

Synaptic wiring motifs in posterior parietal cortex support decision-making

<https://doi.org/10.1038/s41586-024-07088-7>

Received: 13 April 2022

Accepted: 17 January 2024

Published online: 21 February 2024

 Check for updates

Aaron T. Kuan^{1,5,11}, Giulio Bondanelli^{2,3,11}, Laura N. Driscoll^{1,6}, Julie Han^{1,7}, Minsu Kim¹, David G. C. Hildebrand^{1,8}, Brett J. Graham^{1,9}, Daniel E. Wilson¹, Logan A. Thomas^{1,10}, Stefano Panzeri^{2,3}, Christopher D. Harvey¹ & Wei-Chung Allen Lee^{1,4}

The posterior parietal cortex exhibits choice-selective activity during perceptual decision-making tasks^{1–10}. However, it is not known how this selective activity arises from the underlying synaptic connectivity. Here we combined virtual-reality behaviour, two-photon calcium imaging, high-throughput electron microscopy and circuit modelling to analyse how synaptic connectivity between neurons in the posterior parietal cortex relates to their selective activity. We found that excitatory pyramidal neurons preferentially target inhibitory interneurons with the same selectivity. In turn, inhibitory interneurons preferentially target pyramidal neurons with opposite selectivity, forming an opponent inhibition motif. This motif was present even between neurons with activity peaks in different task epochs. We developed neural-circuit models of the computations performed by these motifs, and found that opponent inhibition between neural populations with opposite selectivity amplifies selective inputs, thereby improving the encoding of trial-type information. The models also predict that opponent inhibition between neurons with activity peaks in different task epochs contributes to creating choice-specific sequential activity. These results provide evidence for how synaptic connectivity in cortical circuits supports a learned decision-making task.

Decision-making is a critical component of behaviour and cognition, and understanding how it is implemented has been a long-standing goal in neuroscience. Experiments in primates have revealed the importance of the neocortex for perceptual decision-making, including the posterior parietal cortex (PPC), in which neuronal activity is predictive of upcoming behavioural choices^{1,2}. Such choice-selective activity has also been found in the rodent PPC^{3–10}. However, it remains unclear how this selective neuronal activity arises. Models of cortical decision-making circuits propose that choice alternatives are represented by pools of recurrently connected excitatory neurons, which compete through inhibitory connectivity^{11,12}. Pioneering work often assumed that inhibitory activity was non-selective and inhibitory connectivity non-specific^{13,14}, but recent research has shown that inhibitory activity is as selective as excitatory activity, which suggests that inhibitory connectivity may follow choice-selective rules¹⁰. Indeed, recent models have shown that selective inhibition can crucially alter circuit function by stabilizing network activity or maximizing competition between opposing excitatory pools¹⁵. However, activity measurements in the PPC are consistent with multiple possible circuit architectures^{10,15}, and measurement of the underlying synaptic connectivity is lacking.

Until recently, direct measurements of synaptic connectivity within large neuronal populations have not been technically feasible. However,

advances in high-throughput electron microscopy (EM) have now made it possible to comprehensively map synaptic connectivity within circuits^{16–21}. Such connectomic approaches in the cortex have focused mainly on sensory areas such as the visual cortex^{22–27}, in which inhibitory activity and connectivity are generally less selective than excitatory^{22,28} (but also see refs. 29,30). As a result, little is known about synaptic connectivity in association areas such as the PPC and how it may differ from the sensory cortex.

Here we combined a decision-making task, two-photon calcium imaging and automated serial-section EM^{3,19,31} to measure how synaptic connectivity of hundreds of cortical neurons relates to their functional selectivity in the PPC. We found selective excitatory-to-inhibitory (E-to-I) and inhibitory-to-excitatory (I-to-E) connectivity: excitatory neurons preferentially targeted inhibitory neurons with the same selectivity, whereas inhibitory neurons preferred excitatory targets with opposite selectivity. Together, these preferences form an opponent inhibition motif, in which neurons associated with one choice suppress the activity of neurons associated with the alternative choice. The opponent-inhibition motif was present even between neurons with activity peaks in different task epochs. To investigate the functional implications of this connectivity motif, we modelled recurrent circuits with excitatory and inhibitory populations. The models predict that

¹Department of Neurobiology, Harvard Medical School, Boston, MA, USA. ²Neural Computation Laboratory, Istituto Italiano di Tecnologia, Genoa, Italy. ³Department of Excellence for Neural Information Processing, University Medical Center Hamburg-Eppendorf (UKE), Hamburg, Germany. ⁴FM Kirby Neurobiology Center, Boston Children's Hospital, Boston, MA, USA. ⁵Present address: Department of Neuroscience, Yale School of Medicine, New Haven, CT, USA. ⁶Present address: Allen Institute for Neural Dynamics, Allen Institute, Seattle, WA, USA. ⁷Present address: Khoury College of Computer Sciences, Northeastern University, Seattle, WA, USA. ⁸Present address: Laboratory of Neural Systems, The Rockefeller University, New York, NY, USA. ⁹Present address: Space Telescope Science Institute, Baltimore, MD, USA. ¹⁰Present address: Biophysics Graduate Group, University of California Berkeley, Berkeley, CA, USA. ¹¹These authors contributed equally: Aaron T. Kuan, Giulio Bondanelli. ✉e-mail: s.panzeri@uke.de; christopher_harvey@hms.harvard.edu; wei-chung_lee@hms.harvard.edu

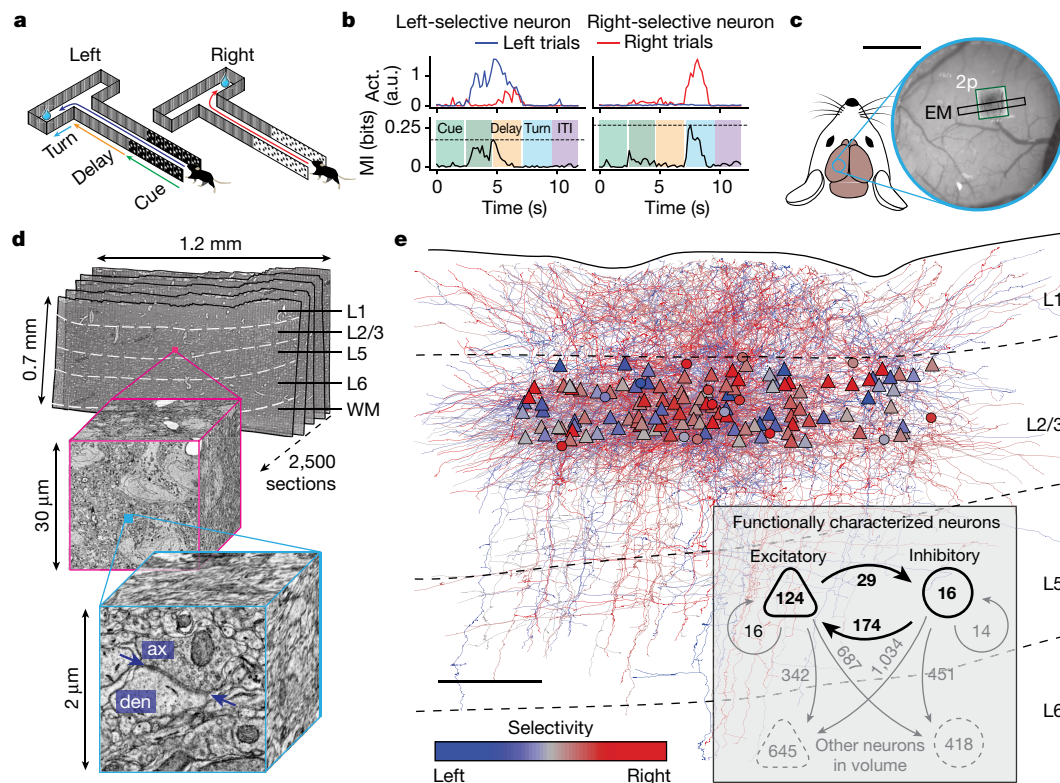


Fig. 1 | Behaviour, functional imaging and EM. **a**, Schematic of decision-making behaviour consisting of a navigational two-alternative forced-choice memory task performed in virtual reality. **b**, The trial-averaged activity (act.) for left (blue) and right (red) trials from two example selective neurons, plotted along with the MI with the trial type (bottom). The magnitude of the maximum selectivity was defined as the maximum value of the MI (dotted horizontal line). a.u., arbitrary units. **c**, Image of the cranial window of the selected mouse, showing the location of overlapping calcium imaging (2p, green) and EM (black) datasets within the PPC. Scale bar, 1 mm. **d**, Schematic of

the volumetric EM dataset consisting of 2,500 serial sections. Insets: images at a progressively higher resolution, highlighting cell bodies (magenta) and an individual synapse (cyan; the arrows indicate the PSD). L, layer; ax, axon; den, dendrite. **e**, Reconstructed circuit in the PPC, consisting of 124 excitatory neurons (triangles) and 16 inhibitory neurons (circles) colour coded by selectivity (key). Inset: summary of the reconstructed circuit, indicating the number of neurons of each type (within shapes) and the number of synaptic connections (next to arrows). The direct E-to-I and I-to-E connections that are analysed in detail in Figs. 2 and 3 are shown in bold. Scale bar, 100 μ m.

opponent inhibition supports amplification of input selectivity and promotes reliable encoding of choices corresponding to each trial type, and that opponent inhibition between neurons of which the activity peaks in different epochs contributes to the creation of choice-specific sequential activity.

Behaviour, imaging and EM

We trained mice to perform a two-alternative forced-choice task in a virtual reality T-maze and used two-photon calcium imaging to measure the activity of layer 2/3 neurons in the left-hemisphere PPC during task performance^{3,31} (Fig. 1a and Extended Data Fig. 1a–c). These behavioural and functional imaging data were included in a previous study³¹. Consistent with previous results^{3,31}, many PPC neurons exhibited temporal activity peaks that were selective for trial type (left or right turn trials; Fig. 1b) and staggered relative to one another in time, forming sequences of neuronal activation that spanned the length of a task trial^{3,31,32} (Extended Data Figs. 1d–f and 2). To quantify trial-type selectivity, we defined the maximum selectivity as the maximum value of the mutual information (MI) between the neuronal activity and trial type across time within the trial (Fig. 1b and Methods). Away from the timepoint of maximum selectivity, neurons often had non-zero activity that typically shared the same left/right preference as the maximum, such that the maximum was generally representative of selectivity throughout the trial (Extended Data Figs. 1e, f and 2). Over the population of neurons, there was a diversity of maximum-selectivity values (Extended Data Figs. 1g, h and 2).

We preserved the brain of one animal immediately after the conclusion of behavioural experiments and used EM to generate a high-resolution structural map of the same neurons of which the activity was previously measured in vivo (Fig. 1c and Methods). We used the GridTape automated transmission EM pipeline¹⁹ to collect and image 2,500 serial 40 nm thin sections and aligned them to form a three-dimensional volume spanning all six cortical layers with about 1.2 mm extent (medial–lateral) and around 100 μ m depth (anterior–posterior). This dataset encompasses approximately 0.1 mm³ at 4.3 nm \times 4.3 nm \times 40 nm per voxel resolution (Fig. 1d). We next co-registered the in vivo and EM data to match calcium-imaging regions of interest to cell bodies in the EM volume (Methods and Extended Data Fig. 1i, j). We were therefore able to relate behaviour, neuronal activity and network anatomy in the PPC.

Using the EM data, we reconstructed the axons and dendrites of the functionally characterized cells within the volume (Fig. 1e) and classified them as excitatory pyramidal cells or inhibitory interneurons (non-pyramidal) on the basis of their morphology ($n = 124$ (pyramidal) and $n = 16$ (non-pyramidal); Extended Data Figs. 3 and 4). Non-pyramidal cells in the PPC were generally as selective as pyramidal cells (Extended Data Fig. 1h; $P = 0.20$, Kolmogorov–Smirnov test), consistent with recent functional imaging experiments¹⁰.

