nature communications

Article

https://doi.org/10.1038/s41467-025-58698-2

Multiplexed subspaces route neural activity across brain-wide networks

Received: 6 May 2024

Accepted: 28 March 2025

Camden J. MacDowell[1], Alexandra Libby[1,3], Caroline I. Jahn[1,3], Sina Tafazoli[1,3], Adel Ardalan[1,3] & Timothy J. Buschman[1,2] ✉

Cognition is flexible, allowing behavior to change on a moment-by-moment basis. Such flexibility relies on the brain's ability to route information through different networks of brain regions to perform different cognitive computations. However, the mechanisms that determine which network of regions is active are not well understood. Here, we combined cortex-wide calcium imaging with high-density electrophysiological recordings in eight cortical and subcortical regions of mice to understand the interactions between regions. We found different dimensions within the population activity of each region were functionally connected with different cortex-wide 'subspace networks' of regions. These subspace networks were multiplexed; each region was functionally connected with multiple independent, yet overlapping, subspace networks. The subspace network that was active changed from moment-to-moment. These changes were associated with changes in the geometric relationship between the neural response within a region and the subspace dimensions: when neural responses were aligned with (i.e., projected along) a subspace dimension, neural activity was increased in the associated regions. Together, our results suggest that changing the geometry of neural representations within a brain region may allow the brain to flexibly engage different brain-wide networks, thereby supporting cognitive flexibility.

Cognition arises from the dynamic interaction of brain regions[1,2]. Broad networks of regions are involved in sensory processing[3], decision making[4], and motor actions[5,6]. Which network is engaged changes over time, both during spontaneous behavior[7–10] and in a task-dependent manner[2,4,11–13]. This is thought to reflect changes in how neural activity propagates through networks of regions which, in turn, can change how information is processed by the brain[14]. For this reason, changing how neural activity propagates between regions is thought to be a fundamental component of cognitive control. Several mechanisms have been proposed for determining how neural activity flows between regions[15], including changes in the gain of neural response[14] or synchrony of brain regions[16,17]. Here, we test a geometric model of cognitive control.

This model is built on two recent observations. First, recordings in early visual cortex found only a part of the neural response within one brain region was functionally connected (i.e., co-varied/correlated) with a second region[18]. In other words, within the high-dimensional space of all possible patterns of neural activity in a region, there was a subspace that was 'shared' with another region. Second, recent work found the geometry of neural representations can change from moment-to-moment, moving between subspaces of neural activity in a task-dependent manner[19]. Together, these results suggest a geometric mechanism for cognitive control—if a neural representation is changed such that it becomes aligned with (i.e., projected along) the subspace that is shared between regions, then this could selectively route that representation to the region(s) associated with that subspace.

---

[1] Princeton Neuroscience Institute, Princeton University, Washington Rd, Princeton, NJ, USA. [2] Department of Psychology, Princeton University, Washington Rd, Princeton, NJ, USA. [3] These authors contributed equally: Alexandra Libby, Caroline I. Jahn, Sina Tafazoli, Adel Ardalan. ✉e-mail: tbuschma@princeton.edu

However, key hypotheses of this mechanism remain untested. First, it is unknown how multiple brain regions interact through subspaces. Second, it is unknown whether aligning a neural representation with a specific subspace predicts the dynamics of neural activity across regions.

To begin to test these hypotheses, we combined large-scale electrophysiology and widefield calcium imaging as mice engaged in spontaneous behaviors. First, we found that subspace dimensions were functionally connected with a network of brain regions (i.e., neural activity projected on that dimension was correlated with activity in other brain regions). Different dimensions were associated with different, independent, but overlapping networks of regions, potentially allowing interactions between regions to be multiplexed. Second, we found that the neural response within a region changed over time to become more (or less) aligned with different subspace dimensions. These changes were related to the propagation of neural activity across regions—when the neural response was more aligned with a subspace dimension, then neural activity was greater in the brain regions associated with that dimension. Altogether, our results support the hypothesis that aligning neural responses to different subspace networks may provide a geometric mechanism for dynamically controlling the flow of information across the brain.

Results

Using four Neuropixels probes[20], we recorded neural activity in eight cortical and subcortical brain regions simultaneously (Fig. 1A-D; see Methods). Overall, >6500 neurons were recorded in hippocampus (HPC), thalamus (TH), prelimbic cortex (PL), frontal motor regions

