

# Vectorized instructive signals in cortical dendrites

<https://doi.org/10.1038/s41586-026-10190-7>

Received: 2 December 2022

Accepted: 23 January 2026

Published online: 25 February 2026

Open access

 Check for updates

Valerio Francioni<sup>1,2</sup>, Vincent D. Tang<sup>1,2,5</sup>, Enrique H. S. Toloza<sup>1,3,4,5</sup>, Zilan Ding<sup>1,2,5</sup>, Norma J. Brown<sup>1,2</sup> & Mark T. Harnett<sup>1,2</sup>✉

Vectorization of teaching signals is a key element of almost all modern machine learning algorithms, including backpropagation, target propagation and reinforcement learning. Vectorization allows a scalable and computationally efficient solution to the credit assignment problem by tailoring instructive signals to individual neurons. Recent theoretical models have suggested that neural circuits could implement single-phase vectorized learning at the cellular level by processing feedforward and feedback information streams in separate dendritic compartments<sup>1–5</sup>. This presents a compelling, but untested, hypothesis for how cortical circuits could solve credit assignment in the brain. Here we used a neurofeedback brain–computer interface task with an experimenter-defined reward function to test for vectorized instructive signals in dendrites. We trained mice to modulate the activity of two spatially intermingled populations (four or five 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 and 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 depended on the causal role of individual neurons in the task and predicted changes in overall activity over the course of learning. Furthermore, targeted optogenetic perturbation of these signals disrupted learning. These results demonstrate a vectorized instructive signal in the brain, implemented via semi-independent computation in cortical dendrites, unveiling a potential mechanism for solving credit assignment in the brain.

Learning is the product of changes in the strength of synaptic connections between neurons<sup>6–13</sup>. Synaptic modifications can have difficult-to-predict effects on network output, particularly in complex hierarchical networks such as the brain. The challenge of determining how individual synapses should be altered to improve task performance is known as the credit assignment problem<sup>14–18</sup>. Whereas this problem is effectively solved in artificial neural networks (ANNs) by the backpropagation-of-error algorithm<sup>19</sup>, how credit assignment is solved in the brain remains unknown<sup>14,15</sup>.

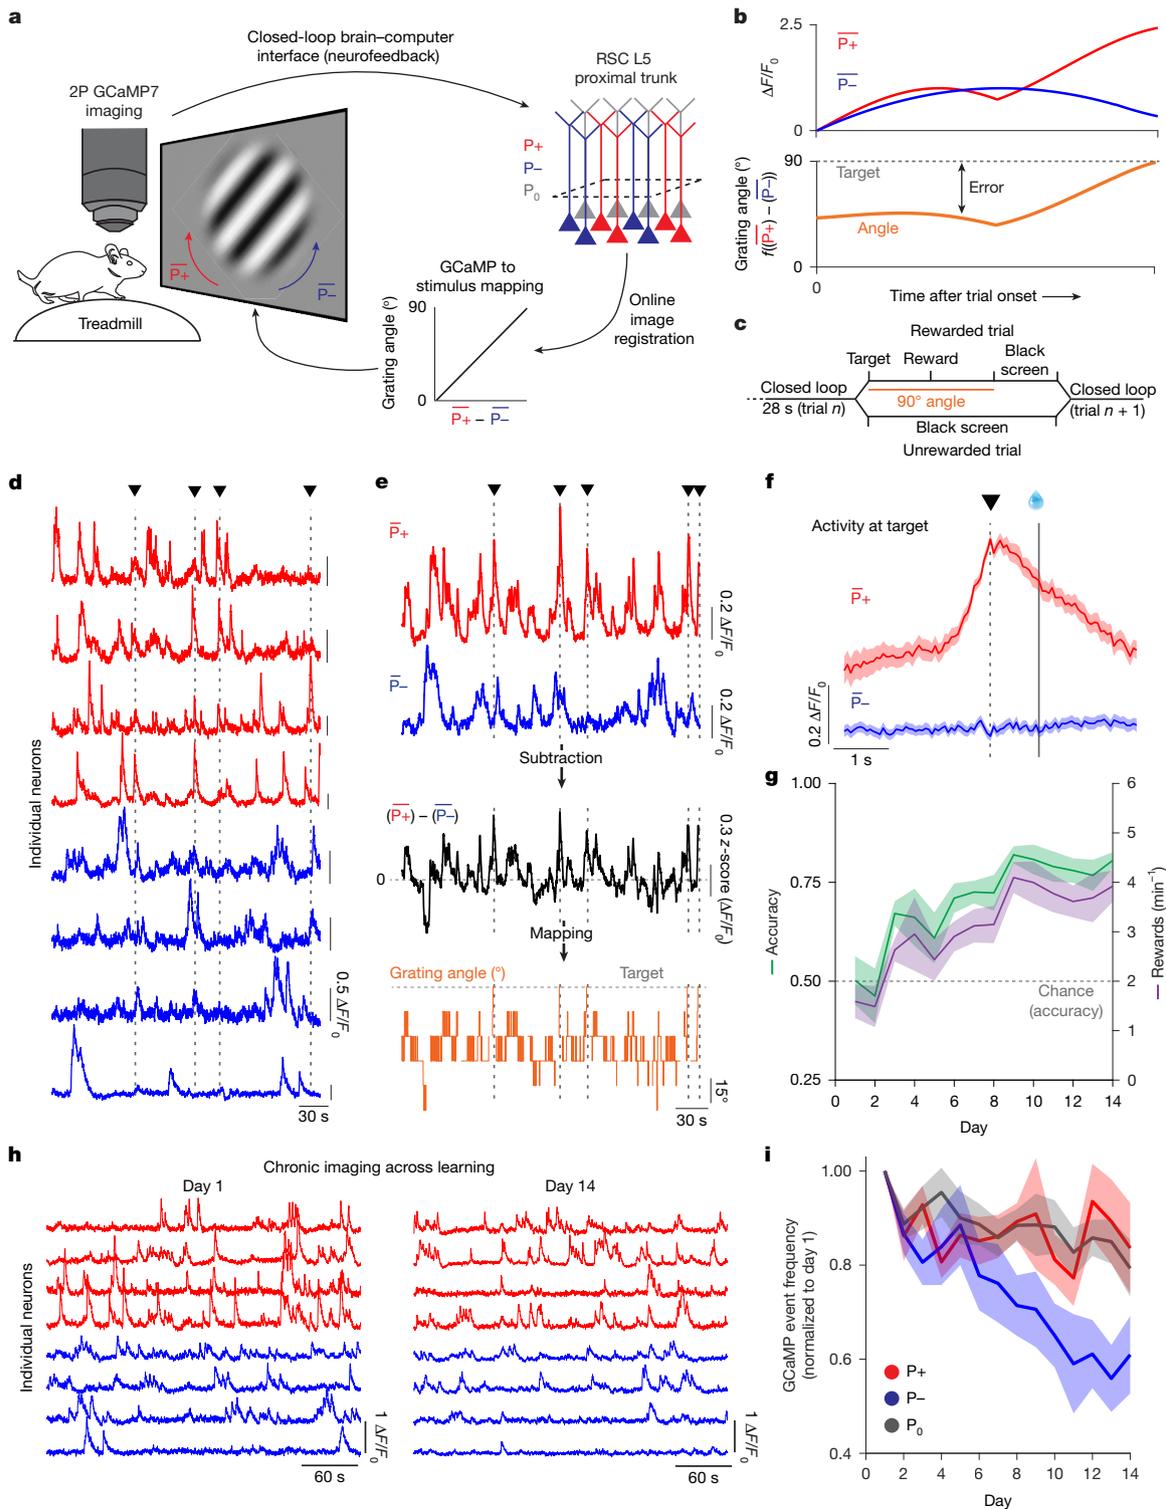
Recent theoretical work has proposed several models by which biological circuits could solve credit assignment, including target learning and backpropagation-like algorithms<sup>1–5,20,21</sup>. Central to both artificial and biologically inspired solutions to credit assignment is the vectorization of instructive signals, as opposed to the broadcasting of a single scalar teaching signal<sup>14</sup>. Effective learning requires, in addition to vectorization, instructive signals to be separable from feedforward inputs to prevent interference<sup>15</sup>. In ANNs, this is achieved via temporal separation, which has long been thought to be biologically implausible.

One hypothesis is that in cortex, credit-related information is spatially, rather than temporally, segregated in the apical dendrites of pyramidal neurons<sup>15</sup>. This aligns with anatomical and circuit evidence that feedforward inputs are received perisomatically and feedback inputs are received in the distal dendrites<sup>22–31</sup>. However, direct evidence regarding the subcellular mechanisms of credit assignment is lacking.

Vectorized teaching signals at the dendritic level should meet four experimentally testable conditions. First, dendritic activity should contain information that is not present in somatic activity alone (although somas could theoretically transmit gradients using qualitatively different spiking patterns<sup>2,4,32</sup>, 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, dendritic activity should reflect the contribution of that neuron to task performance (that is, the reward function). Fourth, disrupting vectorized instructive dendritic signals should impair learning.

<sup>1</sup>McGovern Institute for Brain Research, MIT, Cambridge, MA, USA. <sup>2</sup>Department of Brain and Cognitive Sciences, MIT, Cambridge, MA, USA. <sup>3</sup>Department of Physics, MIT, Cambridge, MA, USA.

<sup>4</sup>Harvard Medical School, Boston, MA, USA. <sup>5</sup>These authors contributed equally: Vincent D. Tang, Enrique H. S. Toloza, Zilan Ding. ✉e-mail: harnett@mit.edu



**Fig. 1** | See next page for caption.

### Specifying a reward function using a BCI task

Evaluating credit assignment in biological neural networks has thus far proved impossible<sup>14,15</sup>. Teaching signals can only be defined relative to a reward function that maps neural activity to task performance. It is unclear whether such functions are explicitly represented in the brain. Even if they are, experimenters are blind to their specific formulation in terms of neural activity<sup>15</sup>. Neurofeedback brain-computer interface (BCI) tasks present a potential solution to this problem by

directly coupling neural activity to task performance, thereby allowing the experimenter to specify the reward function to be optimized<sup>14,20,21</sup>. Previous studies have shown that mice are able to learn BCI tasks using a variety of feedback stimuli and brain areas and that learning induces changes in the activity of the neurons controlling the BCI, including in the hippocampus and various sensory and motor cortices<sup>33-39</sup>. Here we leveraged a visually guided neurofeedback BCI task in cortical pyramidal neurons to test subcellular mechanisms for error and reward-related signalling (Fig. 1a-c and Supplementary Figs. 1 and 2). We trained head-fixed

**Fig. 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 (2P) microscope and free to run on a cylindrical treadmill. Two user-defined populations of GCaMP7f-labelled layer 5 (L5) pyramidal neurons in RSC were imaged at the proximal apical trunk: P+ (red) and P- (blue) were selected to control the rotation of a Gabor patch. P<sub>0</sub> 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°. Activity in P- neurons rotated the Gabor patch stimulus counter-clockwise, towards a 0° 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° apart, from 0° to 90°) linear mapping between the mean activity in P+ neurons minus the activity in P- neurons. **c**, Trial structure: mice had 28 s to reach target activity and receive a reward, delivered 1 s later. In successful trials, the 90° Gabor patch was shown for 2 s, followed by 1 s of black screen presentation. In unsuccessful trials, a 3 s black screen was presented before the onset of the next trial. **d**,  $\Delta F/F_0$  traces as recorded live for P+ (red) and P- (blue) neurons. Vertical dashed lines and triangles represent timepoints where the mouse reached target activity. **e**, Mean activity for the

red (P+) and blue (P-) traces shown in **d**. The black trace shows the arithmetic subtraction of P+ and P- neurons (z-scored). The 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 mouse reached target activity (dashed, vertical line and black triangle) for the session highlighted in **d,e**. Reward was delivered 1 s later (solid vertical line with water reward). Shaded areas represent s.e.m. **g**, Mean performance over days quantified as the fraction of successful trials over the total number of trials and as the number of rewards per minute (one-way repeated measures ANOVA,  $P = 5 \times 10^{-4}$  (accuracy) and  $P = 0.002$  (rewards per minute);  $n = 6$  mice). The dashed horizontal red line represents chance level for accuracy performance (Methods). Shaded areas represent s.e.m. **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<sub>0</sub> neurons 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  $P = 9.3 \times 10^{-4}$  for the effect of population identity, days and an interaction between population identity and days. After Tukey's multiple comparison,  $P = 0.027$  (P+ versus P- neurons),  $P = 0.95$  (P+ versus P<sub>0</sub> neurons) and  $P = 0.01$  (P- versus P<sub>0</sub> neurons).  $n = 6$  mice. Shaded areas represent s.e.m.

mice under a 2-photon microscope to control the activity of two spatially intermingled sets of GCaMP7f-labelled layer 5 pyramidal neurons, in the retrosplenial cortex (RSC), designated P+ and P- (selection criteria in Extended Data Figs. 1 and 4b and Methods). The difference in mean somatic GCaMP activity of P+ versus P- neurons was coupled to rotation of a visual grating relative to a rewarded target angle<sup>33–36,38,39</sup> (Fig. 1d–f and Supplementary Data Fig. 1). We selected RSC owing to the optical accessibility of layer 5 and previous demonstration of independent dendritic events in this area<sup>40</sup>. 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 reducing signal contamination owing to the more precise spatial footprint and faster signal kinetics of the apical trunk<sup>41–43</sup>. We measured task performance with two metrics: accuracy, which represented the fraction of rewarded trials; and speed, which represented the number of rewards obtained per minute. Mice ( $n = 6$ ) learned the task by both metrics (Fig. 1g and Extended Data Fig. 2 and 3).

We compared activity levels of P+ and P- populations, as well as the population of surrounding neurons that were not directly involved in the rotation of the stimulus (termed P<sub>0</sub>), 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), with P+ neurons maintaining their activity levels while P- neurons were downregulated. Whereas, on average, changes in activity in P<sub>0</sub> neurons resembled changes in P+ neurons (Fig. 1i), selecting the subpopulation of P<sub>0</sub> neuron with matching activity levels of P+ and P- neurons on day 1 revealed that changes in activity in P<sub>0</sub> neurons fell in between those of P+ and P- neurons (Extended Data Fig. 4). As the most active neurons on day 1 were also those that were most strongly downregulated (Extended Data Fig. 4c), our results are consistent with a model of learning by sparsification, an energy-efficient solution to the task<sup>44</sup>. Increases in task performance were not correlated with changes in locomotion across days (Extended Data Fig. 3). Moreover, the P+ and P- populations were spatially intermingled, and had the same GCaMP transient frequency on day 1 (Extended Data Figs. 1 and 4a), ruling out the possibility of learning the task by simply engaging a non-specific gain modulation mechanism.

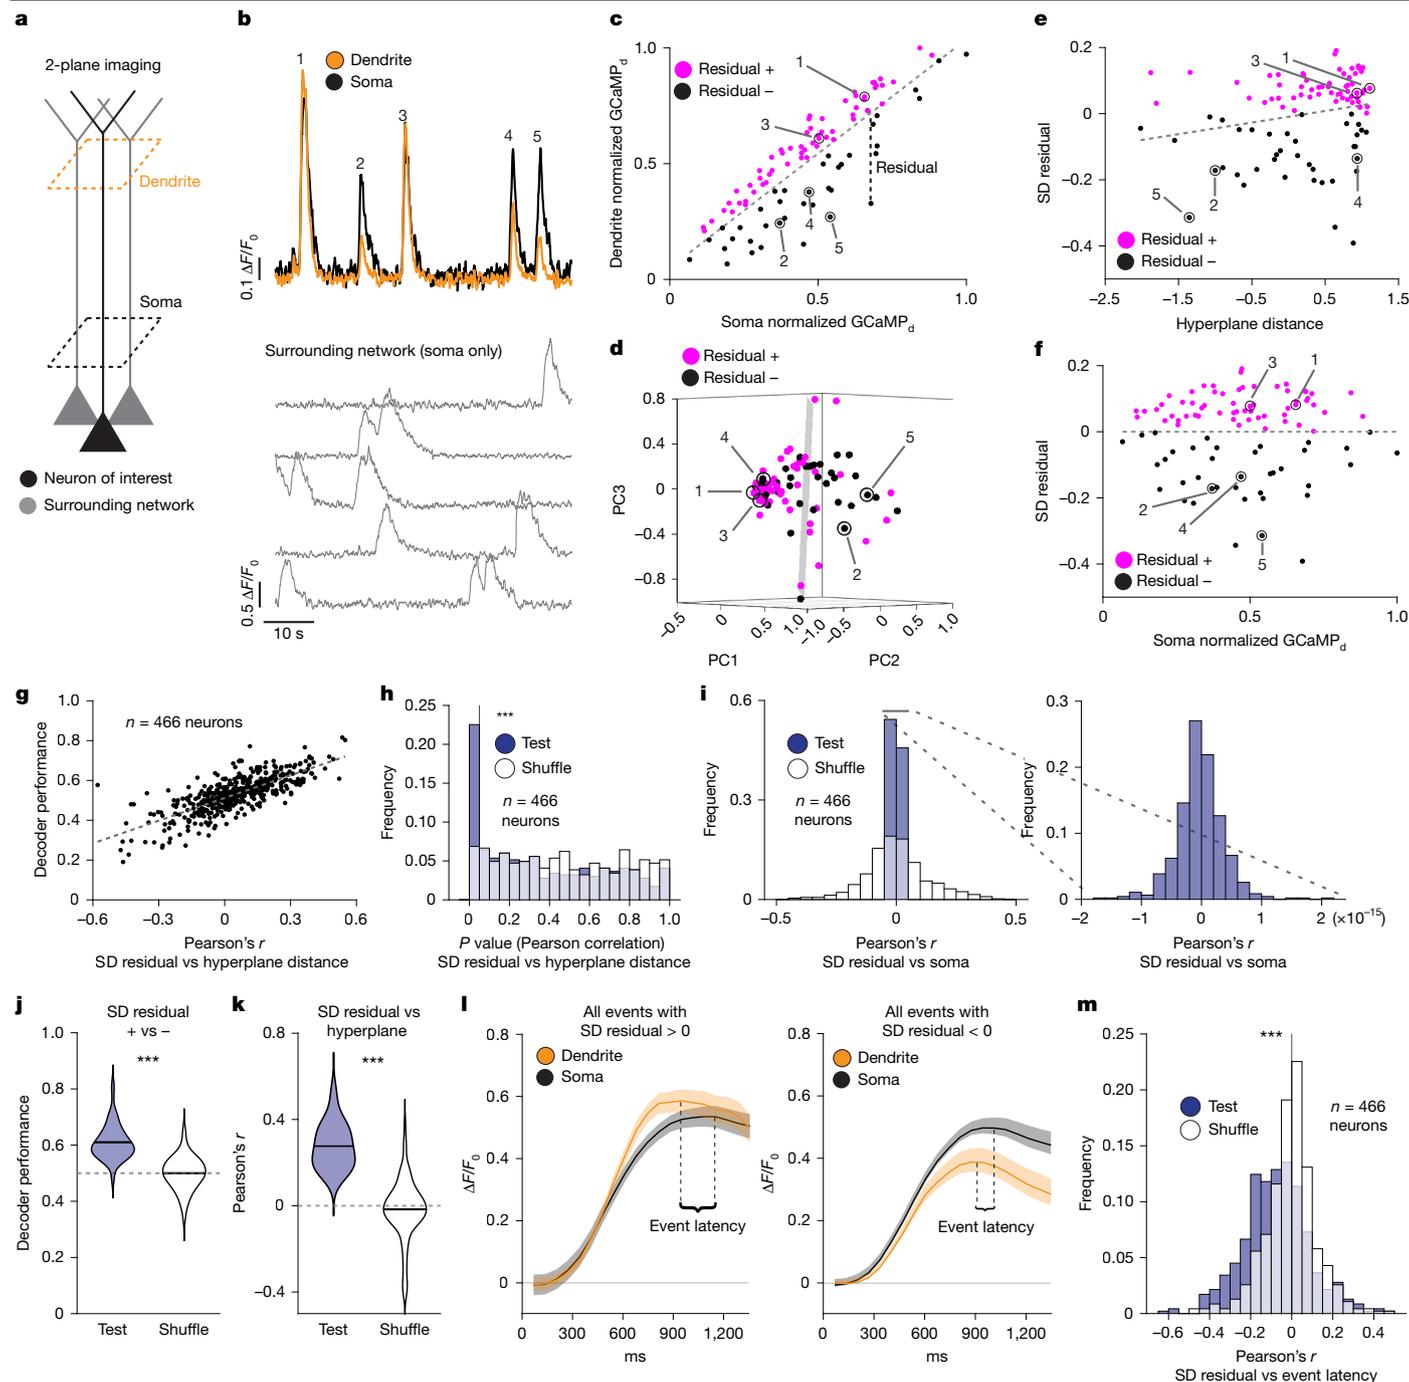
### Dendrites contain information not found in their somas