To map the connectivity of the functionally characterized neurons, we annotated all of the outgoing synapses from their axons within the EM volume and traced the corresponding post-synaptic dendrites back to their cell bodies. We identified 233 synapses in which the post-synaptic cell was also a functionally characterized neuron (Fig. 1e

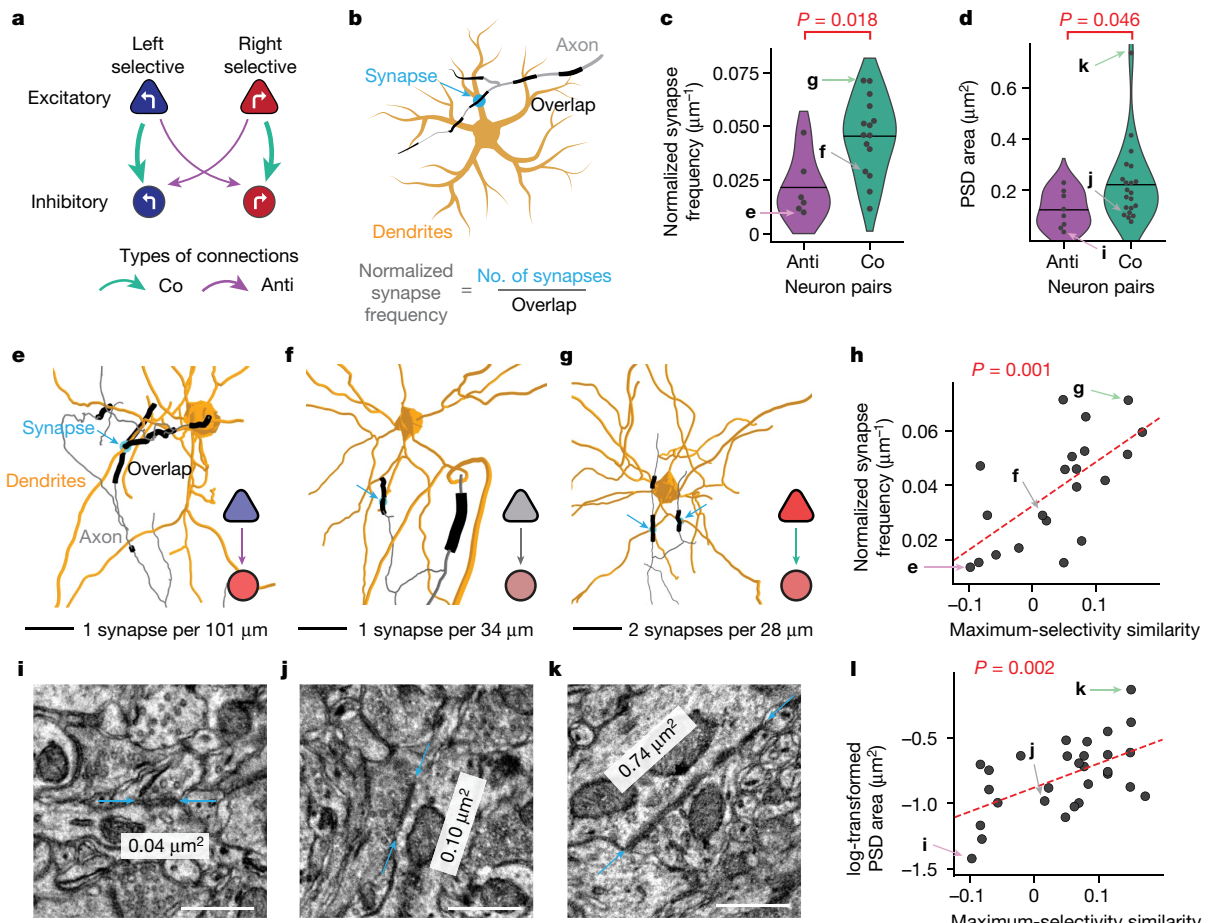


Fig. 2 | Co-selective E-to-I connectivity. **a**, Schematic of E-to-I connections among functionally characterized neurons. Connections are classified as co-selective (co; green) or anti-selective (anti; purple). **b**, Schematic of the axon–dendrite overlap and normalized synapse frequency metric, which quantifies the likelihood of synapses between a specific axon–dendrite pair per μm of overlap (axon path length within $5\ \mu\text{m}$ of the dendrite). Grey, presynaptic axon; orange, post-synaptic dendrites; black, axon–dendrite overlap; cyan arrows, synaptic connections. **c**, The normalized synapse frequency between co-selective neurons (green) is more than twice as frequent as between anti-selective neurons (purple) ($0.021 \pm 0.006\ \mu\text{m}^{-1}$, $n = 6$ connections (anti-selective); $0.045 \pm 0.005\ \mu\text{m}^{-1}$, $n = 15$ (co-selective); $P = 0.018$, Mann–Whitney U -test). **d**, Synapses between co-selective neurons (green) have PSDs almost twice as large as those between anti-selective neurons ($0.12 \pm 0.02\ \mu\text{m}^2$, $n = 8$ synapses (anti-selective); $0.22 \pm 0.03\ \mu\text{m}^2$, $n = 21$ (co-selective); $P = 0.046$, Mann–Whitney U -test). **e–g**, Example connections,

including strongly anti-selective (**e**), weakly co-selective (**f**) and strongly co-selective (**g**) neuron pairs, coloured as in **b**. The left/right selectivity of pre- and post-synaptic neurons is indicated by the coloured icons. Neuron pairs correspond to data points indicated by arrows in **c** and **h**. For **e–g**, scale bars, $20\ \mu\text{m}$. **h**, The normalized synapse frequency is correlated with the maximum-selectivity similarity. $n = 21$ connections. The dotted line is the linear fit ($r = 0.65$, $P = 0.001$, Pearson correlation test). **i–k**, EM images showing example synapses between strongly anti-selective (**i**), weakly co-selective (**j**) and strongly co-selective (**k**) neuron pairs. PSDs are indicated by cyan arrows. Synapses correspond to data points indicated by arrows in **d** and **l** and are from the same connections shown in **e–g**, respectively. For **i–k**, scale bars, $500\ \text{nm}$. **l**, The log-transformed PSD area is correlated with the maximum-selectivity similarity index. $n = 29$ synapses. The dotted line is the linear fit ($r = 0.56$, $P = 0.002$, Pearson correlation test). Data are mean \pm s.e.m. Two-tailed significance tests were performed.

(inset)). The majority of these direct connections were I-to-E (74%), followed by E-to-I (12%), so we focused our analysis on these types of connections. We also quantified the area of the post-synaptic density (PSD area) associated with each synapse (Fig. 1d (inset) and Methods), which correlates with functional synaptic strength³³.

E-to-I connectivity

We first investigated how E-to-I connectivity in the PPC is related to selective activity. We compared the number of synapses between neurons that preferred the same trial type at their maximum selectivity (co-selective) versus those with opposite preferences (anti-selective) (Fig. 2a). As the opportunities for neurons to make synaptic connections are limited to locations where their axons and dendrites come into close proximity, we also quantified the axon–dendrite overlap^{23,34} between

all pairs of neurons (Fig. 2b and Extended Data Fig. 5a). Although the axon/dendrite overlaps were similar for co-selective and anti-selective pairs ($P = 0.56$, Mann–Whitney U -test), the number of synapses per μm overlap (normalized synapse frequency) was over two times higher for co-selective pairs (Fig. 2c; $0.021 \pm 0.006\ \mu\text{m}^{-1}$ (anti-selective) and $0.045 \pm 0.005\ \mu\text{m}^{-1}$ (co-selective); $P = 0.018$, Mann–Whitney U -test; Fig. 2e–g and Extended Data Fig. 5b).

To account for the continuous distribution of selectivity strengths across neurons (Extended Data Fig. 1g,h), we defined for each cell pair a maximum-selectivity similarity, the magnitude of which quantifies how selective the cells are, and the sign of which indicates whether the cells are co-selective or anti-selective (Methods and Extended Data Fig. 5c). The maximum-selectivity similarity was strongly correlated with the normalized synapse frequency ($r = 0.65$, $P = 0.001$, Pearson correlation test; Fig. 2h, Extended Data Fig. 5d and Supplementary Table 1). As the




Fig. 3 | Anti-selective I-to-E connectivity. **a**, Schematic of I-to-E connections among functionally characterized neurons. Connections are classified as co-selective (green) or anti-selective (purple). **b**, The normalized synapse frequency is greater for anti-selective (purple) than co-selective (green) I-to-E connections ($0.020 \pm 0.002 \mu\text{m}^{-1}$, $n = 40$ connections (anti-selective); $0.015 \pm 0.001 \mu\text{m}^{-1}$, $n = 63$ (co-selective); $P = 0.028$, Mann–Whitney *U*-test). **c**, The normalized synapse frequency is negatively correlated with the maximum-selectivity similarity, $n = 103$ connections. The dotted line is the

linear fit ($r = -0.22$, $P = 0.023$, Pearson correlation test). **d**, PSD areas are not significantly different for anti-tuned (purple) and co-tuned (green) I-to-E connections ($0.15 \pm 0.01 \mu\text{m}^2$, $n = 60$ connections (anti-tuned); $0.16 \pm 0.01 \mu\text{m}^2$, $n = 96$ (co-tuned); $P = 0.69$, Mann–Whitney *U*-test). **e**, The PSD area is not significantly correlated with the maximum-selectivity similarity, $n = 156$ synapses. The dotted line is the linear fit ($r = 0.08$, $P = 0.29$, Pearson correlation test). Data are mean \pm s.e.m. Two-tailed significance tests were performed.

maximum selectivity can occur at different times in the trial for different neurons, we also defined a ‘simultaneous-selectivity similarity’, which compares the selectivity of the pre-synaptic and post-synaptic neurons at the same timepoints (Methods and Extended Data Fig. 5e) and therefore relates more directly to synaptic interactions between neurons. The simultaneous-selectivity similarity was also strongly correlated with the normalized synapse frequency (Extended Data Fig. 5f; $r = 0.58$, $P = 0.006$, Pearson correlation test). These results indicate that excitatory neurons preferentially target inhibitory partners with similar selectivity.

We next examined whether the size of individual synapses was correlated with the selectivity of the connected neurons. The PSD area of cortical excitatory synapses is known to correlate with functional strength³³. Here we found that co-selective E-to-I synapses had PSD areas nearly two times larger than those of anti-selective synapses (Fig. 2d; $0.12 \pm 0.02 \mu\text{m}^2$ (anti-selective) and $0.22 \pm 0.03 \mu\text{m}^2$ (co-selective); $P = 0.046$, Mann–Whitney *U*-test). Synapse size was also strongly correlated with maximum-selectivity similarity (Fig. 2i–l; $r = 0.56$, $P = 0.002$, Pearson correlation test) and simultaneous-selectivity similarity (Extended Data Fig. 5g; $r = 0.50$, $P = 0.006$, Pearson correlation test).

Consistent with previous research³¹, we observed that the functional selectivity of PPC neurons can change over time even after the animal attains stable, expert performance of the task (Extended Data Fig. 5h). Concordantly, selectivity similarity between pairs of neurons also drifts on the timescale of days (Extended Data Fig. 5i). Comparing earlier with later behavioural sessions (Extended Data Fig. 1a) revealed that structure–function correlations are weaker on days further from when the brain was preserved (Extended Data Fig. 5j,k). This suggests that synaptic connections in the PPC may also change over timescales of days, which is consistent with turnover rates of some types of axon boutons and dendritic spines³⁵.