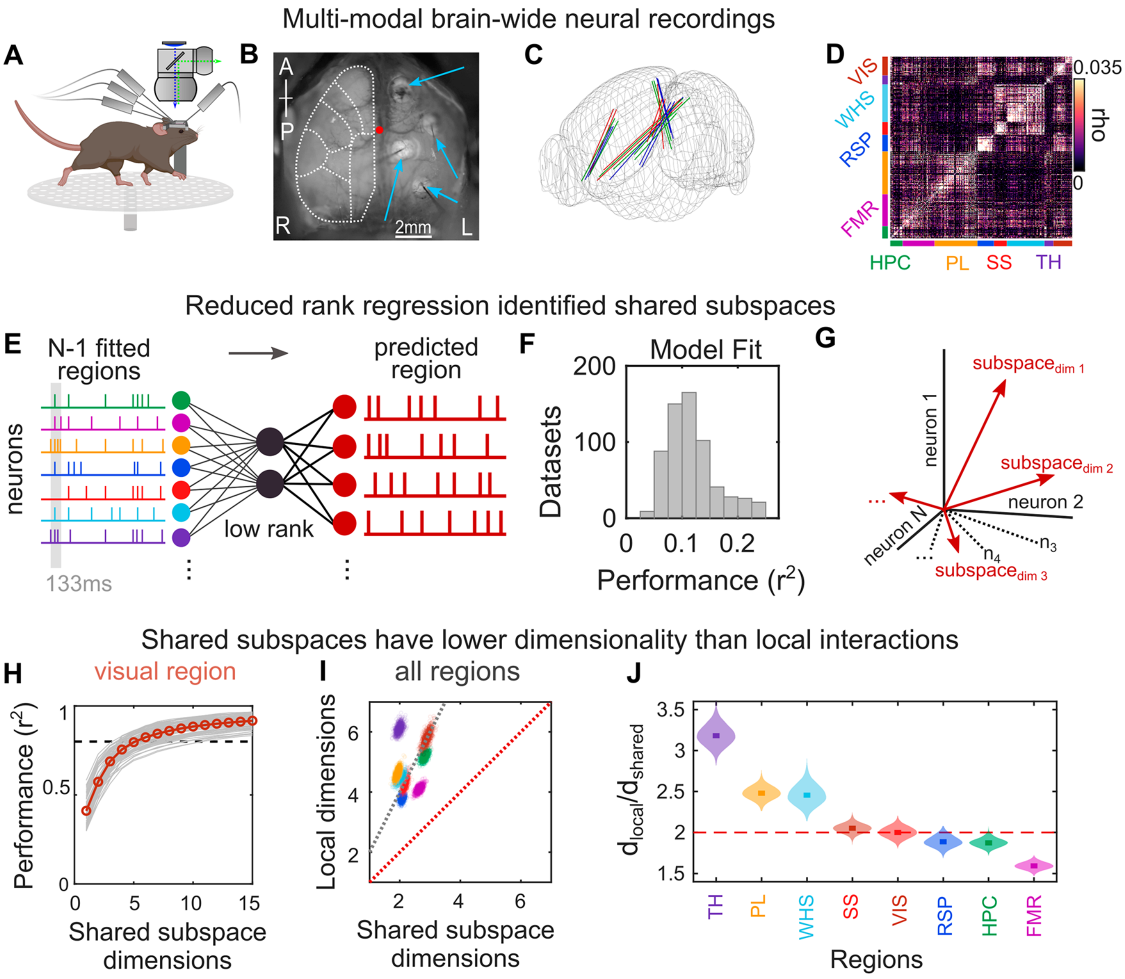

Fig. 1 | Interactions between brain regions are organized into subspaces. **A** Schematic of experiments combining recordings from four Neuropixels probes across cortical and subcortical regions and simultaneous widefield calcium imaging across dorsal cortex. Created in BioRender. Tafazoli, S. (2025) https://BioRender.com/y38t274. **B** Field-of-view of widefield imaging showing the dorsal cortex and four craniotomies with implanted electrodes (blue arrows). White lines outline cortical regions (see Fig. S1). R, L, A and P denote animal's right, left, anterior and posterior. **C** Reconstructed probe locations from six recordings in three mice; colors indicate mice. **D** Correlogram showing correlation between pairs of cells for an example recording. Colored bars along axes indicate brain regions. **E** Schematic of reduced rank regression (RRR) to predict activity of each neural population. Spontaneous spiking activity was binned to 133 ms windows, and the average response was removed so that moment-to-moment spiking variability of one region was predicted using spiking variability of all other regions. **F** Cross-validated performance of all RRR models ($n = 630$ datasets, see Methods). **G** Schematic showing that a shared subspace is spanned by a set of dimensions within the neural activity of the predicted region that are correlated with neural activity in other regions. **H**, **I** A small subspace of neural activity within a region was shared with other regions. **H** Cumulative percent of the explainable variance captured by each dimension of the RRR model (see Methods). Red and gray lines show mean performance and cross-validated performance from $n = 84$ datasets, respectively. Dashed line shows 80% of explainable variance. **I** Bootstrapped distribution showing the number of 'shared' subspace dimensions and the number of 'local' dimensions within a region. Dimensionality was estimated as the number of dimensions of a RRR model that reliably predicted variance in neural activity in withheld data (using a subset of $n = 25$ neurons to balance across areas, decreasing the overall dimensionality estimate, see Methods). Dots show bootstrapped mean ($n = 1000$). Colors correspond to regions. Dotted lines show 1:1 and 2:1 ratio (red and grey, respectively). **J** Violin plots of estimated ratio of local-to-subspace dimensionality for each recorded region. Full distribution shown. See also Figs. S1-S4.

(FMR), retrosplenial cortex (RSP), visual-associated cortex (VIS), primary somatosensory cortex (SS), and whisker somatosensory cortex (WHS; see Methods for detailed parcellation). Simultaneous widefield calcium imaging captured the dynamics of neural activity across dorsal cortex. All recordings were done while mice spontaneously engaged in different behaviors, ensuring a wide variety of cortex-wide neural dynamics[7,8].

To quantify how the eight recorded brain regions were functionally connected, we used reduced rank regression[21] to predict the pattern of neural activity in one brain region as a linear function of the activity in all other regions (Fig. 1E; note: models were fit to moment-by-moment co-variance in neural activity between regions, after removing the average response across similar epochs of cortex-wide dynamics, see Methods and Fig. S1). All eight brain regions were interconnected. The regression models explained a significant amount of the variance in neural activity in all eight regions (Fig. 1F, cross-validated $r^2 = 0.12$, CI: 0.11–0.12, $p < 0.001$ versus shuffled controls, see Methods). Furthermore, each region contributed to the prediction of other regions; excluding any region decreased the cross-validated performance of >99% of the models (mean decrease of 7.6% of the explained variance, CI: 7.3–7.9%; this was robust to changes in model parameters, Fig. S2).

Reduced rank regression (RRR) identified the set of orthogonal dimensions of neural activity within each brain region that were influenced by other regions (Fig. 1G; $Y_i \sim \beta_i X$, where X and Y are the neural activity in the source and predicted regions, respectively, β are linear weights, and $i \in [1..N]$ is the dimension, sorted by variance explained; see Methods). It is important to note that RRR only captures the correlation between populations of neurons and does not measure the directionality of interactions (i.e., neural activity could be propagating from the source region to the target region, vice-versa, or bidirectionally).

Brain regions were connected through a 'shared subspace'. Consistent with previous work[18,22], only a few dimensions of the neural population within a region were correlated with other regions. For example, within visual cortex, 80.0% of the variance in neural activity that could be explained by other regions was captured in the first 5 dimensions and 88% was captured by 10 dimensions (Fig. 1H). A similar pattern was seen across all brain regions; ≥80% of the explainable variance was captured by an average of 5.5 dimensions (CI: 5.3–5.7) and 10 dimensions captured an average of 90.0% of the variance (CI: 89.3–90.1%). Together, these dimensions encapsulated the shared subspace of population activity within each region that was correlated with other regions[18,23,24].

While different dimensions of the shared subspace were orthogonal to one another, they did not engage different populations of neurons. Rather, dimensions engaged the same population of neurons in the target region (Fig. S3), consistent with anatomical evidence that neurons projecting between regions branch to innervate several different regions[25,26].