To determine whether apical dendritic activity contained information that was not encoded in parent somatic activity alone, we used an electrically tunable lens 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 on the basis of the Pearson correlation of their GCaMP signals, thresholded at  $r = 0.6$  as in previous studies<sup>41–43</sup>. 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<sup>41</sup> (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<sup>45</sup>. Deconvolution allowed us to correct for the well-described problem of different signal kinetics across dendritic compartments<sup>46</sup>. Next, we used an area-under-the-curve approach to quantify the magnitude of individual transients (all main results were also validated using a  $\Delta F/F_0$ -based approach to estimation magnitude of transients; Methods and Supplementary Fig. 3) and defined events as coincident whenever they occurred within 500 ms of each other. As these coincident events represent the vast majority of GCaMP transients<sup>40–43,46–52</sup>, 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 varied substantially, despite event timing correlation being very high (Fig. 2b; consistent with prior studies<sup>40,41,43,46,47,49</sup>). As event magnitudes at soma and dendrites were best described by a linear relationship (Extended Data Figs. 5 and 6b), we assessed the relative degree of dendritic amplification versus attenuation with a best-fit line through all events and then calculated the somato-dendritic residual (SD residual) associated with individual transients<sup>43</sup> (Fig. 2b,c). 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 test whether SD residuals contain information that is biologically meaningful, we used activity from all the somas in our field of view in the 2 s 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 ability of the decoder to capture the magnitude of dendritic amplification or attenuation in the classification confidence (Fig. 2e,g,h and Extended Data Figs. 6c,d and 7a,b). This was an emergent property, as the decoder was trained for binary classification only and had no information about the magnitude of dendritic



**Fig. 2** | See next page for caption.

amplification or attenuation. Among 466 neurons, approximately 20% showed a significant correlation between classification confidence and the magnitude of SD residual (Fig. 2h and Extended Data Figs. 6c, d and 7a, b). We found that in these neurons, we could accurately decode 61% of the events as being either amplified or attenuated, well above the 50% chance level (Fig. 2j and Extended Data Figs. 6e and 7c). 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 and Extended Data Figs. 6f and 7d). Of note, our analysis approach completely decorrelates somatic event magnitude from SD residuals (Fig. 2f, l and Extended Data Fig. 6a), 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<sub>0</sub> neurons could be decoded at the same level as P+ and P- neurons (Extended Data Fig. 8), and that decoding does not depend on somatic responses to visual stimuli across the three subpopulations (Extended Data Fig. 9).

We further found that dendritically amplified events consistently peaked earlier than dendritically attenuated events compared with the soma (Fig. 2l, m and Extended Data Figs. 6g and 8e), congruent with results in brain slices<sup>41</sup>.

### Experimental perturbation of SD residuals

Previous studies indicate that anaesthesia reduces top-down input and/or inhibits apical tuft dendrites in layer 5 pyramidal neurons<sup>22,53–55</sup>. We

**Fig. 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 and dendrite for a single neuron of interest (top; P+ and P- neurons across days 1–14) and corresponding activity in 5 surrounding neurons (bottom). Numbers 1–5 indicate identified GCaMP events. **c**, Relationship between integrals of somatic and dendritic transients for the example neuron shown in **b**. Data points represent individual events that were simultaneously detected in soma and dendrite (Methods). A least-squares linear model (dashed grey line) defined events as dendritically amplified (residual +) versus dendritically attenuated (residual -). Events 1–5 correspond to the transients shown in **b**. Subscript d denotes deconvolved. **d**, For each coincident event in the neuron of interest shown in **b, c**, we estimated the network activity vector in the 2 s before using all other neurons in the field of view. Here, the network activity vector was projected onto the first three 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, 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–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 three principal components (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–e**. The 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  $r = 0.74$ ;  $P = 1.4 \times 10^{-84}$ ,  $n = 466$  neurons). Data points represent 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  $P = 1.3 \times 10^{-9}$ ;  $n = 466$  neurons) as estimated in **e, i**. Left, for all neurons, the Pearson's  $r$  for SD residuals and somatic event magnitude as characterized in **f**. The residual-based approach perfectly decorrelates SD residual from somatic activity alone. Right, for test data, a zoomed in version of the same histogram shown on the left. **j**, Decoding performance for neurons with a statistically significant correlation between SD residual and distance from the hyperplane (paired  $t$ -test,  $P = 8.6 \times 10^{-9}$ ;  $0.61 \pm 0.006$  (test) and  $0.50 \pm 0.007$  (shuffle), mean  $\pm$  s.e.m.;  $n = 82$ ). 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,  $P = 3.35 \times 10^{-25}$ ;  $0.28 \pm 0.01$  (test) and  $-7.2 \times 10^{-4} \pm 0.01$  (shuffle), mean  $\pm$  s.e.m.;  $n = 82$ ). **l**, Mean  $\Delta F/F_0$  events for soma and dendrites for all dendritically amplified (left) and dendritically attenuated (right) 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. Dendritically amplified events peaked earlier compared with dendritically attenuated events. **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,  $P = 8 \times 10^{-13}$ ;  $-0.075 \pm 0.007$  (test) and  $-0.005 \pm 0.006$  (shuffle), mean  $\pm$  s.e.m.;  $n = 466$  neurons).

therefore hypothesized that the SD residual should be reduced during anaesthesia compared with wakefulness. To test this, we simultaneously recorded somatic and dendritic activity of layer 5 pyramidal neurons in RSC during these two conditions (Fig. 3a–c). Consistent with previous findings<sup>22</sup>, we observed a marked effect of anaesthesia on the frequency of GCaMP transients (Fig. 3d). For each neuron, we used all events detected during wakefulness to establish the distribution of SD residuals during awake periods. We then measured the effect of anaesthesia on the SD residual using the best-fit somato-dendritic line calculated during wakefulness. Anaesthesia strongly reduced the SD residual (Fig. 3c,e), consistent with previous observations of decreased top-down input<sup>22,55</sup>.

Prior work has also demonstrated that NDNF-positive layer 1 inhibitory interneurons can inhibit the apical dendrites of pyramidal neurons<sup>53,54</sup>. We therefore tested whether NDNF-mediated inhibition reduced SD residuals, indicative of a preferential effect on apical dendritic activity. To do so, we co-injected NDNF-Cre mice ( $n = 4$ ) with both a Cre-dependent version of ChRmine in layer 1 and GCaMP7f expressed under the control of the synapsin promoter in layer 5 (Fig. 3f,g). We then recorded somatic and dendritic GCaMP activity of individual layer 5 neurons, in the presence and absence of layer 1 NDNF<sup>+</sup> interneuron activation via an LED light (Fig. 3h and Extended Data Fig. 10). Similar to our approach during anaesthesia, we first established the control relationship between somatic and dendritic event amplitudes for each neuron and then compared this to the SD residuals of activity during optogenetic activation. NDNF<sup>+</sup> interneuron activation reduced the frequency of GCaMP transients (Fig. 3i) and, consistent with a number of previous *ex vivo* and *in vivo* studies<sup>53,54,56</sup>, strongly decreased the SD residual in individual layer 5 pyramidal neurons (Fig. 3h,j). LED illumination alone (conducted in a control cohort of five mice) was not responsible for these results (Extended Data Fig. 10). Together, these results demonstrate that SD residual is predictably affected by two independent experimental manipulations *in vivo*, establishing it as a robust metric of dendritic versus somatic activity.

### SD residuals decode reward and trial outcome

Next, we evaluated whether the SD residual contained information about task-related variables that could serve as putative teaching

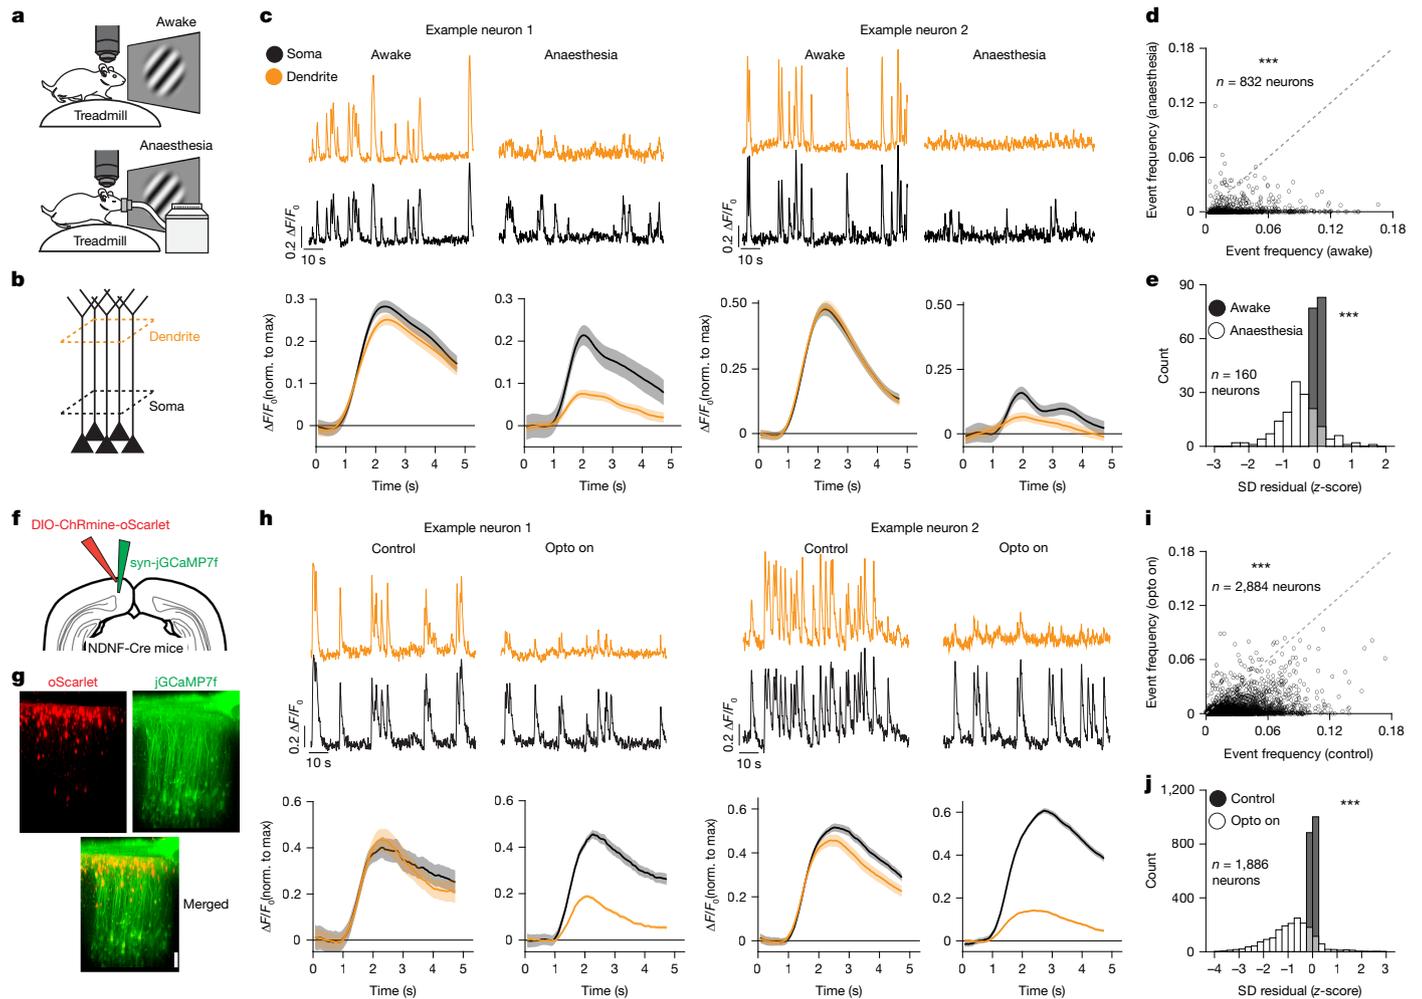
signals. We first tested whether changes in SD residual at the population level contained reward-related information. For each imaging session, we decoded rewarded versus unrewarded trials by comparing the 2 s following neural activity reaching target activation on rewarded trials with the analogous 2 s timeout period during unrewarded trials (Fig. 4a–c). Using a linear SVM trained on SD residuals (Methods), we were able to decode at 63% accuracy on average, above both chance and shuffle performance (Fig. 4d,e and Extended Data Figs. 6 and 11a,b).

Next, we tested whether 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. 4c,f,g and Extended Data Figs. 6 and 11c,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. As the trial time we analysed is pre-outcome, our results indicate that the SD residuals encode instructive signals based on the task-associated reward function.

Finally, we tested the role of layer 1 inhibition in controlling dendritic signals encoding reward and trial outcome (Fig. 4h). To do this, we performed experiments on a second set of four mice expressing ChRmine in NDNF<sup>+</sup> layer 1 interneurons. Optogenetic activation of layer 1 NDNF<sup>+</sup> neurons, but not LED illumination alone (conducted in a control group of five mice), abolished task and reward-related information in the apical dendrites of layer 5 pyramidal neurons (Fig. 4i–l and Extended Data Fig. 10), highlighting a potential role for local cortical inhibition in dendritic processing of task-related variables.