I-to-E connectivity

We next examined whether I-to-E connectivity was also functionally selective (Fig. 3a). We found that the normalized synapse frequency (synapses per μm of axon/dendrite overlap; Fig. 2b) was higher for anti-selective I-to-E pairs compared with co-selective pairs (Fig. 3b and Extended Data Fig. 6a,b; $0.020 \pm 0.002 \mu\text{m}^{-1}$ (anti-selective) and $0.015 \pm 0.001 \mu\text{m}^{-1}$ (co-selective); $P = 0.028$, Mann–Whitney *U*-test), and was negatively correlated with the maximum-selectivity similarity (Fig. 3c; $r = -0.22$, $P = 0.023$, Pearson correlation test, Extended Data Fig. 6c and Supplementary Table 2). These correlations were weaker on days further from when the brain was preserved (Extended Data Fig. 6d). The normalized synapse frequency was also negatively

correlated with simultaneous-selectivity similarity for I-to-E connections (Extended Data Fig. 6e; $r = -0.20$, $P = 0.041$, Pearson correlation test). These results indicate that inhibitory neurons preferentially target excitatory partners with opposite selectivity.

In contrast to E-to-I connectivity, we did not detect a significant difference in the PSD area of I-to-E synapses between co-selective and anti-selective pairs (Fig. 3d; $P = 0.69$, Mann–Whitney *U*-test), and PSD area was not significantly correlated with the maximum-selectivity similarity (Fig. 3e; $r = 0.084$, $P = 0.29$, Pearson correlation test), suggesting that selective I-to-E connectivity may be mediated more by the number of synapses than the strength of individual synapses.

Activity peaks of connected neurons

We considered how connectivity relates to neurons that have differences in the timing of their peak activity within a trial. The majority of connected E-to-I and I-to-E neuron pairs had activity peaks in different temporal epochs in the trial (Extended Data Fig. 7a–c), suggesting that trial-type selectivity may be a stronger determinant of connectivity than the timing of the peak activity. This finding was somewhat surprising because connectivity might be expected to be strongest among neurons with activity peaks at similar timepoints in the trial. However, cells exhibit non-zero activity away from their peak activity times (Extended Data Figs. 1e,f and 2), which enables a presynaptic cell to influence the activity of a post-synaptic cell even if they have activity peaks in different epochs. We tested this idea by computing noise correlations (a measure of simultaneous trial-to-trial co-fluctuations in activity after regressing away covariations due to similar tuning to cue and/or behavioural variables such as running patterns³²; Methods) and found that noise correlations were positively correlated with E-to-I connectivity and negatively correlated with I-to-E connectivity (Extended Data Fig. 7d,e). Furthermore, as noted above, the simultaneous-selectivity similarity was positively correlated with E-to-I synaptic connectivity (Extended Data Fig. 5f) and negatively correlated with I-to-E connectivity (Extended Data Fig. 6e). Together, these results indicate that opponent inhibition patterns of connectivity can affect activity and computation on timescales relevant to direct synaptic transmission.

Circuit modelling

Together, co-selective E-to-I (Fig. 2) and anti-selective I-to-E connectivity (Fig. 3) comprise an opponent inhibition motif (Fig. 4a (top)). We used network modelling to investigate how opponent inhibition may support decision-making computations. We first studied a linear rate model^{10,14} comprising two excitatory and two inhibitory units.

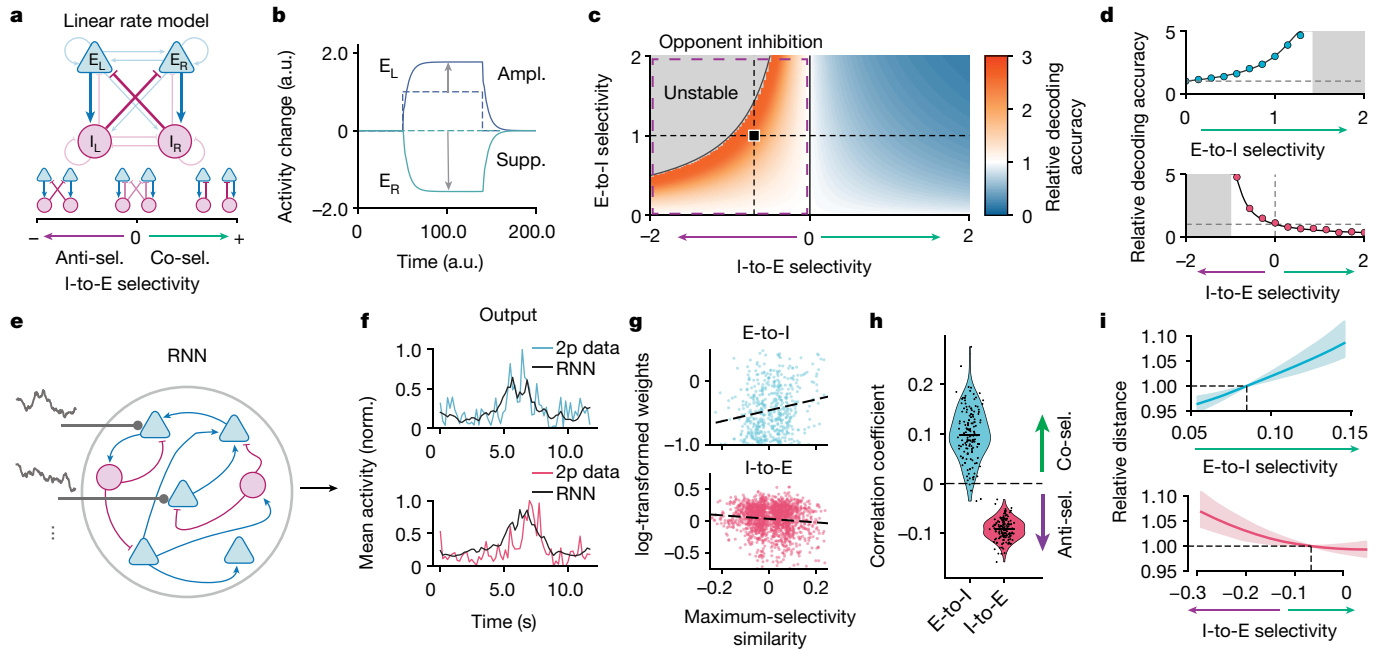


Fig. 4 | Opponent inhibition connectivity motif enhances trial-type signal encoding. **a**, Illustration of the linear rate model comprising two excitatory and two inhibitory units. Left (or right) trial-type input is fed to E_L (or E_R) (top). Bottom, variants in which I-to-E selectivity differs. The purple and green arrows indicate stronger anti-selectivity (anti-sel.) and co-selectivity (co-sel.), respectively. **b**, E_L and E_R activity (solid lines) in response to a left trial-type input (dotted lines) for a network with opponent inhibition. E_L is amplified and E_R is suppressed. **c**, The relative decoding accuracy (ratio of output to input decoding accuracy) as a function of E-to-I and I-to-E selectivity. The black square indicates the parameters used for **b**. **d**, The relative decoding accuracy as a function of E-to-I (top) and I-to-E (bottom) selectivity, corresponding to the dashed lines in **c**. **e**, Illustration of an RNN fit to the population activity. The networks are trained to reproduce experimental trial-averaged activity. **f**, Examples of the PPC activity (coloured lines) and RNN fits (black lines) for one

excitatory (top) and one inhibitory (bottom) neuron. **g**, Correlations between connectivity strength and the maximum-selectivity similarity for E-to-I (top) and I-to-E (bottom) connections for a single RNN. E-to-I connections are positively correlated (co-selective), whereas I-to-E connections are negatively correlated (anti-selective). **h**, The correlation between connection strengths and the maximum-selectivity similarity. $n = 147$ randomly initialized RNNs. E-to-I connections are co-selective, whereas I-to-E connections are anti-selective. **i**, The distance between left and right RNN activity (averaged across time) as a function of selectivity perturbations, normalized to its value in the unperturbed network (dashed lines). E-to-I (top) and I-to-E (bottom) connection weights were perturbed in a manner that increases anti-selectivity or co-selectivity without changing the average connection weight (Supplementary Methods). The solid lines show the median and the error bars show the interquartile range. $n = 147$ randomly initialized RNNs.

Left- or right-selective excitatory neurons (E_L , E_R) receive elevated external input during left or right trials and interact with left- and right-selective inhibitory neurons (I_L , I_R) (Fig. 4a, Methods and Supplementary Methods). In networks with opponent inhibition, input onto E_L decreases E_R activity through feed-forward inhibition, which amplifies E_L activity through feedback disinhibition^{13–15,36–38} (Fig. 4b). In left trials, both suppression of E_R and amplification of E_L increased the difference between neural activity on left and right trials, and this difference was therefore larger for networks with stronger opponent inhibition (Methods, Extended Data Fig. 8a–c and Supplementary Methods). As a consequence, networks with stronger opponent inhibition supported more accurate decoding of trial type in the presence of readout noise (Fig. 4c,d and Extended Data Fig. 8d,e). Opponent inhibition also improved decoding accuracy over a broad range of values of E-to-E selectivity, and even without recurrent excitatory connections (Extended Data Fig. 8f), as well as in networks in which both E and I neurons received external selective input (Extended Data Fig. 9a,b). When time-dependent input noise was included, opponent inhibition amplified the signal more than it amplified the noise, therefore enhancing trial-type encoding (Extended Data Fig. 8g–j).

Although the linear rate model explains how opponent inhibition can affect network coding, it does not include heterogeneity of connection weights, nor does it produce sequential activity peaks as observed in the experimental data. To determine whether its predictions hold for more biologically constrained models, we built a recurrent neural network (RNN) model with the same number of excitatory and inhibitory

neurons as the experimentally reconstructed circuit (Fig. 4e), and trained the connection weights of the RNNs to reproduce the measured calcium activity for left and right trials^{39,40} (Methods). After training, the RNNs generated dynamics that accurately reproduced PPC activity³⁹ (Fig. 4f). Although we did not constrain the selectivity of the RNN connections, the trained RNNs exhibited co-selective E-to-I and anti-selective I-to-E motifs, similar to those found experimentally (Fig. 4g,h and Extended Data Fig. 9c). To investigate whether these motifs supported signal amplification, as predicted by the linear rate model, we systematically manipulated the RNN connectivity⁴¹ by perturbing the E-to-I or I-to-E selectivity around the trained values, and regenerated the dynamics using the new connections (Supplementary Methods). Stronger opponent inhibition (stronger E-to-I co-selectivity or I-to-E anti-selectivity) amplified the separation between left and right population responses (Fig. 4i and Extended Data Fig. 9d), further suggesting that opponent inhibition may enhance the coding of trial-type signals in the PPC.

In the RNNs, as with the experimental data, connections comprising the opponent inhibition motif also include many neuron pairs with peaks in different trial epochs (Extended Data Figs. 7b,c and 10a). By perturbing the connections in the RNNs, we found that the selectivity of connections both between neurons with peaks in the same or different trial epochs promotes the separation between left- and right-choice population responses (Extended Data Fig. 10b–d). An analysis of the modelled currents during network dynamics revealed strong off-peak inhibition specific to connections between neurons peaking in different

epochs (Extended Data Fig. 10e), suggesting that I-to-E connections are critical for the formation of sequential activity peaks. Weakening or removing connections between neurons with peaks in different trial epochs disrupted trial-type selective sequential activity dynamics (Extended Data Fig. 10f–m and Supplementary Methods). Taken together, these results suggest that selective connections between neurons with activity peaks at different times are crucial for generating trial-type-selective and sequential population dynamics.