We refer to the shared subspace as a 'subspace' because it was significantly smaller than the 'local' space of that region's population activity. To estimate the local dimensionality, we used RRR to measure the number of dimensions in one half of the neural population that predicted the activity in the other half of neurons (see Methods). We found that, on average, 2.1 times more dimensions were needed to explain the same amount of variance in neural activity within a region than between regions (CI: 2.0–2.2; Fig. 1I). This suggests only a portion of the space of neural activity within a region was functionally connected to other regions. This was true in all recorded regions (Fig. 1J); although there was variability between regions, which could reflect differences in how each region was functionally connected with the broader network of recorded regions. Importantly, similar results were seen regardless of the total number of regions used to estimate dimensionality (Fig. S4).

Subspace dimensions interact with a distributed network of brain regions

Given these findings, we aimed to understand how shared subspaces integrated neural activity from other brain regions. One hypothesis is that each dimension of the shared subspace reflects an exclusive connection between a pair of regions, essentially acting as a 'channel' between the two regions. In this model, networks of regions would be formed from a set of pairwise interactions (Fig. 2A, upper). Alternatively, each dimension of the subspace may be functionally connected to multiple brain regions, acting as a 'web' to integrate their activity into a single representation (Fig. 2A, lower). In this model, network interactions are inherent in the functional connectivity between regions.

To discriminate these hypotheses, we examined the beta-weights for each subspace dimension (Fig. 2B). The beta-weight of a neuron captures the relative influence of that neuron on neural activity along a subspace dimension within the target region. Therefore, if a subspace dimension reflects the exclusive coupling between two regions, then neurons with high magnitude beta-weights will be concentrated in one source region. Alternatively, if a subspace dimension integrates activity from multiple regions, then neurons with high magnitude beta-weights should be distributed across brain areas. We found evidence for both: some subspace dimensions integrated across distributed networks while others were more exclusive channels between regions.

For example, the largest weights for the first subspace dimension of the visual region were from neurons in retrosplenial, hippocampal and somatosensory regions (Fig. 2C; neural activity was normalized to allow for direct comparison of weights, see Methods). To quantify this, we measured the cumulative distribution of weights for each region and found that retrosplenial, hippocampal, and somatosensory regions contributed a greater proportion of the neurons with larger weights than would be expected by chance (Fig. 2D; area under the curve, AUC, was significantly greater than 0.5, all $p < 0.001$ by permutation, see Methods). Across all recorded regions, the majority of the first two subspace dimensions within each region integrated neural activity from multiple other areas (57%, CI: 38–69%).

However, while many subspace dimensions were correlated with multiple regions, other dimensions were more exclusive. For example, somatosensory neurons were most predictive of activity in the first subspace dimension of whisker cortex (Fig. S5A). In general, higher subspace dimensions were correlated with fewer regions. Overall, 71% of subspace dimensions 3 and 4 and 97% of dimensions 9 and 10 were predominantly correlated with only one other recorded region (CI: 56–81% and 88–100%, respectively; Fig. S5B). These results did not depend on the exact model used; similar results were seen when comparing subspaces across independent models (Fig. S6; see Methods).

Different degrees of integration for different dimensions may allow the brain to control how broadly information is shared. Neural activity along early dimensions will tend to be shared broadly, while higher dimensions will be more specific. Previous work suggests that representations of motor movements and behavioral state are distributed across broad networks of regions[4,7,8,27,28]. To test this, we investigated the correlation between the animal's movements (i.e., nose, whisker pad, and fore-limb motion energy) and neural activity projected along each subspace dimension. Across all regions, activity along early subspace dimensions was significantly more correlated with motor activity (Fig. 2E, F), consistent with the idea that motor information is shared widely[6,8].

In general, the hippocampus and visual cortex contributed the largest weights to the other recorded regions. To quantify the strength of coupling between two regions, we measured the average AUC of the cumulative distribution of weights (as in Fig. 2D) across all subspace dimensions. Hippocampus and visual cortex had the first or second strongest weights for almost all of the other regions (Figs. 3A and S7).