### SD 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 each neuron's causal role in the task (Fig. 5a,b). We reasoned that a scalar error signal would manifest as

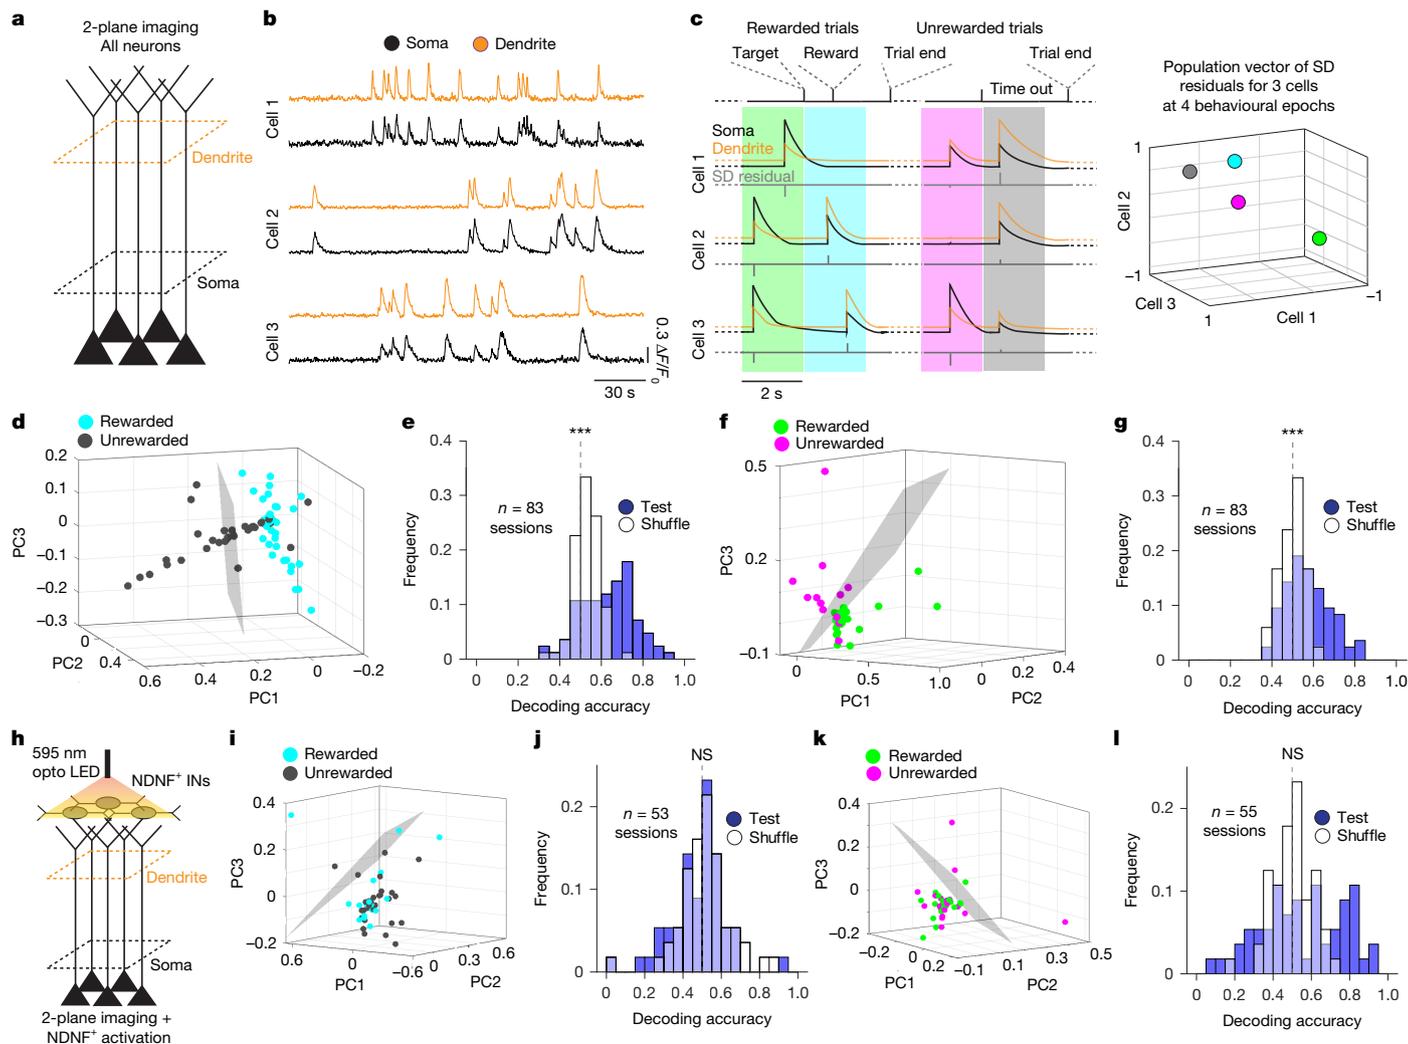


**Fig. 3 | Experimental manipulation of SD residuals.** **a**, Schematic of the experimental approach. First, we imaged neurons in RSC while the mouse was exposed to rotating stimuli identical to those presented during baseline estimation in the BCI task. Next, we anaesthetized the mice (using isoflurane) and imaged the same neurons. **b**, Schematic of the imaging approach. **c**, Top,  $\Delta F/F_0$  GCaMP traces simultaneously recorded in the soma and its corresponding dendrite for two representative neurons during wakefulness and anaesthesia. Bottom, the mean  $\Delta F/F_0$  signal in the somas and in the dendrites of two example neurons, for all events that occurred during wakefulness and anaesthesia. Compared to somatic activity, dendrites are preferentially inhibited during anaesthesia. Shading represents s.e.m. Max, maximum; norm., normalized. **d**, Mean somatic event rate (GCaMP7f) during wakefulness and anaesthesia. The dashed grey line represents the identity line. Paired *t*-test,  $P < 1.9 \times 10^{-126}$ ;  $0.002 \pm 6.3 \times 10^{-5}$  (wakefulness) and  $0.0002 \pm 1.8 \times 10^{-5}$  (anaesthesia);  $n = 832$  neurons from 8 imaging sessions in 4 different mice. **e**, Distribution of SD residual during wakefulness and anaesthesia. Paired *t*-test,  $P = 3.3 \times 10^{-20}$ ;

$0 \pm 2.2 \times 10^{-18}$  (wakefulness) and  $-0.52 \pm 4.9 \times 10^{-2}$  (anaesthesia), mean  $\pm$  s.e.m.;  $n = 160$  neurons. **f**, Targeting strategy schematic. A Cre-dependent version of ChRmine tagged with oScarlet was injected into layer 1 of the RSC of NDNF-Cre mice. GCaMP7f under the control of the synapsin promoter was injected into layer 5 of the RSC in the same mice. **g**, z-stack reconstruction of an imaged field of view. The image was acquired in vivo using 2-photon microscopy (1,000 nm laser wavelength). Scale bar, 50  $\mu$ m. **h**, Same as **c**, for opto on and control conditions. Compared to somatic activity, dendrites are preferentially inhibited during opto on. **i**, Mean somatic event rate (GCaMP7f) during opto on and off. The dashed grey line represents the identity line (paired *t*-test,  $P < 2.2 \times 10^{-308}$ ;  $2.8 \times 10^{-2} \pm 4.5 \times 10^{-4}$  (control) and  $5.9 \times 10^{-3} \pm 2 \times 10^{-4}$  (opto on), mean  $\pm$  s.e.m.;  $n = 2,884$  neurons). Shading represents s.e.m. **j**, Distribution of SD residual during control and opto on. Paired *t*-test,  $P = 9 \times 10^{-214}$ ;  $n = 1,886$  neurons from 56 sessions recorded from 4 mice;  $0 \pm 7.7 \times 10^{-19}$  (control) and  $-0.87 \pm 2.4 \times 10^{-2}$  (opto on).

amplified dendritic activity during periods of error reduction for both P+ and P- neurons and as attenuated dendritic activity during times of error increase. However, a vectorized error signal would exhibit selective P+ versus P- dendritic activation, as the activity of each group is causally mapped 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, given that the angle of the visual stimuli presented to the mice represented the instantaneous task-associated error (Fig. 5a). 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. As our analysis was restricted to time bins with coincident somato-dendritic events in P+ and/or P- neurons, any potential

noise-driven flickering was not present in our analysis. We found that the dendrites of P+ neurons were relatively amplified during error reduction compared with error increase epochs (Fig. 5c). Dendrites in P- neurons exhibited the converse relationship: relative dendritic attenuation and amplification occurred during error reduction and error increase, respectively (Fig. 5d-f and Extended Data Figs. 12 and 6). This relationship could be observed in six out of six mice trained in the task (Extended Data Fig. 13e) 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. 13). Additionally, the same inverted relationship between dendritic signals and task-associated errors was found in the dendrites of P<sub>0</sub> neurons that were functionally correlated to P+ and P- neurons (Extended Data



**Fig. 4 | A population vector of SD 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 and its corresponding dendrite for three representative neurons. **c**, Schematic of SD residual population vector for four different behavioural epochs. Left, somatic, dendritic and residual traces for three cartoon neurons. Green, cyan, purple and black boxes represent 4 different behavioural epochs: 2 s before and after target is reached for rewarded trials and 2 s before and after the end of unrewarded trials. The SD residual trace of each neuron was estimated by collapsing coincident soma–dendrite events into point values at the time of event onset. Resultant SD residual traces for all neurons in each behavioural 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 (Fig. 2 and Methods). Right, 3D plot of SD residuals for the four behavioural epochs shown on the left, where  $x$ ,  $y$  and  $z$  correspond to SD residual of neurons 1–3 (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. Cyan dots are vectors resulting from the 2 s following the reaching of target activity in rewarded trials. Dark grey dots are vectors calculated in the 2 s following the end of an unsuccessful trial (same as **c**). The shaded black hyperplane represents the decision boundary for binary classification calculated using a linear SVM. **e**, Decoding accuracy for test versus shuffle data for 83 sessions (paired  $t$ -test,  $P = 9.8 \times 10^{-9}$ ;  $0.63 \pm 0.01$  (test) and  $0.52 \pm 0.01$  (shuffled), mean  $\pm$  s.e.m.;  $n = 83$ ). **f, g**, Same as in **d, e**; a 264-dimensional vector of SD residual collapsed onto the first 3 principal components for visualization only. Green represents the last 2 s of a rewarded trial and purple represents the last 2 s of an unrewarded trial (paired  $t$ -test,  $P = 7.1 \times 10^{-8}$ ;  $0.57 \pm 0.01$  (test) and  $0.49 \pm 0.01$  shuffled, mean  $\pm$  s.e.m.;  $n = 83$ ). **h**, Schematic of the experimental approach: layer 1 NDNF<sup>+</sup> interneurons (INs) were optogenetically activated during BCI task performance. **i–l**, Same as **d–g**, but during optogenetic activation of NDNF<sup>+</sup> interneurons. **j**, Paired  $t$ -test,  $P = 0.18$ ;  $0.48 \pm 0.02$  (test) and  $0.51 \pm 0.02$  (shuffled), mean  $\pm$  s.e.m.;  $n = 53$  sessions from 4 mice. **l**, Paired  $t$ -test,  $P = 0.13$ ;  $0.54 \pm 0.03$  (test) and  $0.50 \pm 0.01$  (shuffled);  $n = 55$  sessions. NS, not significant.

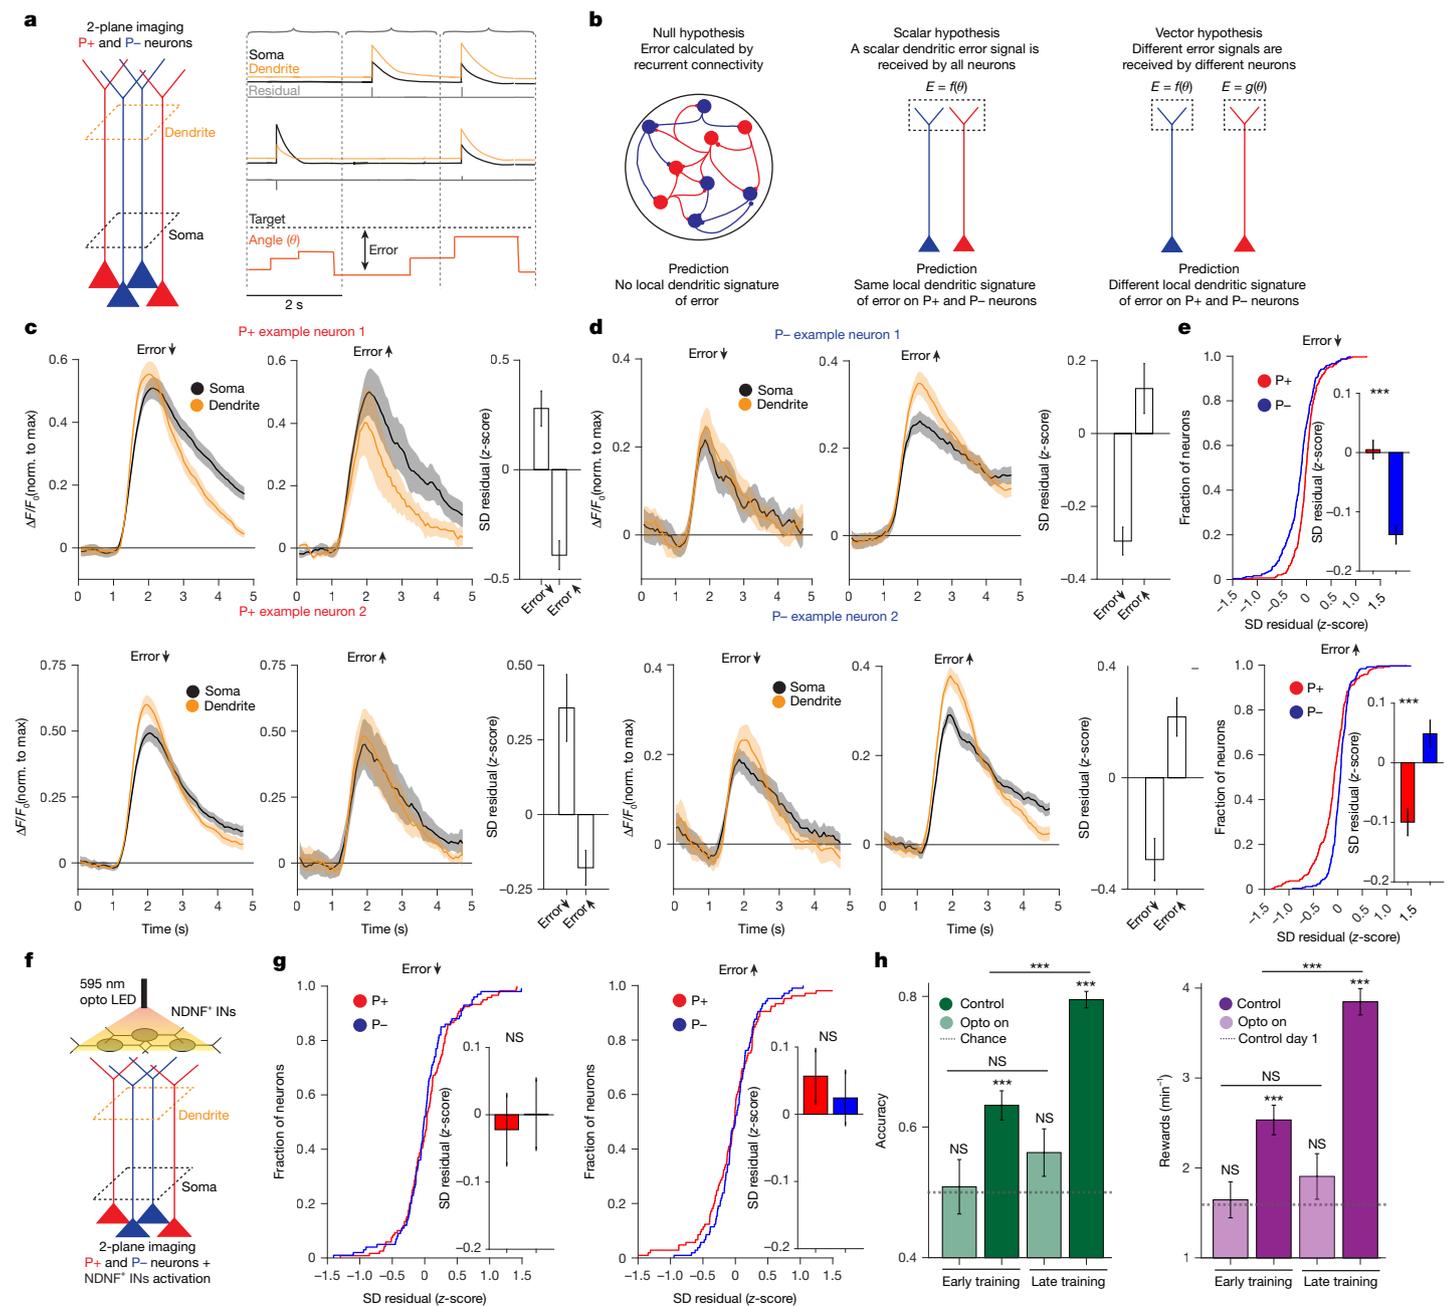
Fig. 14). Of note, SD residuals represented error derivatives, not errors (Extended Data Fig. 13), in contrast to instructive signals found in the classical implementations of backpropagation.

Next, we tested whether vectorized error-related dendritic signals were necessary for learning by optogenetically activating NDNF<sup>+</sup> layer 1 interneurons throughout the BCI task. This abolished vectorized error-related signalling in the apical tuft of layer 5 pyramidal neurons (Fig. 5f, g) and disrupted learning (Fig. 5h) but not in our LED illumination control group (Extended Data Fig. 15). This demonstrates that local

computation in the apical dendritic tuft is necessary for performance improvements in the BCI task.

## Discussion

Here we demonstrate the use of neurofeedback brain–computer interfaces to study the mechanisms of biological credit assignment at the subcellular level. Our results provide—to our knowledge—the first biological evidence of a vectorized solution to the credit assignment



**Fig. 5** | See next page for caption.

problem in the brain via cortical dendrites. Our data are consistent with a model of credit assignment in which learning is instructed by instantaneous, vectorized teaching signals received onto the distal dendrites of pyramidal neurons<sup>1–5</sup>. This spatial segregation mechanism allows cortical circuits to overcome the biologically implausible temporal separation of feedforward and feedback streams conventionally used for computing teaching signals during vectorized learning in ANNs.

The data presented here reveal magnitude differences in coincident somato-dendritic events that can be predicted using activity in the surrounding network of neurons. At the population level, differences in somato-dendritic coupling encode de novo information relative to somatic activity. 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. Finally, optogenetic activation of NDNF<sup>+</sup> layer 1 interneurons disrupted both dendritic computation and learning, demonstrating that dendritic processing is necessary for learning.