If opponent inhibition is used to compute signals relevant for choice, we would expect inhibitory activity to have a role in determining whether the animal makes the correct choice. To test this, we examined how the activity of I and E neurons in the PPC differed on error trials versus correct trials. Selective activity was generally degraded (reduced or reversed) on error trials (Extended Data Fig. 11a), and this degradation was comparable in magnitude for I and E neurons (Extended Data Fig. 11b,c). Thus, the activity of both I and E neurons was related to the mouse's choice. To confirm this result, we performed additional behavioural and functional imaging experiments on mice in which inhibitory neurons were labelled (Methods and Extended Data Fig. 11d) and found results consistent with those from the mouse used for EM analysis (Extended Data Fig. 11c). To understand how error trial selectivity degradation relates to selective connectivity, we also examined how I and E unit activity in the linear rate model correlate with choice (defined on the basis of which E unit was more active in each trial^{13,14}). Only networks with opponent inhibition exhibited inhibitory selectivity degradation on error trials comparable in magnitude to excitatory selectivity degradation, as observed in the experimental data (Extended Data Fig. 11c,e). Taken together, these analyses suggest that inhibitory activity could have a role in producing correct choices and provides further evidence that opponent inhibition contributes to choice-selective activity and, ultimately, decision-making.

Discussion

We sought to understand the relationships between trial-type-selective neuron activity and synaptic connectivity in the PPC. Although selective activity in the PPC has been reported in many previous studies^{3,5–10,42}, accompanying connectivity data have been lacking. Here we used automated serial-section transmission EM¹⁹ to acquire synapse-resolution images in a volume from the PPC that was functionally imaged during behaviour. As neuronal arborizations extend over large distances in the mammalian cortex, it is critical to image a large enough volume to sample them. The EM volume collected here in the PPC contains a much larger volume compared with previous cortical EM datasets^{21–24,26} (but also see refs. 20,27), enabling the reconstruction of substantial portions of axonal and dendritic arborizations, including synaptic connections made on distal branches. The resulting connectivity data, combined with behavioural and functional imaging data from the same animal, enabled us to reveal circuit motifs that support decision-making. Still, these data include only a modest sample size of functionally characterized neurons and synapses between them, and do not include synapses made outside the EM volume. Future functional connectomic datasets involving larger EM volumes and more neurons will probably reveal additional circuit motifs.

We found that the frequency and size of synaptic connections in the PPC depended significantly on the selectivity of pre-synaptic and post-synaptic neurons. For E-to-I connections, co-selective synapses were larger and more frequent, whereas, for I-to-E connections, anti-selective synapses were more frequent. We did not detect a difference in the synapse size between co-selective and anti-selective I-to-E connections. However, synapse size analysis for both E-to-I and I-to-E connections should be interpreted cautiously, as the correlation between synapse size and functional strength in the cortex has been directly measured only for E-to-E synapses³³.

The combination of co-selective E-to-I and anti-selective I-to-E comprises an opponent inhibition motif, in which the activity of left-selective excitatory neurons suppresses the activity of right-selective ones, and vice versa. This motif has been shown to mediate action selection in zebrafish and *Drosophila*^{43,44}, and a related motif has been reported in the ferret visual cortex⁴⁵, but motifs of this type have not previously been reported in the association cortex. Previous research in the mouse PPC proposed that selective connectivity motifs underlie choice-selective inhibitory activity, but could not rule out models with non-selective inhibition¹⁰. Here, the combination of neuronal activity measurements and EM-based connectomics in the same neurons has enabled the identification of the underlying connectivity motifs.

Selective inhibitory connectivity in the PPC contrasts with the primary visual cortex, in which previous connectomic analysis has suggested that E-to-I connectivity is non-selective in mice²² (but also see refs. 46–48). This suggests that specific inhibitory connectivity may be a distinct feature of PPC relative to the primary visual cortex, which underlies specialized functional roles of different cortical areas. Alternatively, opponent inhibition may be a more general motif that can arise in both sensory and association cortices with task learning⁴⁹, as well as in the motor cortex⁵⁰. Experiments comparing connectivity across multiple cortical areas in the same trained animal will be needed to assess how general opponent inhibition is across the cortex.

Although we observed selective connectivity across a population of neurons, the individual connection probabilities between neuron pairs were quite variable. This connection noise is a source of biological variability⁵¹ that places limits on our ability to detect connectivity motifs given limited experimental sample sizes. Thus, although data presented here were sufficient to reveal the opponent inhibition motif, more subtle motifs may require more data to uncover. For example, we did not find selectivity in E-to-E connections (Extended Data Fig. 6f,j), but recurrent E-to-E co-selectivity (which is found in the primary visual cortex^{23,52}) might be observable with a larger dataset. Thus, we anticipate that future connectomics experiments encompassing even larger volumes and numbers of functionally characterized neurons will uncover additional connectivity motifs and elucidate differences among interneuron subtypes. Nevertheless, our circuit modelling suggests that advantages of opponent inhibition apply over a wide range of E-to-E and I-to-I selectivity (Extended Data Fig. 8f).

In models of decision-making, the formation of categorical choices is typically facilitated by non-selective lateral inhibition¹². Recently, it has been proposed that selective inhibition could have one of two possible roles: promoting competition with anti-selective I-to-E connectivity, or stabilizing dynamics through co-selective I-to-E connectivity¹⁵. These distinct contributions are also present in the linear rate model presented here (Extended Data Fig. 8f). Our anatomical data suggests that PPC lies in the competition regime.

Although some decision-making models focus on the production of categorical choices through winner-take-all dynamics in attractor models^{15,15}, previous research suggests that, during navigational tasks, the PPC produces more complex dynamics in which multiple activity patterns arise for each trial type⁵³. These neural trajectories in the PPC probably represent a wide range of task and behavioural variables, including the mouse's choice, its navigational movements and position, and sensory cues from the environment^{3,31,54–56}. For this reason, the model developed here focuses on a graded encoding of the choice signal, whereby the PPC circuit helps to separate these multifaceted neural trajectories to enhance the encoding of the signals relevant for navigational decision-making.

Temporal sequences of activity in the cortex have become a prominent topic of discussion because they have been identified across many cortical areas and in many contexts in recent years^{3,57,58}. Although the analysis in this study has focused on choice selectivity, choice and temporal selectivity are intermingled within neuronal activity dynamics. Indeed, we observed that choice-specific connections include neuron pairs that

have activity peaks at different times (Extended Data Fig. 7a–c). Using modelling, we showed that recurrent networks need such connections to generate choice-specific sequences (Extended Data Fig. 10), which provides a theoretical intuition for why we find such connections in the experimental connectivity data. Although further studies will be needed to fully understand how synaptic connectivity underlies temporal activity sequences, our study provides a start in this direction.

In summary, we identified an anatomical opponent inhibition motif consisting of functionally selective connectivity between excitatory and inhibitory neurons in the PPC. Using modelling, we showed that this opponent inhibitory motif improves the encoding of trial-type information. Together, these results identify an anatomical connectivity motif in the PPC that supports decision-making.

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/s41586-024-07088-7>.

- Assad, J. A. & Maunsell, J. H. R. Neuronal correlates of inferred motion in primate posterior parietal cortex. *Nature* **373**, 518–521 (1995).
- Shadlen, M. N. & Newsome, W. T. Motion perception: seeing and deciding. *Proc. Natl Acad. Sci. USA* **93**, 628–633 (1996).
- Harvey, C. D., Coen, P. & Tank, D. W. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* **484**, 62–68 (2012).
- Raposo, D., Kaufman, M. T. & Churchland, A. K. A category-free neural population supports evolving demands during decision-making. *Nat. Neurosci.* **17**, 1784–1792 (2014).
- Goard, M. J. et al. Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. *eLife* **5**, 471–477 (2016).
- Funamizu, A., Kuhn, B. & Doya, K. Neural substrate of dynamic Bayesian inference in the cerebral cortex. *Nat. Neurosci.* **19**, 1682–1689 (2016).
- Hwang, E. J., Dahlen, J. E., Mukundan, M. & Komiyama, T. History-based action selection bias in posterior parietal cortex. *Nat. Commun.* **8**, 1242 (2017).
- Song, Y. H. et al. A neural circuit for auditory dominance over visual perception. *Neuron* **93**, 940–954 (2017).
- Zhong, L. et al. Causal contributions of parietal cortex to perceptual decision-making during stimulus categorization. *Nat. Neurosci.* **22**, 963–973 (2019).
- Najafi, F. et al. Excitatory and inhibitory subnetworks are equally selective during decision-making and emerge simultaneously during learning. *Neuron* **105**, 165–179 (2020).
- Shadlen, M. N., Britten, K. H., Newsome, W. T. & Movshon, J. A. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J. Neurosci.* **16**, 1486–1510 (1996).
- Wang, X. J. Decision making in recurrent neuronal circuits. *Neuron* **60**, 215–234 (2008).
- Wang, X. J. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* **36**, 955–968 (2002).
- Wong, K. F. & Wang, X. J. A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* **26**, 1314–1328 (2006).
- Roach, J. P., Churchland, A. K. & Engel, T. A. Choice selective inhibition drives stability and competition in decision circuits. *Nat. Commun.* **14**, 147 (2023).
- Denk, W. & Horstmann, H. Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure. *PLoS Biol.* **2**, e329 (2004).
- Xu, C. S. et al. Enhanced FIB-SEM systems for large-volume 3D imaging. *eLife* **6**, e25916 (2017).
- Eberle, A. L. et al. High-resolution, high-throughput imaging with a multibeam scanning electron microscope. *J. Microsc.* **259**, 114–120 (2015).
- Phelps, J. S. et al. Reconstruction of motor control circuits in adult *Drosophila* using automated transmission electron microscopy. *Cell* **184**, 759–774 (2021).
- Shapson-Coe, A. et al. A connectomic study of a petascale fragment of human cerebral cortex. Preprint at *bioRxiv* <https://doi.org/10.1101/2021.05.29.446289> (2021).
- Kasthuri, N. et al. Saturated reconstruction of a volume of neocortex. *Cell* **162**, 648–661 (2015).
- Bock, D. D. et al. Network anatomy and in vivo physiology of visual cortical neurons. *Nature* **471**, 177–182 (2011).
- Lee, W.-C. A. et al. Anatomy and function of an excitatory network in the visual cortex. *Nature* **532**, 370–374 (2016).
- Motta, A. et al. Dense connectomic reconstruction in layer 4 of the somatosensory cortex. *Science* **366**, eaay3134 (2019).
- Dorkenwald, S. et al. Binary and analog variation of synapses between cortical pyramidal neurons. *eLife* **11**, e76120 (2022).
- Turner, N. L. et al. Reconstruction of neocortex: organelles, compartments, cells, circuits, and activity. *Cell* **185**, 1082–1100 (2022).
- MiCRONS Consortium et al. Functional connectomics spanning multiple areas of mouse visual cortex. Preprint at *bioRxiv* <https://doi.org/10.1101/2021.07.28.454025> (2021).
- Hofer, S. B. et al. Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. *Nat. Neurosci.* **14**, 1045–1052 (2011).
- Runyan, C. A. et al. Response features of parvalbumin-expressing interneurons suggest precise roles for subtypes of inhibition in visual cortex. *Neuron* **67**, 847–857 (2010).
- Ma, W. P. et al. Visual representations by cortical somatostatin inhibitory neurons—selective but with weak and delayed responses. *J. Neurosci.* **30**, 14371–14379 (2010).
- Driscoll, L. N., Pettit, N. L., Minderer, M., Chetih, S. N. & Harvey, C. D. Dynamic reorganization of neuronal activity patterns in parietal cortex. *Cell* **170**, 986–999 (2017).
- Runyan, C. A., Piasini, E., Panzeri, S. & Harvey, C. D. Distinct timescales of population coding across cortex. *Nature* **548**, 92–96 (2017).
- Holler, S., Köstinger, G., Martin, K. A. C., Schuhknecht, G. F. P. & Stratford, K. J. Structure and function of a neocortical synapse. *Nature* **591**, 111–116 (2021).
- Stepanyants, A., Hof, P. R. & Chklovskii, D. B. Geometry and structural plasticity of synaptic connectivity. *Neuron* **34**, 275–288 (2002).
- Holtmaat, A. & Svoboda, K. Experience-dependent structural synaptic plasticity in the mammalian brain. *Nat. Rev. Neurosci.* **10**, 647–658 (2009).
- Cannon, S. C., Robinson, D. A. & Shamma, S. A proposed neural network for the integrator of the oculomotor system. *Biol. Cybern.* **49**, 127–136 (1983).
- Seung, H. S. in *The Handbook of Brain Theory and Neural Networks* 2nd edn (ed. Arbib, M. A.) 94–97 (MIT, 2003).
- Lim, S. & Goldman, M. S. Balanced cortical microcircuitry for maintaining information in working memory. *Nat. Neurosci.* **16**, 1306–1314 (2013).
- Rajan, K., Harvey, C. D. & Tank, D. W. Recurrent network models of sequence generation and memory. *Neuron* **90**, 128–142 (2016).
- Song, H. F., Yang, G. R. & Wang, X. J. Training excitatory-inhibitory recurrent neural networks for cognitive tasks: a simple and flexible framework. *PLoS Comput. Biol.* **12**, e1004792 (2016).
- Kim, R. & Sejnowski, T. J. Strong inhibitory signaling underlies stable temporal dynamics and working memory in spiking neural networks. *Nat. Neurosci.* **24**, 129–139 (2021).
- Orlandi, J. G., Abdolrahmani, M., Aoki, R., Lyamzin, D. R. & Benucci, A. Distributed context-dependent choice information in mouse posterior cortex. *Nat. Commun.* **14**, 192 (2023).
- Koyama, M. et al. A circuit motif in the zebrafish hindbrain for a two alternative behavioral choice to turn left or right. *eLife* **5**, e16808 (2016).
- Jovanic, T. et al. Competitive disinhibition mediates behavioral choice and sequences in *Drosophila*. *Cell* **167**, 858–870 (2016).
- Wilson, D. E., Scholl, B. & Fitzpatrick, D. Differential tuning of excitation and inhibition shapes direction selectivity in ferret visual cortex. *Nature* **560**, 97–101 (2018).
- Wilson, D. E. et al. GABAergic neurons in ferret visual cortex participate in functionally specific networks. *Neuron* **93**, 1058–1065 (2017).
- Keller, A. J. & Martin, K. A. C. Local circuits for contrast normalization and adaptation investigated with two-photon imaging in cat primary visual cortex. *J. Neurosci.* **35**, 10078–10087 (2015).
- Cardin, J. A., Palmer, L. A. & Contreras, D. Stimulus feature selectivity in excitatory and inhibitory neurons in primary visual cortex. *J. Neurosci.* **27**, 10333–10344 (2007).
- Tseng, S. Y., Chetih, S. N., Arlt, C., Barroso-Luque, R. & Harvey, C. D. Shared and specialized coding across posterior cortical areas for dynamic navigation decisions. *Neuron* **110**, 2484–2502 (2022).
- Zagha, E., Ge, X. & McCormick, D. A. Competing neural ensembles in motor cortex gate goal-directed motor output. *Neuron* **88**, 565–577 (2015).
- Bartol, T. M. et al. Nanoconnectomic upper bound on the variability of synaptic plasticity. *eLife* **4**, e10778 (2015).
- Cossell, L. et al. Functional organization of excitatory synaptic strength in primary visual cortex. *Nature* **518**, 399–403 (2015).
- Morcos, A. S. & Harvey, C. D. History-dependent variability in population dynamics during evidence accumulation in cortex. *Nat. Neurosci.* **19**, 1672–1681 (2016).
- Krumin, M., Lee, J. J., Harris, K. D. & Carandini, M. Decision and navigation in mouse parietal cortex. *eLife* **7**, e42583 (2018).
- Mimica, B., Dunn, B. A., Tombaz, T., Srikanth Bojja, V. P. T. N. & Whitlock, J. R. Efficient cortical coding of 3D posture in freely behaving rats. *Science* **362**, 584–589 (2018).
- Minderer, M., Brown, K. D. & Harvey, C. D. The spatial structure of neural encoding in mouse posterior cortex during navigation. *Neuron* **102**, 232–248 (2019).
- Hahnloser, R. H. R., Kozhevnikov, A. A. & Fee, M. S. An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature* **419**, 65–70 (2002).
- Koay, S. A., Charles, A. S., Thiberge, S. Y., Brody, C. D. & Tank, D. W. Sequential and efficient neural-population coding of complex task information. *Neuron* **110**, 328–349 (2022).