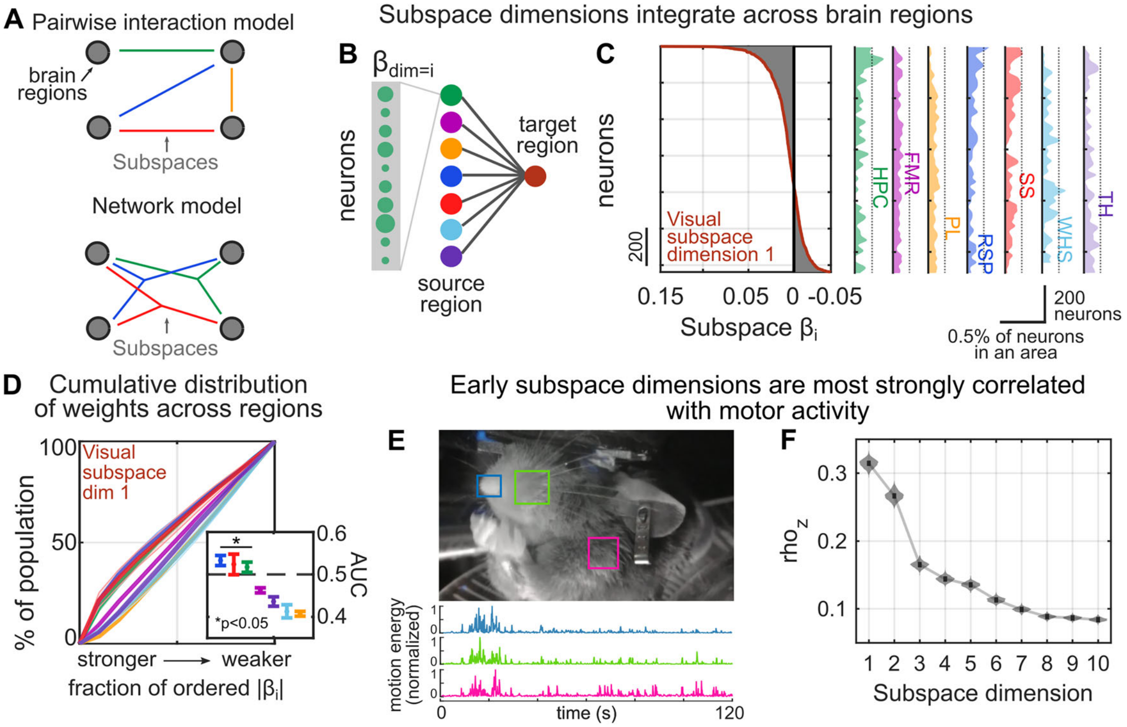

Fig. 2 | Subspaces between areas reflect brain-wide networks. **A** Schematic showing two hypotheses of how shared subspaces could integrate neural activity across brain regions. Subspaces may reflect pairwise interactions between regions (top) or may reflect networks of regions interacting through a single dimension (bottom). **B** Schematic of how beta weights of the regression model estimate the contribution of neurons in a source region to a target region. **C** Example distribution of beta weights (x-axis) across all neurons (y-axis) contributing to predicting activity along the first dimension of the shared subspace in visual cortex. Left plot shows distribution across all regions. Right plots show density histogram of weights within each region (normalized to account for different numbers of neurons: dotted lines added to aid visualization). **D** Cumulative distribution of the percent of neurons within a source area contributing to the beta weights along the first dimension of the visual subspace (as in **C**). X-axis shows the beta weights across all neurons, ordered from strongest to weakest. Lines show mean, shaded region shows bootstrapped 95% confidence interval from $n = 84$ datasets (see Methods). Inset shows the mean and confidence intervals of the area under the curve (AUC) for the contributions from each region. P values were $p < 0.001$; $p = 0.026$; $p < 0.001$ for RSP, SS, and HPC, respectively; one-sided bootstrap test versus an AUC of 0.5, $n = 1000$ permutations. **E** The relationship between subspace activity and motor activity was quantified by correlating the activity along each subspace dimension with the motion energy of animals' nose (blue), whisker-pad (green), and shoulder (pink; see Methods). **F** Bootstrapped distribution showing the average correlation (y-axis, z-transformed) between motor activity and activity along each subspace dimension (x-axis) across datasets. Subspace dimension 1 was significantly more correlated with motor activity than all other dimensions ($p < 0.001$, one-sided bootstrap test, $n = 1000$ bootstraps from $n = 630$ datasets). See also Figs. S5-S6.

This suggests that, during spontaneous behavior, hippocampus and visual cortex may act as 'hubs' that are broadly functionally connected with other regions (Fig. 3B).

Subspace dimensions are functionally connected to cortex-wide subspace networks

So far, our results suggest a) brain regions are functionally connected through a shared subspace and b) that each dimension of this shared subspace is connected with a network of brain regions. However, even with multiple Neuropixels probes, we have a limited view of the network of brain regions that are connected to an individual subspace dimension. Therefore, in order to visualize the broader network, we leveraged simultaneous widefield imaging and electrophysiological recordings.

Widefield calcium imaging captured the activity of populations of pyramidal neurons in the superficial layers of dorsal cortical regions[29]. In order to match timescales across imaging and electrophysiology, we used a shallow feed-forward network to estimate the neural activity underlying the calcium signal (Figs. S8-S10; see Methods). Then, we correlated the moment-by-moment fluctuations in neural activity along each subspace dimension of a brain region, as estimated with the RRR fit to electrophysiological recordings, with fluctuations in neural activity across the entire cortex, as measured by the widefield imaging (Fig. 4A; see Methods; analyses used the first ten subspace dimensions). The resulting maps visualize the network of cortical regions that co-varied with neural activity along each subspace dimension of a region (Fig. 4B; additional examples in Fig. S11; all maps were thresholded at significance, $p < 0.05$, corrected for false-discovery rate, see Methods). These maps reflect the 'subspace network' of cortical regions that are functionally connected to a dimension of a region's shared subspace.

Validating our approach, the subspace networks observed in the imaging data were significantly correlated with the relative strength of interactions observed in electrophysiological recordings (dimension 1: $r_{spearman} = 0.22$, CI: 0.15–0.29, $p = 0.001$; dimension 2: $r_{spearman} = 0.13$, CI: 0.072–0.20, $p = 0.001$; overall: $r_{spearman} = 0.055$, CI: 0.033–0.076, $p = 0.001$; see Methods). Note that Fig. 4B shows subspace networks for a representative dataset; the exact subspace networks (and their order) may vary across datasets. Nevertheless, several observations were consistent across all datasets.

First, we found there was diversity in the anatomical structure of subspace networks. Consistent with our electrophysiology results, the first subspace dimension of most regions was correlated with a broad network that spanned the majority of the dorsal cortex (Fig. 4B, first column; 75.4% of pixels, CI: 73.3–77.2%). In contrast, subsequent subspace dimensions engaged more localized networks of cortical areas (Figs. 4B and S12). Some of these subspace networks followed known functional boundaries, such as anterior-lateral somatomotor networks[30,31] (Fig. 4B, green circles), while others followed anatomical boundaries, such as primary visual cortex (Fig. 4B, yellow circles). Yet,

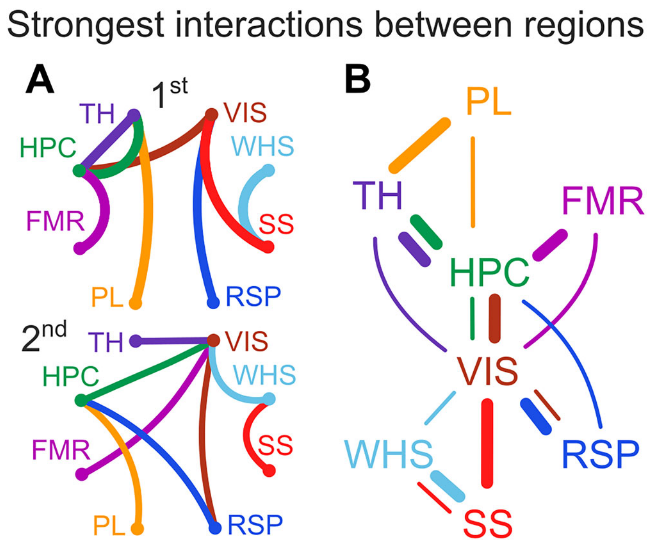

Fig. 3 | Schematic of interactions between regions. A Schematic showing the network of strongest (top) and second strongest (bottom) interactions between areas, across all datasets. Strength of contributions were taken as the AUC of beta weights (as in Fig. 2D, except averaged across all subspace dimensions). Colored lines indicate the subspace target region, e.g., the strongest contribution to FMR is from HPC (purple line). **B** Network graph showing the two strongest interactions between brain areas. Thicker/thinner lines show the first/second strongest interaction, respectively. See also Fig. S7.

other subspace networks were novel, spanning traditional borders (Fig. 4B, red circles).

