely, R. M., Koralek, A. C., Athalye, V. R., Costa, R. M. & Carmena, J. M. Volitional Modulation of Primary Visual Cortex Activity Requires the Basal Ganglia. *Neuron* (2018) doi:10.1016/j.neuron.2018.01.051.
34. Clancy, K. B. & Mrsic-Flogel, T. D. The sensory representation of causally controlled objects. *Neuron* **109**, 677-689.e4 (2021).
35. Jeon, B. B., Fuchs, T., Chase, S. M. & Kuhlman, S. J. Existing function in primary visual cortex is not perturbed by new skill acquisition of a non-matched sensory task. *Nat. Commun.* **13**, 3638 (2022).
36. Mitani, A., Dong, M. & Komiyama, T. Brain-Computer Interface with Inhibitory Neurons Reveals Subtype-Specific Strategies. *Curr. Biol.* **28**, 77-83.e4 (2018).
37. Koralek, A. C., Jin, X., Long II, J. D., Costa, R. M. & Carmena, J. M. Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature* **483**, 331–335 (2012).
38. Voigts, J. & Harnett, M. T. Somatic and Dendritic Encoding of Spatial Variables in Retrosplenial Cortex Differs during 2D Navigation. *Neuron* (2019) doi:10.1016/J.NEURON.2019.10.016.
39. Beaulieu-Laroche, L., Toloza, E. H. S., Brown, N. J. & Harnett, M. T. Widespread and Highly Correlated Somato-dendritic Activity in Cortical Layer 5 Neurons. *Neuron* (2019) doi:10.1016/j.neuron.2019.05.014.
40. Peters, A. J., Lee, J., Hedrick, N. G., O'neil, K. & Komiyama, T. Reorganization of corticospinal output during motor learning. *Nat. Neurosci.* **20**, 1133–1141 (2017).
41. Francioni, V., Padamsey, Z. & Rochefort, N. L. High and asymmetric somato-dendritic coupling of v1 layer 5 neurons independent of visual stimulation and locomotion. *eLife* (2019) doi:10.7554/eLife.49145.
42. Padamsey, Z., Katsanevaki, D., Dupuy, N. & Rochefort, N. L. Neocortex saves energy by reducing coding precision during food scarcity. *Neuron* **110**, 280-296.e10 (2022).

43. Rupprecht, P. *et al.* A database and deep learning toolbox for noise-optimized, generalized spike inference from calcium imaging. *Nat. Neurosci.* **24**, 1324–1337 (2021).
44. Kerlin, A. *et al.* Functional clustering of dendritic activity during decision-making. *eLife* **8**, (2019).
45. Hill, D. N., Varga, Z., Jia, H., Sakmann, B. & Konnerth, A. Multibranch activity in basal and tuft dendrites during firing of layer 5 cortical neurons in vivo. *Proc. Natl. Acad. Sci.* **110**, 13618–13623 (2013).
46. Grienberger, C., Chen, X. & Konnerth, A. NMDA Receptor-Dependent Multidendrite Ca 2+ Spikes Required for Hippocampal Burst Firing In Vivo. *Neuron* **81**, 1274–1281 (2014).
47. Otor, Y. *et al.* Dynamic compartmental computations in tuft dendrites of layer 5 neurons during motor behavior. *Science* **376**, 267–275 (2022).
48. Helmchen, F., Svoboda, K., Denk, W. & Tank, D. W. In vivo dendritic calcium dynamics in deep-layer cortical pyramidal neurons. *Nat. Neurosci.* (1999) doi:10.1038/14788.
49. Francioni, V. & Harnett, M. T. Rethinking Single Neuron Electrical Compartmentalization: Dendritic Contributions to Network Computation In Vivo. *Neuroscience* **489**, 185–199 (2022).
50. Xu, N. *et al.* Nonlinear dendritic integration of sensory and motor input during an active sensing task. *Nature* **492**, 247–51 (2012).
51. Mikulasch, F. A., Rudelt, L., Wibral, M. & Priesemann, V. Where is the error? Hierarchical predictive coding through dendritic error computation. *Trends Neurosci.* **46**, 45–59 (2023).
52. Branco, T. & Häusser, M. The single dendritic branch as a fundamental functional unit in the nervous system. *Curr. Opin. Neurobiol.* (2010) doi:10.1016/j.conb.2010.07.009.
53. Guizar-Sicairos, M., Thurman, S. T. & Fienup, J. R. Efficient subpixel image registration algorithms. *Opt. Lett.* **33**, 156–158 (2008).
54. Brainard, D. H. The Psychophysics Toolbox. *Spat. Vis.* **10**, 433–436 (1997).

55. Pachitariu, M. *et al.* Suite2p: beyond 10,000 neurons with standard two-photon microscopy. 061507 Preprint at https://doi.org/10.1101/061507 (2017).
56. Keemink, S. W. *et al.* FISSA: A neuropil decontamination toolbox for calcium imaging signals. *Sci. Rep.* (2018) doi:10.1038/s41598-018-21640-2.
57. Hsieh, C.-J., Chang, K.-W., Lin, C.-J., Keerthi, S. S. & Sundararajan, S. A dual coordinate descent method for large-scale linear SVM. in *Proceedings of the 25th international conference on Machine learning - ICML '08* 408–415 (ACM Press, 2008). doi:10.1145/1390156.1390208.

Acknowledgements

We thank Ila Fiete and Courtney Yaeger for comments on the manuscript. V.F. is Y. Eva Tan Molecular Therapeutics Fellow. V.D.T. is a Janet and Sheldon (1959) Razin Fellow. M.T.H. was supported by the NIH (RO1NS106031 and R01NS113079) and by the Klingenstein-Simons Fellowship, the Vallee Foundation Scholars, and the McKnight Scholars programs.

Contributions

V.F. conceptualized and designed the experimental approach, designed the BCI, habituated and performed surgery on mice, collected the data, conceptualized and implemented the analyses, prepared the figures and wrote the manuscript. V.D.T. helped in the conceptualization of data analysis, building the data analysis pipeline and writing the manuscript. N.J.B. performed surgeries on mice. E.H.S.T. helped writing the manuscript. M.T.H supervised all aspects of the project.

Ethics declaration

The authors declare no competing interest.

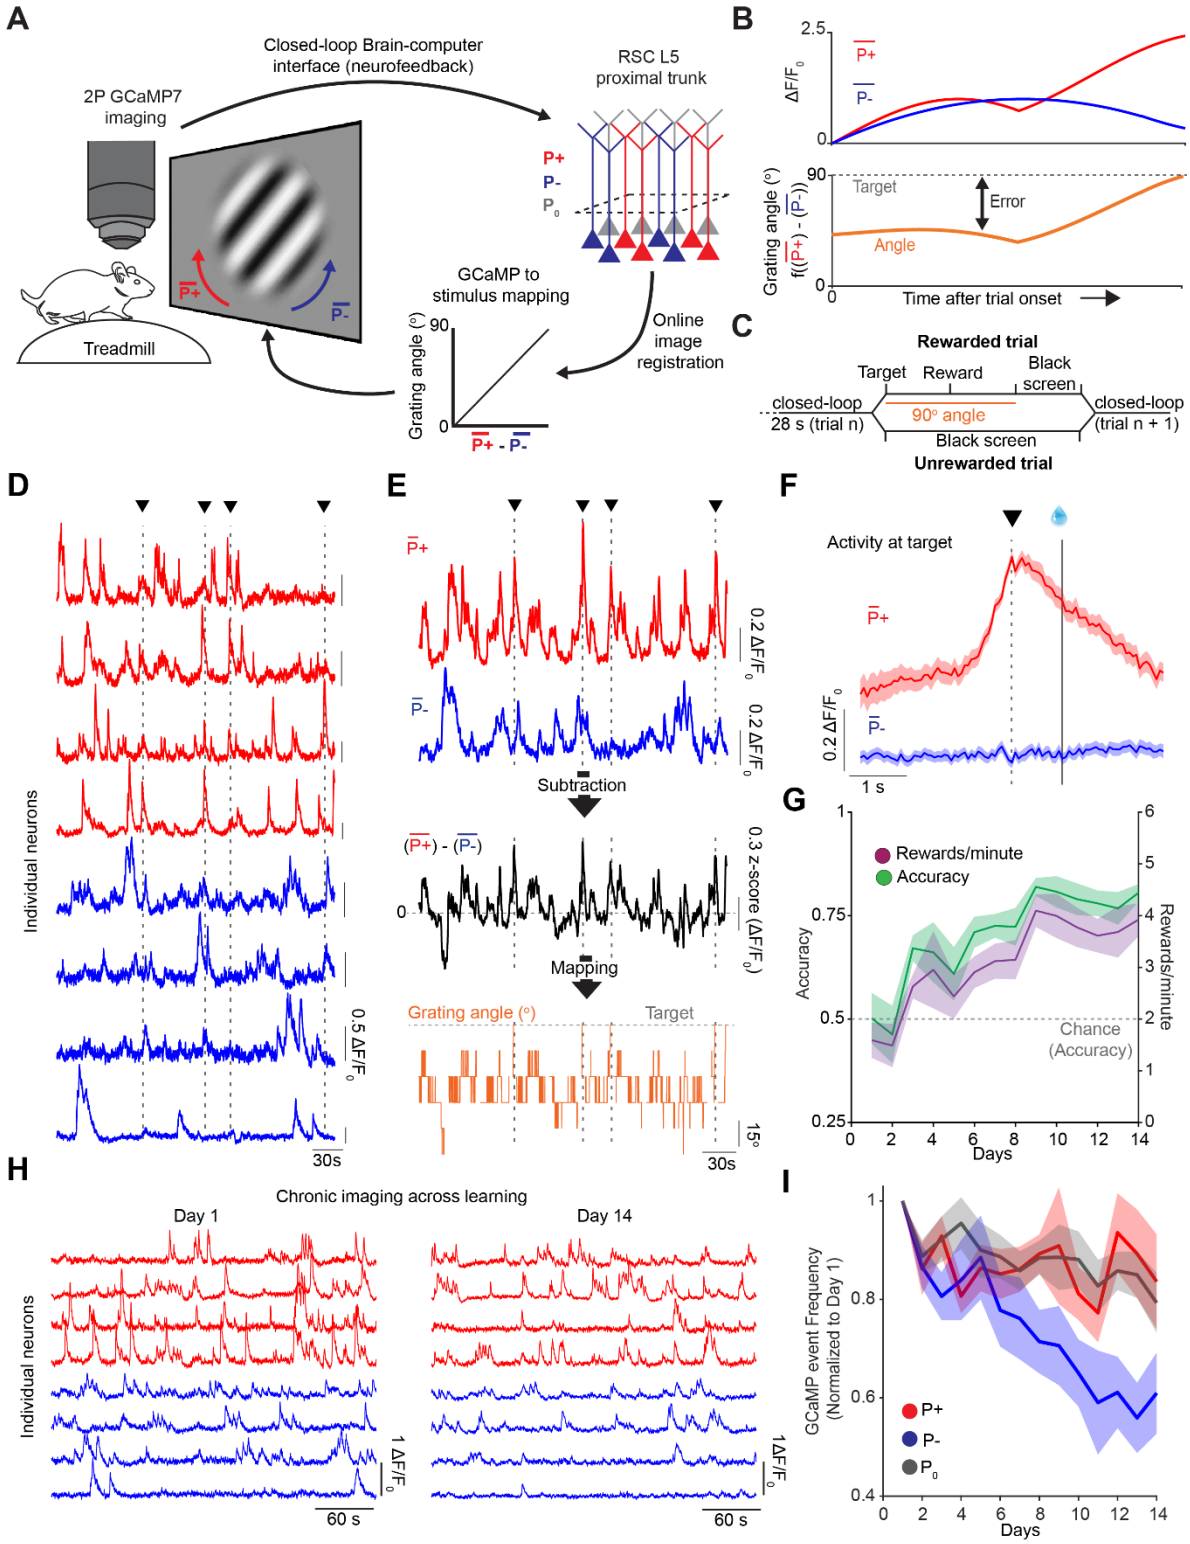
A
Closed-loop Brain-computer interface (neurofeedback)
2P GCaMP7 imaging
RSC L5 proximal trunk
P+
P-
P0
Treadmill
GCaMP to stimulus mapping
Grating angle (°)
90
0
P+ - P-
Online image registration
B
ΔF/F0
2.5
0
P+
P-
Grating angle (°)
90
Target
Error
Angle
0
Time after trial onset
C
Rewarded trial
Target
Reward
Black screen
closed-loop 28 s (trial n)
90° angle
closed-loop (trial n + 1)
Black screen
Unrewarded trial
D
Individual neurons
0.5 ΔF/F0
30s
E
P+
P-
0.2 ΔF/F0
Subtraction
(P+) - (P-)
0.3 z-score (ΔF/F0)
Mapping
Grating angle (°)
Target
15°
30s
F
Activity at target
P+
P-
0.2 ΔF/F0
1 s
G
Rewards/minute
Accuracy
Chance (Accuracy)
Days
H
Chronic imaging across learning
Day 1
Day 14
Individual neurons
1 ΔF/F0
60 s
I
GCaMP event Frequency (Normalized to Day 1)
P+
P-
P0
Days

Figure 1: Mice learn a neurofeedback BCI task through the differential regulation of P+ and P- neurons.

a, Schematic of the BCI setup. Mice were head-fixed and imaged under a 2-photon microscope and free to run on a cylindrical treadmill. Two user-defined populations of GCaMP7f-labeled layer 5 pyramidal neurons in retrosplenial cortex (RSC) were imaged at the proximal apical trunk: P+ (in red) and P- (in blue) were selected to control the rotation of a Gabor patch. P_0 neurons were designated as all other neurons in the field of view. Single frames were online-registered (motion-corrected). Activity in P+ neurons rotated the patch clockwise, towards the target angle of 90-degree. Activity in P- neurons rotated the Gabor patch stimulus counter-clockwise, towards a 0-degree angle. **b,** Schematic of the mapping between P+ and P- activity, stimulus angle, target activity and error. Error was the distance between current and target activation. The angle represents a binned (7 bins, 15 degrees apart, from 0 to 90 degrees) linear mapping between the mean activity in P+ neurons minus the activity in P- neurons. **c,** Trial structure: mice had 28 seconds to reach target activity and receive a reward delivered 1 second later. In successful trials, the 90-degree Gabor patch was shown for 2 seconds, followed by 1 second of black screen presentation. In unsuccessful trials, a 3 second black screen was presented before the onset of the next trial. **d,** $\Delta F/F_0$ traces as recorded live for P+ (in red) and P- (in blue) neurons. Vertical dashed lines and triangles represent timepoints where the animal reached target activity. **e,** Mean activity for the red (P+) and blue (P-) traces shown in d. Black trace shows the arithmetic subtraction of P+ and P- neurons (z-scored). Orange trace shows the corresponding visual stimulus angle as presented to the mouse. **f,** Mean $\Delta F/F_0$ for P+ and P- activity aligned to the time in which the animal reached target activity (dashed, vertical line and black triangle) for the session highlighted in d and e. Reward was delivered 1 second later (solid vertical line with water reward). Shaded areas are standard error of the mean (SEM). **g,** Mean performance over days quantified as the fraction of successful trials over the total number of trials in green, and as the number of rewards per minute in purple (One-way repeated measures ANOVA, $p = 5e^{-4}$ and p = for reward accuracy and p = 0.002 for accuracy and rewards/minute, respectively. n = 6 mice). Dashed horizontal red line represents chance level for accuracy performance (see Methods). **h,** $\Delta F/F_0$ traces for the same P+ and P- neurons on training day 1 and training day 14. **i,** Calcium transient frequency for P+, P- and P_0 neurons (in red, blue, and black, respectively) across the 14 days of training normalized to the activity on day 1. All neurons were tracked over the full 14 days of imaging. (Two-way repeated measures ANOVA, p = 0.012, p = 0.004 and $p = 9.3e^{-4}$ for the effect of population identity, days and an interaction between these 2. After Tukey's multiple

comparison, $p = 0.027$, $p = 0.95$ and $p = 0.01$ for P+ vs. P-, P+ vs P_0 and P- vs P_0 neurons, respectively. n = 6 mice).