Our results demonstrate the existence of a signed, vectorized dendritic input that is tailored in a condition-specific manner to individual neurons. The extent to which this dendritic activity reflects moment-to-moment computational signals—as opposed to teaching signals for synaptic weight changes—remains to be uncovered. This could be explored via manipulations of dendritic activity through different phases of learning<sup>57</sup>. Importantly, vectorized error signals need not be confined to learning. In control-theory-inspired frameworks of credit assignment, error signals are applied during task operation, to steer the system online<sup>58,59</sup>. Further work is needed to assess whether these dendritic signals result from glutamatergic inputs from higher-order

**Fig. 5 | Dendritic error signals are cell-specific and depend on the causal contribution of the neuron to the task.** **a**, Left, experimental schematic. Top right, idealized somatic, dendritic and residual traces, respectively, for three neurons. Bottom right, relationship between stimulus angle, target and error. All closed-loop data from the BCI paradigm (which excludes rewards and timeout periods) were chunked into 2 s 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 behaviour. This model predicts that neurons with opposite causal contribution to behaviour 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 mean  $\pm$  s.e.m. of SD residual value (z-scored) for all events that occurred during error decrease and increase epochs. **d**, Same as **c**, for two P- neurons. In contrast to the P+ neurons, dendritic activity is relatively attenuated for error reduction epochs compared with error increase epochs. **e**, Top, during error reduction epochs, the cumulative distribution function for SD residuals (z-scored across all neurons) for P+ and P- neurons. The bar graph represents mean  $\pm$  s.e.m. 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;  $P = 5.3 \times 10^{-7}$ ;  $0.005 \pm 0.01$  (P+) and  $-0.14 \pm 0.02$  (P-),

mean  $\pm$  s.e.m.;  $n = 292$  (P+) and  $240$  (P-) neurons). Bottom, dendrites in P+ neurons are more attenuated than P- neurons during error reduction epochs (*t*-test;  $P = 1.2 \times 10^{-7}$ ;  $-0.1 \pm 0.02$  (P+) and  $0.05 \pm 0.01$  (P-), mean  $\pm$  s.e.m.;  $n = 267$  (P+) and  $249$  (P-)). **f**, Experimental schematic. GCaMP signals in the soma and dendrites of P+ and P- neurons were recorded during optogenetic activation of layer 1 NDNF+ neurons. **g**, As in **e** but during layer 1 NDNF+ neuron activation. Left, during error reduction epochs, the cumulative distribution function for SD residuals (z-scored across all neurons) for P+ and P- neurons. The bar graph represents mean  $\pm$  s.e.m. for the population distribution shown in the cumulative distribution function for P+ and P- neurons (*t*-test;  $P = 0.58$ ;  $0.06 \pm 0.04$  (P+) and  $0.02 \pm 0.04$  (P-), mean  $\pm$  s.e.m.;  $n = 119$  (P+) and  $100$  (P-) neurons). Right, same as **e**, for error increase epochs. Dendrites in P+ neurons are more attenuated than those in P- neurons during error reduction epochs (*t*-test;  $P = 0.7$ ;  $0.02 \pm 0.05$  (P+) and  $0 \pm 0.03$  (P-);  $n = 105$  (P+) and  $105$  (P-)). **h**, BCI performance from four mice. Left, accuracy in early and late training for control (opto off) and opto on conditions. Dashed grey line represents chance level (paired *t*-test;  $P = 0.83$ ;  $0.51 \pm 0.04$ , mean  $\pm$  s.e.m.;  $n = 32$  for opto on against chance during early training; paired *t*-test;  $P = 3 \times 10^{-7}$ ;  $0.63 \pm 0.03$ ;  $n = 48$  for control against chance during early training;  $P = 0.1$ ;  $0.56 \pm 0.02$ ;  $n = 24$  for opto on against chance during late training;  $P = 5.7 \times 10^{-23}$ ;  $0.8 \pm 0.01$ ;  $n = 36$  for control against chance during late training; *t*-test,  $P = 0.36$  (opto on) and  $1.3 \times 10^{-7}$  (control) for early versus late training). Right, rewards per minute in early and late training for control (opto off) and opto on conditions. Dashed grey line represents rewards per minute for control on day 1 (paired *t*-test;  $P = 0.8$ ;  $1.6 \pm 0.2$ , mean  $\pm$  s.e.m.;  $n = 32$  for opto on against control on day 1, during early training; paired *t*-test;  $P = 7 \times 10^{-7}$ ;  $0.63 \pm 0.26$ ;  $n = 48$  for control against control on day 1, during early training;  $P = 0.23$ ;  $1.9 \pm 0.26$ ;  $n = 24$  for opto on against control on day 1, during late training;  $P = 4.2 \times 10^{-17}$ ;  $3.8 \pm 0.14$ ;  $n = 36$  for control against control on day 1, during late training; *t*-test,  $P = 0.42$  (opto on) and  $1.5 \times 10^{-7}$  (control) for early versus late training).

cortical areas, from neuromodulation, or as a product of recurrent excitatory and inhibitory local computation. Dopaminergic signalling specifically has been causally implicated in both error signalling and in learning neurofeedback BCI tasks in rodents and humans<sup>33,60–62</sup> and thus represents a compelling target for future investigation. Further experiments are also needed to test whether errors signals are calculated locally at each hierarchical layer or are transmitted across layers, as in the classical formulation of backpropagation<sup>19</sup>. Previous neurofeedback BCI studies have demonstrated that degrading the contingency between neuronal activity and feedback stimuli impairs learning<sup>33–35</sup>; future work will have to determine whether external stimuli are always necessary for error representations or whether animals can access the cost function via internal states exclusively, and how the dendritic representation of error might change as a result.

The error signals we observed have appealing connections to the gradient calculations found in the backpropagation algorithm. In contrast to the classical implementation of backpropagation, however, we observed that dendrites received signals that bore signatures of error derivative rather than error itself. Intriguingly, our results could also be consistent with target propagation (specifically, difference target propagation)<sup>14,20,21</sup>. Indeed, our data indicate that dendritic activity contains a target signal for the parent soma in addition to task-related error information. Future approaches, built on the framework we present here, could be used to disentangle the specific learning algorithms used by the brain<sup>14,63</sup>.

Together, our results help to reconcile early findings and theories of dendritic function, which focused on single dendritic branches as the building blocks for independent computation, with later in vivo findings that have demonstrated prevalent co-occurrence of dendritic and somatic events<sup>15,24,51,64</sup>. By demonstrating that apical dendrites locally compute reward and error-related signals, our results present a framework for dendritic computation that does not require fully independent dendrites to perform credit assignment for adaptive behaviour and highlight new directions for the development of biologically inspired ANNs.

## Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-026-10190-7>.

1. Sacramento, J., Bengio, Y., Costa, R. P. & Senn, W. Dendritic cortical microcircuits approximate the backpropagation algorithm. In *Proc. 32nd International Conference on Neural Information Processing Systems* (eds Bengio, S. & Wallach, H. M.) 8735–8746 (ACM, 2018).
2. 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).
3. Guerguiev, J., Lillicrap, T. P. & Richards, B. A. Towards deep learning with segregated dendrites. *eLife* **6**, e22901 (2017).
4. Greedy, W., Zhu, H. W., Pemberton, J., Mellor, J. & Costa, R. P. Single-phase deep learning in cortico-cortical networks. In *Proc. 36th International Conference on Neural Information Processing Systems* (eds Koyejo, S. et al.) 24213–24225 (ACM, 2022).
5. Körding, K. P. & König, P. Supervised and unsupervised learning with two sites of synaptic integration. *J. Comput. Neurosci.* **11**, 207–215 (2001).
6. Magee, J. C. & Grienberger, C. Synaptic plasticity forms and functions. *Annu. Rev. Neurosci.* **43**, 95–117 (2020).
7. 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).
8. Caporale, N. & Dan, Y. Spike timing-dependent plasticity: a Hebbian learning rule. *Annu. Rev. Neurosci.* **31**, 25–46 (2008).
9. 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).
10. Kirkwood, A., Rioult, M. G. & Bear, M. F. Experience-dependent modification of synaptic plasticity in visual cortex. *Nature* **381**, 526–528 (1996).
11. Bliss, T. V. P. & Lomo, 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).
12. 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).
13. Abbott, L. F. & Blum, K. I. Functional significance of long-term potentiation for sequence learning and prediction. *Cereb. Cortex* **6**, 406–416 (1996).
14. Lillicrap, T. P., Santoro, A., Marris, L., Akerman, C. J. & Hinton, G. Backpropagation and the brain. *Nat. Rev. Neurosci.* **21**, 335–346 (2020).

15. Richards, B. A. & Lillicrap, T. P. Dendritic solutions to the credit assignment problem. *Curr. Opin. Neurobiol.* **54**, 28–36 (2019).
16. Richards, B. A. et al. A deep learning framework for neuroscience. *Nat. Neurosci.* **22**, 1761–1770 (2019).
17. Minsky, M. Steps toward artificial intelligence. *Proc. IRE* **49**, 8–30 (1961).
18. Lansdell, B. J., Prakash, P. R. & Kording, K. P. Learning to solve the credit assignment problem. In *Proceedings of the International Conference on Learning Representations (ICLR)* <https://openreview.net/forum?id=ByeUBANtvB> (2020).
19. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).
20. Meulemans, A., Carzaniga, F., Suykens, J., Sacramento, J. & Grewe, B. F. A theoretical framework for target propagation. in *Proc. 34th International Conference on Neural Information Processing Systems* vol. 33 (eds Larochelle, H. et al.) 20024–20036 (ACM, 2020).
21. Lee, D.-H., Zhang, S., Fischer, A. & Bengio, Y. Difference target propagation. In *Machine Learning and Knowledge Discovery in Databases 9284* (eds Appice, A. et al.) 498–515 (Springer, 2015).
22. Makino, H. & Koyama, 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* **398**, 338–341 (1999).
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* **110**, 1532–1546 (2022).
26. Marques, T., Nguyen, J., Fioreze, G. & Petreanu, L. The functional organization of cortical feedback inputs to primary visual cortex. *Nat. Neurosci.* **21**, 757–764 (2018).
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. Fişek, M. et al. Cortico-cortical feedback engages active dendrites in visual cortex. *Nature* **617**, 769–776 (2023).
32. Naud, R. & Sprekeler, H. Sparse bursts optimize information transmission in a multiplexed neural code. *Proc. Natl Acad. Sci. USA* **115**, E6329–E6338 (2018).
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* **97**, 1356–1368.e4 (2018).
34. 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.* **17**, 807–809 (2014).
35. Clancy, K. B. & Mrsic-Flogel, T. D. The sensory representation of causally controlled objects. *Neuron* **109**, 677–689.e4 (2021).
36. 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).
37. Lai, C., Tanaka, S., Harris, T. D. & Lee, A. K. Volitional activation of remote place representations with a hippocampal brain–machine interface. *Science* **382**, 566–573 (2023).
38. Mitani, A., Dong, M. & Komiyama, T. Brain–computer interface with inhibitory neurons reveals subtype-specific strategies. *Curr. Biol.* **28**, 77–83.e4 (2018).
39. 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).
40. Voigts, J. & Harnett, M. T. Somatic and dendritic encoding of spatial variables in retrosplenial cortex differs during 2D navigation. *Neuron* **105**, 237–245.e4 (2019).
41. 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* **103**, 235–241.e4 (2019).
42. 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).
43. 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* **8**, e49145 (2019).
44. 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).
45. 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).
46. Kerlin, A. et al. Functional clustering of dendritic activity during decision-making. *eLife* **8**, e46966 (2019).
47. 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. USA* **110**, 13618–13623 (2013).
48. Grienberger, C., Chen, X. & Konnerth, A. NMDA receptor-dependent multidendrite Ca<sup>2+</sup> spikes required for hippocampal burst firing in vivo. *Neuron* **81**, 1274–1281 (2014).
49. Otor, Y. et al. Dynamic compartmental computations in tuft dendrites of layer 5 neurons during motor behavior. *Science* **376**, 267–275 (2022).
50. Helmchen, F., Svoboda, K., Denk, W. & Tank, D. W. In vivo dendritic calcium dynamics in deep-layer cortical pyramidal neurons. *Nat. Neurosci.* **2**, 989–996 (1999).
51. Francioni, V. & Harnett, M. T. Rethinking single neuron electrical compartmentalization: dendritic contributions to network computation in vivo. *Neuroscience* **489**, 185–199 (2022).
52. Xu, N. et al. Nonlinear dendritic integration of sensory and motor input during an active sensing task. *Nature* **492**, 247–251 (2012).
53. Abs, E. et al. Learning-related plasticity in dendrite-targeting layer 1 interneurons. *Neuron* **100**, 684–699.e6 (2018).
54. Palmer, L. M. et al. The cellular basis of GABA<sub>B</sub>-mediated interhemispheric inhibition. *Science* **335**, 989–993 (2012).
55. Keller, A. J., Roth, M. M. & Scanziani, M. Feedback generates a second receptive field in neurons of the visual cortex. *Nature* **582**, 545–549 (2020).
56. Jiang, X., Wang, G., Lee, A. J., Stornetta, R. L. & Zhu, J. J. The organization of two new cortical interneuronal circuits. *Nat. Neurosci.* **16**, 210–8 (2013).
57. Reinhold, K. et al. Striatum supports fast learning but not memory recall. *Nature* **643**, 458–467 (2025).
58. Meulemans, A. et al. Credit assignment in neural networks through deep feedback control. In *Advances in Neural Information Processing Systems 34* (eds Ranzato, M. et al.) 4674–4687 (Curran Associates, 2021).
59. Todorov, E. & Jordan, M. I. Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* **5**, 1226–1235 (2002).
60. Athalye, V. R., Carmena, J. M. & Costa, R. M. Neural reinforcement: re-entering and refining neural dynamics leading to desirable outcomes. *Curr. Opin. Neurobiol.* **60**, 145–154 (2020).
61. Gershman, S. J. et al. Explaining dopamine through prediction errors and beyond. *Nat. Neurosci.* **27**, 1645–1655 (2024).
62. Kasahara, K., DaSalla, C. S., Honda, M. & Hanakawa, T. Basal ganglia–cortical connectivity underlies self-regulation of brain oscillations in humans. *Commun. Biol.* **5**, 712 (2022).
63. Mikulasch, F. A., Rudelt, L., Wibrall, M. & Priesemann, V. Where is the error? Hierarchical predictive coding through dendritic error computation. *Trends Neurosci.* **46**, 45–59 (2023).
64. Branco, T. & Häusser, M. The single dendritic branch as a fundamental functional unit in the nervous system. *Curr. Opin. Neurobiol.* **20**, 494–502 (2010).

**Publisher’s note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2026

## 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 and NDNF-Cre heterozygous mice were maintained on a 12 h:12 h 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. Except for anaesthesia experiments, mice 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. Mice were injected with dexamethasone (4 mg kg<sup>-1</sup>), carprophen (5 mg kg<sup>-1</sup>) and buprenorphine (slow release, 0.5 mg kg<sup>-1</sup>) 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 3 mm-wide craniotomy was performed. In Rbp4-Cre mice, at 3–4 different sites, we injected 100 nl of AAV1-syn-FLEX-jGCaMP7f-WPRE (Addgene, 04492-AAV1, 2–5 × 10<sup>12</sup> viral genomes (vg) ml<sup>-1</sup> 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 RSC (2.5 mm caudal of bregma). The same labelling approach was used to perform anaesthesia experiments. In NDNF-Cre mice, we injected 100 nl of AAV8-nEF-Con/Foff2.0-ChRmine-oScarlet (Addgene, 137161-AAV8, 7 × 10<sup>12</sup> vg ml<sup>-1</sup> after a 1:5 dilution from the original concentration) 150 μm from the surface of the brain and 75 nl of AAV1-syn-jGCaMP7f-WPRE (Addgene, 104488-AAV1, 2–5 × 10<sup>12</sup> vg ml<sup>-1</sup> concentration after a 1:10 dilution from the original concentration) 500 μm from the surface of the brain. The dura was left intact. Cranial windows consisted of two stacked 3 mm 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 to perform imaging under head-fixed conditions. At the end of the procedure, a single dose of 25 ml kg<sup>-1</sup> of Ringer's solution was injected subcutaneously to rehydrate the mouse. 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 4× pulse splitter to reduce photodamage and bleaching. For excitation and photon collection we used a 16× Nikon objective with 0.8 numerical aperture. Bidirectional scanning was performed (512 × 796 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 to avoid light contamination. Mice were habituated to human handling for 5–10 min every day and to head-fixation for 15 min 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 behavioural experiments. The locomotion of the mouse was recorded using an optical encoder (E6, US Digital, 2500 cycles per revolution) tracking the rotation of a cylindrical treadmill with a radius of 19 cm 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 five pyramidal neurons at two different planes: proximal to the soma and right below the nexus (tuft bifurcation point).

### Optogenetic stimulation

For optogenetic stimulation we used a Cyclops LED driver from open ephys (OEPS-6602) triggered using a direct 6 ms TTL pulse delivered via the NeuroLabware Dual PSOC box. The driver controlled a fibre-coupled 595 nm LED (8.7 mW, 100 mA Thorlabs M595F2). LED illumination was synchronized with the photomultiplier tubes (PMTs) of the imaging system using custom-made Matlab scripts. In brief, for every new frame acquired by the 2-photon microscope, the LED was activated for the initial 6 ms of the frame, while the PMTs were kept shut off for an additional 1 ms (7 ms total of PMT off time). PMTs would then reactivate to collect calcium data for the remaining approximately 24 ms of the approximately 31 ms frame.

### BCI task

Similar to previous implementations of BCI learning paradigms<sup>34,35,38</sup>, mice were trained so that they obtained rewards by modulating the activity of eight or ten layer 5 pyramidal neurons in the RSC 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° (horizontal) and P- neurons whose activity rotated the stimulus away from the target angle, towards a 0° (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° and 90°, passive viewing) for approximately 13 min (12,000 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 12,000-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 s.d. 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, and used 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° angle corresponded to the minimum value in the subtracted  $\Delta F/F_0$  signal distribution while target, or 90° 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° interval between 0° and 90°. 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. The screen refreshed every time a 2-photon frame at the soma was acquired (at 15 Hz). In line with previous studies performing neurofeedback BCI in rodents<sup>33–36,38</sup>, we binned the visual stimulus to avoid noise-driven, frame-by-frame stimulus updates at the screen refresh rate, which is beyond the perceptual threshold in mice<sup>65,66</sup>. To avoid introducing a second, orthogonal dimension to our task that would disrupt the straightforward mapping between neuronal activity and task error, we did not introduce any requirement on the number of P+ or P- neurons required to be simultaneously active to trigger a reward or a stimulus update. In each trial, mice had 28 s to reach target

# Article

activity. If they did, a reward, consisting of 4  $\mu$ l of 10% sucrose water was delivered 1 s after. Additionally, after reaching target activity, the stimulus froze to a 90° angle for 2 s. After that, mice saw a black screen for an additional 1 s and a new trial was initiated. All new trials were initiated by a 0.5 s isoluminant grey stimulus. If a mouse did not reach target activity within the 28 s of the trial, a 3 s 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 s. For the optogenetics experiments, we recorded 2 different baselines (13 min each, during passive viewing of the Gabor patch—same as described above). The first one with the LED off (control) was used for post hoc analysis only of the data shown in Fig. 3. The second baseline was recorded during LED stimulation (opto on) and was used to map neuronal activity to angle and target during the closed-loop part of the BCI task. This mapping was crucial to ensure that the decoder was calibrated consistently for the closed-loop BCI task, which was recorded during the opto on condition only. Early and late training were defined as days 1–8 and 9–14 respectively, based on average performance (accuracy) remaining above 0.75 in the control (opto off) condition. For anaesthesia experiments, we recorded two passive viewing sessions (awake and anaesthetized, 24 min each) where we presented the same set of stimuli presented when recording the baseline session for the BCI task. We anaesthetized the mice in between these two sessions by initially administering (via inhalation) 4% isoflurane that subsequently decreased to 1% for the duration of the imaging session.