Interestingly, we found subspace networks associated with different dimensions and different brain regions often had similar topographies (Fig. 4B, highlighted by colored circles). For example, the subspace network associated with dimensions 6 and 9 in frontal motor were similar to one another (Fig. 4B, networks shared 67.5% of pixels, $p < 0.001$ above chance, one-sided binomial test). This was seen in all regions and animals: overall, 38% of dimensions within a region had similar subspace networks (Fig. S13E; quantified as sharing the same cluster, $p <= 2e{-}5$, permutation test, see Methods). This suggests these subspace networks are multi-dimensional, which could increase the information capacity of the subspace network, allowing more complex representations to be communicated between regions.

While subspace networks were similar in different dimensions within a region, they were also similar across regions. For example, the subspace network associated with dimension 6 and 9 in frontal motor cortex was similar to those associated with dimension 5 in thalamus and dimensions 5 and 10 in retrosplenial (Fig. 4B, green circles). Likewise, the subspace networks associated with dimensions 2 and 3 in thalamus and dimension 3 in retrosplenial were similar to one another (Fig. 4B, blue circles; shared 86.7% of pixels, $p < 0.001$). Clustering across all datasets and regions found 13 distinct categories of subspace networks (with the categories reflected in color circles in Fig. 4B, see Fig. 5 for clusters and Methods). Importantly, similar subspace networks were observed across multiple brain regions and recordings, suggesting that they are generalized networks and not exclusive to the exact set of brain regions we recorded (Fig. S13).

As with the example subspace networks from one animal (Fig. 4B), the clustered subspace networks showed a diversity of structures (Fig. 5). This includes canonical networks observed in previous studies[32–34], such as a somatomotor network (cluster 9 in Fig. 5., also highlighted in green in Fig. 4B), lateral cortical network (clusters 6), and visual network (cluster 2), suggesting that these networks are co-activated in support of specific cognitive computations such as somatosensation, motor control, vision, and executive function[32,35,36]. Other clustered networks were novel and distributed. For example, cluster 8 (also highlighted in blue in Fig. 4B) shows a network spanning motor, somatosensory, and retrosplenial regions and cluster 11 shows a network inclusive of all regions except parietal cortex. These subspace networks were, in general, bilaterally symmetric[9,30] and included multiple anatomical regions. Collectively these canonical subspace networks covered the entire cortex and can be conceived as providing a basis for supporting interactions across cortical regions. We discuss this next.

Subspace networks are multiplexed across the cortex

Across clusters, subspace networks were distinct from one another. For example, in frontal motor, the subspace networks for dimensions 4 and 8 were significantly non-overlapping with dimensions 6 and 9 (Fig. 3B; average overlap between pairs = 2.3%, all $p < 0.001$, below chance, one-sided binomial test). Overall, any two subspace networks overlapped on an average of 28.9% of their significant pixels. However, there was a wide distribution of overlap across different pairs of subspace networks (STD = 28.9%) and a considerable fraction (37.9%) of pairs overlapped in less than 10% of their pixels, suggesting they were distinct, independent, subspace networks (Fig. 4C). The diversity of subspace networks is also reflected in the diversity of clusters, many of which were restricted to a few regions and had relatively little overlap with other subspace networks (Fig. 5).

Each individual subspace network tended to involve a few distinct regions (Figs. 4B and 5), but the set of all subspace networks covered the cortical surface (Fig. 6A). On average, each cortical area was a significant component of approximately three subspace networks from each target region. To visualize this, Fig. 6B shows the 'barcode' of subspace networks that involves each cortical area for a given target region. For example, in frontal motor cortex, dimension 3 is shared broadly, including both anterior and posterior regions (blue triangles, Fig. 6B, top left), while dimension 4 is restricted to posterior regions (green triangles) and dimensions 2, 5, and 6 are restricted to anterior regions (red, purple, and orange triangles). In this way, different dimensions of the neural response within frontal motor cortex can engage different networks of cortical areas. In other words, some parts of the neural response in frontal motor cortex will be functionally connected with both anterior and posterior cortex (e.g., dimension 3) while other parts of the neural response will be distinct to each region (e.g., dimensions 2, 4, 5, and 6). Similar patterns were observed in other regions (Fig. 6B). Altogether, these results suggest subspace networks may allow communication between regions to be multiplexed, as different dimensions of neural activity within a brain region are functionally connected with distinct, yet overlapping, networks of brain regions. The fact that subspace dimensions are orthogonal suggests that information can be independently represented and shared with each subspace network.

Interactions between brain areas are associated with alignment of local activity to subspace networks

So far, our results show different dimensions of neural activity within a region are functionally connected with different networks of regions. As outlined above, this could provide the basis for controlling interactions between brain regions. Specifically, the hypothesis is that aligning neural responses with a specific subspace dimension predicts how that neural activity propagates between cortical regions. To test this, we related the geometry of neural responses within a region to the spatiotemporal dynamics of neural activity across the cortex.

As expected[6,7,9], cortex-wide imaging found neural activity was dynamic: at different moments in time, different networks of brain regions were engaged (Fig. S1A-E). Using a convolutional factorization approach[37], we identified 14 'motifs' of cortex-wide neural dynamics[7]. Each motif reflected a unique, ~1 second-long spatiotemporal pattern of neural activity across brain regions (likely reflecting different