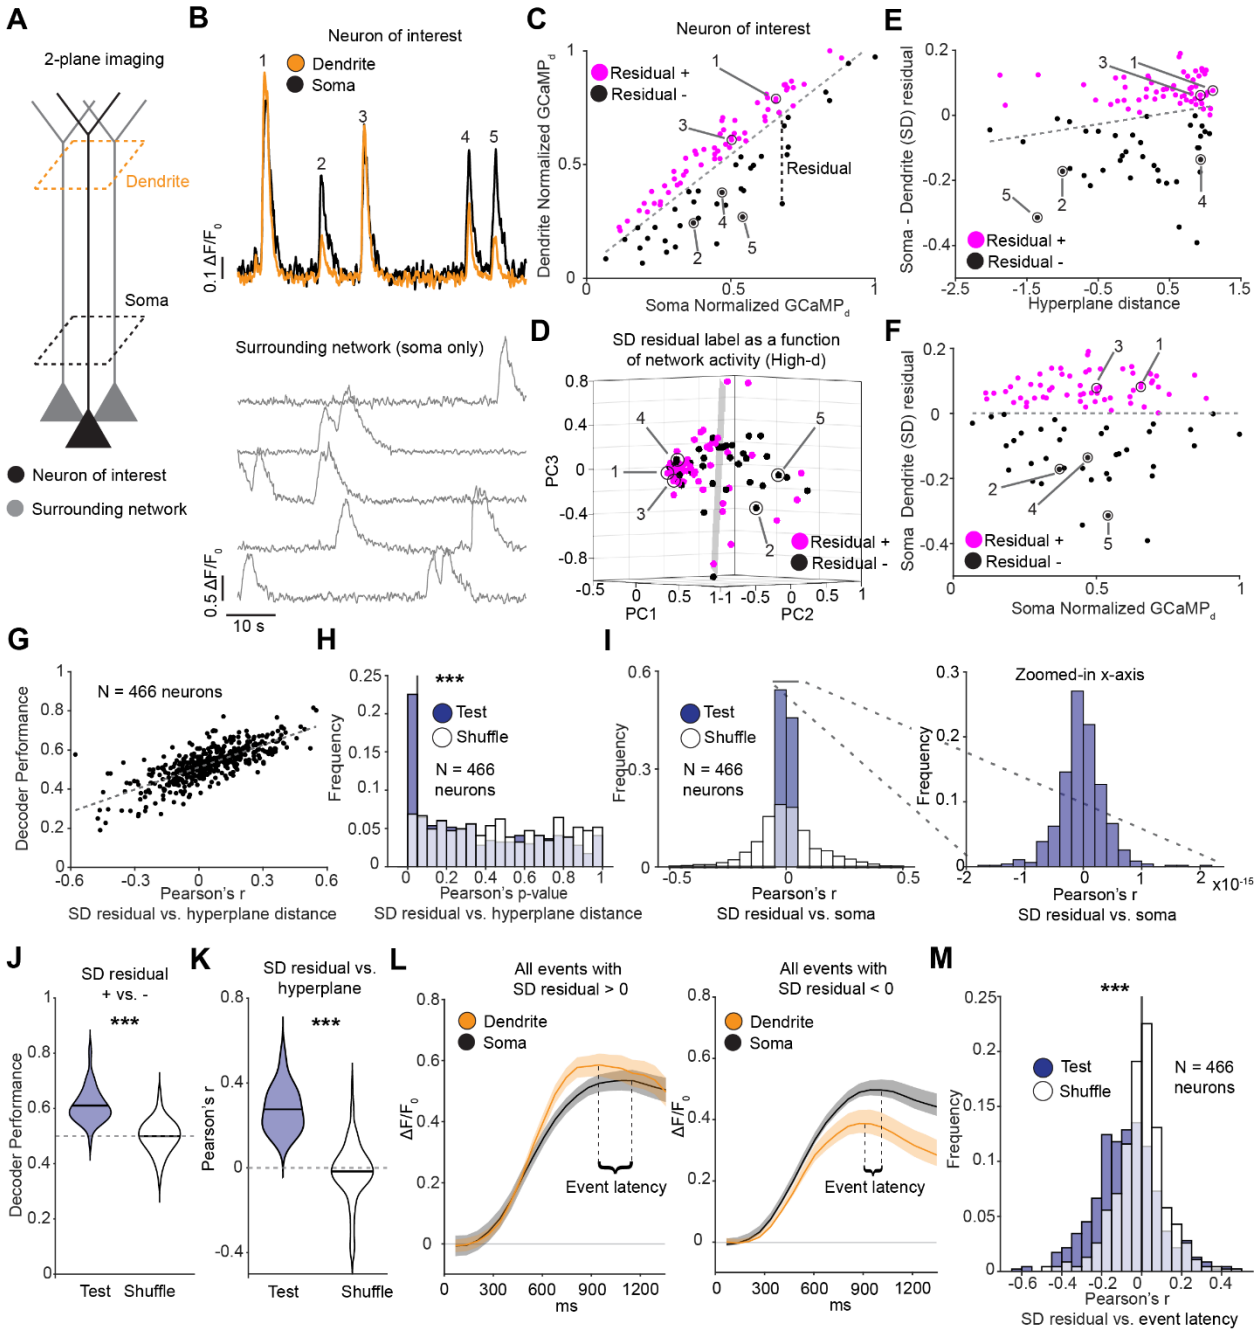
A
2-plane imaging
Dendrite
Soma
Neuron of interest
Surrounding network
B
Neuron of interest
Dendrite
Soma
0.1 ΔF/F0
Surrounding network (soma only)
0.5 ΔF/F0
10 s
C
Neuron of interest
Residual +
Residual -
Residual
Dendrite Normalized GCaMPd
Soma Normalized GCaMPd
D
SD residual label as a function of network activity (High-d)
PC1
PC2
PC3
Residual +
Residual -
E
Soma - Dendrite (SD) residual
Hyperplane distance
Residual +
Residual -
F
Soma Dendrite (SD) residual
Soma Normalized GCaMPd
Residual +
Residual -
G
N = 466 neurons
Decoder Performance
Pearson's r
SD residual vs. hyperplane distance
H
***
Test
Shuffle
N = 466 neurons
Frequency
Pearson's p-value
SD residual vs. hyperplane distance
I
Test
Shuffle
N = 466 neurons
Frequency
Pearson's r
SD residual vs. soma
Zoomed-in x-axis
x10-15
J
SD residual + vs. -
***
Decoder Performance
Test
Shuffle
K
SD residual vs. hyperplane
***
Pearson's r
Test
Shuffle
L
All events with SD residual > 0
Dendrite
Soma
ΔF/F0
Event latency
ms
All events with SD residual < 0
Dendrite
Soma
ΔF/F0
Event latency
ms
M
***
Test
Shuffle
N = 466 neurons
Frequency
Pearson's r
SD residual vs. event latency

605

Figure 2: Differences in somatic and dendritic magnitudes for coincident events are predicted by local network dynamics.

a, Schematic of two-plane 2-photon calcium imaging of a network of neurons at the proximal and distal trunk. **b,** $\Delta F/F_0$ traces recorded simultaneously in the soma (black) and dendrite (orange) for a single neuron of interest (top, P+ and P- neurons across days 1-14) and corresponding activity in 5 surrounding neurons. Numbers 1-5 indicate identified GCaMP events. **c,** Relationship between somatic and dendritic transients' integral for the example neuron shown in b. Datapoints represent individual events simultaneously detected in soma and dendrite (see Methods). A least-squares linear model (dashed grey line) defined events as dendritically amplified (in magenta, residual +) versus dendritically attenuated (in black, residual -). Events (1-5) correspond to the transients shown in b. **d,** For each coincident event in the neuron of interest shown in b and c, we estimated the network activity vector in the 2 seconds before, using all other neurons in the field of view. Here, network activity vector was projected onto the first 3 principal components for visualization only. The shaded black hyperplane represents the decision boundary for binary classification (dendritically amplified versus dendritically attenuated) calculated using a linear SVM. Events 1-5 correspond to the network activity vector associated with transients 1-5 shown in b and c. **e,** The relationship between SD residuals estimated in c and the distance from the decision boundary (hyperplane distance) estimated in d for all coincident somato-dendritic events in the neuron of interest. Events 1-5 correspond to those shown in b, c and d. The dashed grey line represents the least-squares best-fit line. To maintain visual consistency with d, the distance from the hyperplane was estimated on the first 3 principal components. This is for visualization only. **f,** The relationship between SD residual as estimated in c, and somatic event magnitude. Highlighted events 1-5 correspond to those shown in b, c, d and e. Dashed grey line represents the least-squares best fit line. **g,** Decoder performance as a function of the correlation between SD residuals and hyperplane distance (Pearson's r = 0.74; p-value = $1.4e^{-84}$, n = 466 neurons). Datapoints represents individual neurons. **h,** Distribution of p-values for test data and a control randomly shuffled distribution, testing the correlation between SD residuals and distance from the hyperplane (or classification confidence, Wilcoxon signed rank test = p = $1.3e^{-9}$. N = 466 neurons) as estimated in e. **i,** Left panel, for all neurons, the Pearson's r for somato-dendritic residuals and somatic event magnitude as characterized in f. The residual-based approach perfectly decorrelates SD residual from somatic activity alone. Right panel, for test data, a zoomed in version of the same histogram shown in the left panel. Note the 10^{-15} scale on the x-axis. **j,** Decoding performance for neurons with a statistically significant correlation between SD residual

and distance from the hyperplane (paired t-test, $p = 8.6e^{-9}$. Mean = 0.61 and 0.50; SEM = 0.006 and 0.007; n = 82). Dashed grey line indicates chance level. **k,** Pearson's r for neurons with a statistically-significant correlation between SD residual and the distance between population vector and hyperplane (paired t-test, $p = 3.35e^{-25}$. Mean = 0.28 and $-7.2e^{-4}$; SEM = 0.01 and 0.01; n = 82). **l,** Mean $\Delta F/F_0$ events for soma (black) and dendrites (orange) for all dendritically amplified (left panel) and dendritically attenuated (right panel) events in a single neuron. $\Delta F/F_0$ traces are aligned to somatic peak time. Event latency is defined as the time between the somatic and dendritic peaks. Compared to dendritically attenuated events, dendritically-amplified events peaked earlier. **M,** Pearson correlation value between the SD residual and the event latency between soma and its corresponding dendrite indicating that the larger the SD residual, the earlier the dendritic peak is compared to the somatic one (paired t-test, $p = 8e^{-13}$. Mean = -0.075 and -0.005; SEM = 0.007 and 0.006. n = 466 neurons).

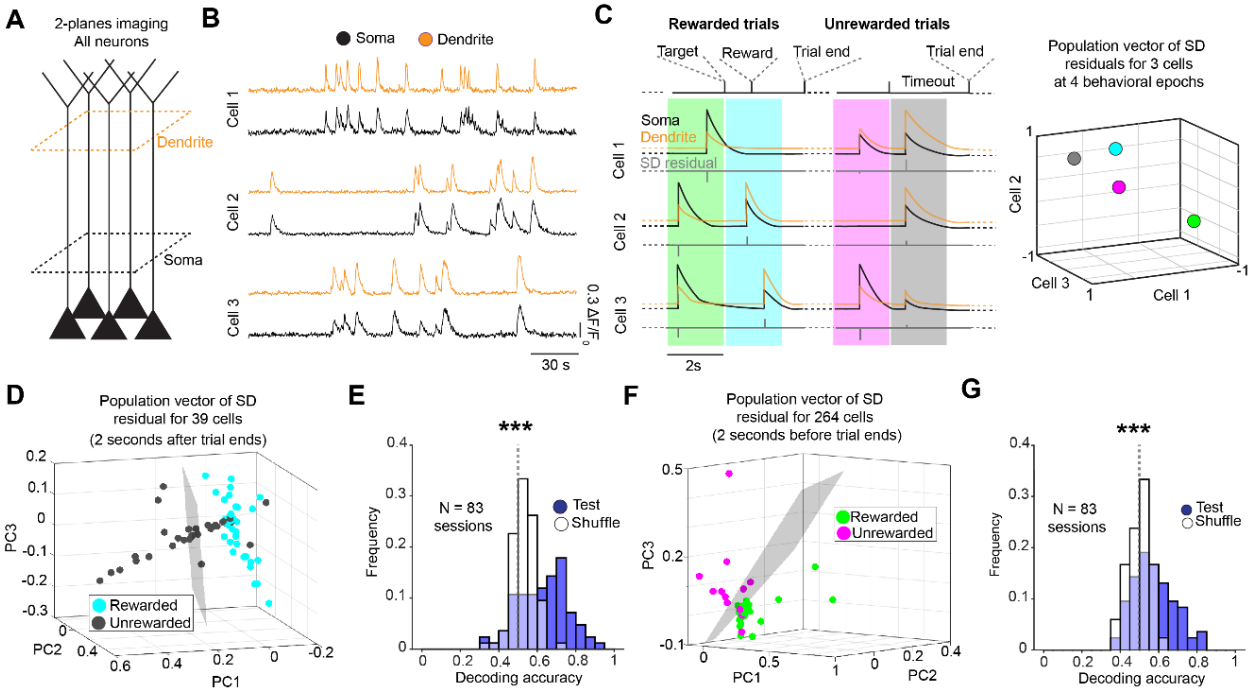
A
2-planes imaging
All neurons
Dendrite
Soma
B
Soma
Dendrite
Cell 1
Cell 2
Cell 3
30 s
C
Rewarded trials
Unrewarded trials
Target
Reward
Trial end
Trial end
Timeout
Soma
Dendrite
SD residual
2s
Population vector of SD
residuals for 3 cells
at 4 behavioral epochs
Cell 1
Cell 2
Cell 3
D
Population vector of SD
residual for 39 cells
(2 seconds after trial ends)
Rewarded
Unrewarded
PC1
PC2
PC3
E
***
N = 83
sessions
Test
Shuffle
Frequency
Decoding accuracy
F
Population vector of SD
residual for 264 cells
(2 seconds before trial ends)
Rewarded
Unrewarded
PC1
PC2
PC3
G
***
N = 83
sessions
Test
Shuffle
Frequency
Decoding accuracy