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

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© The Author(s), under exclusive licence to Springer Nature Limited 2024

Methods

Behaviour and calcium imaging

All of the experimental procedures were approved by the Harvard Medical School Institutional Animal Care and Use Committee. Mice were housed under a reversed 12 h–12 h light–dark light cycle. We trained mice to perform a two-alternative forced-choice delay task in a virtual reality T-maze and performed *in vivo* calcium imaging as previously described^{3,31} (Fig. 1a). In brief, the task consisted of cue, delay and turn phases. During the cue phase, the mice were presented with one of two visual cues on the walls of the T-maze. In the second half of the T-maze (the delay period), the cues were replaced by a neutral grey. At the T-intersection, the mice turned left or right based on the presented cue to earn a water reward. The mice performed the task by running on a spherical treadmill and were presented with visual stimuli projected onto a screen using the ViRMen software engine⁵⁹ (forked at <https://github.com/lauradriscoll/virmen>).

We used raw calcium imaging and behavioural data from a single mouse originally from a previous study³¹. Owing to the extensive effort required to generate a large enough EM volume (~6 months for imaging, ~6 months for alignment, ~7,000 annotation hours for neuron tracing), it was not feasible to generate multiple such datasets from multiple animals in this study. In brief, the mouse was a male C57BL/6J mouse (The Jackson Laboratory) aged 8 weeks at the start of behavioural training, 14–18 weeks during imaging and 18–19 weeks when the brain was preserved for EM imaging. GCaMP6m was expressed in left-hemisphere PPC layer 2/3 neurons by viral injection (AAV2/1-synapsin-1-GCaMP6m) and neuronal activity was recorded during behaviour using a two-photon microscope (5.3 Hz volume imaging rate) controlled by ScanImage (v.4; Vidrio Technologies). Behaviour and functional imaging were recorded every day for around 30 consecutive days (with a few 1 day breaks). However, for the main analysis in this study, we focused on data from the last four sessions. For additional analyses investigating how structure-function relationships evolve over time, we used sessions 8–10 days before (early), 4–7 days before (middle) and the last 4 days before (late) euthanasia (Extended Data Figs. 1a, 5j, k and 6d).

For each session, we performed source extraction using Suite2p⁶⁰ (<https://www.suite2p.org>) and manually screened the resulting regions of interest (ROIs) to obtain putative cell bodies and associated calcium signals. The dF/F calcium signals were deconvolved using the constrained FOOPSI algorithm⁶¹ (<https://github.com/epnev/constrained-foopsi>) to obtain an event rate that estimates the relative firing rate of each neuron over time. We synchronized the event rates to the behavioural trials structure to obtain trial-wise event rates for each neuron. As the length of trials can vary based on the how quickly the mouse runs, we synchronized based on three landmarks during the trial: the beginning of the trial, the start of the delay period and the end of the trial. Trial aligned event rates had 63 timepoints (each 188 ms) according to the following scheme: timepoints 1–13 correspond to the cue (beginning) epoch and are aligned to the landmark timepoint: 1, the start of trial (running onset); timepoints 14–25 correspond to the cue (end) epoch and are aligned to the landmark timepoint: 26, the start of delay period (cue offset); timepoints 26–38 correspond to the delay epoch and are aligned to the landmark timepoint: 26, the start of delay period (cue offset); timepoints 39–51 correspond to the turn epoch and are aligned to the landmark timepoint: 52, the end of the trial (reward given or omitted); and timepoints 52–63 correspond to the inter-trial interval and are aligned to the landmark timepoint: 52, the end of the trial (reward given or omitted).

Thus, synchronized trials were not completely continuous in time (there are two discontinuities). We next calculated trial-averaged event rates by averaging over left and right trials separately.

For experiments analysing the activity of excitatory and inhibitory neurons during error trials, additional mice were trained to perform

the virtual reality behavioural task. We injected the left-hemisphere PPC with AAV9-hSyn-jGCaMP7f, AAV9-mDlx-NLS-mRuby2 and AAV8-CaMKIIa-ChRmine-mCherry-Kv2.1 and installed a cranial window and headplate^{62–64}. This set of viruses labelled both excitatory and inhibitory neurons with jGCaMP7f, excitatory neurons with a somatic membrane-localized red fluorophore and inhibitory neurons with a nuclear-localized red fluorophore (Extended Data Fig. 11d). We recorded neural activity in the PPC using two-photon microscopy (~30 Hz frame rate) and synchronized behaviour and imaging using frame triggers from the microscope and iteration triggers from ViRMen. We also collected a z stack spanning $\pm 10 \mu\text{m}$ around the imaging plane to record structural data from both the red and green channels. ROIs were automatically extracted using Suite2p and non-cell sources were discarded. After aligning the calcium imaging field-of-view to the structural z stack, for each source, we measured the red fluorescence in the somatic membrane and in the nucleus. To measure the nuclear red fluorescence, we eroded the masks for each source and took the mean grey value of the eroded ROI in the aligned red channel image; the somatic membrane red fluorescence was the mean grey value of all of the pixels in the ROI that were not included in the nuclear region. To identify neurons with enriched nuclear red fluorescence, we computed the ratio of nuclear to membrane fluorescence. We took cells with the top 10% of these values to be putative inhibitory interneurons labelled by AAV9-mDlx-NLS-mRuby2 and considered the bottom 70% of these values to be putative excitatory neurons with red fluorescence in the somatic membrane. These thresholds were chosen to be conservative in our labelling of inhibitory and excitatory neurons.

Information-theoretic selectivity index

For each neuron in each session, we calculated an information-theoretic trial-type selectivity index as follows. The event rates were converted to binary values by setting all non-zero values to 1, measuring when the neuron is active versus inactive. Then, as done previously³², we calculated the instantaneous MI⁶⁵ between this binarized event rate and the identity of the sensory cue (which has two possible values for right and left trials) for each timepoint in the trial using the Information Breakdown Toolbox⁶⁶. This measures the instantaneous trial-type information, with trial type intended as the type of cue presented to the mouse. All left- and right-turn trials were included in this analysis, including error trials in which the mouse turned the wrong direction. Alternative analyses, such as including only correct trials or defining the trial type based on the mouse's choice, lead to slightly different values of MI, but did not significantly change the results. We subtracted the limited-sampling bias from the information estimates using the Panzeri–Treves bias correction⁶⁷, which improved the accuracy of information estimates by removing the confounding effect of differences in trial numbers across sessions (see the 'Method selection for estimating trial type selectivity and noise correlation' section in the Supplementary Methods; Extended Data Fig. 12). We computed information in each available session (among the last four sessions) and then averaged the value across sessions to provide an information estimate for each neuron, as this optimized the SNR of the calculation (Extended Data Fig. 12). We estimated the significance of the information value at each timepoint ($P < 0.05$) using the 95th percentile of the null-hypothesis distribution obtained by randomly permuting neural responses across trials⁶⁶ ($n = 1,000$ permutations), and we set to zero non-significant information values to avoid attributing selectivity to timepoints that were not. The MI has units of bits and is bounded between 0 and 1.

We defined the magnitude of the maximum selectivity to be the maximum value of the MI across timepoints. The timepoint of peak MI was different for different neurons. We defined the sign of the maximum selectivity based on the trial-averaged event rates at the timepoint of peak MI. The maximum selectivity was given a negative sign if the trial-averaged binarized event rate was greater on left trials

than right. Thus, the maximum selectivity ranges from -1 to $+1$, with the sign indicating preferred trial type and the magnitude indicating MI with trial type, and 0 indicating non-significant information. Maximum selectivity values were calculated separately for each behavioural session. To obtain the overall maximum selectivity for each neuron, these values were averaged across the last four behavioural sessions before the mouse was euthanized (see the ‘Co-registration between in vivo and EM data’ section below for alignment of multiple sessions and EM data). For additional analyses investigating how structure–function relationships evolve over sessions, we calculated maximum selectivity separately for early, middle and late sessions (Extended Data Figs. 1a, 5j,k and 6d).