## P+ and P– neuron selection

On day 1, we drew 20–40 regions of interest (ROIs) in a single field of view prior to starting our baseline recording. Next, we recorded a session of passive visual stimuli that we would later use as our baseline recording for day 1. At the end of this recording, all  $\Delta F/F_0$  traces for all drawn ROIs were plotted and visually inspected using a custom Matlab script. The experimenter would then select either eight or ten of these traces based on event frequency, signal to noise ratio (SNR; determined as the ratio between noise band width and maximum event size), baseline stability and calcium transient dynamics (with a clear rise, peak, and exponential decay—as opposed to plateau-looking events). The best 8–10 neurons would then be selected from the available pool of neurons on which ROIs were drawn. No arbitrary parameter cut-off (for example, minimum event frequency or SNR) was introduced. The subdivision of these eight to ten neurons into the P+ and P– population would then be determined by a random number generator. Once selected, P+ and P– neurons would remain the same for the entire duration of the experiment.

## Online motion-correction

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 s (300–600 frames) collected onto our field of view. To motion-correct each subsequent frame, we selected four smaller central areas to register independently from one another (2D rigid translation) against the corresponding four areas in the reference image<sup>67</sup>. We finally rigid-translate the entire 2D image by taking the average translation in  $x$  and  $y$  for these four subregions.

## Visual stimuli

Visual stimuli were generated using the Psychophysics Toolbox package for MATLAB (MathWorks)<sup>68</sup> 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° 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<sup>69</sup>. For each field of

view, we removed duplicates by excluding ROIs whose signal correlation was above 0.6 and whose centre was within 20  $\mu$ m of distance. In order to separate trunk signals from potential neuropil contamination, fluorescence signals of our ROIs were processed using FISSA<sup>70</sup> 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  $(F_i - F_0)/F_0$ .  $F_0$  is defined as the tenth percentile of a 120-s-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<sup>45</sup>.

## 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-designed 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 used a 2 s 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 s 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 corresponding distal trunk whenever their  $\Delta F/F_0$  signal correlation was equal above 0.6. For optogenetics and anaesthesia experiments, we matched proximal trunks to their corresponding distal trunk using activity during control (opto off) and wakefulness, respectively. Whenever we found more than one 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. To estimate the SD residual during optogenetic stimulation and anaesthesia, first we estimated the best-fit line using somato-dendritic activity during light off and wakefulness conditions, respectively. Next, we calculated the residual for all events detected during opto on and anaesthesia as the distance between these events and the previously calculated best-fit line.

## Decoders

To decode whether individual transients would be amplified or attenuated, we trained a SVM binary classifier (linear kernel) using stochastic gradient descent<sup>71</sup> (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 preceding 2 s, 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 datasets. 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 s before and 2 s after the reach of target activity, respectively for successful trials, and 2 s before and 2 s after the end of a trial for unsuccessful trials. To decode successful from unsuccessful trials, we generated a  $n$ -dimensional SD residual vector by taking the residual for each neuron for which we identified a somato-dendritic pair (see above) in these 2 s epochs. Neurons inactive in the 2 s 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 s.e.m. Statistical tests and independent samples are described in figure legends. All  $t$ -tests in the Article are two-sided.

## Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

## Data availability

Data are available upon request. Source data are provided with this paper.

## Code availability

The BCI code is available at <https://github.com/harnett/CLOOSE>. Analysis code is available upon request.

65. Marques, T. et al. A role for mouse primary visual cortex in motion perception. *Curr. Biol.* **28**, 1703–1713.e6 (2018).
66. You 游文愷, W.-K. & Mysore, S. P. Dynamics of visual perceptual decision-making in freely behaving mice. *eNeuro* **9**, 0161-21.2022 (2022).
67. Guizar-Sicairos, M., Thurman, S. T. & Fienup, J. R. Efficient subpixel image registration algorithms. *Opt. Lett.* **33**, 156–158 (2008).
68. Brainard, D. H. The Psychophysics Toolbox. *Spat. Vis.* **10**, 433–436 (1997).
69. Pachitariu, M. et al. Suite2p: beyond 10,000 neurons with standard two-photon microscopy. Preprint at *bioRxiv* <https://doi.org/10.1101/061507> (2017).
70. Keemink, S. W. FISSA: a neuropil decontamination toolbox for calcium imaging signals. *Sci. Rep.* **8**, 3493 (2018).
71. 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 *Proc. 25th International Conference on Machine Learning* 408–415 (ACM, 2008).

**Acknowledgements** We thank I. Fiete, R. Naud and C. Yaeger for comments on the manuscript. V.F. was supported by the Y. Eva Tan Postdoctoral Fellowship. V.D.T. was supported by the MathWorks Science Fellowship and the Janet and Sheldon Razin Fellowship. E.H.S.T. was supported by the National Institute of General Medical Sciences (T32GM007753), the Paul and Daisy Soros Fellowship, and the Yang ICoN Graduate Fellowship. M.T.H. was supported by the NIH (R01NS106031, R01NS113079 and R01MH135141) and by the Klingenstein–Simons Fellowship, the Vallee Foundation Scholars and the McKnight Scholars programmes.

**Author 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. E.H.S.T. helped to write the manuscript. Z.D. helped in performing surgeries on mice and with data acquisition. N.J.B. performed surgeries on mice. M.T.H. supervised all aspects of the project.

**Competing interests** The authors declare no competing interests.

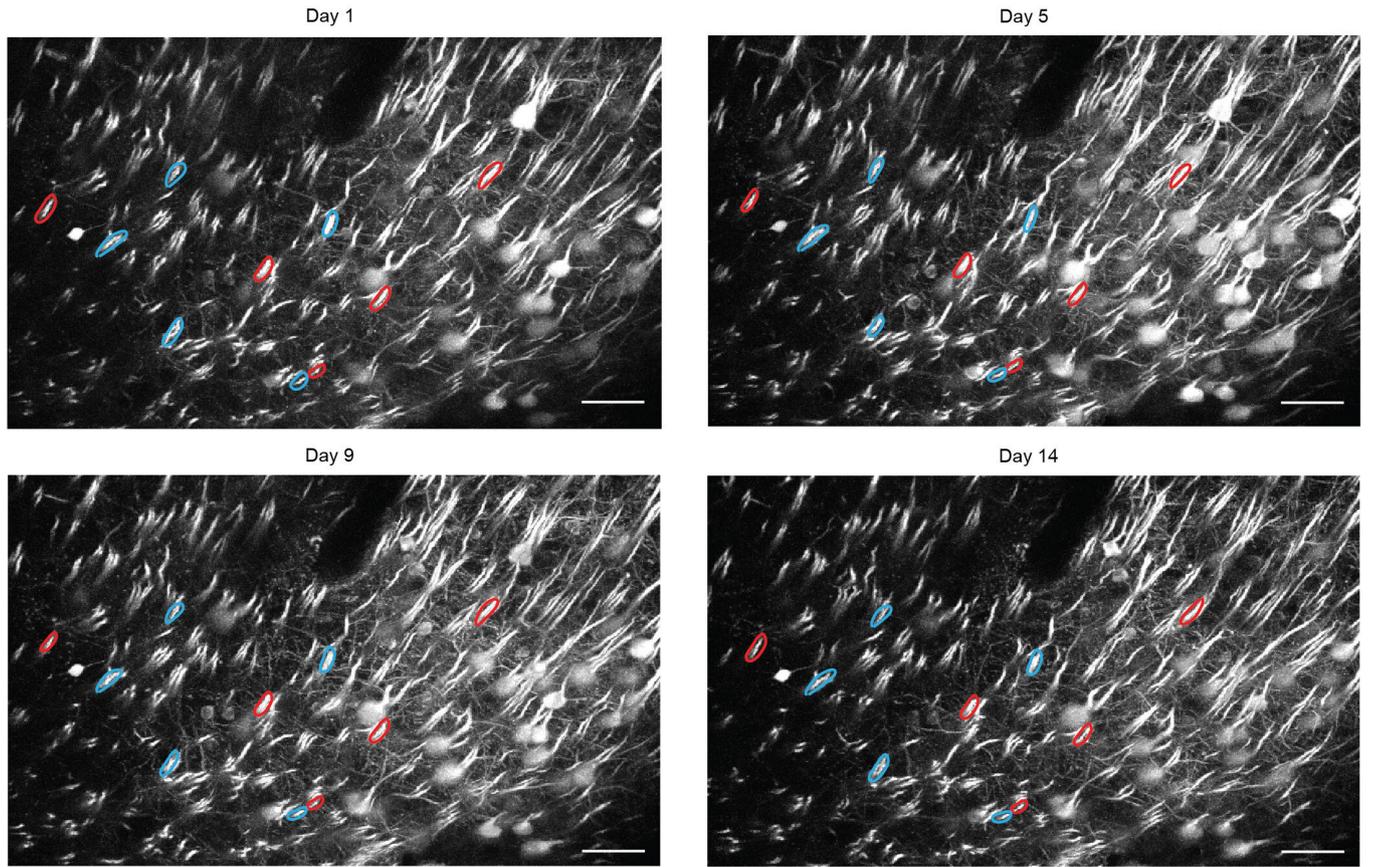
## Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41586-026-10190-7>.

**Correspondence and requests for materials** should be addressed to Mark T. Harnett.

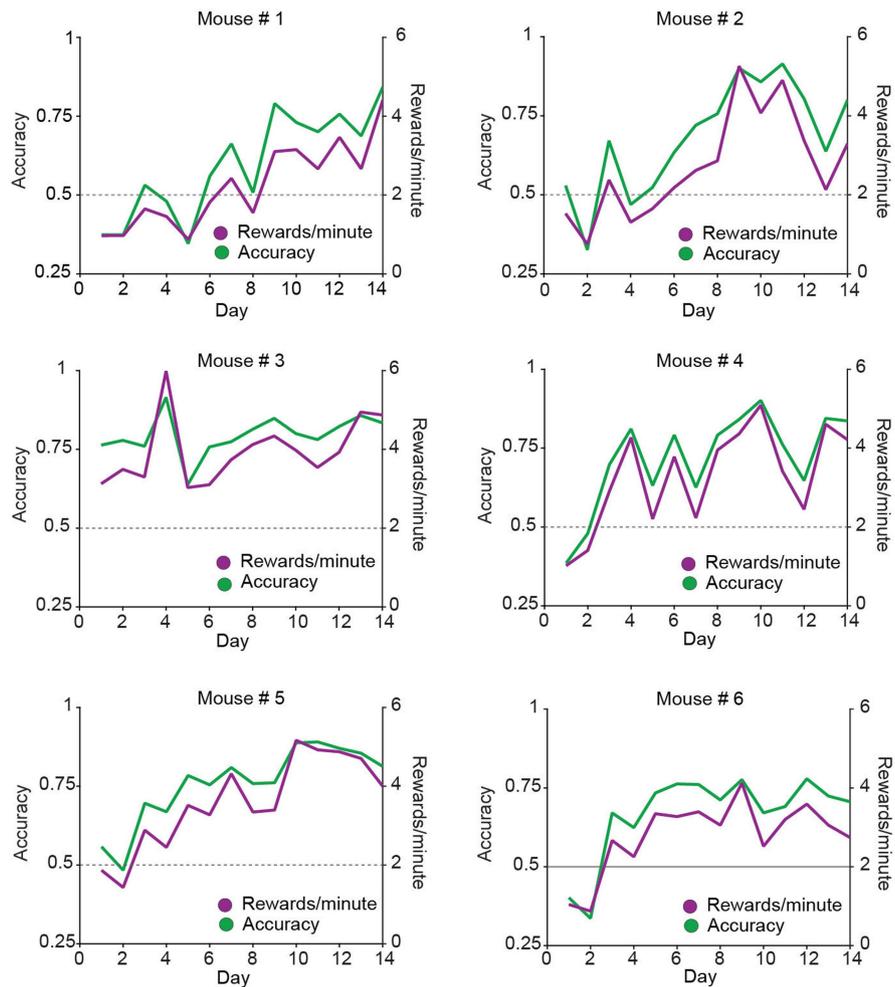
**Peer review information** *Nature* thanks the anonymous reviewer(s) for their contribution to the peer review of this work.

**Reprints and permissions information** is available at <http://www.nature.com/reprints>.

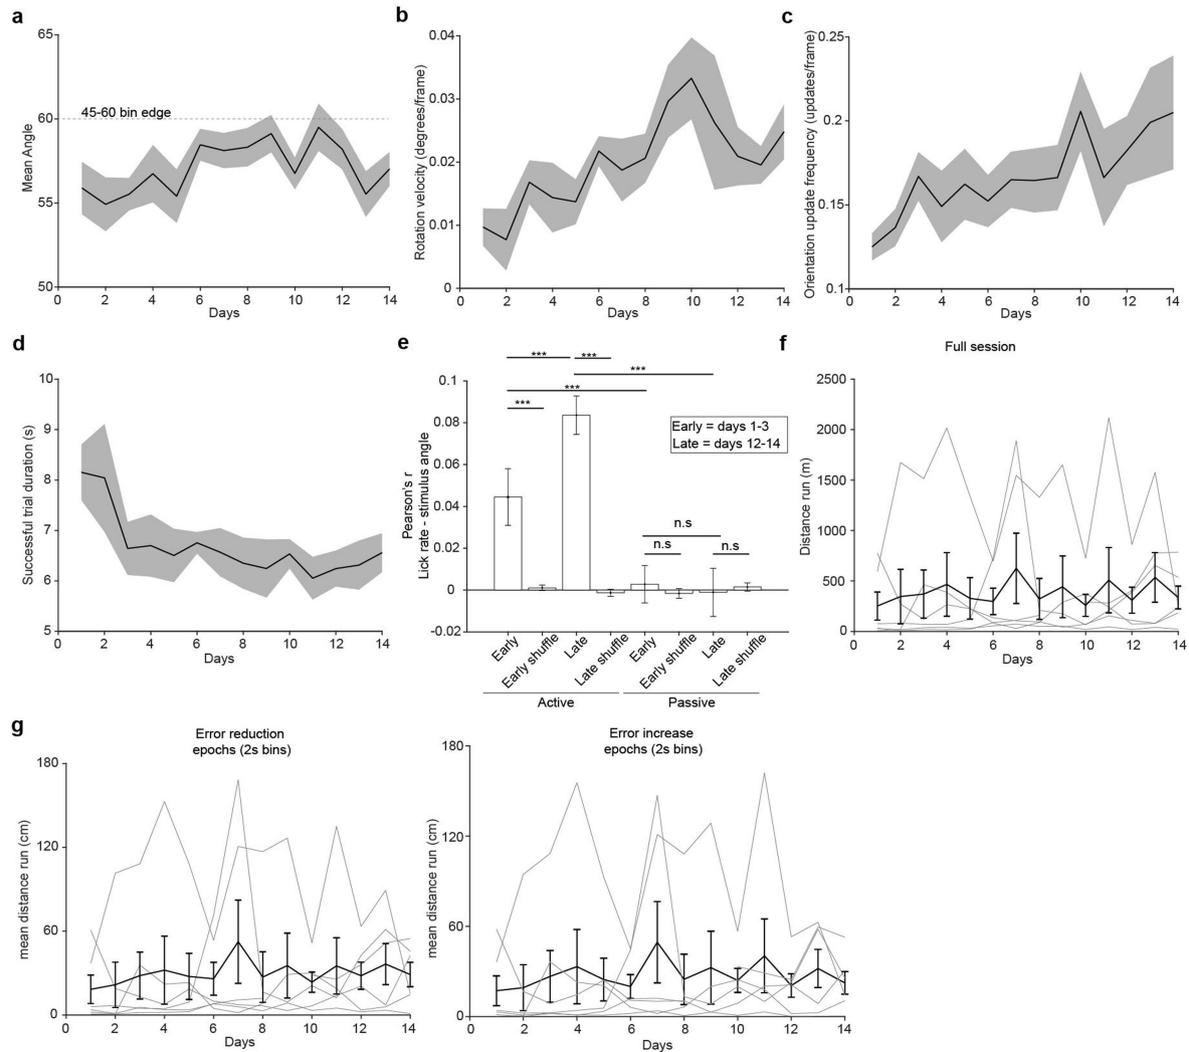


**Extended Data Fig. 1 | Examples of a field of view with P+ and P- neurons.** Examples of a field of view with the same, GCaMP7f-labelled, chronically-tracked P+ and P- neurons (5 each, shown in red and blue, respectively), imaged

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  $\mu\text{m}$ .



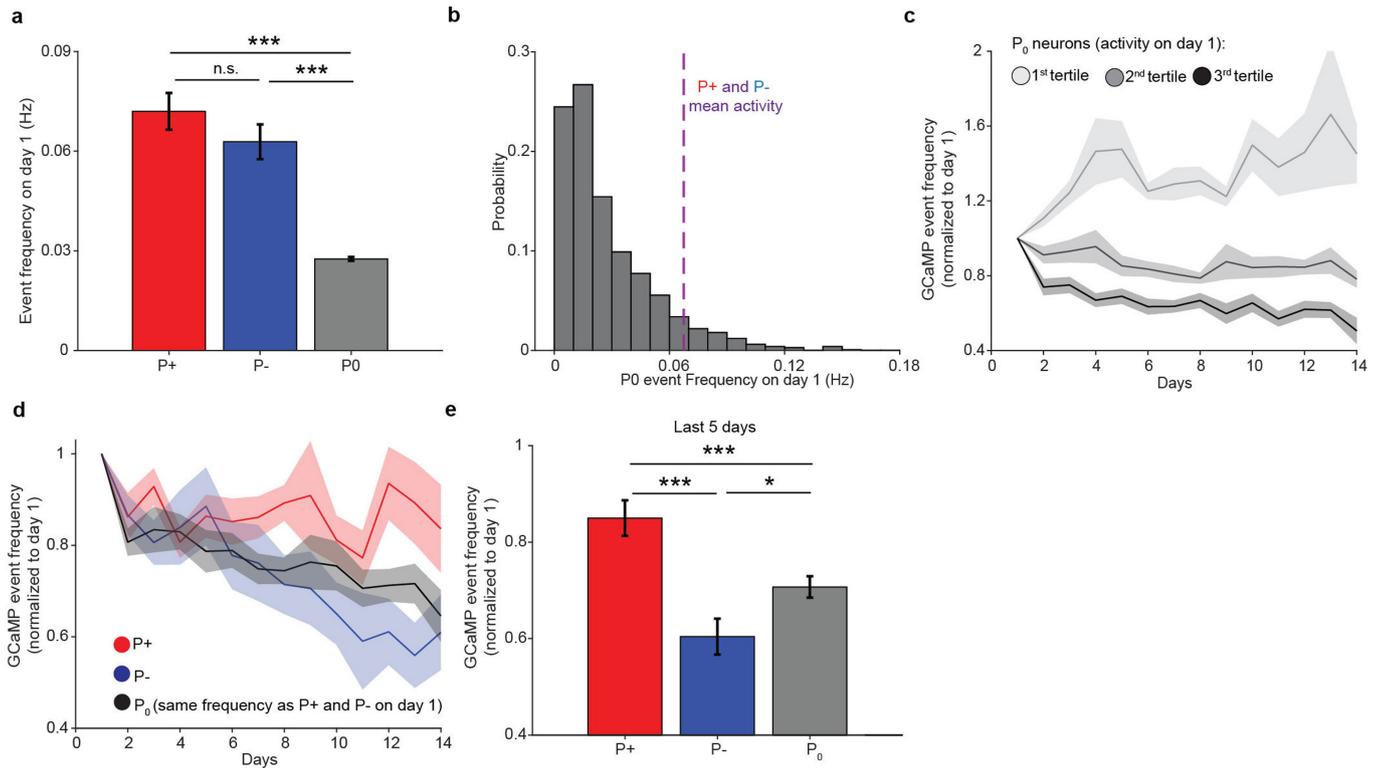
**Extended Data Fig. 2 | Task performance for individual animals.** Task performance evaluated using two metrics: accuracy (the fraction of successful trials) and rewards per minute for each of 6 mice. Summary graph and statistics in Fig. 1g. The grey dashed line represents chance accuracy.