Figure 3: A population vector of somato-dendritic residuals contains reward-related information.

a, Schematic of the experimental approach. We isolated all neurons in the field of view with a paired soma and dendrite and used the SD residual population vector to decode task-relevant variables. **b,** $\Delta F/F_0$ GCaMP traces simultaneously recorded in the soma (black) and its corresponding dendrite (orange) for three representative neurons. **c,** Schematic of SD residual population vector for 4 different behavioral epochs. Left panel, in black, orange and grey the somatic, dendritic and residual traces for 3 cartoon neurons. The green, cyan, purple and black boxes represent 4 different behavioral epochs: 2 seconds before and after target is reached for rewarded trials and 2 seconds before and after the end of unrewarded trials. Each neuron's SD residual trace was estimated by collapsing coincident soma-dendrites events into point values at the time of event onset. Resultant SD residual traces for all neurons in each behavioral epoch were used to estimate the n-dimensional vector of SD residuals where n corresponded to the number of neurons for with paired somas and dendrites (see method and figure 2). Right panel, 3-D plot of SD residuals for the four behavioral epochs shown in the left panel, where x, y and z correspond to SD residual of neurons 1-3 from left. **d,** A 39-dimensional vector of SD residuals collapsed onto the first 3 principal components for visualization purposes only. Each dot corresponds to a vector of SD residuals. In cyan, vectors resulting from the two seconds following the reaching of target activity in rewarded trials. In dark grey, vectors calculated in the 2 seconds following the end of an unsuccessful trial (same as C). Shaded black hyperplane represents decision boundary for binary classification calculated using a linear SVM. **e,** Decoding accuracy for test vs. shuffle data for 83 sessions (paired t-test, $p = 9.8e^{-9}$. Mean = 0.63 and 0.52; SEM = 0.01 and 0.01; n = 83 for test and shuffled data, respectively. **f, g,** Same as in d and e: a 264-dimensional vector of SD residual collapsed onto the first 3 principal components for visualization only. Green represents the last two seconds of a rewarded trial while purple represents the last two seconds of an unrewarded trial (paired t-test, $p = 7.1e^{-8}$. Mean = 0.57 and 0.49; SEM = 0.01 and 0.01; n = 83 for test and shuffled data, respectively).

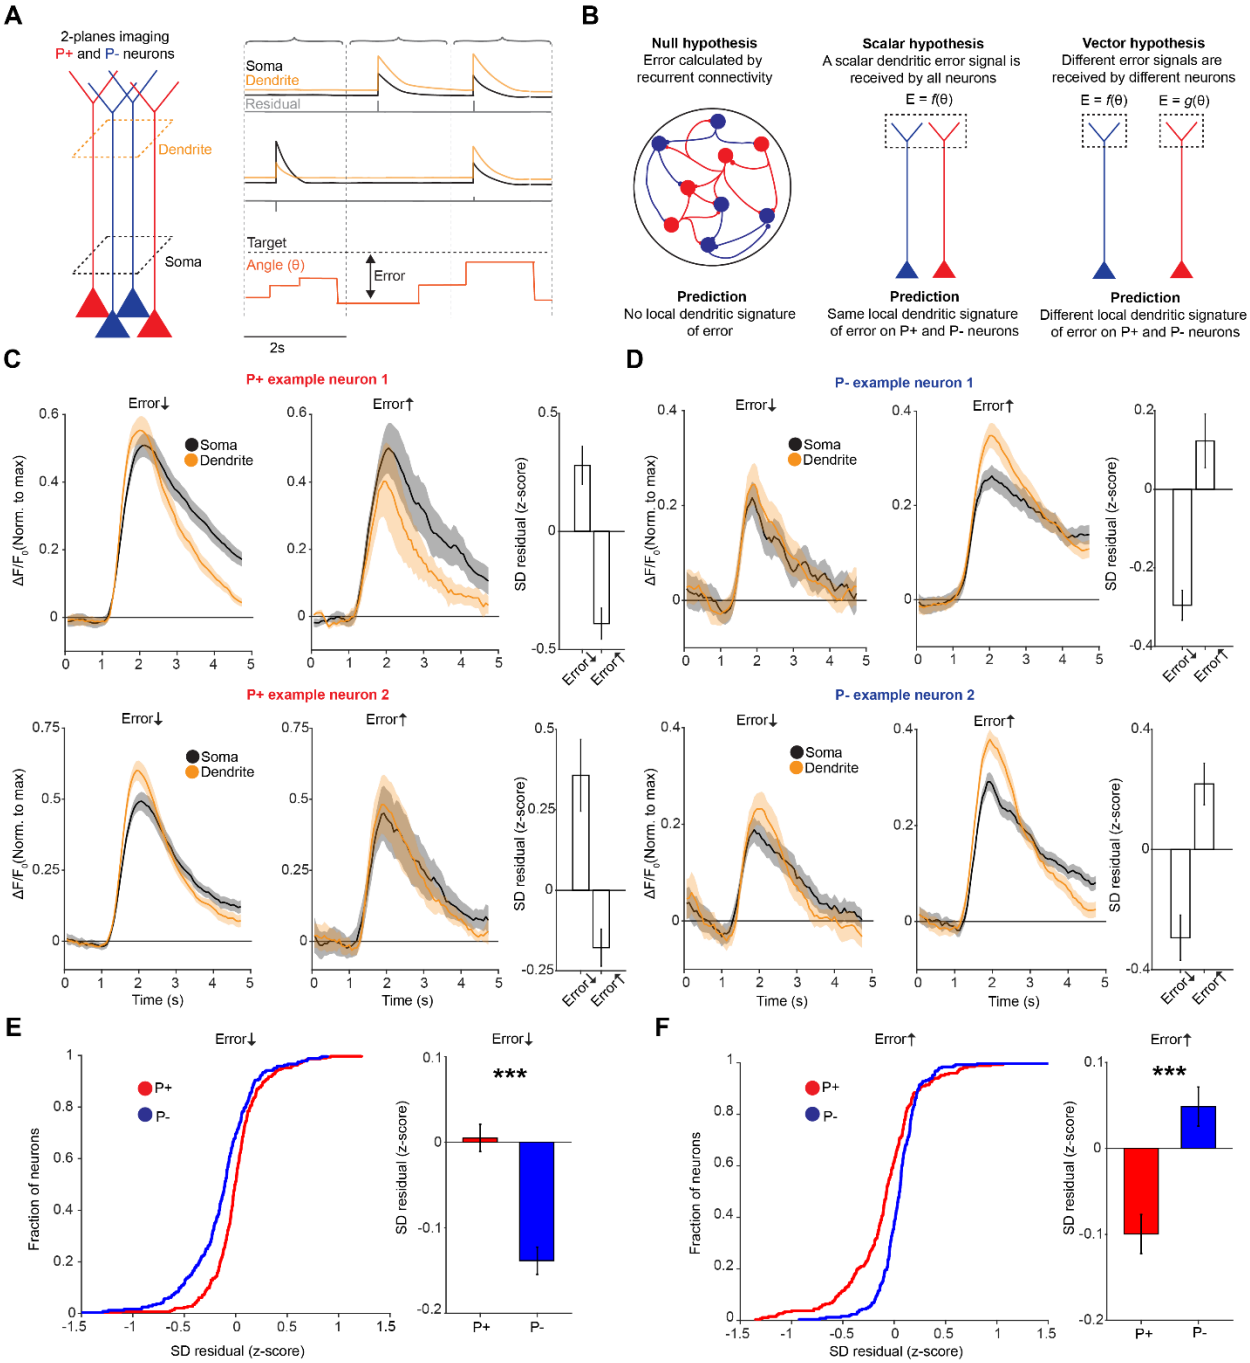
A
2-planes imaging
P+ and P- neurons
Dendrite
Soma
Soma
Dendrite
Residual
Target
Angle (θ)
Error
2s
B
Null hypothesis
Error calculated by recurrent connectivity
Prediction
No local dendritic signature of error
Scalar hypothesis
A scalar dendritic error signal is received by all neurons
E = f(θ)
Prediction
Same local dendritic signature of error on P+ and P- neurons
Vector hypothesis
Different error signals are received by different neurons
E = f(θ)
E = g(θ)
Prediction
Different local dendritic signature of error on P+ and P- neurons
C
P+ example neuron 1
Error↓
Error↑
Soma
Dendrite
ΔF/F0 (Norm. to max)
SD residual (z-score)
P+ example neuron 2
Time (s)
D
P- example neuron 1
P- example neuron 2
E
Error↓
P+
P-
Fraction of neurons
SD residual (z-score)
***
F
Error↑

Figure 4: Dendritic error signals are cell-specific and depend on the causal contribution of the neuron to the task.

a, Experimental schema. Black, orange and grey indicate idealized somatic, dendritic and residual traces for 3 neurons. Relationship between stimulus angle, target, and error shown below. All closed-loop data from the BCI paradigm (which excludes rewards and timeout periods) were chunked into 2 seconds bins. Epochs of error decrease and increase were defined as bins in which the mean derivative of the angle increased and decreased, respectively. **b,** Three possible hypotheses: In the null hypothesis scenario, error is calculated at the population level through recurrent dynamics independent of dendrites. In the scalar hypothesis, a single error signal is broadcasted through the dendrites of all neurons in the network. This model predicts that all neurons receive the same error signal. In the vector hypothesis, error signals received on the dendrites of individual neurons are tailored according to the causal involvement of individual neuron to behavior. This model predicts that neurons with opposite causal contribution to behavior will receive different error signals onto their dendrites. **c,** For two individual P+ neurons, the mean $\Delta F/F_0$ signal in the somas (black) and in the dendrites (orange) for all events that occurred during epochs of error reduction and error increase respectively. Compared to somatic activity, dendrites are relatively amplified during error reduction compared to error increase epochs. The bar graph represents the mean SD residual value (z-scored) for all events occurred during error decrease and increase epochs. Error bars represent SEM. **d,** Same as c for two P- neurons. Contrary to the P+ neurons, dendritic activity is relatively attenuated for error reduction epochs compared to error increase epochs. **e,** Left panel, during error reduction epochs, the cumulative distribution function for SD residuals (z-scored across all neurons) for P+ (in red) and P- (in blue) neurons. The bar graph represents the mean and SEM for the population distribution shown in the cumulative distribution function for P+ and P- neurons. Dendrites of P+ neurons are relatively amplified compared to the dendrites of P- neurons during error reduction epochs (t-test; $p = 5.3e^{-7}$; mean = 0.005 and -0.14; SEM = 0.01 and 0.02; n = 292 and 240 for P+ and P- neurons respectively). **f,** Same as e, for error increase epochs. Dendrites in P+ neurons are more attenuated than P- neurons during error reduction epochs (t-test, $p = 1.2e^{-7}$; mean = - 0.1 and 0.05; SEM = 0.02 and 0.01; n = 267 and 249 for P+ and P- neurons respectively).

Supplemental Materials

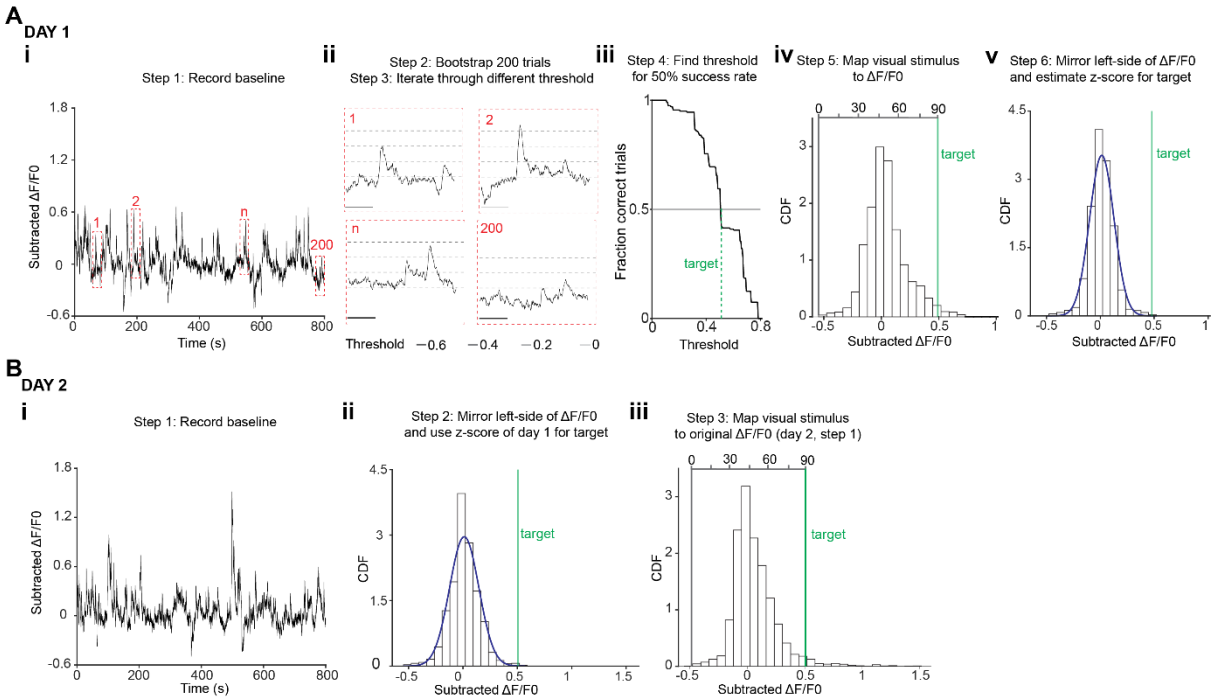

Supplementary Figure 1: Determining target activity and mapping visual feedback to ΔF/F0 within and across days.

a, Mapping strategy on day 1. **ai**, A baseline of approximately 13 minutes (12,000 frames imaged at ~15Hz) was recorded each day for P+ and P- neurons, and the mean ΔF/F0 for each subpopulation was subtracted from one another (also see Fig. 1e). **aii,** The subtracted signal was used to bootstrap 200 trials (28 seconds each resampling with replacement 200 start times, red boxes). Next, we iteratively searched across different thresholds. **aiii,** For each threshold (in incremental steps of 0.005 ΔF/F0), we selected the one which was crossed 50% of the times in the 200 resampled trials. This value was set as the target value (i.e. 90-degree angle). **aiv**, Next we mapped the subtracted ΔF/F0 to a visual feedback stimulus. 90-degree was mapped as the ΔF/F0 at which the animal reached target on 50% of the resampled trials (see aii), 0-degrees was mapped as the minimum value in the ΔF/F0 distribution. The space in between was linearly spaced into 6 bins whose edges marked the boundary between 15-degrees angles. **av**, The left side of the subtracted ΔF/F0 was mirrored on the right and fitted using a Gaussian function. The z-score of the target value (see aii), was recorded for mapping on subsequent days. **b**, Target activity and mapping strategy on day 2-14. **bi**, We recorded baseline activity in the same way as

shown in ai. **bii**, After estimating the distribution of the subtracted ΔF/F0 signal, we mirrored the left part of the distribution and fitted a Gaussian function as described in aiv. Our target was defined as the same z-score of the target value on day 1 (see iv). **biii**, On the original subtracted ΔF/F0 distribution, we mapped 90-degrees as the z-score of the target value on day 1 (see bii) and 0-degrees as the minimum value in the subtracted ΔF/F0 distribution.