This information-theoretic definition of selectivity is distinct from metrics based on receiver-operator characteristic analysis¹⁰ or trial-averaged activity^{3,31} used in other studies. In contrast to these other metrics, the MI metric quantifies information available in individual trials and can capture any linear or nonlinear tuning. However, these different metrics for selectivity are highly correlated with each other, so the main conclusions of this study also hold when using alternative definitions of selectivity.

Note that, in our dataset, the number of right-selective neurons outnumbered the left-selective (Extended Data Figs. 1g and 2). As a result, the number of functionally characterized left-selective inhibitory neurons was small. Thus, our findings on trial-type selective connectivity primarily comes from connections involving right-selective inhibitory neurons.

Error trial analysis

To investigate the relationship between neuronal activity and the animal’s choice, we compared correct and error trials (Extended Data Fig. 11). Error trials were defined as trials in which the animal turned the wrong direction, causing a mismatch between the cue and the choice. As the rate of error trials was low, we pooled trials from 11 sessions (including early, middle and late sessions indicated in the Extended Data Fig. 1a). However, each individual neuron was typically detected in several, but not all 11 sessions. To account for experimental differences across sessions, we normalized the deconvolved activity of each neuron in each session to the average activity level of that neuron across the whole session.

For each neuron, we defined the preferred and non-preferred cue based on the functional selectivity of that neuron. For example, for a left-selective neuron, the preferred cue was the one that indicates a left turn (white cue), and the non-preferred was the one that indicates a right turn (black cue). By definition, neurons are generally more active during preferred-cue trials than non-preferred.

To analyse how activity changes on error trials compared to correct trials, we quantified for each neuron the relative change in activity for preferred ($\Delta_{\text{preferred}}$) and non-preferred ($\Delta_{\text{non-preferred}}$) trials:

$$\Delta_{\text{preferred}} = \frac{\text{act}(\text{preferred correct trials}) - \text{act}(\text{preferred error trials})}{\text{act}(\text{all preferred trials})}$$

$$\Delta_{\text{non-preferred}} = \frac{\text{act}(\text{non-preferred error trials}) - \text{act}(\text{non-preferred correct trials})}{\text{act}(\text{all non-preferred trials})}$$

where $\text{act}(x)$ refers to the mean activity over trials of type x , with activity averaged over both trial types and across time for each trial (except in Extended Data Fig. 11a, in which activity is shown as a function of time for visualization purposes). Generally, when the animal makes an error, neuronal activity is lower on preferred trials and higher on non-preferred trials, which is consistent with the activity encoding to some extent the (erroneous) choice. Thus, the functional selectivity of

neurons is degraded (or even switched) on error trials. We defined the signs of $\Delta_{\text{preferred}}$ and $\Delta_{\text{non-preferred}}$ to both be positive in this case.

To quantify the overall selectivity degradation, we defined the ‘error trial selectivity degradation’ for a given neuron i (Extended Data Fig. 11b) as:

$$D_i = \Delta_{\text{preferred}} + \Delta_{\text{non-preferred}}$$

For a neural population, the error trial selectivity degradation was defined as:

$$D = \text{median}(\Delta_{\text{preferred}}) + \text{median}(\Delta_{\text{non-preferred}}),$$

where the median is taken over neurons belonging to the neural population considered (Extended Data Fig. 11c). The median, rather than the mean, was used to reduce the effect of outliers that can result due to the small number of error trials.

Positive values of D_i indicate that the neuron’s selective activity is degraded in error trials compared with correct trials. In fact, if a neuron’s activity is selective for the wrong choice on error trials, $\Delta_{\text{preferred}}$, $\Delta_{\text{non-preferred}}$ and D_i are all positive. Null values of D_i indicate that the selective activity does not change between correct and error trials. If a neuron’s activity is the same on error and correct trials, $\Delta_{\text{preferred}}$, $\Delta_{\text{non-preferred}}$ and D_i are all 0 . Negative values of D_i (rarely seen in our data) would indicate that selective activity is enhanced in error trials.

To assess the statistical significance of D for different populations of neurons (such as excitatory or inhibitory), we performed a permutation test in which we shuffled the identity of error trials (keeping the total number of error trials the same) and recalculated D for each shuffle. The P value of D was determined by comparing the measured D to this null distribution.

Note that the value of the degradation index D is sensitive to the number of trials and the number of error trials. If the number of error trials is small, results will become noisier and values of D will be closer to zero. To mitigate this, we excluded neurons with fewer than four total error trials of a given type (preferred or non-preferred). However, trial numerosity was far greater in the EM mouse (because it was based on many co-registered sections) compared with the additional mice with only 1–2 sessions. Thus, the magnitude of D cannot be meaningfully compared between the EM and additional mice, but can be compared between I and E neurons within each dataset, which share the same number of trials.

EM dataset

On the last day of in vivo imaging, we injected the tail vein with a fluorescent dye to label blood vessels (rhodamine B isothiocyanate–dextran (molecular mass, 70 kDa), 5% (v/v), Sigma-Aldrich) and acquired an anatomical 2p reference stack of the imaging ROI in the PPC (green channel, GCaMP6m for cells; red channel, rhodamine for blood vessels). After performing behaviour and calcium imaging, the mouse was perfused transcardially (2% formaldehyde/2.5% glutaraldehyde in 0.1 M cacodylate buffer with 0.04% CaCl_2) and the brain was prepared for EM imaging as previously described²³. In brief, 200- μm -thick coronal vibratome sections were cut, post-fixed and en bloc stained with 1% osmium tetroxide/1.5% potassium ferrocyanide followed by 1% uranyl acetate, dehydrated with a graded ethanol series and embedded in resin (TAAB 812 Epon, Canemco).

We cut serial 1- μm -thick sections from regions a few sections away from the PPC, stained them with toluidine blue (EMS), imaged them with light microscopy and aligned them to large blood vessels on the surface of the brain (from photos taken through the cranial window) to estimate which sections overlap with the PPC regions imaged in vivo. We then used micro-CT (Zeiss Versa) to confirm the correct vibratome section by registering corresponding vasculature between the micro-CT volume and 2p reference stack.

Article

We cut a series of 2,500 thin sections (thickness, 40–45 nm) and imaged them using the GridTape system as previously described¹⁹. In brief, we trimmed the tissue block to mesa containing the tissue of interest using an ultramicrotome (Leica UC7) and a diamond trimming knife (EMS-Diatome). Then, using an automated tape ultramicrotome (ATUM), we sectioned and collected 2,500 sections onto a reel of GridTape over a period of about 15 h. During pickup, 37 sections (out of 2,500, comprising 1.4%) did not adhere to the transparent slots on the GridTape and 8 sections (0.3%) had ruptured film supports.

After sectioning and pickup, we post-stained the sections with lead citrate using a semi-automated reel-to-reel system. The sections were imaged using an automated transmission EM system over a period of about 6 months. For each section, an ROI of approximately 1.2 mm × 0.7 mm was imaged (for the first 800 sections, a larger ROI of about 1.5 × 0.7 mm was imaged). During staining and imaging, 38 (1.5%) additional sections were damaged in a way that precluded successful imaging. Overall, 2,427 (97.1%) of the intended 2,500 sections were successfully imaged. The missed sections included one four-section gap, two three-section gaps and ten two-section gaps. The dataset consisted of around 400 TB of raw data (16-bit images). The raw images were converted to 8-bit format and stitched together to form a contiguous three-dimensional volume using an elastic spring mesh algorithm (AlignTK).

Co-registration between in vivo and EM data

On the last day of in vivo imaging, we recorded a volumetric reference stack of the ROI in the PPC with 1 μm³ voxel size, which was used as a bridge between the calcium imaging planes for each session (4 planes separated by 25 μm in z) and the EM volume. First, we co-registered the reference stack to a down-sampled version of the EM volume (Extended Data Fig. 1i). Then, for each session, we co-registered the imaging planes to the reference stack. Both registrations were calculated by manually identifying a moderate number of correspondence points (~10–30) using the ImageJ (v.1.53) plugin BigWarp⁶⁸ (release 6.0.0; <https://imagej.net/BigWarp>) to calculate an affine transformation matrix (custom MATLAB code).

Using these affine transformations, we overlaid the extracted source ROIs from each calcium imaging session onto the EM space, and manually inspected each ROI for matching cell bodies in the EM volume (Extended Data Fig. 1j). Some ROIs were not associated with any cell bodies, presumably because they were other objects such as large dendrites. To increase confidence, we matched ROIs to EM for multiple sessions simultaneously, and used the trial-aligned activity to help to determine whether ROIs in different sessions were from the same neuron. Activity was usually, but not always, similar for the same neuron over several sessions. At the conclusion of this correspondence process, we identified 140 functionally characterized neurons within the EM volume with matching calcium imaging ROIs. For most neurons, matching ROIs were found in multiple, but not all, sessions. This might be because they were not identified by the source extraction algorithm, or because they were excluded due to uncertainty in the co-registration and manual matching procedure. The co-registration results were reviewed by a 2nd expert annotator, and only cell matches that were agreed upon by the two independent annotators were included in analysis.

Neuronal circuit reconstruction

The morphology and connectivity of the functionally characterized neurons were reconstructed through manual tracing by a team of annotators using the CATMAID^{23,69,70} (release 2018.11.09) collaborative annotation software. Starting from the cell body (which was previously co-registered to calcium imaging source ROIs), all branches (including axons and dendrites) were traced completely until they either ended or reached the boundary of the EM volume. In some cases, data quality issues such as missing sections or poor image quality prevented further tracing. The functionally characterized neurons were classified as pyramidal or non-pyramidal (inhibitory) on the basis of

their reconstructed morphology. Pyramidal cells were classified as such on the basis of characteristic features that included a prominent, pial-projecting apical dendrite, outward/downward projecting basal dendrites and downward-projecting axon (Extended Data Fig. 3). Non-pyramidal cells had a variety of morphologies probably corresponding to distinct interneuron subtypes (Extended Data Fig. 4). In total, 124 excitatory and 16 inhibitory neurons with co-registered calcium imaging were traced within the EM volume. For the main analyses, 116 excitatory and 15 inhibitory neurons, which were detected in the calcium imaging data from the last four sessions (see the ‘Early, middle and late behavioural sessions’ section below) were used.

Synaptic connections were identified by characteristic ultrastructural features⁷¹, including concentrated synaptic vesicles and a post-synaptic density (PSD), which is a darkening/thickening of the post-synaptic membrane at synapses. We annotated synapses in a way that also encodes an estimate of the area of the PSD. Connector objects were annotated on the section in which the PSD appeared longest, and the pre- and post-synaptic nodes were placed in a way such that the distance between them was equal to the length of the PSD (Figs. 1d and 2i–k). This length was taken to be the diameter of circular PSD to estimate PSD area. Although PSDs are not exactly circular, the errors inherent to this estimation are probably small compared with the large variations in size from synapse to synapse^{25,72}.

Starting from the outgoing synapses on the axons of functionally characterized neurons, each post-synaptic neuron was traced until either the cell body was found or the neuron left the edge of the EM volume. Those post-synaptic partners with cell bodies within the volume were classified as pyramidal or non-pyramidal as described above. Some of the post-synaptic partners were other functionally characterized neurons (direct connections). Thus, the first-order output connectivity of the functionally characterized neurons was traced completely within the EM volume (Fig. 1e (inset)).

All tracing (including neuronal morphologies and synapse size estimations) was reviewed a second time by an independent reviewer. In general, we took a conservative approach to terminate ambiguous continuations to avoid merge errors. Neuron tracing required approximately 7,000 h, including around 5,000 h of tracing and about 2,000 h of reviewing.