**Extended Data Fig. 3 | Visual stimulus and behavioral correlates of learning.**

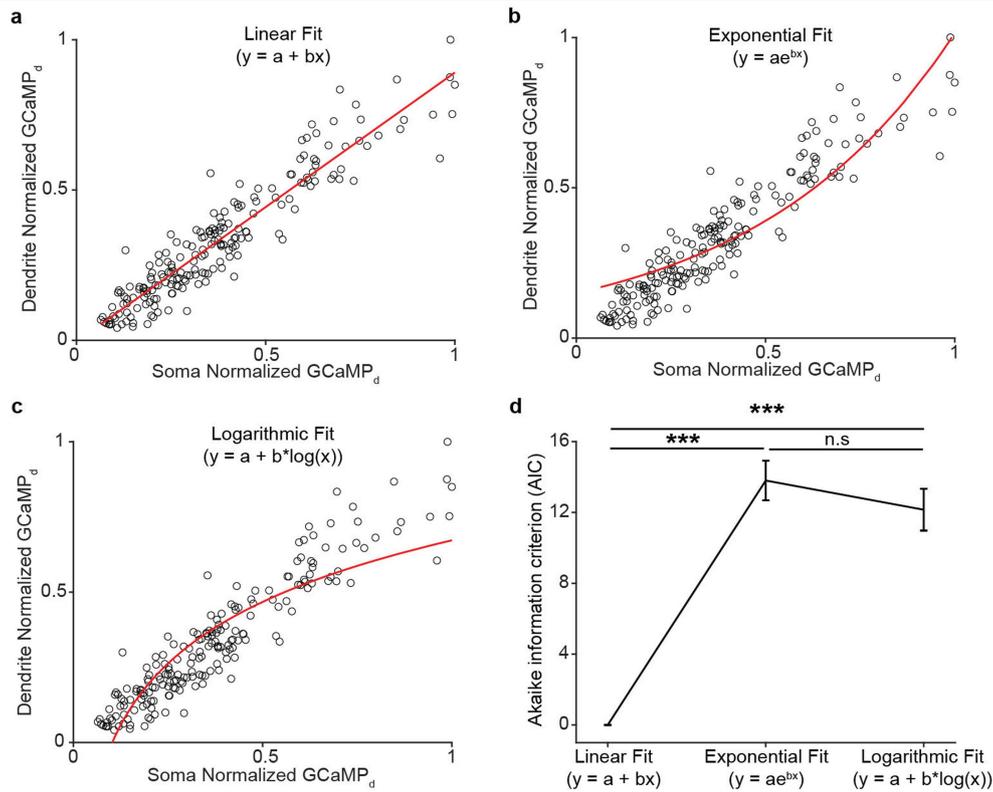
**a**, Mean visual stimulus angle across 14 days of learning. Due to binning, stimuli below the 60-degree threshold were presented as 45-degrees. **b**, Rotation speed across the 14 days of learning. Rotations towards and away from a reward were defined as positive and negative rotations, respectively. **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,  $p = 0.13, 1.4 \times 10^{-11}$  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,  $p = 8.5 \times 10^{-4}$  for early-active vs late-active,  $3.4 \times 10^{-4}$  for early-active vs early-active (shuffle),  $8.4 \times 10^{-11}$  for late-active vs late-active (shuffle),  $1.7 \times 10^{-4}$  for early-active vs early passive,  $5 \times 10^{-11}$  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).  $N = 15$  for all conditions). Error bars and shaded areas are s.e.m.. **f**, 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 min or 76500 imaging frames (2 planes at 30 Hz)). In grey, the distance run by individual mice (One-way repeated measures ANOVA,  $p = 0.58, n = 6$  across 14 days). **g**, 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 (left 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). Error bars represent s.e.m.



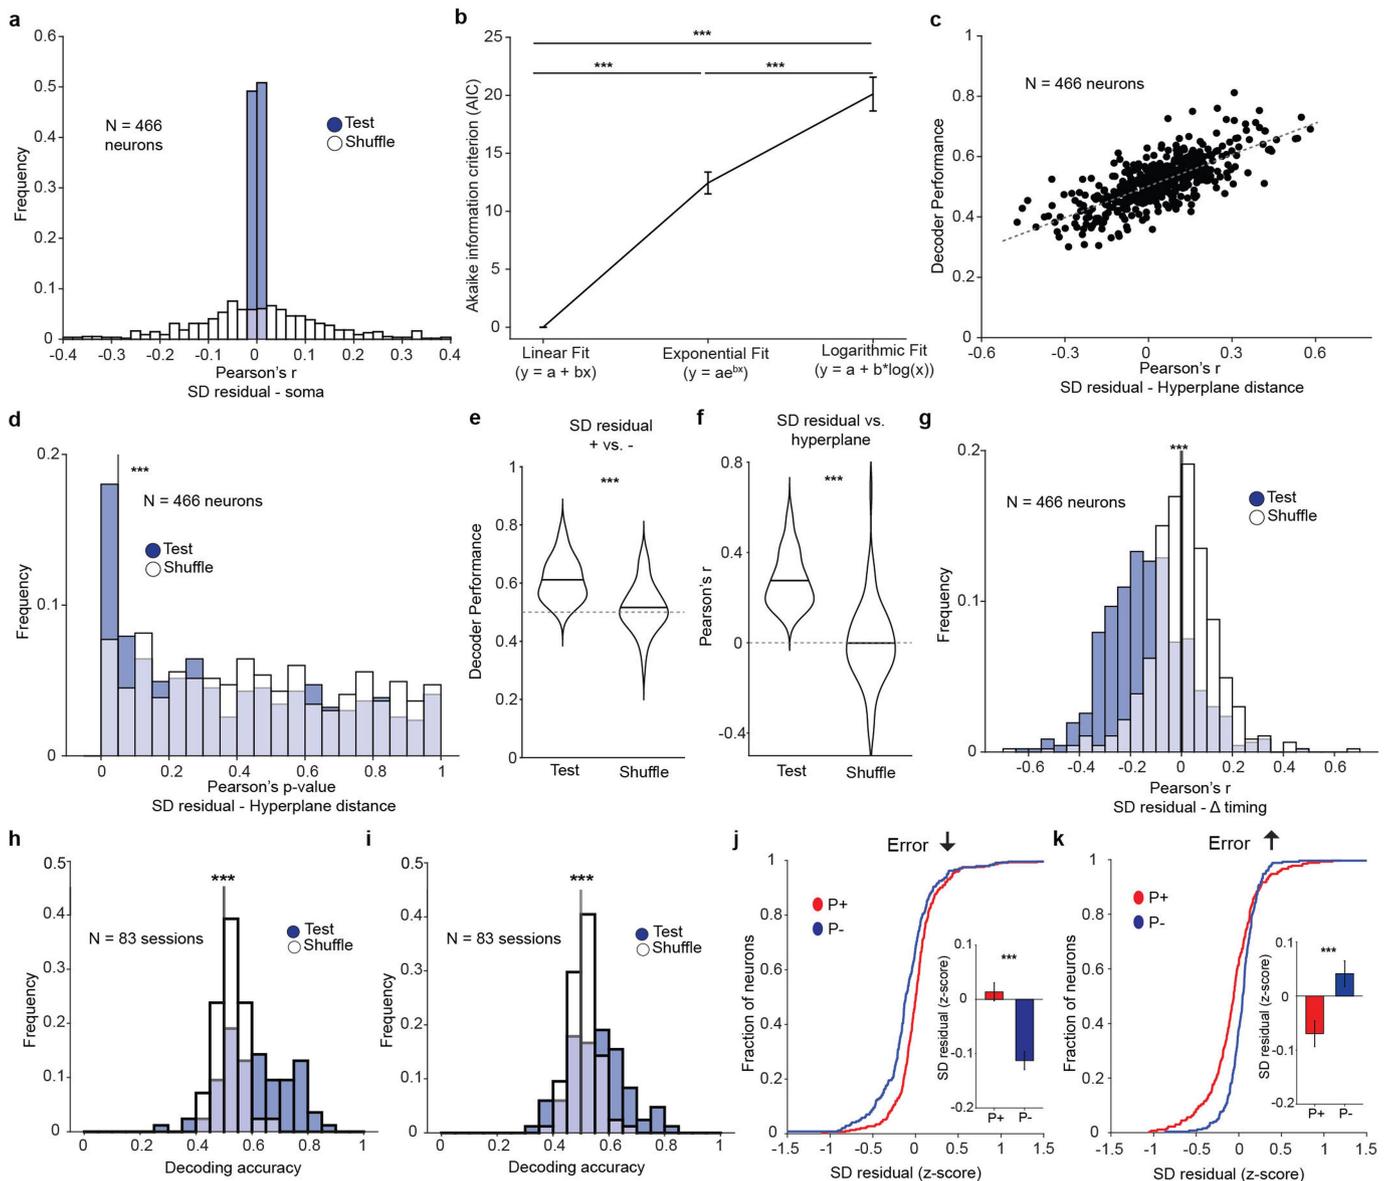
**Extended Data Fig. 4 | Changes in GCaMP transient frequency over days depend on day 1 frequency.** **a**, Event frequency for P+, P- and P<sub>0</sub> neurons on day 1. P+ and P- neurons were selected to be of higher activity on day 1 compared to other neurons in the field of view (One-way ANOVA,  $p = 2.52e^{-30}$ . After Tukey's multiple comparisons correction,  $p = 0.35$ ,  $9.57e^{-10}$  and  $9.56e^{-10}$  for P+ vs P-, P+ vs P<sub>0</sub> and P- vs P<sub>0</sub>, respectively, mean = 0.072, 0.063 and 0.028; s.e.m. = 0.006, 0.005 and  $5.6e^{-4}$ ;  $n = 27$ , 27 and 1839 for P+, P- and P<sub>0</sub>, respectively). **b**, Distribution of event frequency for P<sub>0</sub> neurons on day 1. The mean event frequency for P+ and P- neurons (purple, dashed line) fall on the 8<sup>th</sup> percentile of the P<sub>0</sub> neurons distribution. **c**, Calcium transient frequency for P<sub>0</sub> 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). Shading represents s.e.m. **d**, Calcium transient frequency for P+, P- and P<sub>0</sub> neurons. Only P<sub>0</sub> 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. Mean values and s.e.m. are shown). **e**, Calcium transient frequency for P+, P- and P<sub>0</sub> 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<sub>0</sub> and P- vs P<sub>0</sub>, respectively; mean = 0.85, 0.60 and 0.71, s.e.m. = 0.037, 0.037 and 0.02 for P+, P- and P<sub>0</sub>, respectively.  $n = 6$  mice).



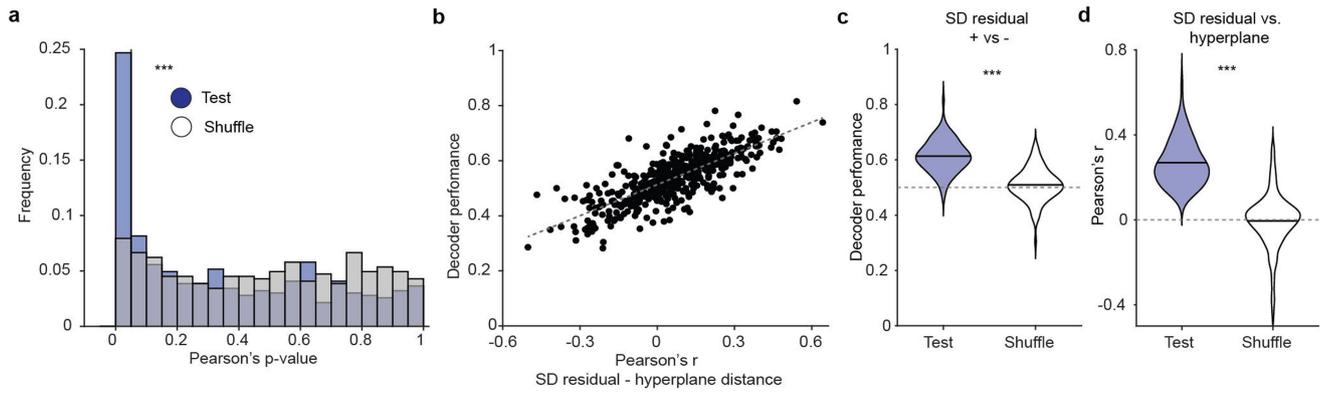
**Extended Data Fig. 5 | Testing the assumption of linearity between somatic and dendritic event magnitude.** **a**, Representative example of a single neuron whose events were isolated and their magnitude estimated in the soma and dendrites, plotted against each other. Axes are normalized to the largest 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; s.e.m. = 0, 1.11 and 1.17;  $n = 543$  neurons for Linear, Exponential and logarithmic fit, respectively).



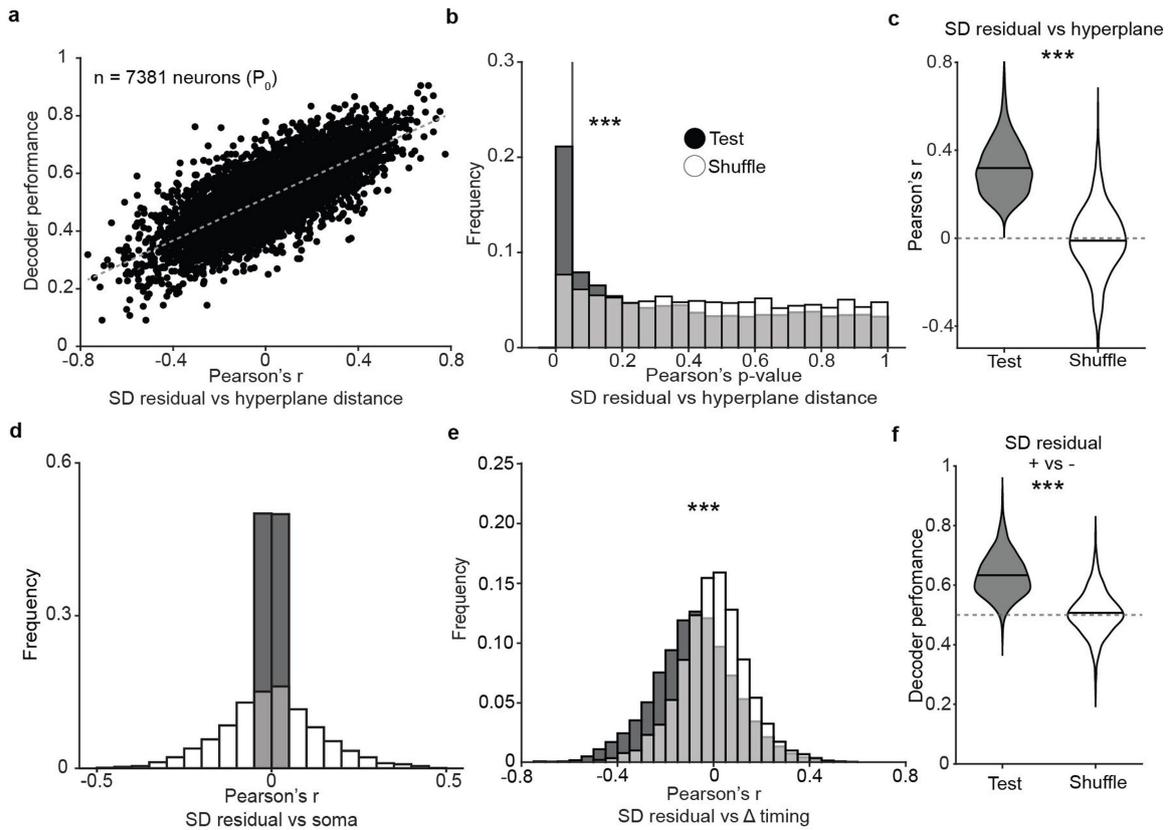
**Extended Data Fig. 6 | SD residual analysis using  $\Delta F/F_0$ -based event size estimation.** **a**, Same as in Fig. 2i. Our  $\Delta F/F_0$ -based approach to estimate event size (see Methods and Supplementary Fig. 2) decorrelated the SD residual from somatic magnitude. **b**, Same as Supplementary Fig. 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; s.e.m. = 0, 0.9 and 1.4;  $n = 543$  neurons for Linear, Exponential and logarithmic fit, respectively). **c**, Same as Fig. 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 Fig. 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 Fig. 2j. Decoder performance (linear SVM) for statistically-significant neurons in d (paired t-test,  $p = 8.3e^{-10}$ .