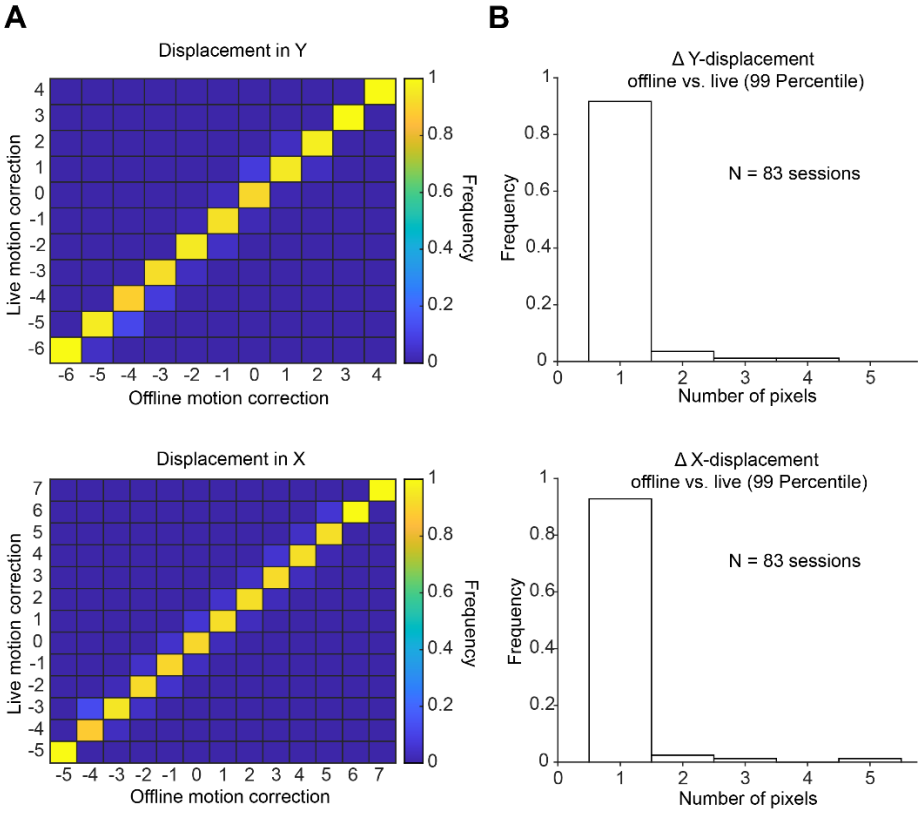

Supplementary Figure 2: Validation of live motion correction

a, Live motion correction performance for a given imaging session. We validated our live motion-correction algorithm by comparing its performance against a commonly used offline motion correction algorithm[55]. To be able to compare performance, we motion-corrected against the same reference image used for offline motion correction. For a given pixel displacement in X and Y in the offline algorithm, we calculated the proportion of frames which our online algorithm estimated to be displaced by the same amount. Each column is normalized to the total number of frames for a given displacement value. **b,** For X and Y displacements, in all our imaging sessions, the 99th percentile of the difference between live and offline motion correction. Our data show that in more than 90% of our imaging sessions, the estimated pixel displacement using live and offline motion correction differ by 1 pixel at the 99th percentile level.

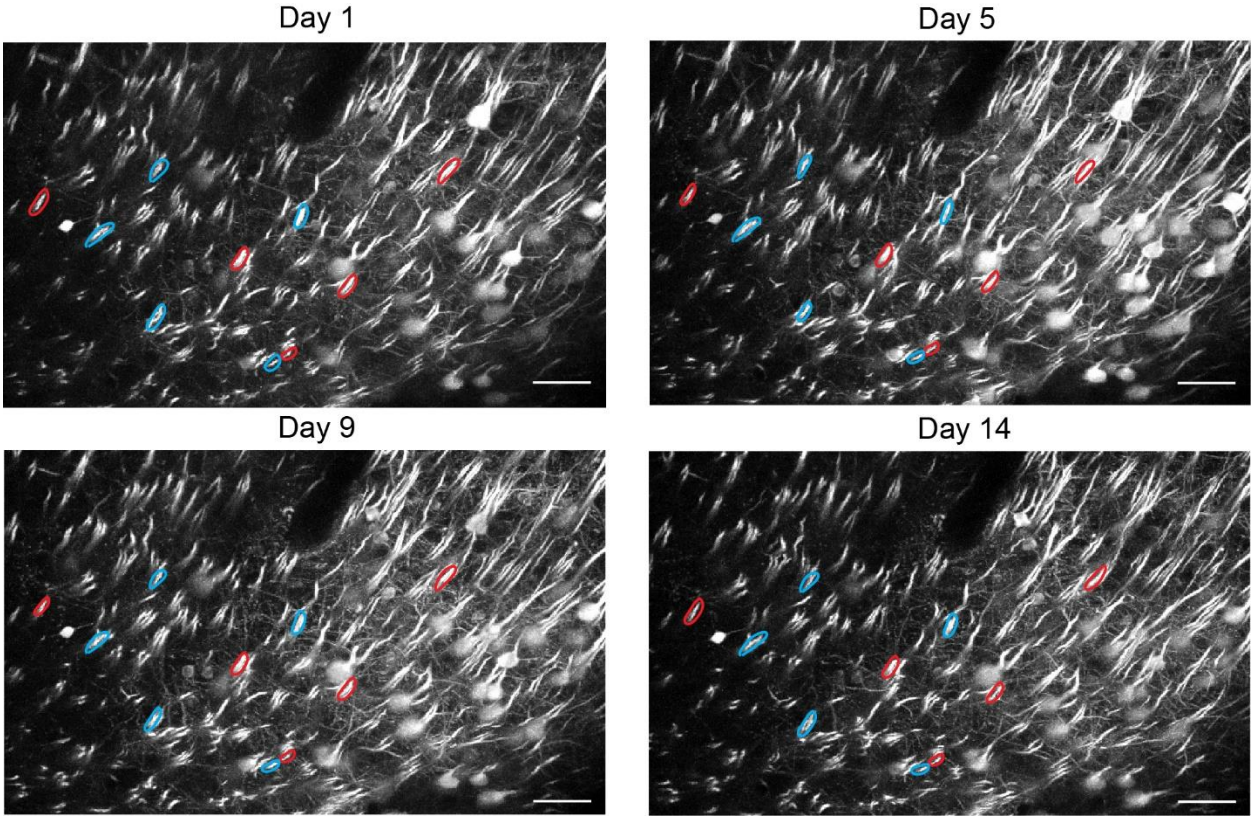

Supplementary Figure 3: Representative field of view with P+ and P- neurons

A representative field of view with the same, GCaMP7f-labelled, chronically-tracked P+ and P- neurons (5 each, shown in red and blue, respectively), imaged at the proximal trunk (as a proxy for soma) during 14 days of learning, shown at day 1, 5, 9 and 14. P+ and P- neurons were selected on day 1. Scale bar represents 50 µm.

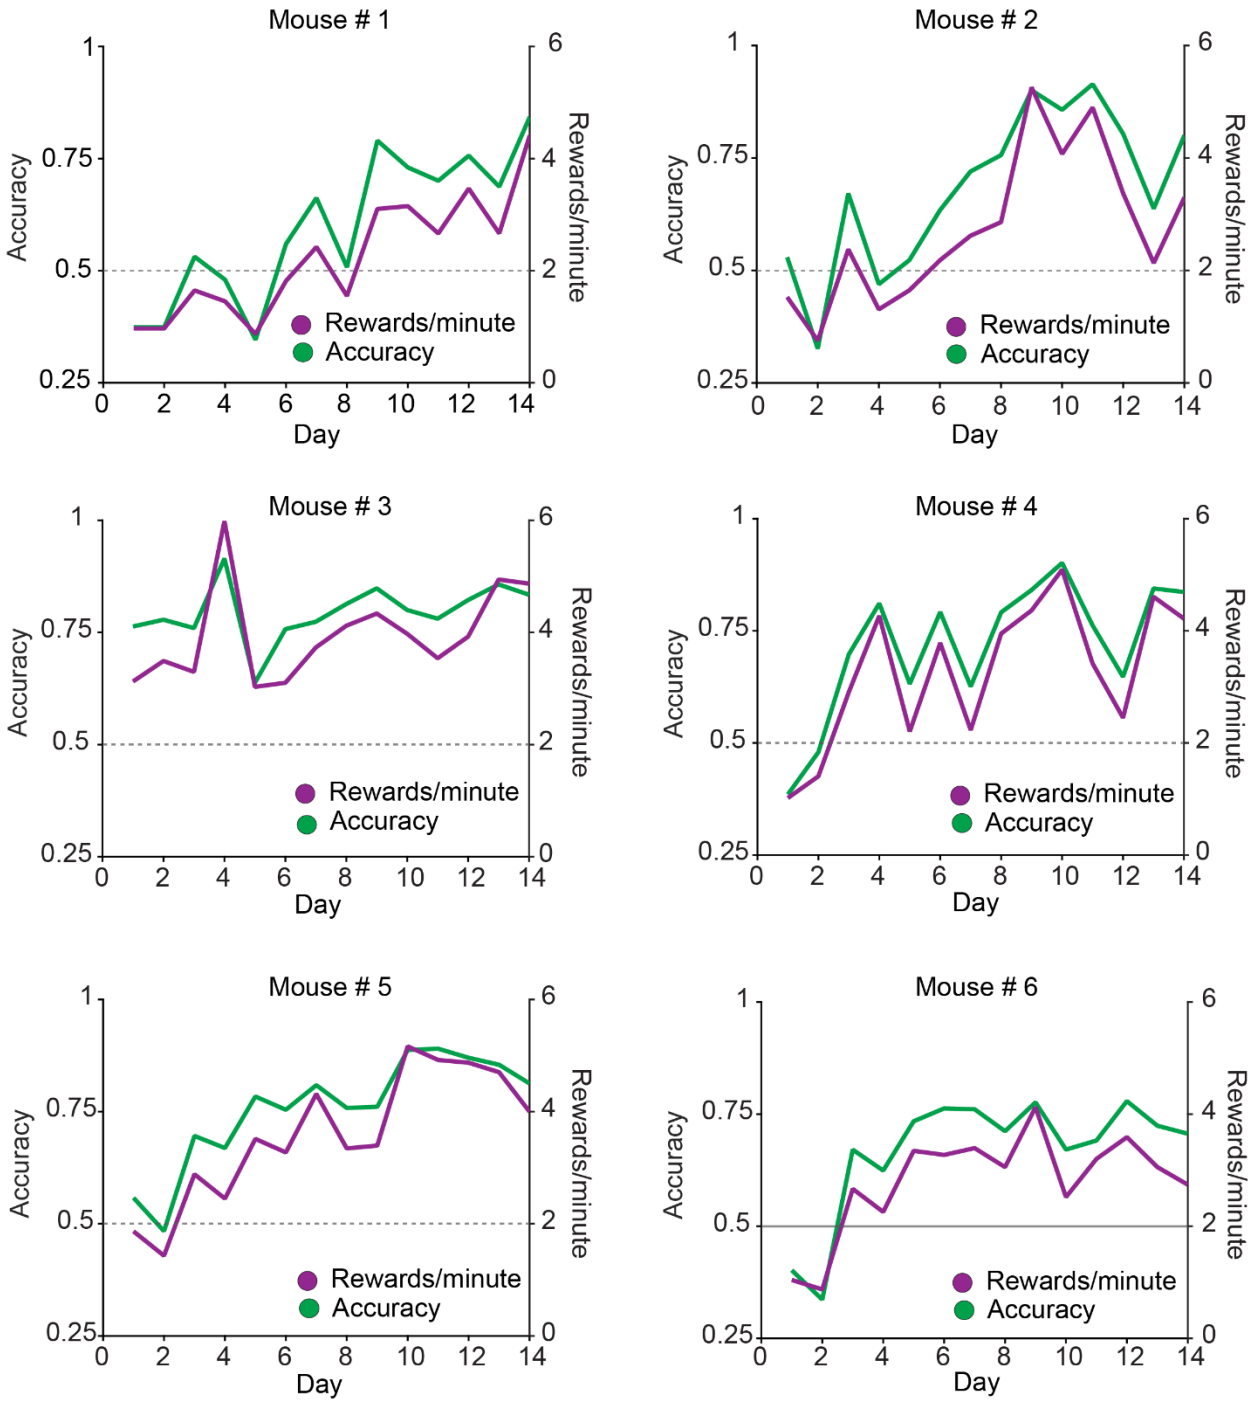

Supplementary Figure 4: Task performance for individual animals

Task performance evaluated using two metrics: accuracy (the fraction of successful trials) and rewards per minute for each of 6 mice. Summary graph and statistics in Figure 1g. The grey dashed line represents chance accuracy.