Connectivity analysis

Analysis of neuron morphology and connectivity was performed by querying the CATMAID database and performing calculations using custom Python code based on the navis (v.1.1.0; <https://github.com/navis-org/navis>) and pymaid (v.2.4.0; <https://github.com/navis-org/pymaid>) libraries. Specific details for particular analyses are provided below.

Axon–dendrite overlap

To quantify how many opportunities pairs of neurons have to make synaptic connections, we calculated the axon–dendrite overlap, which is the length of pre-synaptic axon that comes within 5 μm of the post-synaptic dendrites. All neuron pairs with synaptic connections had non-zero axon–dendrite overlap, because the axons and dendrites need to come into close proximity to form a synapse. Neuron pairs that have a high frequency of synapses per overlap were interpreted to have a high affinity for forming synaptic connections. The axon–dendrite overlap was calculated using the `cable_overlap` function in navis. To mitigate tracing irregularities, the neuron skeletons were first resampled at 100 nm inter-node distance and then smoothed with a radius of 1 μm. We also examined the effect of changing the 5 μm distance threshold in Extended Data Fig. 5d and 6c.

Maximum-selectivity similarity

The maximum-selectivity similarity s quantifies in a continuous manner how similar the selectivity is between two neurons, and was defined as (Extended Data Fig. 5c):

$$s = \text{sign}(c_1 c_2) \sqrt{|c_1| |c_2|},$$

where c_1 and c_2 are the maximum selectivity values (across timepoints) for the two neurons, and the sign of s is positive if c_1 and c_2 have the same sign (co-selective) and negative if they have opposite sign (anti-selective). Thus, the maximum-selectivity similarity encodes both directionality and strength of selectivity.

A pair of neurons was considered to be co-selective if their selectivity indices (see the 'Information-theoretic selectivity index' section above) had the same sign, and anti-selective for opposite sign.

Simultaneous-selectivity similarity

The maximum-selectivity similarity used in the main analysis combines the maximum selectivity (over timepoints) of the pre- and post-synaptic neurons separately. This metric does not guarantee that the same timepoint is used to calculate selectivity for the pre- and post-synaptic neurons. Thus, a pair of neurons can have high maximum-selectivity similarity even if they are never active at the same time or during the same temporal epoch.

To quantify selectivity similarity of neurons restricted to timepoints at which they are both active, we developed the simultaneous-selectivity similarity (Extended Data Figs. 5e–g and 6e). For this metric, the simultaneous-selectivity similarity $s(t)$ is first calculated from the simultaneous trial-type MI values $c_1(t)$ and $c_2(t)$ for each timepoint (see the 'Information-theoretic selectivity index' section above):

$$s(t) = \text{sign}(c_1(t) c_2(t)) \sqrt{|c_1(t)| |c_2(t)|}.$$

Finally, we defined the simultaneous-selectivity similarity as the value of $s(t)$ at the timepoint with the maximum absolute value across timepoints within the trial length. Therefore, this metric considers only selectivity of simultaneous activity in the pre- and post-synaptic neurons.

Synapse size

The PSD area was calculated for each individual synaptic connection based on synapse annotations that estimated the PSD size (see the 'Neuronal circuit reconstruction' section above). The PSD area has been shown to be a correlate of synapse strength (at least for E-to-E synapses in sensory cortex³³).

Normalized synapse frequency

The normalized synapse frequency f quantifies how many synapses are made between a source and target neuron, normalized to the length of overlap between the pre-synaptic axon and post-synaptic dendrite.

$$f = N/L_a,$$

where N is the number of synaptic connections and L_a is the axon-dendrite overlap, defined as the path length of the source neuron's axon that is within 5 μm (maximal spine length²³, Extended Data Figs. 5d and 6c) of dendrites of the target neuron (Fig. 2b). The L_a value quantifies how many opportunities there are for the two neurons to connect based on their fine-scale morphology²³.

Non-connected neuron pairs

In addition to neuron pairs connected by synapses, there were also many non-connected pairs of neurons that have non-zero axon-dendrite overlap but do not form synaptic connections within the EM dataset. However, we cannot rule out the possibility that they form connections outside the bounds of the dataset (due to the limited volume) or that some connections were missed due to tracing errors. Nevertheless, the non-connected pairs can be informative of connection selectivity; in particular, pairs that have a large amount of axon-dendrite overlap without making a connection suggest a low affinity for forming synapses.

To include non-connected neuron pairs in the selectivity analysis, we calculated a pooled synapse frequency f_{pool} as:

$$f_{\text{pool}} = \frac{N_{\text{cn}} + N_{\text{non}}}{L_{\text{cn}} + L_{\text{non}}},$$

where N_{cn} is the total number of synapses among connected pairs, N_{non} is the number of non-connected pairs, L_{cn} is the total axon-dendrite overlap among connected pairs, and L_{non} is the total overlap among non-connected pairs. To avoid outliers and noise from pairs with very little cable overlap (which can vary strongly based on small differences in local tracing), only non-connected pairs with an axon-dendrite overlap of greater or equal to the average overlap per synapse ($L_{\text{cn}}/N_{\text{cn}}$) were included in this calculation. As non-connected pairs are weighted as if they have one synapse each, f_{pool} should be considered to be an estimate of the upper bound on synapse frequency. For Extended Data Figs. 5b and 6b, f_{pool} was calculated separately for co-selective and anti-selective pairs. Confidence intervals and P values (two-tailed) were calculated by bootstrapping over pairs of non-connected neurons.

Structure-function correlations

To quantify structure-function relationships (in Figs. 2h,l and 3c,e and Extended Data Figs. 5f,g,j,k and 6d,e, and across all cell types in Extended Data Fig. 6f–m), we calculated Pearson (linear) correlation coefficients between the functional (maximum-selectivity similarity or simultaneous-selectivity similarity) and anatomical (normalized synapse frequency or PSD area) measures. For normalized synapse frequency, each datapoint was a connection between two neurons, which can involve one or more synapses. The P values reported for these correlations indicate the probability that the correlation deviates from zero (Pearson correlation test with Student's t -statistics⁷³, using the `scipy.stats.pearsonr` or `MATLAB corr` functions).

Early, middle and late behavioural sessions

For analyses investigating how structure-function relationships evolve over time, functional selectivity indices were calculated separately from early (8–10 days before sacrifice), middle (4–7 days before) and late (0–3 days before) sessions. Correlations between similarity indices and connectivity (normalized synapse frequency or PSD area) were then calculated separately for early, middle and late sessions.

Noise correlations

In Extended Data Fig. 7d,e, we compute noise correlations between all pairs of simultaneously recorded neurons within the last four sessions. To discount contributions to the neurons' activity due to shared tuning of neurons to behavioural variables such as running patterns, we implemented a partial noise correlation calculation (adapted from a previous study³²). For each neuron, and separately for each trial type, we performed at each timepoint a linear regression using as dependent variables the single-trial neuron's activity and as independent variables the single-trial values of the x and y positions and velocities and the heading angle. We then regressed away the linear contribution of the behavioural variables to neural activity as well as the trial-type mean activity, and computed noise correlations using the average over time during the trial of the residual activity. We pooled all available data from the four last sessions to compute the estimate of noise correlations for each pair. The details of the calculation procedure were set to optimize the signal-to-noise ratio of calculations on simulated activity of pairs of neurons (Extended Data Fig. 12; see the 'Method selection for estimating trial type selectivity and noise correlation' section of the Supplementary Methods). As regressions are meaningless in the absence of variability of firing across trials, for each pair of neurons, only timepoints at which both neurons had significant non-zero activity (corresponding to timepoints in which the P values of a one-sample t -test against zero activity was lower

than 0.05 for both neurons) were included in the noise correlation calculation.

Same-epoch and different-epoch connections

For each neuron, we defined the peak epoch as the epoch (see the ‘Behaviour and calcium imaging’ section for definitions of the epochs) in which the trial average activity (defined as the mean over trials and over left and right trial types) reached its maximum value. Connections between neuron pairs that shared the same peak epoch were referred to as same-epoch connections, whereas connections between pairs with different peak epochs were referred to as different-epoch connections.

Network models

We examined a simple linear network model with recurrent connections (Fig. 4 and Extended Data Figs. 8, 9a,b and 11c,e). The network comprises excitatory and inhibitory units organized in two trial-type selective subnetworks, each of them including one excitatory and one inhibitory unit (respectively E_L, I_L and E_R, I_R for left and right subnetworks). The dynamics of the network follows the differential equation:

$$\dot{r}_i = -r_i + \sum_{j=1}^{N=4} J_{ij} r_j + I_{\text{ext},i} + \eta_i(t), \quad (1)$$

where $\mathbf{r} = (r_{E_L}, r_{I_L}, r_{E_R}, r_{I_R})$ represents the firing rate deviation from the baseline activity level. The term J_{ij} represents the connectivity weight between the presynaptic unit j and the post-synaptic unit i . We model left and right trial types through modulations of the external trial-type selective input to the excitatory units, $I_{\text{ext},i}$, assumed to be constant in time. The term $\eta_i(t)$ represents a source of zero-mean Gaussian input noise to unit i , which in general depends on time.

For each pre-synaptic and post-synaptic types $X, Y \in E, I$, we denoted by w_{YX}^{in} and w_{YX}^{out} the synapses that connect units belonging to the same or to different subnetworks, respectively, and we assumed that these connections were symmetric with respect to the left and right subnetworks L and R . We denote by $\Delta_{YX} = w_{YX}^{\text{in}} - w_{YX}^{\text{out}}$ the connection selectivity for the X-to-Y synapses.

We examined the response of the network to external inputs. We defined the network response as the steady state of the network with external input I_{ext} and zero external noise. When the dynamics is linear, the steady-state \mathbf{r}^* can be written as a function of the connectivity matrix \mathbf{J} and of the external input \mathbf{I}_{ext} as:

$$\mathbf{r}^* = (\mathbf{I} - \mathbf{J})^{-1} \mathbf{I}_{\text{ext}}. \quad (2)$$

We therefore examined the network dynamics as a function of Δ_{YX} at fixed average connection strength $S_{YX} = w_{YX}^{\text{in}} + w_{YX}^{\text{out}}$. In Fig. 4b–d and Extended Data Figs. 8 and 9a,b, we set the overall magnitude of the connectivity weights to $S_{IE} = S_{EI} = 2$, while varying the E-to-I selectivity in the range $\Delta_{IE} \in [0, 2]$ and the I-to-E selectivity in the range $\Delta_{EI} \in [-2, 2]$ (Fig. 4b–d and Extended Data Fig. 8).

Assuming that the external selective inputs are symmetric for left and right trials (that is, for left trials, $\mathbf{I}_L = (c_1, 0, c_2, 0)$ with $c_1 > c_2$, while, for right trials, $\mathbf{I}_R = (c_2, 0, c_1, 0)$ with $c_1 > c_2$), we computed the trial-type encoding dimensions, defined as the vectors connecting the trial-specific mean responses of the excitatory units on left and right trials, both at the input and output stages, given by $\mathbf{d}_{\text{in}} = \mathbf{I}_L - \mathbf{I}_R$ and $\mathbf{d}_{\text{out}} = \mathbf{r}_L^* - \mathbf{r}_R^*$, where $\mathbf{r}_{L/R}^*$ are the responses of the network to inputs $\mathbf{I}_{L/R}$. The input encoding dimension is the dimension connecting trial-specific responses in the absence of recurrent connections. The input and output encoding dimensions are related through the relationship:

$$\mathbf{d}_{\text{out}} = \frac{1}{\delta} \mathbf{d}_{\text{in}}, \quad (3)$$

where $\delta = (1 - \Delta_{EE})(1 + \Delta_{II}) + \Delta_{EI}\Delta_{IE}$. When the value of δ decreases, the separation between the output mean activities $\|\mathbf{d}_{\text{out}}\|$ increases. Moreover, when $\delta < 1$, the separation between trial-specific activities at the output is larger than at the input stage. Importantly, under the assumption of symmetric inputs for left and right trial types, the separation of mean responses depends only on the connection selectivity Δ_{XY} , but not on the average connection strength S_{XY} . In the absence of E-to-E and I-to-I connections, increased separation of trial-specific activities occurs when $\Delta_{EI}\Delta_{IE} < 0$, corresponding to co-selective E-to-I and anti-selective I-to-E (Fig. 4c,d and Extended Data Figs. 8 and 9a,b).