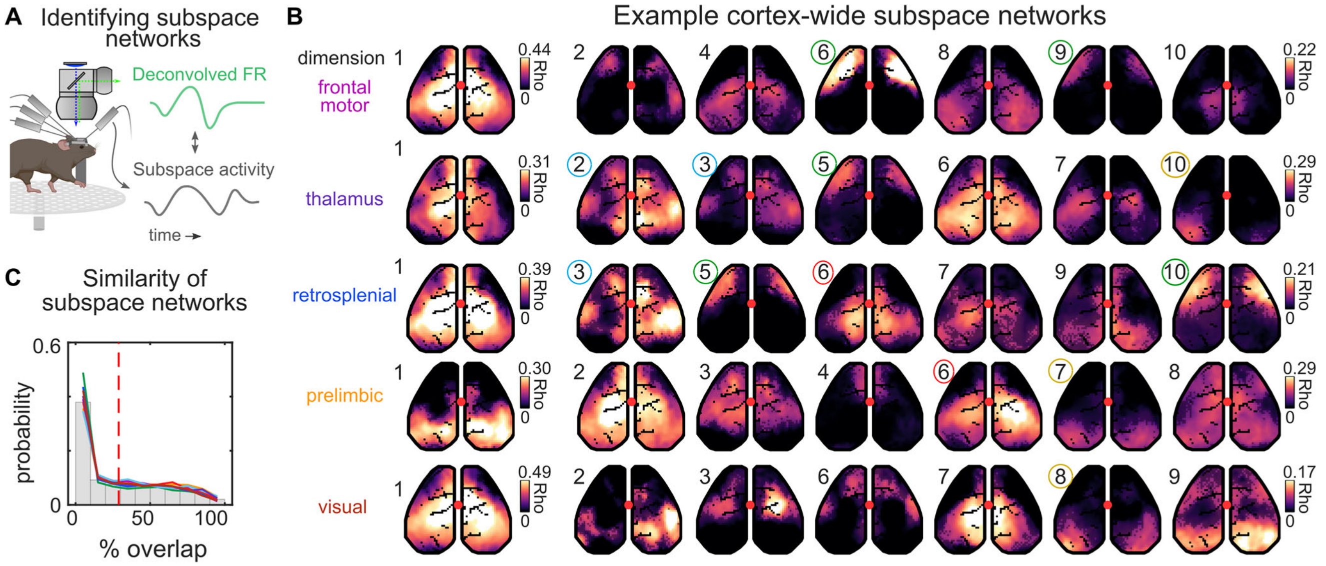

Fig. 4 | Subspaces engage independent but overlapping cortical networks. **A** Electrophysiology was combined with widefield calcium imaging of population-level neural activity across cortex. Variance in neural activity projected along each subspace dimension was correlated with imaging signal to identify the subspace-associated cortical networks. To better match the time constants of imaging and electrophysiology, we used a feedforward neural network to estimate the neural activity underlying the imaged calcium signal (see Methods). Created in BioRender. Tafazoli, S. (2025) https://BioRender.com/y38t274. **B** Representative cortical maps for different subspace dimensions. Rows are different regions. Numbers indicate the subspace dimension corresponding to each map. Color intensity of map indicates the strength of correlation between fluorescence activity at that pixel and the spiking activity along a given subspace (see Methods). Transparency of the color in each image is thresholded at significance such that non-significant pixels are more transparent. For visualization, color is scaled independently for first network versus subsequent networks. Colored circles around numbers indicate cluster identity of subspace networks as determined through clustering analyses (Figs. 5 and S13). Note, presented data is from an example dataset and ordering may differ between recordings. **C** Similarity of cortical networks. Similarity (x-axis) was computed as the percentage of overlap in significant pixels between cortical maps across all subspace dimensions for a region (see Methods). Y-axis shows the fraction of all compared subspaces. Gray bars show all data (across regions, $n = 630$ dataset). Colored lines show mean distribution of each brain region ($n = 8$) across datasets. See also Figs. S8-S13.

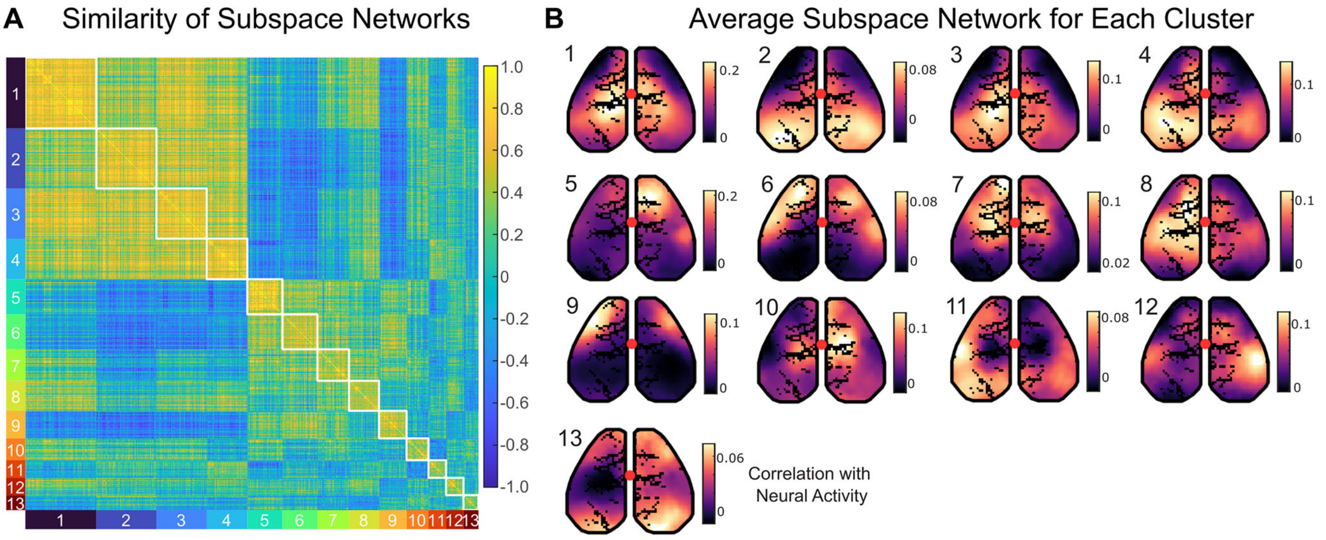

Fig. 5 | Canonical cortical subspace networks. **A** Similarity matrix for all subspace networks, sorted by cluster identity shown in **B**. Cluster groupings are shown along the x and y axes. Colorbar shows correlation of subspace network across the entire spatial map. White boxes indicate the within-cluster pairs of subspace networks. **B** Median subspace network for each cluster. Color axis shows strength of correlation. Only pixels that were valid across all recordings were used (i.e., the map of vasculature is combined across all recordings). Clusters were labeled according to the number of subspace networks in each cluster (descending). See also Fig. S13.

cognitive processes[7,38]). These motifs occurred multiple times per session and reliably generalized across sessions, animals, and environments[7,10,38]. Altogether, the 14 motifs captured 65.8% of the variance in cortex-wide neural activity in withheld data (CI: 65.0-66.6%, see Fig. S1 and Methods). Importantly, which motif was expressed changed on a moment-by-moment basis[7,10], reflecting the flexibility of interactions between brain regions during spontaneous behaviors.