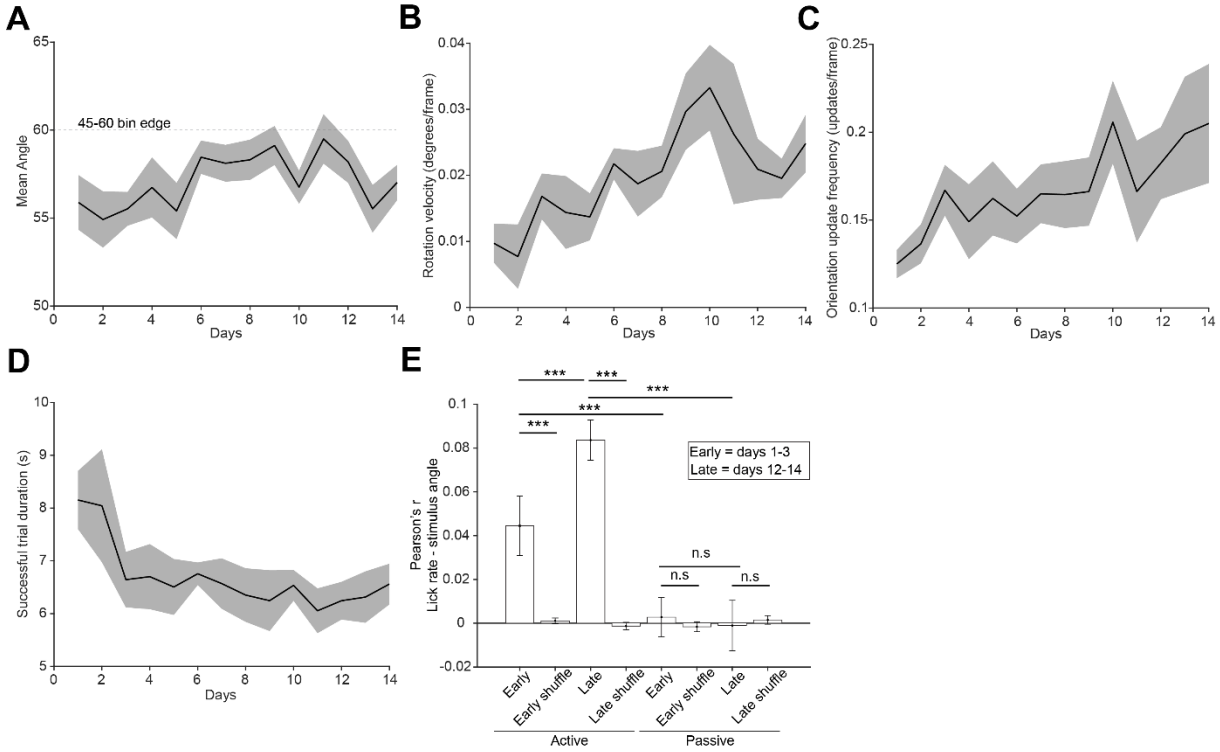

Supplementary Figure 5: Visual stimulus and behavioral correlates of learning.

a, Mean visual stimulus angle across 14 days of learning. Due to binning, stimuli below the 60-degrees threshold were presented as 45-degrees. **b,** Rotation speed across the 14 days of learning. Rotations towards and away from a reward were defined as positive and negative rotations, respectively. **c,** Mean frequency at which a new orientation was presented. **d,** Mean duration of successful trials across the 14 days of learning. **e,** Pearson's correlation between licking frequency and stimulus angle (outside reward periods) during active (closed-loop) and passive stimulus presentation at the beginning and the end of learning (Two-way repeated measures ANOVA, $p = 0.13$, $1.4e^{-11}$ and 0.028 for the effect of state (active vs passive), days, and interaction between state and days, respectively. After two-stage Benjamini, Krieger and Yekutieli correction for false discovery rate, $p = 8.5e^{-4}$ for early-active vs late-active, $3.4e^{-4}$ for early-active vs early-active (shuffle), $8.4e^{-11}$ for late-active vs late-active(shuffle), $1.7e^{-4}$ for early-active vs early passive, $5e^{-11}$ for late-active vs late-passive, 0.83 for early-passive vs late-passive, 0.8 for early-passive vs early-passive(shuffle), 0.81 for late-passive vs late-passive (shuffle)).

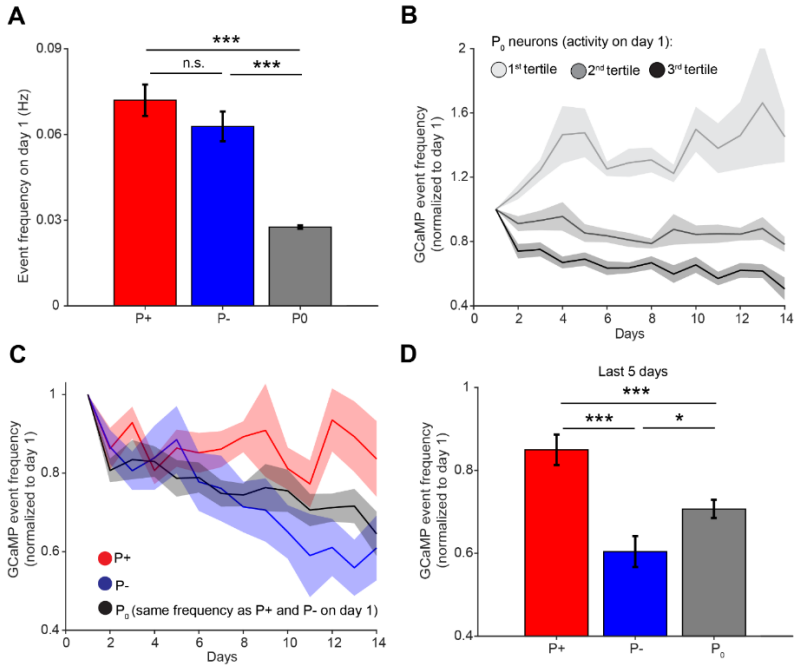

Supplementary Figure 6: Changes in GCaMP transient frequency over days depend on day 1 frequency.

a, Event frequency for P+, P- and P_0 neurons on day 1. P+ and P- neurons were selected to be of higher activity on day 1 compared to other neurons in the field of view (One-way ANOVA, p = $2.52e^{-30}$. After Tukey's multiple comparisons correction, p = 0.35, $9.57e^{-10}$ and $9.56e^{-10}$ for P+ vs P-, P+ vs P_0 and P- vs P_0, respectively. mean = 0.072, 0.063 and 0.028; SEM = 0.006, 0.005 and $5.6e^{-4}$; n = 27, 27 and 1839 for P+, P- and P_0, respectively). **b,** Calcium transient frequency for P_0 neurons divided in tertiles based on activity on day 1, across the 14 days of training normalized to the activity on day 1. All neurons were tracked over the full 14 days of imaging (n = 6 mice). **c,** Calcium transient frequency for P+, P- and P_0 neurons. Only P_0 neurons with the same calcium transient frequency as P+ and P- neurons on day 1 (neurons within the 95% confidence interval of the joint P+ and P- distribution) were selected. Neurons were tracked across the 14 days of imaging and activity was normalized to the activity on day 1 (Two-way repeated measures ANOVA, p = 0.008, p = 0.006 and p = 0.003 for the effect of population identity, days and an interaction between these 2, respectively. n = 6 mice across 14 days). **d,** Calcium transient frequency for P+, P- and P_0 on days 10 to 14, normalized to day 1 (One-way repeated measures ANOVA, p = $1e^{-7}$. After Tukey's multiple comparison, p = $3.3e^{-6}$, $5.47e^{-4}$ and 0.02 for P+ vs P-, P+ vs P_0 and P- vs P_0, respectively; mean = 0.85, 0.60 and 0.71, SEM = 0.037, 0.037 and 0.02 for P+, P- and P_0, respectively. n = 6 mice).

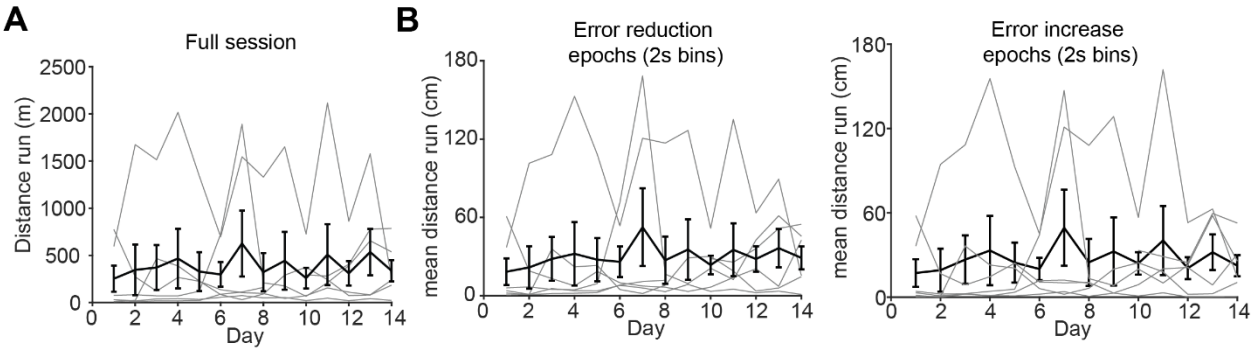

Supplementary Figure 7: Learning is not due to changes in locomotion across days.

a, In black, the mean and standard error of the mean for the distance run by all animals in individual imaging sessions. All sessions were of equal length (41.26 minutes or 76500 imaging frames (2 planes at 30Hz)). In grey, the distance run by individual mice (One-way repeated measures ANOVA, p = 0.58, n = 6 across 14 days). **b,** After excluding all timeouts, grey, black and reward-related periods, we divided our trials into 2 seconds bins and estimated whether the net rotation of the stimulus (the mean of the derivative angle) was positive (lefts panel, towards the reward, error reduction) or negative (right panel, away from the reward, error increase) in each bin and calculated the average distance run during error reduction and error increase epochs. We found no differences neither within condition across days, nor among conditions (Two-way repeated measures ANOVA, p = 0.15, 0.86 and 0.14 for the effects of error type, days and error type and days interaction n = 6 animals across 14 days).

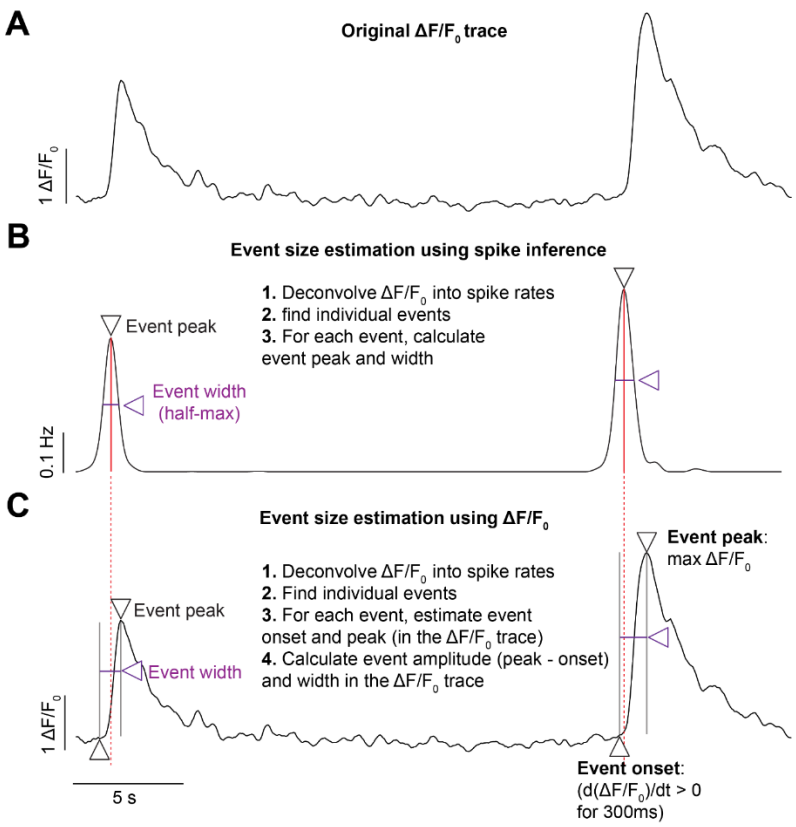

Supplementary Figure 8: Event detection.

a, $\Delta F/F_0$ trace for an individual neuron. **b,** For the same $\Delta F/F_0$ trace shown in A, calcium signal was deconvolved into a spike rate using CASCADE[43]. Event size was calculated by first detecting individual peaks (downward facing triangle, *findpeaks* function in MATLAB, see Method) and then calculating the area under the curve by estimating the amplitude (red, vertical line) and width at half maximum (purple line). **c,** Same $\Delta F/F_0$ trace as in a. To obtain $\Delta F/F_0$ metrics of event size, we first deconvolved $\Delta F/F_0$ signals into spike rates using CASCADE and detected individual transients (same as steps 1-2 as in b). For each detected event we then found its start (upward facing triangle) and peak (downward facing triangle) as the point at which the derivative of $\Delta F/F_0$ trace became positive for 300 ms in a row in a 2 seconds backward window from the detected spike inference peak, and the maximum $\Delta F/F_0$ value in the following 2 seconds (unless constrained by the presence of another event). The amplitude of the event was defined as $\Delta F/F_0$ value at peak – $\Delta F/F_0$ at onset. The width was defined as peak – onset frame.