In Fig. 4c,d and Extended Data Figs. 8 and 9a,b, we computed the decoding accuracy of an optimal linear decoder trained to classify trial types from the neural activity of the excitatory units. Given the difference in mean activity between left and right trial types $\Delta\mu$, and the noise covariance matrix C , we computed the signal-to-noise ratio as $\text{SNR} = \Delta\mu^T C^{-1} \Delta\mu$.

The decoding accuracy of an optimal linear classifier can then be computed as:

$$\text{Decoding accuracy} = \Phi\left(\sqrt{\frac{\text{SNR}}{4}}\right), \text{ where } \Phi(x) = \int_{-\infty}^x \frac{e^{-z^2/2}}{\sqrt{2\pi\sigma^2}} dz. \quad (4)$$

In Fig. 4c,d and Extended Data Figs. 8 and 9a,b, we defined the relative decoding accuracy as:

$$\text{Relative decoding accuracy} = \frac{\text{Decoding accuracy}(\text{output}) - 0.5}{\text{Decoding accuracy}(\text{input}) - 0.5}$$

In Extended Data Fig. 11, we examined the predictions of the network model on how the activity of the excitatory and inhibitory units affect the correctness of the behavioural choices. As in previous work^{13,14}, the behavioural choice was modelled by comparing the firing rate of the two excitatory units. The model expressed a left choice when the left excitatory unit fired more strongly than the right excitatory unit. The model choices were determined from the excitatory responses alone because generally only excitatory neurons project to downstream areas. Correct/error trials were those in which the trial-type identity and choice coincided/differed. We considered the version of the model with input on excitatory units. For only the analyses of Extended Data Fig. 11, we added a non-selective background input term on top of the trial-type selective input component. The common background input determined the average firing across the preferred and non-preferred trial-types. The inclusion of a background input does not affect any of the results of the other analyses (Fig. 4b–d and Extended Data Figs. 8 and 9a,b), but aided the comparison between the model and the data in this specific analysis as, in the comparisons of activity on error versus correct trials in PPC data, the activity differences were normalized to the average activity over all trials (Extended Data Fig. 11; see the ‘Error trial analysis’ section above).

In Fig. 4e–i and Extended Data Figs. 9c,d and 10, we trained high-dimensional RNN models⁷⁴ to reproduce the trial-averaged PPC activity traces on left and right trials. We then analysed the connectivity weights obtained after training to evaluate the predictions of the linear rate model on the trained networks. Before training, the PPC traces were normalized to the peak of each cell’s activity on preferred trials. The RNN is described by the differential equation

$$\tau \dot{x}_i = -x_i + \sum_{j=1}^N J_{ij} \phi(x_j) + I_{\text{ext},i}(t) + \sigma \xi_i(t), \quad (5)$$

where x_i represent the synaptic current of neuron i and $r_i = \phi(x_i)$ is the firing rate of neuron i , $I_{\text{ext},i}$ represents the trial-type selective external input, $\sigma \xi_i(t)$ is a source of time-dependent external noise (Gaussian noise with zero mean and standard deviation σ , uncorrelated across neurons and time) and $\tau = 0.025$ is the time constant of the neurons.

The network is constrained to have the same number of neurons and the same type (116 E neurons and 15 I neurons) as the experimentally reconstructed circuit. Accordingly, the connectivity matrix J_{ij} is constrained to satisfy Dale's law, that is, with elements belonging to the same column having the same sign according to the type of the pre-synaptic neurons⁴⁰. The external input $I_{ext,i}$ consisted of temporally correlated white noise (see the 'The recurrent neural network model' section of the Supplementary Methods). In Fig. 4e–i and Extended Data Fig. 10, the input was fed only to the excitatory neurons. In Extended Data Fig. 9c,d, both excitatory and inhibitory neurons selective for the same trial type received selective external input with the same amplitude. Only RNNs that fit well with PPC activity ($R^2 > 0.3$) were considered in the analyses of Fig. 4f–i and Extended Data Figs. 9c,d and 10.

Reporting summary

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

Data availability

Directions for accessing the EM dataset, reconstructed neurons and calcium imaging data are available at GitHub (https://github.com/htem/PPC_inhibitoryMotifs).

Code availability

Software and analysis code is available at Zenodo (<https://zenodo.org/doi/10.5281/zenodo.10310186>)⁷⁵. Code used to perform all information-theoretic analyses was published previously⁶⁶ and is available online (<https://doi.org/10.1186/1471-2202-10-81>). Code for the network model analyses is available at Zenodo (<https://doi.org/10.5281/zenodo.10200999>)⁷⁶.

59. Aronov, D. & Tank, D. W. Engagement of neural circuits underlying 2D spatial navigation in a rodent virtual reality system. *Neuron* **84**, 442–456 (2014).
60. Pachitariu, M. et al. Suite2p: beyond 10,000 neurons with standard two-photon microscopy. Preprint at *bioRxiv* <https://doi.org/10.1101/061507> (2016).
61. Pnevmatikakis, E. A. et al. A structured matrix factorization framework for large scale calcium imaging data analysis. Preprint at <https://doi.org/10.1101/061507> (2016).
62. Chan, K. Y. et al. Engineered AAVs for efficient noninvasive gene delivery to the central and peripheral nervous systems. *Nat. Neurosci.* **20**, 1172–1179 (2017).
63. Dimidschstein, J. et al. A viral strategy for targeting and manipulating interneurons across vertebrate species. *Nat. Neurosci.* **19**, 1743–1749 (2016).
64. Marshel, J. H. et al. Cortical layer-specific critical dynamics triggering perception. *Science* **365**, eaaw5202 (2019).
65. Shannon, C. E. A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 379–423 (1948).
66. Magri, C., Whittingstall, K., Singh, V., Logothetis, N. K. & Panzeri, S. A toolbox for the fast information analysis of multiple-site LFP, EEG and spike train recordings. *BMC Neurosci.* **10**, 81 (2009).
67. Panzeri, S. & Treves, A. Analytical estimates of limited sampling biases in different information measures. *Network* **7**, 87–107 (1996).
68. Bogovic, J. A., Hanslovsky, P., Wong, A. & Saalfeld, S. Robust registration of calcium images by learned contrast synthesis. In *Proc. 2016 IEEE 13th International Symposium on Biomedical Imaging* 1123–1126 (ISBI, 2016).

69. Saalfeld, S., Cardona, A., Hartenstein, V. & Tomancak, P. CATMAID: collaborative annotation toolkit for massive amounts of image data. *Bioinformatics* **25**, 1984–1986 (2009).
70. Schneider-Mizell, C. M. et al. Quantitative neuroanatomy for connectomics in *Drosophila*. *eLife* **5**, 1133–1145 (2016).
71. Peters, A., Palay, S. L. & Webster, H. deF. *The Fine Structure of the Nervous System: Neurons and Their Supporting Cells* (Oxford Univ. Press, 1991).
72. Loewenstein, Y., Kuras, A. & Rumpel, S. Multiplicative dynamics underlie the emergence of the log-normal distribution of spine sizes in the neocortex in vivo. *J. Neurosci.* **31**, 9481–9488 (2011).
73. Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. *Numerical Recipes: The Art of Scientific Computing* (Cambridge Univ. Press, 2007).
74. Perich, M. G. et al. Inferring brain-wide interactions using data-constrained recurrent neural network models. Preprint at *bioRxiv* <https://doi.org/10.1101/2020.12.18.423348> (2021).
75. Kuan, A. T. Data analysis code for 'synaptic wiring motifs in posterior parietal cortex support decision-making'. *Zenodo* <https://doi.org/10.5281/zenodo.10310186> (2023).
76. Bondanelli, G. Circuit modeling code for 'synaptic wiring motifs in posterior parietal cortex support decision-making'. *Zenodo* <https://doi.org/10.5281/zenodo.10200999> (2023).
77. Wang, Q. et al. The Allen Mouse Brain Common Coordinate Framework: a 3D reference atlas. *Cell* **181**, 936–953 (2020).

Acknowledgements We thank M. Eroles, M. Larson, K. Ng, L. Hulshof, T. Xie, O. Sato, T. Khanna, E. DiPietro, E. Phalen, L. Guo, S. Kushner, L. Decoursey, S. Xie, C. Zharyy, E. Zboinski, L. Sadilina, M. Narwani, R. Xu, K. Molloy, D. Patel, B. Reicher, L. Walsh, Z. Diaks, Z. Wing Fan, N. Byrne, J. Shin, T. Pedersen and A. Buckner for neuron tracing; X. Chen for providing image annotation software; W. Copeland, J. Kiperman, B. Sanders and T. Yusuf for assisting with EM imaging; W. Lou, S. Gerhard, J. Phelps and R. Zheng for assistance with image alignment; A. Buckner, R. Xu and S. Ah Kim for assistance with mouse training; N. Pettit for assistance with calcium imaging; S. Kira, A. Emanuel and A. Wang for providing mice; G. Hood for providing alignment software (AlignTK); P. Schlegel for providing analysis software (pymaid and navis); K. Hayworth for help with X-ray microCT processing; and M. Andermann, R. Born, T. Nguyen, X. Chen and D. Hidalgo for comments on the manuscript. AAV8-CaMKIIa-ChRmine-mCherry-Kv2.1 was a gift from the Deisseroth laboratory. pAAV-mDlx-NLS-mRuby2 was a gift from V. Gradinaru (Addgene, 99130-AAV9). This work was supported by the NIH (R01NS108410 to C.D.H., S.P. and W.-C.A.L.; DP1MH125776 and R01NS089521 to C.D.H.; RF1MH114047 to W.-C.A.L.; F32MH118698 to D.E.W.; and K99EB032217 to A.T.K.), the Bertarelli Program in Translational Neuroscience and Neuroengineering, Edward R. and Anne G. Lefler Center, Stanley and Theodora Feldberg Fund to W.-C.A.L. Portions of this research were conducted on the Orchestra High Performance Compute Cluster at Harvard Medical School partially provided through NIH NCCR (S10RR028832).

Author contributions A.T.K., G.B., L.N.D., S.P., C.D.H. and W.-C.A.L. conceptualized the project and designed experiments. L.N.D. and D.E.W. performed mouse behaviour and calcium imaging experiments. W.-C.A.L., D.G.C.H. and A.T.K. prepared tissue samples for EM. D.G.C.H. and A.T.K. performed GridTape sectioning. B.J.G. and A.T.K. developed automated EM techniques and performed EM imaging. A.T.K. and L.A.T. performed EM image processing and alignment. M.K. and A.T.K. performed co-registration between calcium imaging and EM datasets. J.H. and A.T.K. performed and managed neuron tracing and data annotation. A.T.K. performed analysis of circuit connectivity. A.T.K., G.B. and S.P. analysed functional data. G.B. performed circuit modelling. A.T.K., G.B., S.P., C.D.H. and W.-C.A.L. wrote the paper, and all of the authors assisted in reviewing and revising the manuscript.

Competing interests Harvard University filed a patent application regarding GridTape (WO2017184621A1) on behalf of the inventors, including W.-C.A.L., D.G.C.H. and B.J.G., and negotiated licensing agreements with interested partners. The other authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41586-024-07088-7>.

Correspondence and requests for materials should be addressed to Stefano Panzeri, Christopher D. Harvey or Wei-Chung Allen Lee.

Peer review information Nature thanks the anonymous reviewers for their contribution to the peer review of this work.

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