Although motifs were measured using widefield imaging data, they reflected the activity of individual neurons (Fig. S1F). For example, during motifs A and B, neural activity was increased above baseline in whisker, somatosensory, and retrosplenial cortex (all $p < 0.001$, paired signed-rank test; Fig. 7A, insets). However, the magnitude of neural response differed between the two motifs: somatosensory and whisker cortices were more active during motif A, while retrosplenial cortex

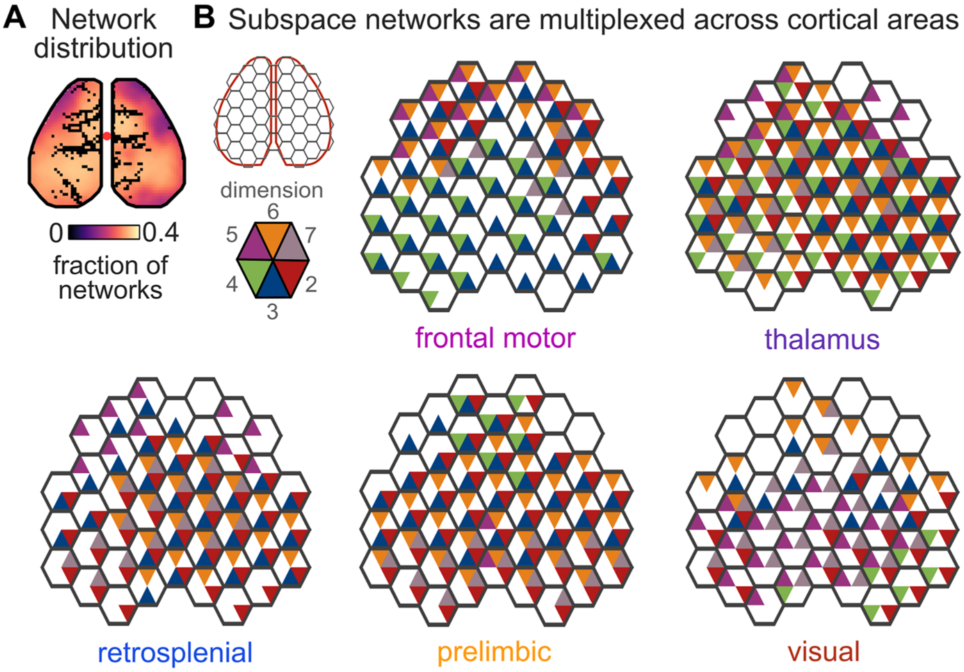

Fig. 6 | Cortex-wide subspace networks are multiplexed. **A** Subspace networks uniformly involved regions across cortex. Map shows the average fraction of subspace networks that were significantly correlated with each pixel (out of 10 possible maps per subspace, averaged across datasets from all target regions, $n = 630$). **B** Honeycomb plots showing multiplexing of subspace networks across cortical areas. Each hexagon shows the dimensions of a subspace that significantly engaged over 25% of pixels within that hexagonal parcel of the cortical surface. For visualization, only a subset of dimensions (2-7) is shown. As in Fig. 4B, data is from an example dataset and so numbering and exact maps may vary across recordings. See also Figs. S11-S13.

was more active during motif B (both $p < 0.05$, paired signed-rank test). This suggests the two motifs involved different responses across the three regions. Indeed, as seen in Fig. 7B, the co-evolution of neural activity across all three regions differed between motifs. Motif A mainly involved neural activity in whisker and somatosensory cortex (Fig. 7B, red), while motif B largely involved activity in whisker and retrosplenial cortex (Fig. 7B, blue). Reflecting the different dynamics across regions, the average neural activity during motif A and B were restricted to two independent, nearly-orthogonal, planes within the three-region neural space (Fig. 7B, inset; see Methods).

Here, we wanted to test whether the different neural responses observed during motif A and B reflected the engagement of different subspace networks. In particular, if a network was engaged when the neural response within a region was aligned with the associated subspace dimension (Fig. 7C; i.e., the neural response had a high magnitude projection onto the subspace dimension). For example, since somatosensory (SS) and whisker (WHS) cortex are both more active during motif A, then we expect the response in somatosensory cortex should be more aligned to (i.e., have a greater projection along) the subspace that functionally connects somatosensory and whisker cortex (i.e., the SS-WHS subspace) during Motif A. In contrast, because somatosensory and whisker cortex are less active during motif B, we expect the response in somatosensory cortex to be less aligned to (i.e., project less along) the SS-WHS subspace during motif B.

To measure alignment of the neural response and the subspace dimension, we calculated the angle between the neural response and subspace dimension. When the response is aligned to the subspace, then the angle will be lower (i.e., more parallel) while, when not aligned, the angle will be greater (i.e., closer to orthogonal at 90°). Consistent with the idea that the neural response is more aligned with a subspace dimension when the associated network is engaged, the angle between the neural response in somatosensory cortex and the SS-WHS subspace dimension was 73.2° during motif A but 85.8° during motif B (Fig. 8A, left, $p = 0.03$, bootstrap test, see Methods). A similar pattern was seen in whisker cortex: the neural response was more aligned with the SS-WHS subspace during Motif A than Motif B (Fig. 8A, second left).

Reflecting the dynamic nature of this alignment, a different pattern of results was seen during Motif B. Neural responses in retrosplenial (RSP) and whisker (WHS) cortex were more aligned with the RSP-WHS subspace during motif B than motif A (Fig. 8A, two right plots). This highlights how aligning the neural response in whisker cortex to either the SS-WHS or RSP-WHS subspaces is associated with changes in the moment-to-moment interactions between those regions, as captured by motifs A and B, respectively.

This phenomenon was highly consistent. Similar results were observed across all recordings of the trio whisker, somatosensory, and retrosplenial cortex during Motifs A and B (Fig. 8B; $p < 0.001$, bootstrap test). Furthermore, a similar pattern of results was seen across six different sets of triads of regions and pairs of motifs (Figs. 8C and S14; $p < 0.001$, bootstrap test; region triads collectively included all 8 recorded regions, see Methods). Control analyses confirmed this was not due to changes in the subspace networks or how the inter-regional subspace was estimated (see Methods). Finally, alignment was observed regardless of whether subspaces were defined using neural activity during either motif (see Methods).