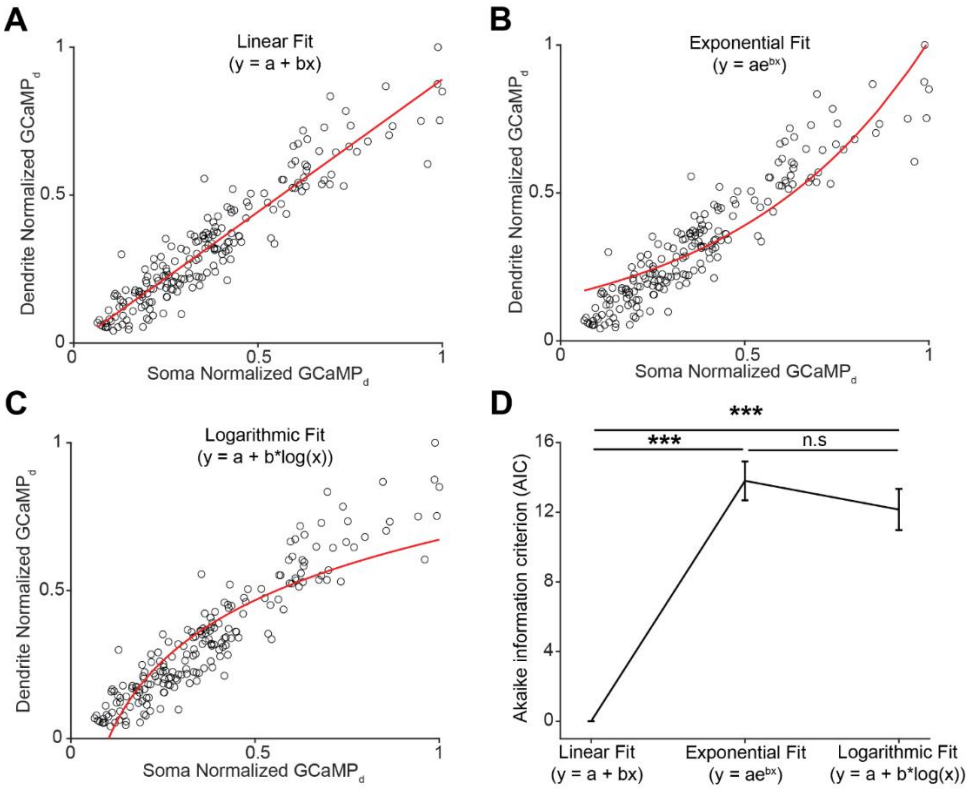

Supplementary Figure 9: Testing the assumption of linearity between somatic and dendritic event magnitude.

a, Representative example of a single neuron whose events were isolated and their magnitude estimated in the soma and dendrites, plotted against each other. Axes are normalized to the largest event. To describe the relationship between somatic and dendritic event magnitude, we fit a linear model. **b,** Same as in a, except this time we fit an exponential model to the data. **c,** Same as in a and b, this time however we fit a logarithmic model to the data. **d,** For all neurons, we calculated the Akaike information criteria (AIC) to describe the goodness of a model (while penalizing for the number of parameters used). Lower values of AIC mean better fit. AIC values were zeroed compared to the linear fit model. Error bars represent standard error of the mean (One-way repeated measures ANOVA; $p < 1e^{-15}$, $p < 1e^{-15}$, $p = 0.56$ for Linear vs Exponential, Linear vs Logarithmic and Exponential vs. Logarithmic after Tukey's multiple comparison correction. Mean = 0, 13.8 and 12.1; SEM = 0, 1.11 and 1.17; n = 543 neurons for Linear, Exponential and logarithmic fit, respectively).

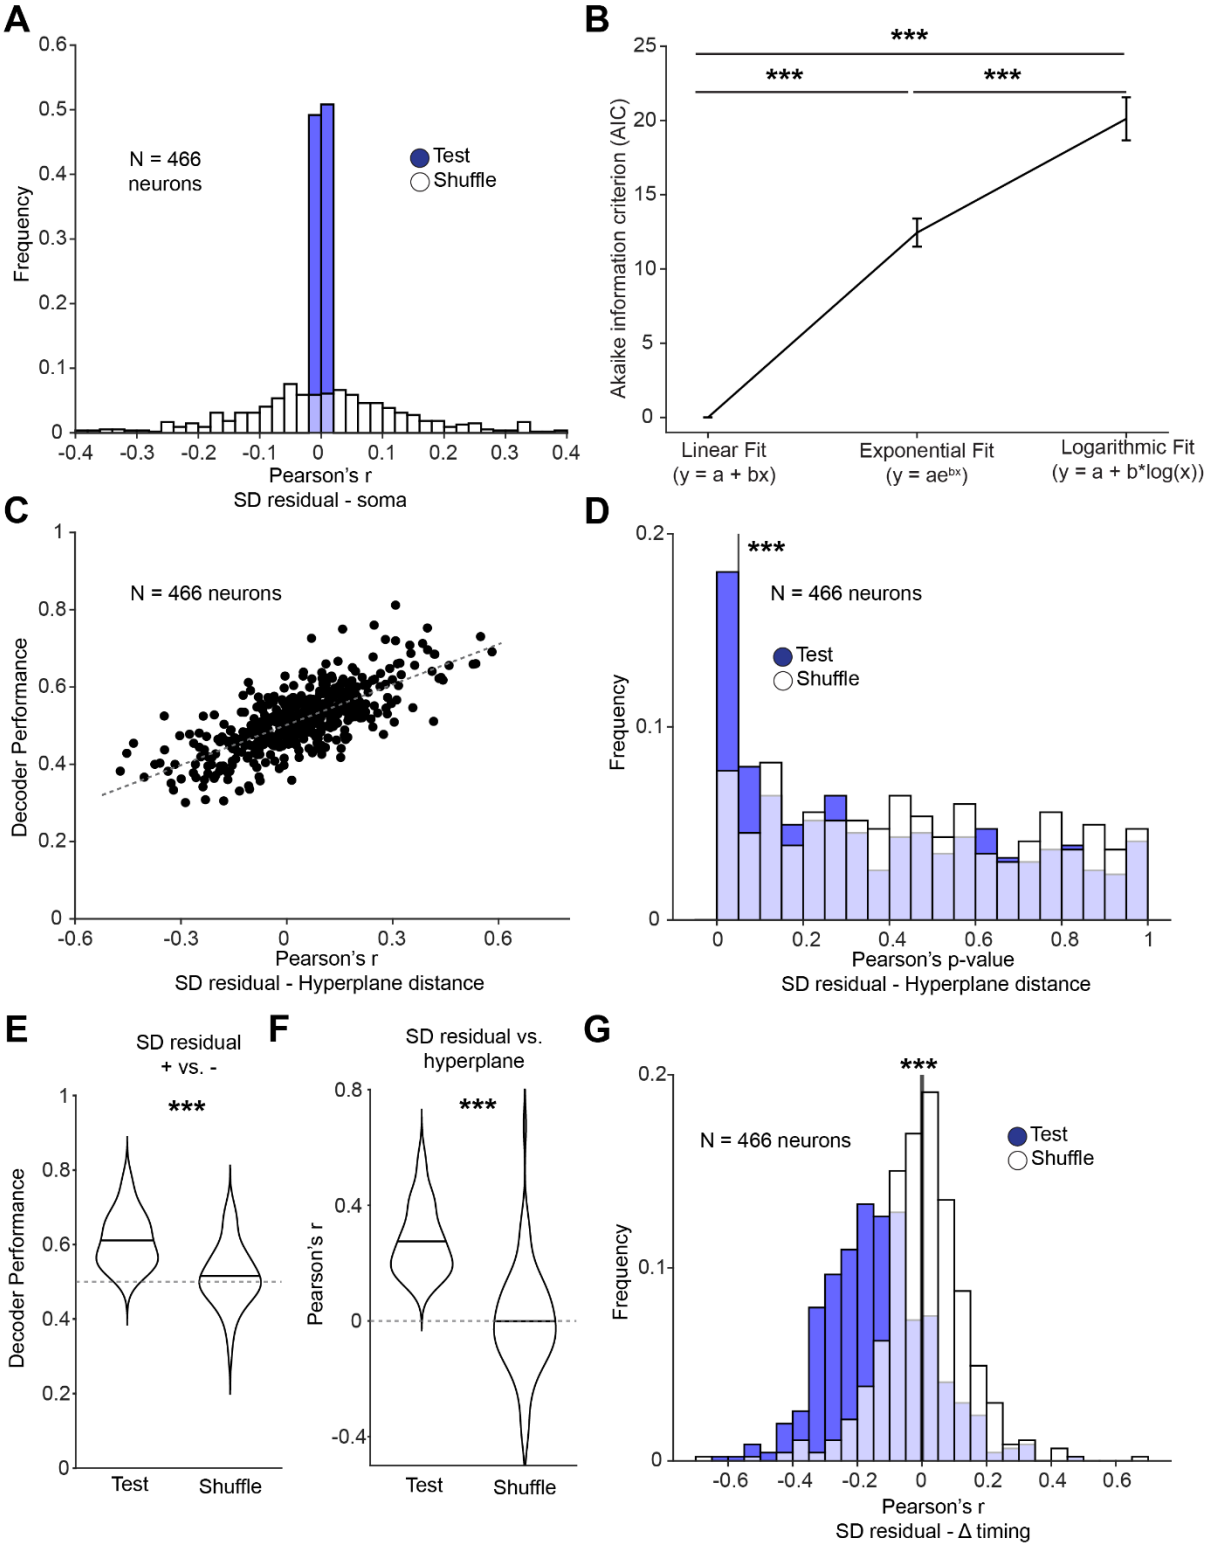
A
N = 466
neurons
Test
Shuffle
Frequency
Pearson's r
SD residual - soma
B
***
***
***
Akaike information criterion (AIC)
Linear Fit
(y = a + bx)
Exponential Fit
(y = ae^bx)
Logarithmic Fit
(y = a + b*log(x))
C
N = 466 neurons
Decoder Performance
Pearson's r
SD residual - Hyperplane distance
D
***
N = 466 neurons
Test
Shuffle
Frequency
Pearson's p-value
SD residual - Hyperplane distance
E
SD residual
+ vs. -
***
Decoder Performance
Test
Shuffle
F
SD residual vs.
hyperplane
***
Pearson's r
Test
Shuffle
G
***
N = 466 neurons
Test
Shuffle
Frequency
Pearson's r
SD residual - Δ timing

Supplementary Figure 10: Decoding somato-dendritic residuals using $\Delta F/F_0$-based event size estimation.

a, Same as in Figure 2i. Our $\Delta F/F_0$-based approach to estimate event size (see Methods and Supplementary Figure 2) decorrelated the SD residual from somatic magnitude. **b,** Same as Supplementary Figure 1. Linear models best describe the relationship between somatic and dendritic event magnitudes. For all neurons, we calculated the Akaike information criteria (AIC). Lower values of AIC mean a better fit. AIC values were zeroed compared to the linear fit model. Error bars represent standard error of the mean (One-way repeated measures ANOVA; $p = 3e^{-10}$, $p = 1e^{-10}$, $p = 1e^{-10}$ for Linear vs Exponential, Linear vs Logarithmic and Exponential vs. Logarithmic after Tukey's multiple comparison correction. Mean = 0, 12.4 and 20.1; SEM = 0, 0.9 and 1.4; n = 543 neurons for Linear, Exponential and logarithmic fit, respectively). **c,** Same as Figure 2g. Pearson correlation between decoder performance (linear SVM) and the correlation between SD residual and the distance from the hyperplane (or SVM classification confidence). Each dot represents one neuron (Pearson correlation = 0.73. $p = 2.5e^{-78}$). **d,** Same as Figure 2h, the proportion of neurons with a statistically-significant (alpha = 0.05) correlation between the SD residual and the distance from the hyperplane (or classification confidence, Wilcoxon signed rank test = $p = 7.3e^{-4}$. N = 466 neurons). **e,** Same as Figure 2j. Decoder performance (linear SVM) for statistically-significant neurons in d (paired t-test, $p = 8.3e^{-10}$. Mean = 0.61 and 0.51; SEM = 0.008 and 0.01; n = 66). **f,** Same as Figure 2k. Pearson correlation between the SD residual and the distance from hyperplane for statistically-significant neurons (paired t-test, $p = 1.6e^{-15}$. Mean = 0.27 and $-9.6e^{-4}$; SEM = 0.01 and 0.02; n = 66). **g,** Same as Figure 2m. Pearson correlation between relative event timing in soma and dendrites and SD residual (residual from best fit). A negative correlation means that the larger the residual (more dendritically-amplified) the earlier the peak timing in the dendrite compared to the soma, paired t-test, $p = 9.7e^{-40}$. Mean = 0.13 and 0.005; SEM = 0.007 and 0.006; n = 466).

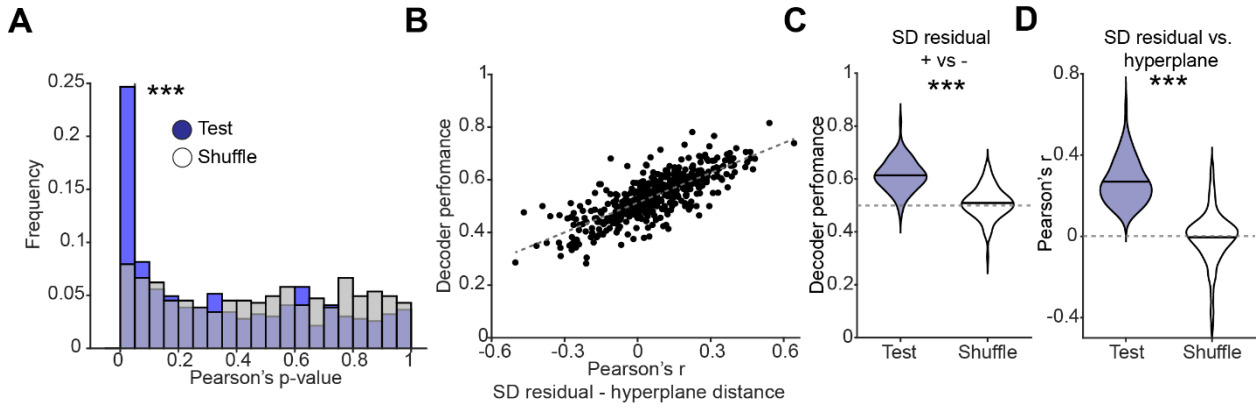

Supplementary Figure 11: Decoding SD residuals using network activity in a single frame preceding event onset.

a, Decoder performance as a function of the correlation between SD residuals and hyperplane distance (Pearson's r = 0.77; p-value = $7.1e^{-94}$, n = 466 neurons). Datapoints represents individual neurons. **b,** Distribution of p-values for test data and a control randomly shuffled distribution, testing the correlation between SD residuals and distance from the hyperplane (or classification confidence, Wilcoxon signed rank test = p = $2.8e^{-9}$. N = 466 neurons). **c,** Decoding performance for neurons with a statistically significant correlation between SD residual and distance from the hyperplane (paired t-test, p = $8.6e^{-28}$. Mean = 0.61 and 0.49; SEM = 0.006 and 0.007; n = 99). Dashed grey line indicates chance level. **d,** Pearson's r for neurons with a statistically-significant correlation between SD residual and the distance between population vector and hyperplane (paired t-test, p = $7.9e^{-30}$. Mean = 0.27 and -0.005; SEM = 0.01 and 0.01; n = 99).