Mean = 0.61 and 0.51; s.e.m. = 0.008 and 0.01;  $n = 66$ ). **f**, Same as Fig. 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  $-9.6e^{-4}$ ; s.e.m. = 0.01 and 0.02;  $n = 66$ ). **g**, Same as Fig. 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; s.e.m. = 0.007 and 0.006;  $n = 466$ ). **h, i**, Same as Fig. 4e (left panel, paired t-test,  $p = 4.3e^{-9}$ , Mean = 0.61 and 0.52; s.e.m. = 0.01 and 0.01, respectively for test and shuffle data.  $N = 83$  sessions) and 4g (right panel, paired t-test,  $p = 2.4e^{-4}$ , Mean = 0.55 and 0.50; s.e.m. = 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. **j**, Same as Fig. 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  $-0.11$ ; s.e.m. = 0.02 and 0.02;  $n = 292$  and 240 for P+ and P- neurons, respectively). **k**, Same as Fig. 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; s.e.m. = 0.02 and 0.01;  $n = 267$  and 249 for P+ and P- neurons, respectively).



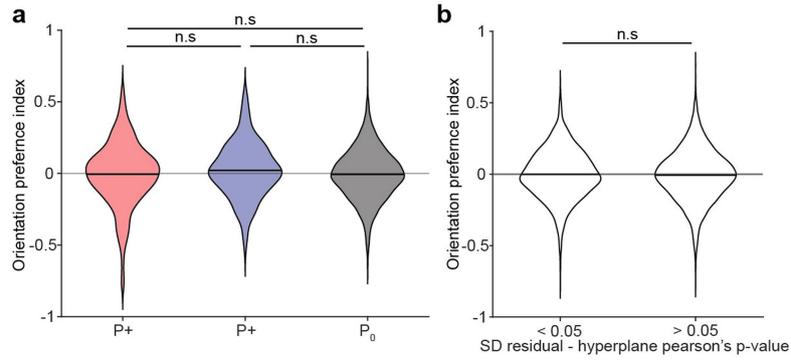
**Extended Data Fig. 7 | Decoding SD residuals using network activity in a single frame preceding event onset. a**, Decoder performance as a function of the correlation between SD residuals and hyperplane distance (Pearson's  $r = 0.77$ ;  $p$ -value =  $7.1e^{-94}$ ,  $n = 466$  neurons). Data points represent 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; s.e.m. = 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  $-0.005$ ; s.e.m. = 0.01 and 0.01;  $n = 99$ ).



**Extended Data Fig. 8 | 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  $< 2.2e^{-308}$ ,  $n = 7381$  neurons). Data points represent 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

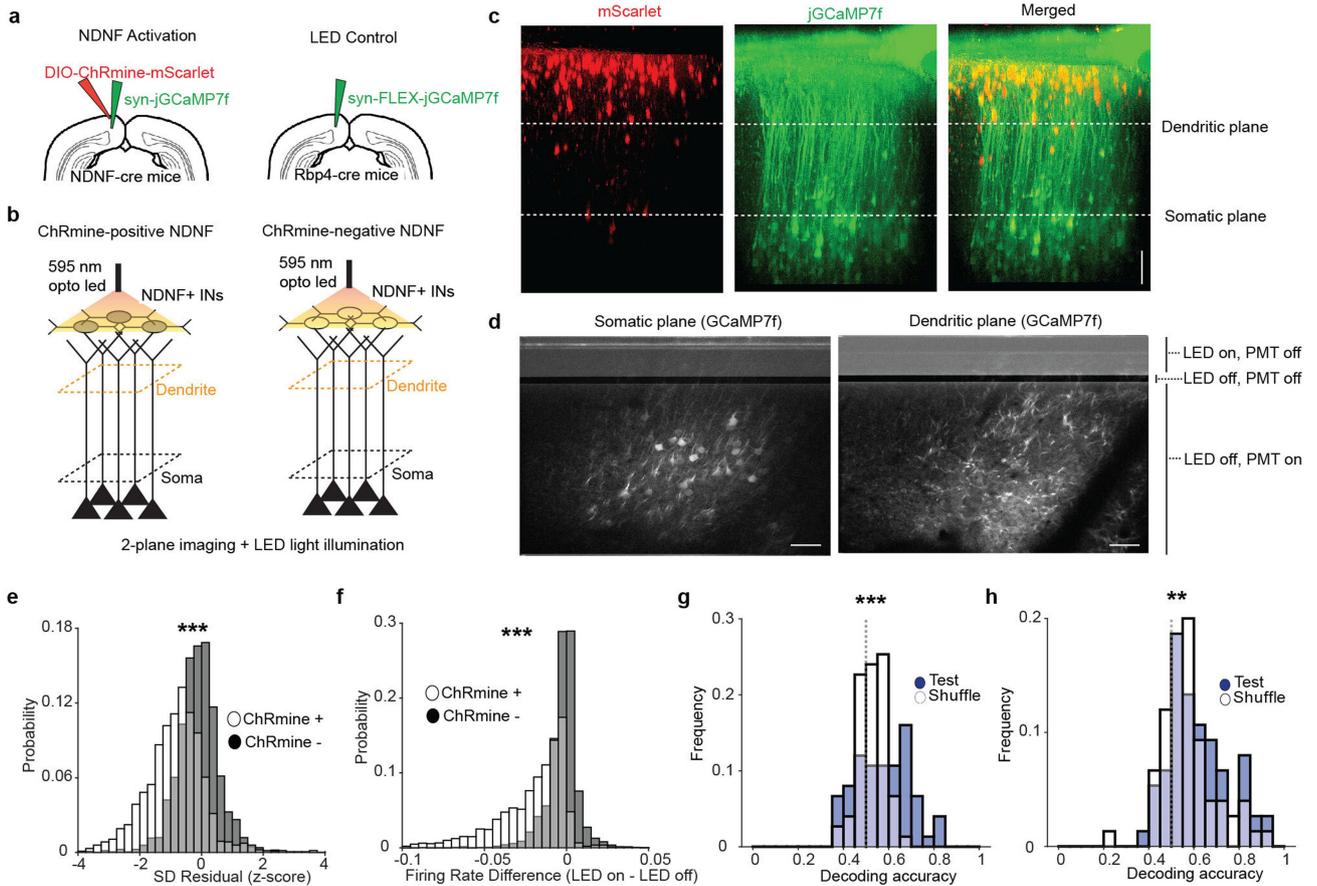
(paired t-test,  $p < 2.2e^{-308}$ , Mean = 0.31 and  $-0.01$ ; s.e.m. = 0.003 and 0.005). **d**, For all neurons, the Pearson's  $r$  for somato-dendritic residuals and somatic event magnitude. **e**, Pearson correlation value between the SD residual the event latency between soma and its corresponding 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  $-6.9e^{-4}$ ; s.e.m. = 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; s.e.m. = 0.002 and 0.002). Dashed grey line indicates chance level.



**Extended Data Fig. 9 | 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<sub>0</sub> (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<sub>0</sub> and P- vs P<sub>0</sub>, respectively. Mean =  $-0.005, 0.02$  and  $-0.007$ ; s.e.m. =  $0.01, 0.01$

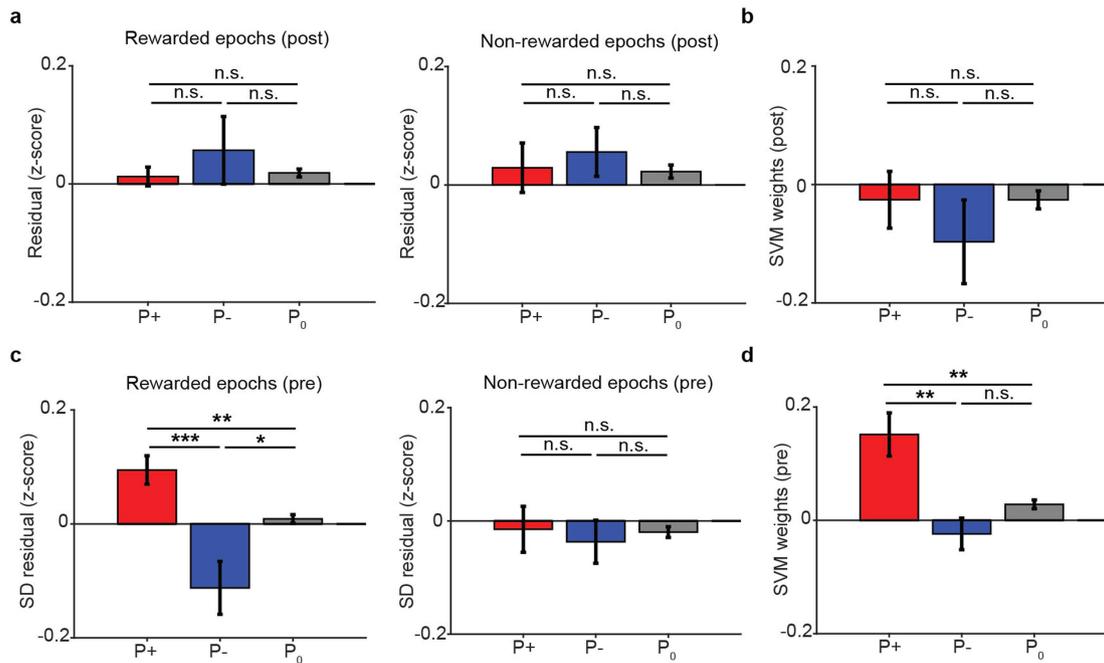
and  $0.02$ ;  $n = 268, 198$  and  $7381$  for P+, P- and P<sub>0</sub>, respectively). **b**, Orientation preference index for all neurons with a statistically-significant Pearson's correlation between the SD residual and the network distance from the maximally-separating hyperplane (t-test,  $p = 0.31$ . Mean =  $-0.002$  and  $-0.007$ ; s.e.m. =  $0.004$  and  $0.002$ ;  $n = 1298$  and  $6547$  for significant and non-significant neurons, respectively).



**Extended Data Fig. 10 | LED illumination alone cannot explain changes to somatic event rate, SD residual, or decoding of task-related variables.**

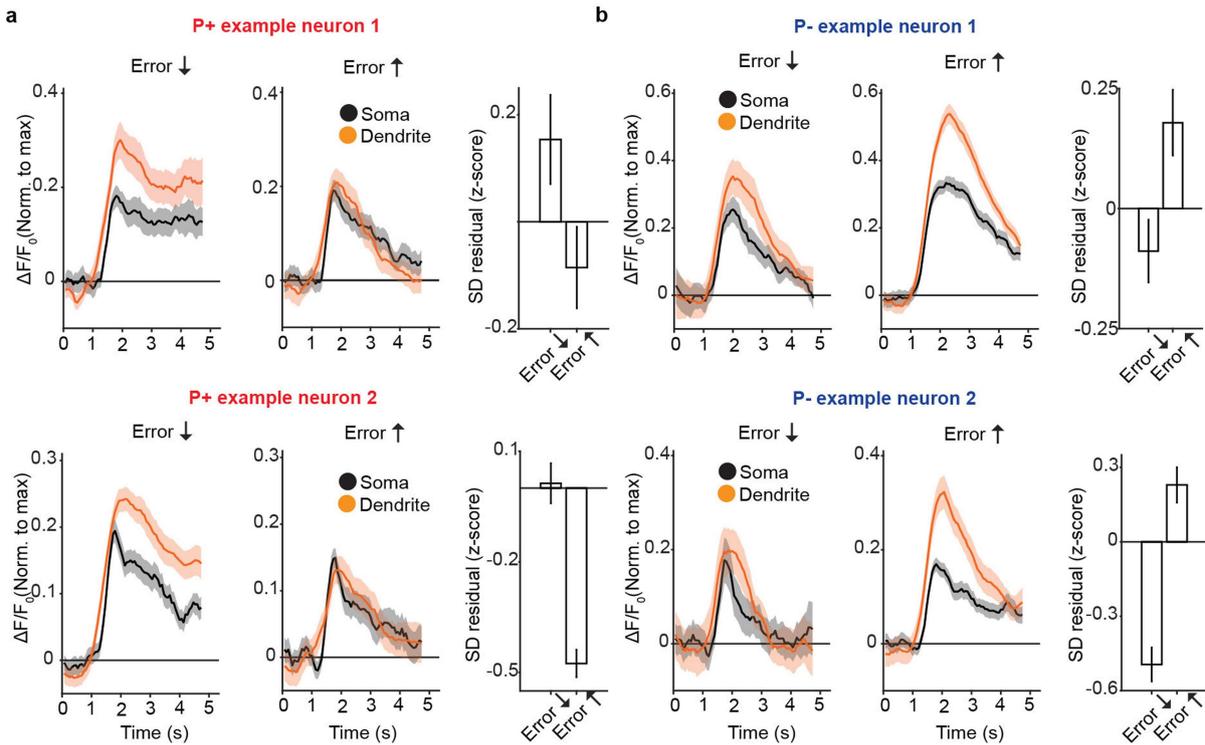
**a,b**, Viral targeting and imaging/LED stimulation strategy schematic for the activation of NDNF+ interneurons in 4 mice (56 imaging sessions), and for a control group of 9 Rbp4 mice (42 imaging sessions) expressing GCaMP7f in layer 5 neurons and exposed to LED light stimulation without expressing ChRmine in NDNF+ interneurons. **c**, Same as Fig. 3g, but for NDNF activation experiments. White dashed lines indicate the somatic and dendritic field of view as shown in **d**. Scale bar = 50  $\mu\text{m}$ . **d**, Somatic and dendritic field of view as indicated by the dashed lines in **a**. Layer 5 neurons labelled with GCaMP7f. The two images show the artefact produced by the simultaneous imaging and LED illumination for optogenetic stimulation. LED light activates for 6 ms at the beginning of each frame while PMTs stay off for 1 extra ms (7 ms total PMT off time). After the first 7 ms, the LED light turns off while PMTs engage to record GCaMP7f activity for the remainder of the frame. Scale bar = 100  $\mu\text{m}$ . **e**, Probability distributions of the SD residual for the two groups of mice. For each neuron, the relationship between somatic and dendritic activity was first established

during the opto-off condition, and then all events during opto-on were superimposed onto this relationship. An SD residual of 0 means that there is no difference in SD residual during LED light on vs LED light off (t-test,  $p = 3.37e^{-136}$ . Mean = -0.1 and -0.9; s.e.m. = -0.02 and 0.02;  $n = 2087$  and 1886 neurons for ChRmine + and ChRmine - groups, respectively). **f**, Somatic event rate in the groups mice estimated as event rate (Hz) during opto on - opto off (t-test,  $p = 1.6e^{-297}$ . Mean = -0.003 and -0.02; s.e.m. = -0.0002 and 0.0004;  $n = 3481$  and 2884 neurons for ChRmine +ve and ChRmine -ve groups, respectively). **g**, Same as Fig. 4j (left panel, paired t-test,  $p = 0.001$ . Mean = 0.62 and 0.57; s.e.m. = 0.02 and 0.02, respectively for test and shuffle data.  $N = 64$  sessions from 5 mice) and for a control group of 5 Rbp4 mice expressing GCaMP7f in layer 5 neurons and exposed to LED stimulation without expressing ChRmine in NDNF+ interneurons. **h**, Same as Fig. 4l (right panel, paired t-test,  $p = 0.001$ . Mean = 0.58 and 0.52; s.e.m. = 0.01 and 0.001, respectively for test and shuffle data.  $N = 65$  sessions in 5 mice) for a control group of 5 Rbp4 mice expressing GCaMP7f in layer 5 neurons and exposed to LED stimulation without expressing ChRmine in NDNF+ interneurons.



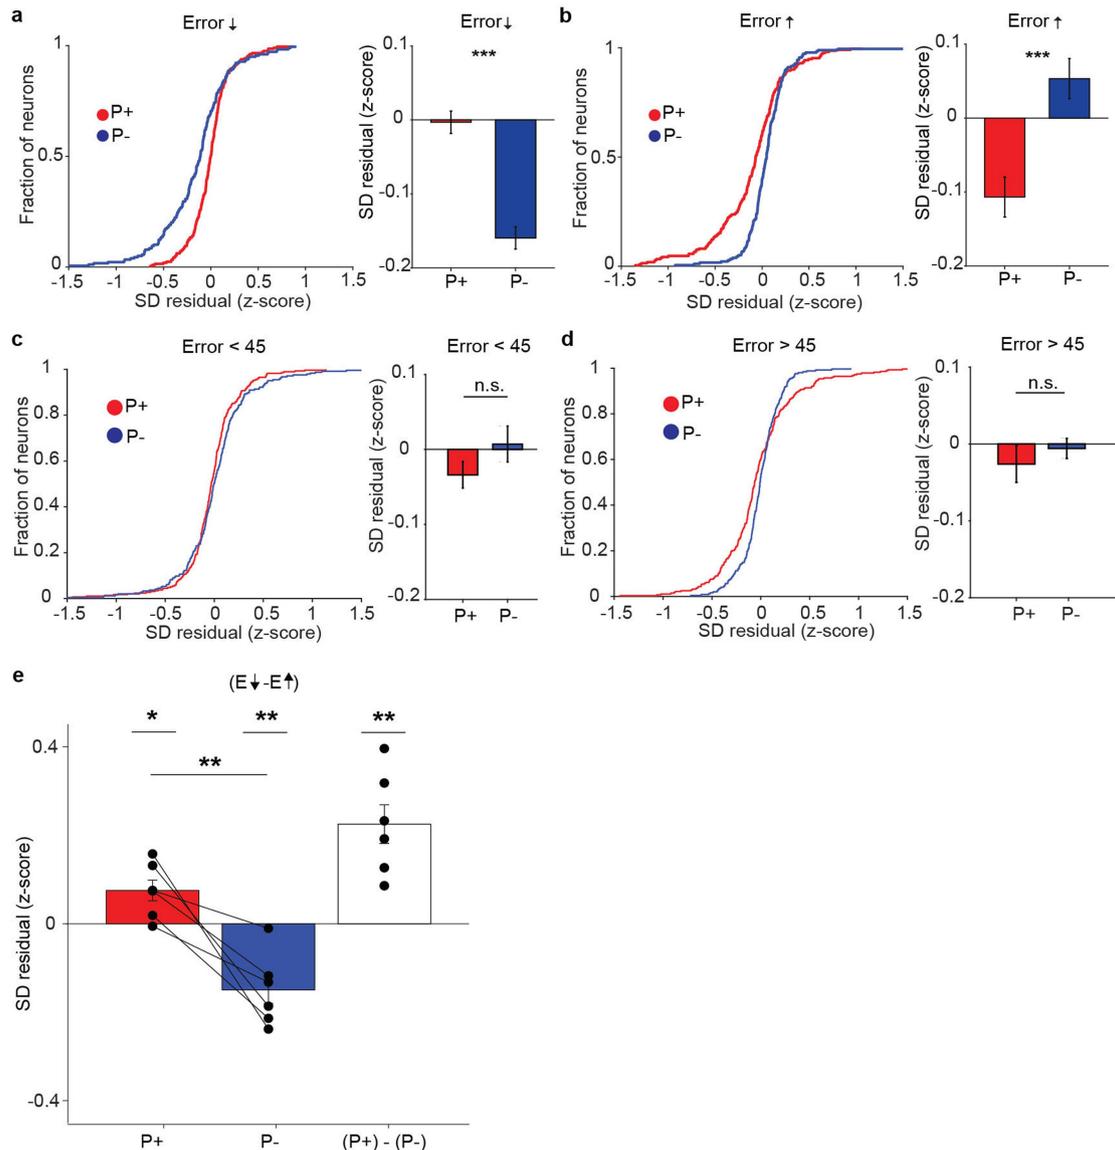
**Extended Data Fig. 11 | Dendritic reward signals are cell-specific during periods of reward anticipation but not during reward consumption.** **a**, SD residuals in the two seconds following target reach for rewarded trials (left) and the end of a trial for rewarded and unrewarded trials (right), respectively, for P+, P- and P<sub>0</sub> neurons (rewarded epochs, One-way repeated measures ANOVA,  $p = 0.5$ . After Tukey's multiple comparison  $p = 0.75$ ,  $p = 0.93$  and  $p = 0.78$  for P+ vs P-, P+ vs P<sub>0</sub> and P- vs P<sub>0</sub>, respectively. Mean = 0.01, 0.06 and 0.02; s.e.m. = 0.02, 0.06 and 0.01 for P+, P- and P<sub>0</sub>, 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<sub>0</sub> and P- vs P<sub>0</sub>, respectively. Mean = 0.03, 0.06 and 0.02, s.e.m. = 0.04, 0.04 and 0.01 for P+, P- and P<sub>0</sub>, respectively.  $n = 84$  sessions). **b**, SVM weights for decoding rewarded vs unrewarded trials after the end of a trial for P+, P- and P<sub>0</sub> 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<sub>0</sub> and P- vs P<sub>0</sub>, respectively. Mean = -0.03, -0.1 and -0.03; s.e.m. = 0.05, 0.07