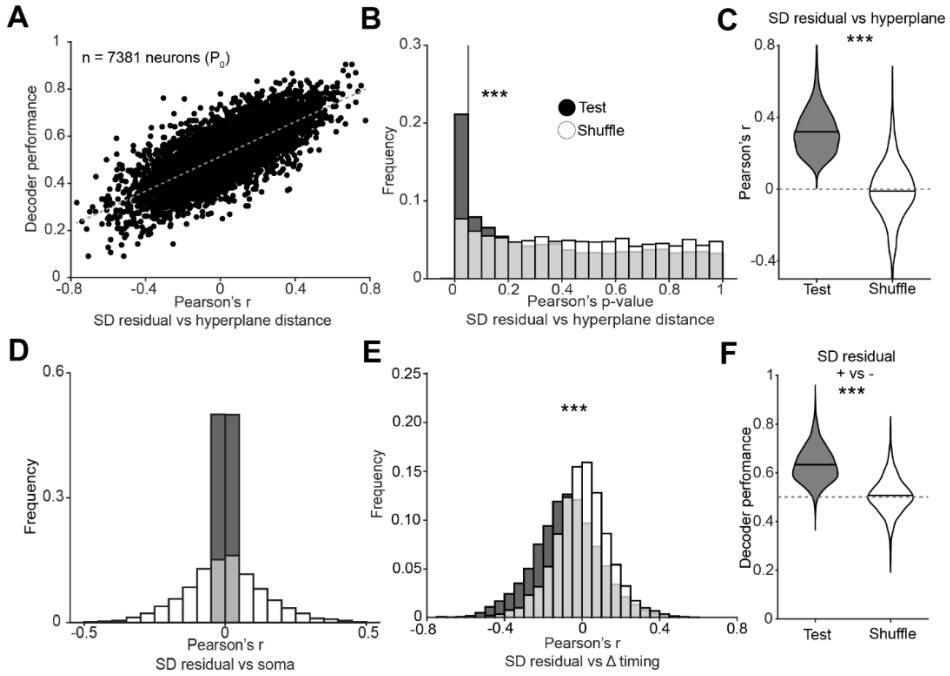

Supplementary Figure 12: Differences in somatic and dendritic magnitudes for coincident events are predicted by local network dynamics in P_0.

a, Decoder performance as a function of the correlation between SD residuals and hyperplane distance (Pearson's r = 0.75; p-value < $2.2e^{-308}$ (smallest computable value using Matlab 2020a), n = 7381 neurons). Datapoints represents individual neurons. **b,** Distribution of p-values for test data and a control randomly shuffled distribution, testing the correlation between SD residuals and distance from the hyperplane (or classification confidence, Wilcoxon signed rank test = p = $5e^{-96}$. N = 7381 neurons). **c,** Pearson's r for neurons with a statistically-significant correlation between SD residual and the distance between population vector and hyperplane (paired t-test, < $2.2e^{-308}$ (smallest computable value using Matlab 2020a). Mean = 0.31 and -0.01; SEM = 0.003 and 0.005). **d,** For all neurons, the Pearson's r for somato-dendritic residuals and somatic event magnitude. **e,** Pearson correlation value between the SD residual the event latency between soma and its corresponding dendrite indicating that the larger the SD residual, the earlier the dendritic peak is compared to the somatic one (paired t-test, p = $1.48e^{-169}$. Mean = -0.077 and -$6.9e^{-4}$; SEM = 0.002 and 0.001; n = 7381 neurons). **f,** Decoding performance for neurons with a statistically significant correlation between SD residual and distance from the hyperplane (paired t-test, p = $1.1e^{-254}$. Mean = 0.63 and 0.50; SEM = 0.002 and 0.002). Dashed grey line indicates chance level.

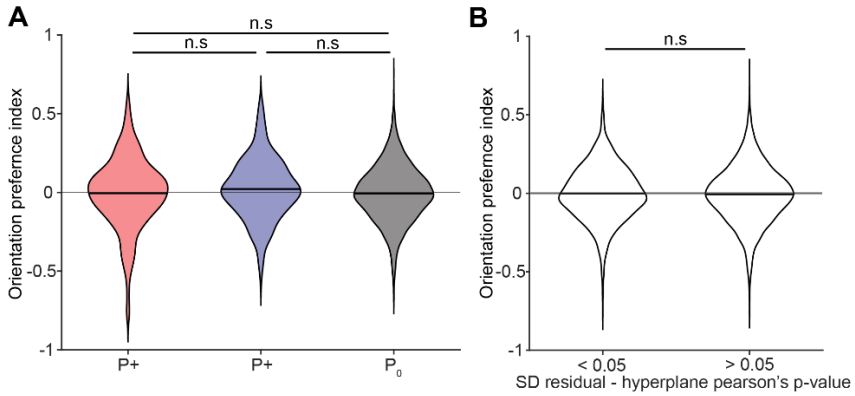

Supplementary Figure 13: Orientation preferences of decoded neurons.

a, Orientation preference index, defined as $(R_{90} – R_0)/ (R_{90} + R_0)$ where R_{90} and R_0 are the inferred spike rates at 90- and 0-degrees angles, respectively during passive visual stimulus presentation for P+, P- and P_0 (One-way ANOVA, p = 0.09; After Tukey's multiple comparisons p = 0.25, 0.98 and 0.07 for P+ vs P-, P+ vs P_0 and P- vs P_0, respectively. Mean = -0.005, 0.02 and -0.007; SEM = 0.01, 0.01 and 002; n = 268, 198 and 7381 for P+, P- and P_0, respectively). **b,** Orientation preference index for all neurons with a statistically-significant Pearson's correlation between the SD residual and the network distance from the maximally-separating hyperplane (t-test, p = 0.31. Mean = -0.002 and -0.007; SEM = 0.004 and 0.002; n = 1298 and 6547 for significant and non-significant neurons, respectively).

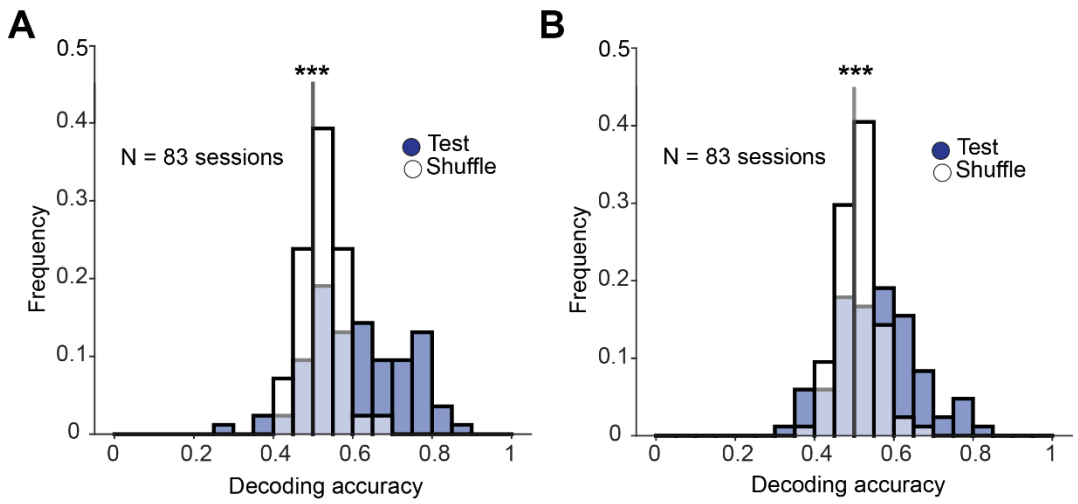

Supplementary Figure 14: Decoding task related variables using $\Delta F/F_0$-based event size estimation.

a, b, Same as Figure 3e (left panel, paired t-test, $p = 4.3e^{-9}$. Mean = 0.61 and 0.52; SEM = 0.01 and 0.01, respectively for test and shuffle data. N = 83 sessions) and 3g (right panel, paired t-test, $p = 2.4e^{-4}$. Mean = 0.55 and 0.50; SEM = 0.01 and 0.01, respectively for test and shuffle data. N = 83 sessions) respectively, using a $\Delta F/F_0$-based metrics for estimating the size of the SD residual.

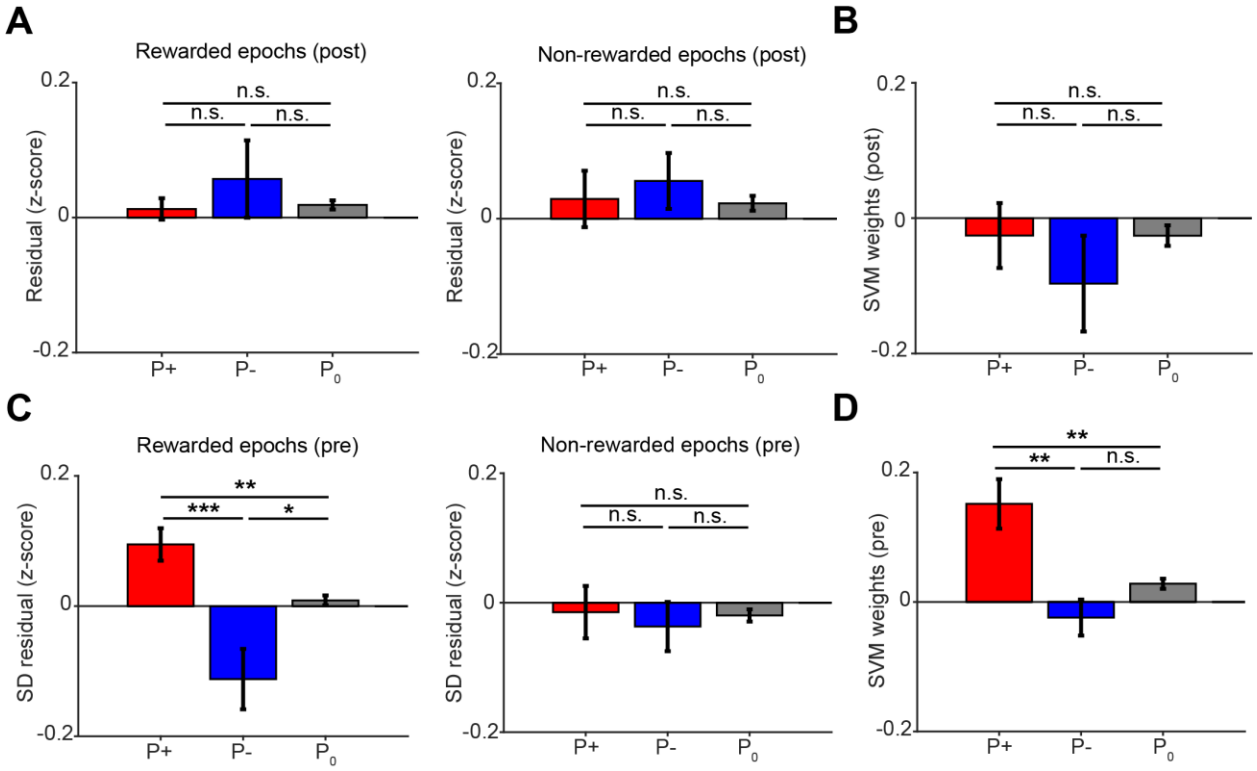

Supplementary Figure 15: Dendritic reward signals are cell-specific during periods of reward anticipation but not during reward consumption.

a, SD residuals in the two seconds following target reach for rewarded trials (left) and the end of a trial for rewarded and unrewarded trials (right), respectively, for P+, P- and P_0 neurons (rewarded epochs, One-way repeated measures ANOVA, p = 0.5. After Tukey's multiple comparison p = 0.75, p = 0.93 and p = 78 for P+ vs P-, P+ vs P_0 and P- vs P_0, respectively. Mean = 0.01, 0.06 and 0.02; SEM = 0.02, 0.06 and 0.01 for P+, P- and P_0, respectively. Unrewarded epochs, One-way repeated measures ANOVA, p = 0.47. After Tukey's multiple comparison p = 0.65, p = 0.99 and p = 0.57 for P+ vs P-, P+ vs P_0 and P- vs P_0, respectively. Mean = 0.03, 0.06 and 0.02, SEM = 0.04, 0.04 and 0.01 for P+, P- and P_0, respectively. n = 84 sessions). **b,** SVM weights for decoding rewarded vs unrewarded trials after the end of a trial for P+, P- and P_0 neurons. Here, the rewarded side of the hyperplane was arbitrarily assigned as the positive side, while the unrewarded one was assigned as the negative one (One-way repeated measures ANOVA, p = 0.72. After Tukey's multiple comparison p = 0.87, p = 0.99 and p = 0.71 for P+ vs P-, P+ vs P_0 and P- vs P_0, respectively. Mean = -0.03, -0.1 and -0.03; SEM = 0.05, 0.07 and 0.2 for P+, P- and P_0, respectively. n = 84 sessions). **c,** SD residual in the two seconds preceding target reach for rewarded trials (left) and in the two seconds preceding the end of a trial for unrewarded trials (right), for P+, P- and P_0 neurons (rewarded epochs, One-way repeated measures ANOVA,

$p = 2e^{-4}$. After Tukey's multiple comparison $p = 5.6e^{-4}$, $p = 4.8e^{-3}$ and $p = 0.035$ for P+ vs P-, P+ vs P_0 and P- vs P_0, respectively. Mean = 0.09, -0.11 and 0.01, SEM = 0.03, 0.05 and 0.01 for P+, P- and P_0, respectively. Unrewarded epochs, One-way repeated measures ANOVA, $p = 0.83$. After Tukey's multiple comparison $p = 0.92$, $p = 0.99$ and $p = 0.9$ for P+ vs P-, P+ vs P_0 and P- vs P_0, respectively. Mean = -0.01, -0.04 and -0.01, SEM = 0.04, 0.04 and 0.01 for P+, P- and P_0, respectively. n = 84 sessions). **d,** SVM weights for decoding rewarded vs unrewarded trials for P+, P- and P_0 neurons. As in b, the rewarded side of the hyperplane was arbitrarily assigned as the positive side, while the unrewarded one was assigned as the negative one (One-way repeated measures ANOVA, $p = 4.1e^{-4}$. After Tukey's multiple comparison $p = 3.8e^{-3}$, $p = 5.6e^{-3}$ and $p = 0.18$ for P+ vs P-, P+ vs P_0 and P- vs P_0, respectively. Mean = 0.15, -0.02 and 0.03, SEM = 0.04, 0.03 and 0.01 for P+, P- and P_0, respectively. n = 84 sessions).

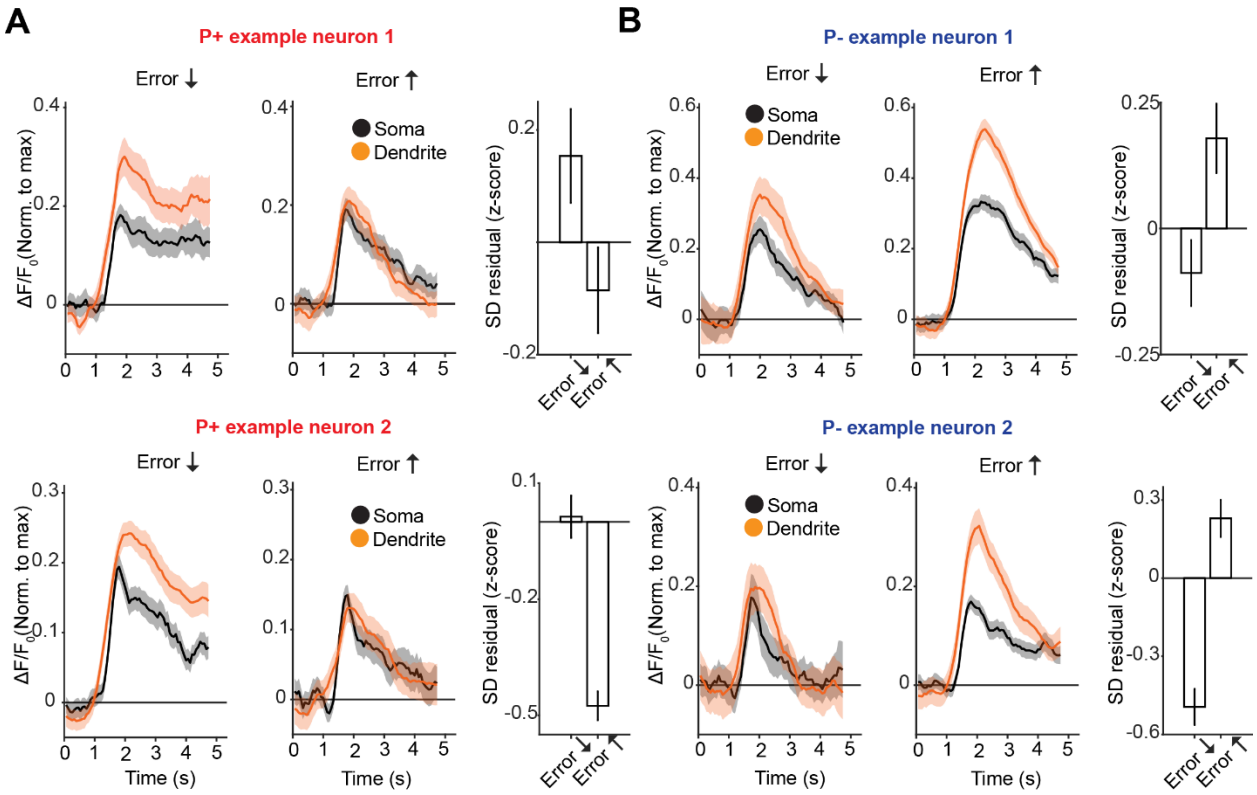

Supplementary Figure 16: Four additional example neurons (2 P+ and 2 P-) showing different somato-dendritic relationships during error decrease and error increase epochs.

For two P+ and two P- neurons, mean ΔF/F0 signal in the soma (black) and in the dendrite (orange) are shown for all events that occurred during epochs of error reduction and error increase, respectively. The bar graph represents the mean SD residual value (z-scored) for all events occurred during error decrease and increase epochs. Shaded areas and error bars represent SEM.

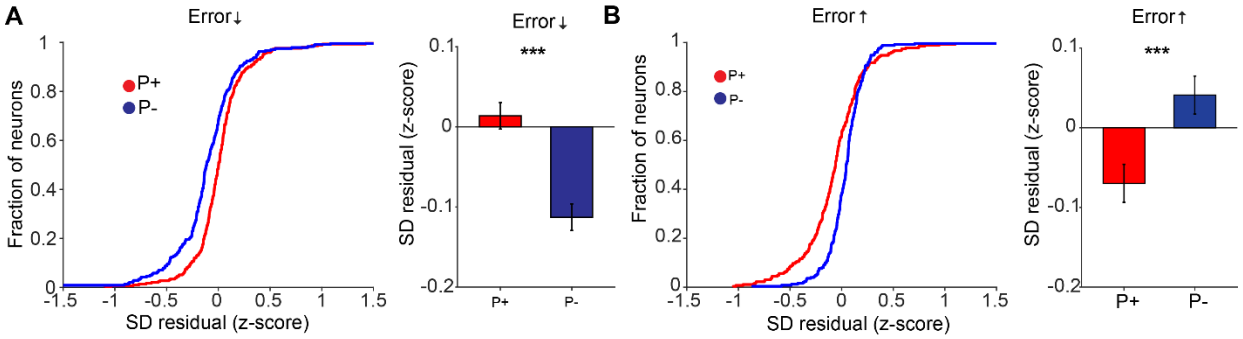

Supplementary Figure 17: Neuron specific error signals using our $\Delta F/F_0$ metric to estimate event size.

a, Same as Figure 4c. During error reduction epochs, dendrites of P+ neurons are relatively amplified compared to the dendrites of P- neurons (t-test, $p = 1.4e^{-5}$; Mean = 0.01 and -0.11; SEM = 0.02 and 0.02; n = 292 and 240 for P+ and P- neurons, respectively). **b,** Same as Figure 4d. During error reduction epochs, dendrites of P+ neurons are relatively attenuated compared to the dendrites of P- neurons (t-test, $p = 3.1e^{-6}$; Mean = -0.07 and 0.04; SEM = 0.02 and 0.01; n = 267 and 249 for P+ and P- neurons, respectively).

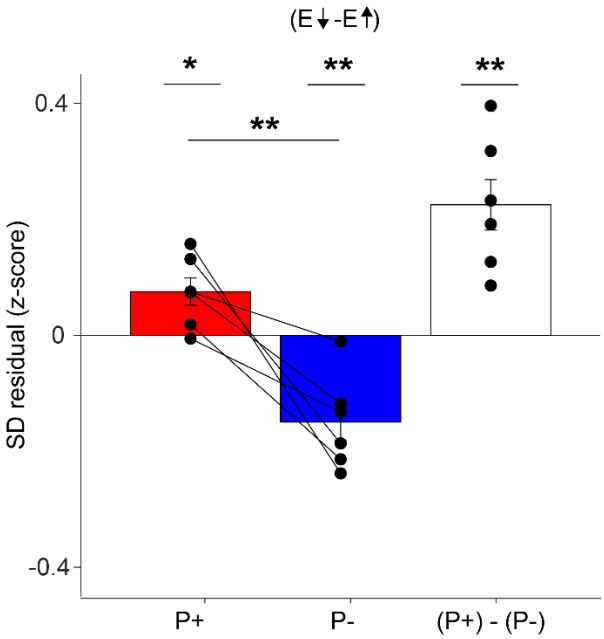

Supplementary Figure 18: Dendritic signals are consistent across all mice.

For all mice (N = 6) the separability of dendritic signals, defined as the SD residual during epochs of error decrease minus the SD residual during epochs of error increase, for P+ (in red), P- (in blue), and overall dendritic separability ((P+) – (P-)) (Paired t-test, p = 0.03, p = 0.006 and p = 0.005; Mean = 0.08, -0.15 and 0.23; SEM = 0.03, 0.03 and 0.05 for P+, P- and (P+) – (P-), respectively. N = 6 animals. Note that comparing P+ vs. P- or (P+) – (P-) vs. 0 is statistically equivalent).

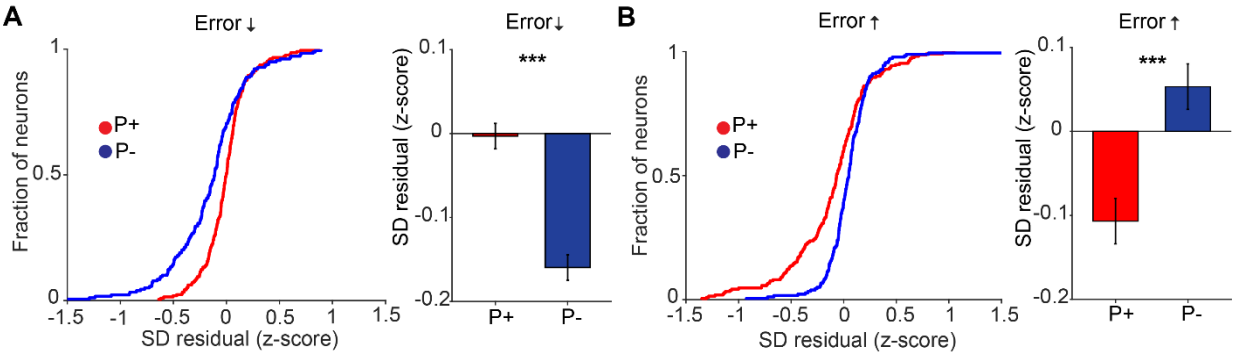

Supplementary Figure 19: Dendritic error signals cannot be explained away by differences in somatic event magnitude.

a, b, The cumulative distribution function for SD residuals (z-scored per neuron) in P+ (red) and P- (blue) neurons during error decrease in a (t-test, $p = 3.5e^{-6}$, mean = -0.003 and -0.16; SEM = 0.015 and 0.03; n = 211 and 179 for P+ and P-) and error increase in b (t-test, $p = 4.6e^{-6}$, mean -0.1 and 0.05; SEM = 0.03 and 0.02; n = 211 and 179 for P+ and P- neurons) epochs. Only neurons with the same somatic response to error decrease and increase were selected for this analysis.

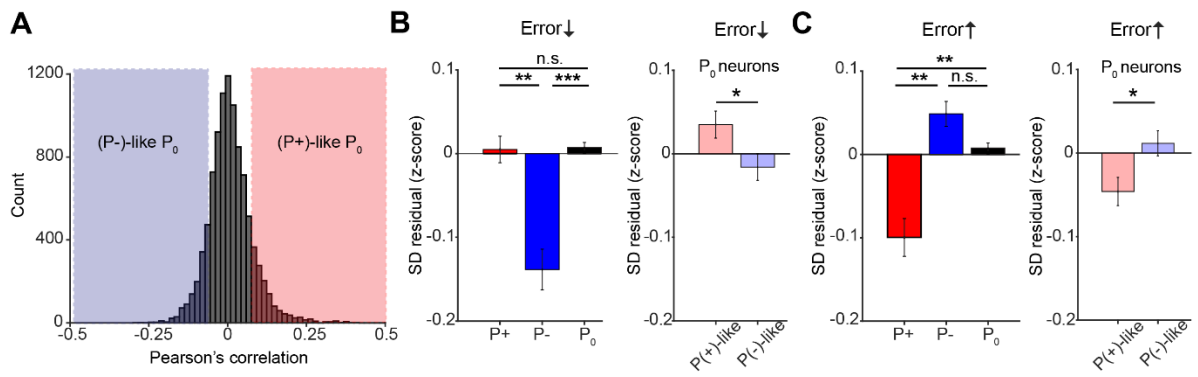

Supplementary Figure 20: P_0 neurons which are functionally-correlated to P+ and P- neurons also receive vectorized dendritic error signals.

a, Across all sessions, we divided P_0 neurons into (P+)- and P(-)-like based on their correlation to the (P+)-(P-) subtracted signal we used to rotate the visual stimulus. Among all P_0 neurons (n = 9796) we defined (P+)- and (P-)-like P_0 neurons as those with a correlation above 1 and below -1 standard deviation. **b,** Left panel, during error reduction epochs, bar graph comparing the SD residual (z-scored) in P+, P- and P_0 neurons (One-way ANOVA, p = $1.6e^{-4}$. After Tukey's multiple comparisons, p = 0.006, p = 1 and p = $8.9e^{-5}$ for P+ vs P-, P+ vs P_0 and P- vs P_0, respectively. Mean = 0.004, -0.13 and 0.007; SEM = 0.02, 0.02 and 0.01; n = 292, 240 and 8335 for P+, P- and P_0, respectively). Right panel, during error reduction epochs, bar graph comparing the SD residual (z-scored) in (P+)- and P(-)-like P_0 neurons (t-test, p = 0.025; mean = 0.04 and -0.02; SEM = 0.02 and 0.02; n = 1059 and 1110 for (P+)- and P(-)-like P_0 neurons, respectively). **c,** Left panel, during error increase epochs, bar graph comparing the SD residual (z-scored) in P+, P- and P_0 neurons (One-way ANOVA, p = 0.002. After Tukey's multiple comparisons, p = 0.004, p = 0.45 and p = 0.003 for P+ vs P-, P+ vs P_0 and P- vs P_0, respectively. Mean = -0.1, 0.05 and 0.01; SEM = 0.02, 0.02 and 0.01; n = 267, 249 and 8356 for P+, P- and P_0, respectively). Right panel, during error increase epochs, bar graph comparing the SD residual (z-scored) in (P+)- and P(-)-like P_0 neurons (t-test, p = 0.010; mean = -0.05 and 0.01; SEM = 0.02 and 0.01; n = 1070 and 1099 for (P+)- and P(-)-like P_0 neurons, respectively).

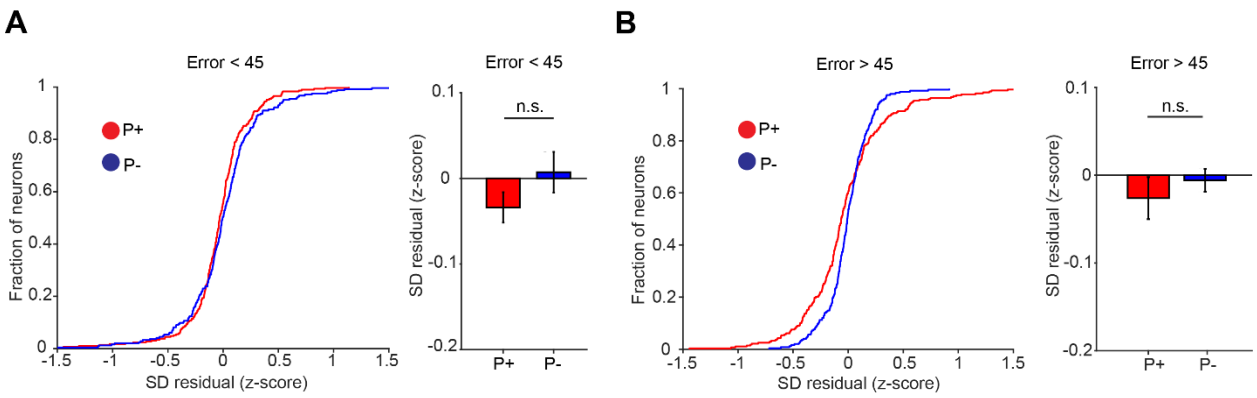

Supplementary Figure 21: Dendritic signals do not represent absolute error values.

a, The cumulative distribution function for SD residuals (z-scored per neuron) in P+ (red) and P- (blue) neurons during epochs where the absolute error was smaller than 45-degrees (t-test, p = 0.16; mean = -0.03 and 0.07; SEM = 0.017 and 0.02; n = 290 and 244 for P+ and P- neurons, respectively). **b**, Same as a, but for epochs in which absolute error was larger than 45-degrees (t-test, p = 0.47; mean -0.03 and -0.006; SEM = 0.02 and 0.01; n = 278 and 249 for P+ and P- neurons, respectively).