and 0.2 for P+, P- and P<sub>0</sub>, 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<sub>0</sub> 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<sub>0</sub> and P- vs P<sub>0</sub>, respectively. Mean = 0.09, -0.11 and 0.01, s.e.m. = 0.03, 0.05 and 0.01 for P+, P- and P<sub>0</sub>, 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<sub>0</sub> and P- vs P<sub>0</sub>, respectively. Mean = -0.01, -0.04 and -0.01, s.e.m. = 0.04, 0.04 and 0.01 for P+, P- and P<sub>0</sub>, respectively.  $n = 84$  sessions). **d**, SVM weights for decoding rewarded vs unrewarded trials for P+, P- and P<sub>0</sub> 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<sub>0</sub> and P- vs P<sub>0</sub>, respectively. Mean = 0.15, -0.02 and 0.03, s.e.m. = 0.04, 0.03 and 0.01 for P+, P- and P<sub>0</sub>, respectively.  $n = 84$  sessions).



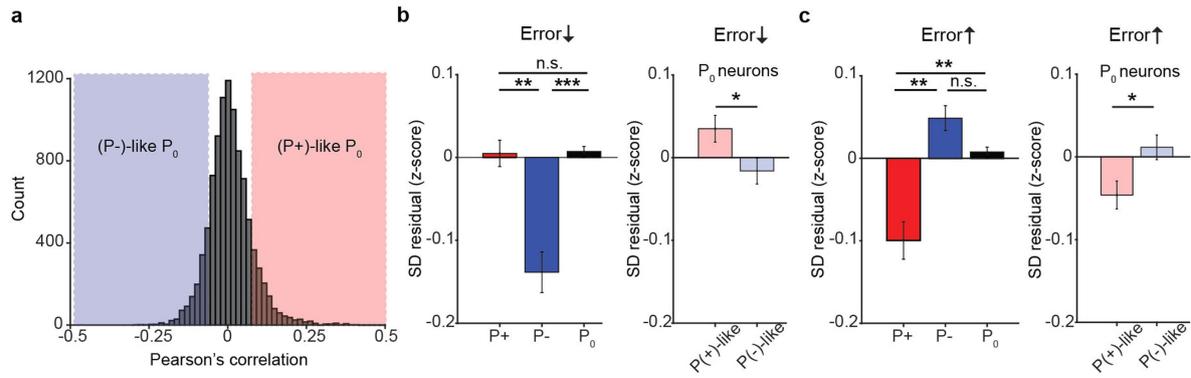
**Extended Data Fig. 12 | Four additional example neurons (2 P+ and 2 P-) showing different somato-dendritic relationships during error decrease and error increase epochs.** For two P+ and two P- neurons, mean  $\Delta F/F_0$  signal in the soma (black) and in the dendrite (orange) are shown for all 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 s.e.m..



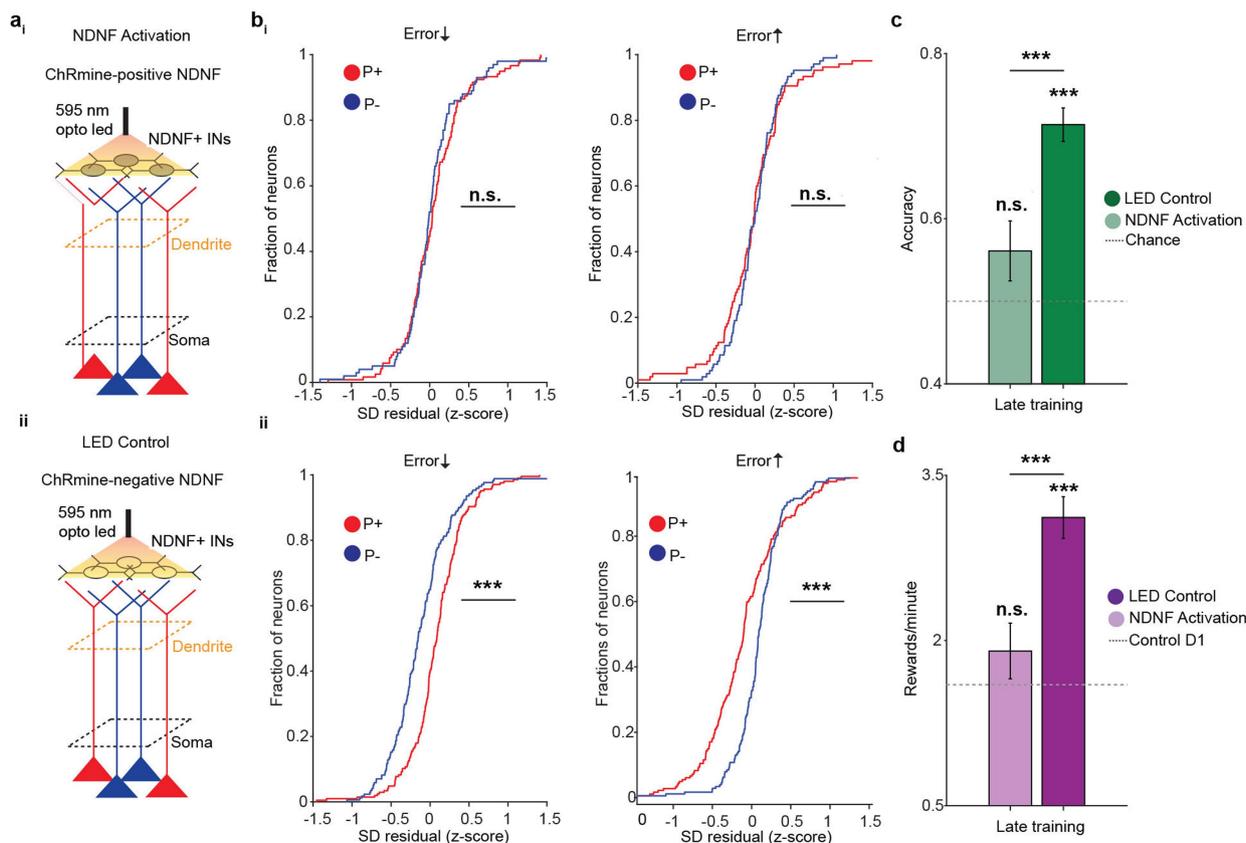
**Extended Data Fig. 13 | Dendritic error signals cannot be explained away by differences in somatic event magnitude, do not represent absolute error values, and are consistent across mice. a, b,** The cumulative distribution function for SD residuals (z-scored per neuron) in P+ (red) and P- (blue) neurons during error decrease in a (t-test,  $p = 3.5 \times 10^{-6}$ , mean = -0.003 and -0.16; s.e.m. = 0.015 and 0.03;  $n = 211$  and 179 for P+ and P-) and error increase in b (t-test,  $p = 4.6 \times 10^{-6}$ , mean -0.1 and 0.05; s.e.m. = 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. **c,** 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; s.e.m. = 0.017 and 0.02;  $n = 290$  and 244 for P+ and P- neurons, respectively). **d,** 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; s.e.m. = 0.02 and 0.01;  $n = 278$  and 249 for P+ and P- neurons, respectively). **e.** 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; s.e.m. = 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).



**Extended Data Fig. 14 | P<sub>0</sub> neurons which are functionally-correlated to P+ and P- neurons also receive vectorized dendritic error signals.** **a**, Across all sessions, we divided P<sub>0</sub> 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<sub>0</sub> neurons (n = 9796) we defined (P+)- and (P-)-like P<sub>0</sub> 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<sub>0</sub> neurons (One-way ANOVA,  $p = 1.6 \times 10^{-4}$ . After Tukey's multiple comparisons,  $p = 0.006$ ,  $p = 1$  and  $p = 8.9 \times 10^{-5}$  for P+ vs P-, P+ vs P<sub>0</sub> and P- vs P<sub>0</sub>, respectively. Mean = 0.004, -0.13 and 0.007; s.e.m. = 0.02, 0.02 and 0.01; n = 292, 240 and 8335 for P+, P- and P<sub>0</sub>, respectively). Right panel, during error reduction epochs, bar

graph comparing the SD residual (z-scored) in (P+)- and P(-)-like P<sub>0</sub> neurons (t-test,  $p = 0.025$ ; mean = 0.04 and -0.02; s.e.m. = 0.02 and 0.02; n = 1059 and 1110 for (P+)- and P(-)-like P<sub>0</sub> neurons, respectively). **c**, Left panel, during error increase epochs, bar graph comparing the SD residual (z-scored) in P+, P- and P<sub>0</sub> 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<sub>0</sub> and P- vs P<sub>0</sub>, respectively. Mean = -0.1, 0.05 and 0.01; s.e.m. = 0.02, 0.02 and 0.01; n = 267, 249 and 8356 for P+, P- and P<sub>0</sub>, respectively). Right panel, during error increase epochs, bar graph comparing the SD residual (z-scored) in (P+)- and P(-)-like P<sub>0</sub> neurons (t-test,  $p = 0.010$ ; mean = -0.05 and 0.01; s.e.m. = 0.02 and 0.01; n = 1070 and 1099 for (P+)- and P(-)-like P<sub>0</sub> neurons, respectively).



**Extended Data Fig. 15 | LED illumination alone is not responsible for the abolishment of vectorized error signals and impaired learning.**

**a**, Experimental schematic: P+ and P- neurons were recorded during the activation of NDNF+ interneurons (N = 4 mice) (a<sub>i</sub>) and in a control group of 5 Rbp4 mice expressing GCaMP7f in layer 5 neurons and exposed to LED stimulation without expressing ChRmine in NDNF+ interneurons (a<sub>ii</sub>). **b**, Upper panel (b<sub>i</sub>). Same as Fig. 5g. The activation of L1 NDNF+ neuron leads to the abolishment of vectorized error signals in P+ and P- neurons. Left, 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 (t-test;  $p = 0.58$ ; mean = 0.06 and 0.02; s.e.m. = 0.04 and 0.04; n = 119 and 100 for P+ and P- neurons respectively from 4 mice). Right panel, for error increase epochs (t-test,  $p = 0.7$ ; mean = -0.02 and 0; SEM = 0.05 and 0.03; n = 105 and 105 for P+ and P- neurons respectively). Lower panel, the activation of the LED light in Rbp4-Cre animals expressing GCaMP7f in layer 5 neurons, but not ChRmine in NDNF+ interneurons (b<sub>ii</sub>). Left, 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 (t-test;  $p = 8.2e^{-7}$ ; mean = 0.08 and -0.1; s.e.m. = 0.03 and 0.03; n = 206 and 171 for P+ and P- neurons respectively from 5 mice). Right panel, for error increase epochs (t-test,  $p = 5.9e^{-4}$ ; mean = -0.06 and 0.1; s.e.m. = 0.04 and 0.03; n = 182 and 177 for P+ and P- neurons respectively). **c, d**, BCI performance (accuracy in c and rewards per minute in d) for the last 6 days of training (late training) in the two groups of animals: the ones expressing ChRmine in NDNF+ interneurons (n = 4) and their ChRmine-negative LED illumination counterparts (n = 5) (for accuracy in c, t-test;  $p = 3.1e^{-4}$ ; mean = 0.56 and 0.71; s.e.m. = 0.04 and 0.02; n = 24 and 30 for NDNF+ neurons activation and LED control, respectively; paired t-test,  $p = 0.1$  for NDNF+ neurons activation vs chance and  $p = 2e^{-11}$  for LED control vs chance. For rewards/minute in d, t-test;  $p = 2.6e^{-4}$ ; mean = 1.9 and 3.1; s.e.m. = 0.2 and 0.3; n = 24 and 30 for NDNF+ neurons activation and LED control, respectively, paired t-test,  $p = 0.22$  for NDNF+ neurons activation vs control day 1 and  $p = 6.8e^{-9}$  for LED control vs control day 1).

Corresponding author(s): Mark HarnettLast updated by author(s): Jan 16, 2026

## Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

### Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a | Confirmed

- The exact sample size ( $n$ ) for each experimental group/condition, given as a discrete number and unit of measurement
- A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- The statistical test(s) used AND whether they are one- or two-sided  
*Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
- A description of all covariates tested
- A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- For null hypothesis testing, the test statistic (e.g.  $F$ ,  $t$ ,  $r$ ) with confidence intervals, effect sizes, degrees of freedom and  $P$  value noted  
*Give  $P$  values as exact values whenever suitable.*
- For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- Estimates of effect sizes (e.g. Cohen's  $d$ , Pearson's  $r$ ), indicating how they were calculated

*Our web collection on [statistics for biologists](#) contains articles on many of the points above.*

### Software and code

Policy information about [availability of computer code](#)

Data collection

Scanbox Version 3  
Psychophysics Toolbox Version 3

Data analysis

Matlab 2020a  
Suite2P (<https://github.com/MouseLand/suite2p>)  
Cascade (<https://github.com/HelmchenLabSoftware/Cascade>)  
ROIMatchPub (<https://github.com/ransona/ROIMatchPub>)  
Custom-built BCI software (Matlab 2020a)  
Custom-built analysis pipeline (Matlab 2020a)  
Graphpad Prism 8.4

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio [guidelines for submitting code & software](#) for further information.

## Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our [policy](#)

Analysis code and data are available upon request. Code for running the brain-computer interface is available at <https://github.com/harnett/CLOSE>.

## Research involving human participants, their data, or biological material

Policy information about studies with [human participants or human data](#). See also policy information about [sex, gender \(identity/presentation\), and sexual orientation](#) and [race, ethnicity and racism](#).

Reporting on sex and gender	NA
Reporting on race, ethnicity, or other socially relevant groupings	NA
Population characteristics	NA
Recruitment	NA
Ethics oversight	NA

Note that full information on the approval of the study protocol must also be provided in the manuscript.

## Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

- Life sciences       Behavioural & social sciences       Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see [nature.com/documents/nr-reporting-summary-flat.pdf](https://www.nature.com/documents/nr-reporting-summary-flat.pdf)

## Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	No sample size calculation was performed. Sample sizes are comparable to or larger than similar studies (Clancy et al 2014; Neely et al. 2018; Jeon et al. 2022)
Data exclusions	A single imaging session was excluded from the study due to a technical problem at the time of data acquisition (water leakage under the objective) as reported in the manuscript. Inclusion criteria for analysis are clearly described in the manuscript.
Replication	Key findings were replicated across two independent experimental groups.
Randomization	Experimental groups were not randomized. Mice were assigned to specific experimental groups based on their genotypes. All groups were trained on the same brain-computer interface task.
Blinding	Blinding was not performed during data collection or analysis. The brain-computer interface task was automated and consistent across experimental series, and analysis of experimental series were performed using fixed analysis pipelines.

## Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

## Materials &amp; experimental systems

## Methods

- n/a  Involved in the study
- Antibodies
- Eukaryotic cell lines
- Palaeontology and archaeology
- Animals and other organisms
- Clinical data
- Dual use research of concern
- Plants

- n/a  Involved in the study
- ChIP-seq
- Flow cytometry
- MRI-based neuroimaging

## Animals and other research organisms

Policy information about [studies involving animals](#); [ARRIVE guidelines](#) recommended for reporting animal research, and [Sex and Gender in Research](#)

Laboratory animals	Male and female Rbp4-Cre and NDNF-cre heterozygous mice were used for experiments at 8-15 weeks of age. Animals were housed on a 12-h light cycle, at 65-70° F and ~60% humidity.
Wild animals	No wild animals were used in this study.
Reporting on sex	Equal numbers of male and female animals were used in this study.
Field-collected samples	No field collected samples were used in this study.
Ethics oversight	All experiments were compliant with guidance and regulation from the NIH and the Massachusetts Institute of Technology Committee on Animal Care.

Note that full information on the approval of the study protocol must also be provided in the manuscript.

## Plants

Seed stocks	NA
Novel plant genotypes	NA
Authentication	NA