While the effect was very consistent, we did find that the change in angle was often relatively small (a few degrees). To understand the effect of these changes, we used a computational model to show how a small change in the angle between a neural response and the shared subspace (~5-10 degrees) can lead to large changes in decodability (Fig. S15). Consistent with the model, we found that during motif B, the neural activity in somatosensory cortex has a small projection onto the SS-WHS subspace (Fig. 9A, blue, lines along axes show the extent of the response within somatosensory cortex and projected along the SS-WHS subspace). In other words, the neural response is expected to have a minor effect on activity in whisker cortex. In contrast, during motif A, the neural response in somatosensory cortex had a larger

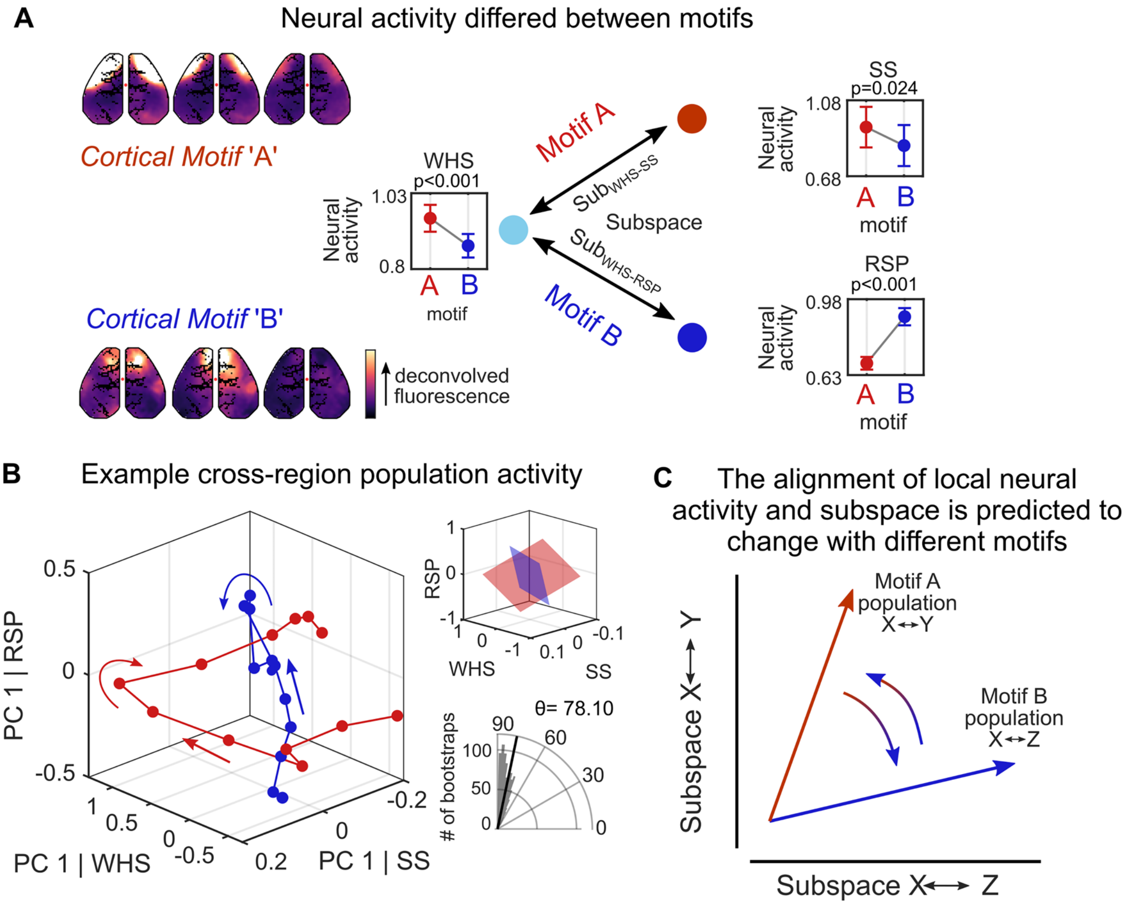

Fig. 7 | Neural responses are aligned to the shared subspace between two regions when those regions are engaged. A Left plots show pattern of cortical fluorescent activity during motif A (top) and motif B (bottom). Right plots show the neural response in whisker (WHS), somatosensory (SS), and retrosplenial (RSP) cortex (from electrophysiological recordings) during each motif. Arrows reflect bidirectional interactions between brain regions since reduced rank regression cannot specify directionality. Dot and error bars indicate mean and standard error of baseline-normalized spiking activity across neurons during each motif for an example recording (SS, $n = 29$; WHS, $n = 230$; RSP, $n = 190$ neurons, see Methods; p-values estimated with two-sided paired signed-rank test). **B** Projection of neural activity in three-region neural space during motif A (red) and B (blue). Average neural activity in each region, during each motif, was projected along the first principal component (PC; see Methods). Arrows show evolution of neural activity over time following onset of each motif (motif A, $n = 336$ occurrences; B, $n = 266$; 133 ms timesteps). Top inset shows planes fit to this PC space projection. Bottom inset shows bootstrapped distribution of angle between projection of activity ($n = 1000$ bootstraps across motif occurrences). **C** Schematic of proposed mechanism wherein aligning a neural response to a subspace dimension propagates activity along the associated subspace network. Arrows reflect the predicted changes in the alignment during different motifs. See also Fig. S14.

projection onto the SS-WHS subspace (Fig. 9A, red). In this way, a small change in the angle between the neural representation and a subspace dimension (Fig. 9B upper) can lead to a large change in the projected response (Fig. 9B, lower).

Discussion

Different dimensions of neural activity are functionally connected with different networks of brain regions

Using a combination of large-scale, multi-region electrophysiology and widefield calcium imaging, we studied how neural populations in multiple brain regions dynamically interact. Consistent with previous research[23,24], we found cortical and subcortical regions were functionally connected through a subspace of their neural activity (Fig. 1). Interestingly, many subspace dimensions were connected, not only with one other region, but rather a network of brain regions (Figs. 2–5). In particular, the subspace dimensions that explained the most variance in neural activity (e.g., subspace dimensions 1 and 2) tended to integrate neural activity across networks that were broadly spatially distributed.

This observation is consistent with previous work showing broad, cortex-wide, networks of functionally connected brain regions in humans[39], monkeys[40,41], and mice[33]. However, conceptually, these networks have been conceived (and modeled) as activating a set of pairwise connections between regions (Fig. 2A, upper). Our results suggest that distributed networks are inherent in the functional connectivity—each dimension of neural activity is shared with a network of brain regions (Fig. 2B, lower). These results argue against a simple directed graph of network connectivity in the brain. Instead, each brain region in the network has multiple outgoing connections that each project to a broader network of brain regions (schematized in Fig. 10A).

A relatively small number of subspace networks were consistently observed across dimensions, regions, and animals (Fig. 5). This is consistent with previous work showing a canonical set of functionally connected networks of brain regions that are consistent across individuals and tasks[33,36]. However, our results suggest these networks co-exist within a region as different dimensions within the neural space, each of which is functionally connected with a different network. The low number of subspace networks is also consistent with previous work showing a relatively low dimensional number of cortical states[7,28], possibly facilitating learning of cognitive control[38].

Communication between brain regions is multiplexed

Our results suggest communication between brain regions may be multiplexed, with different parts of neural activity within a region

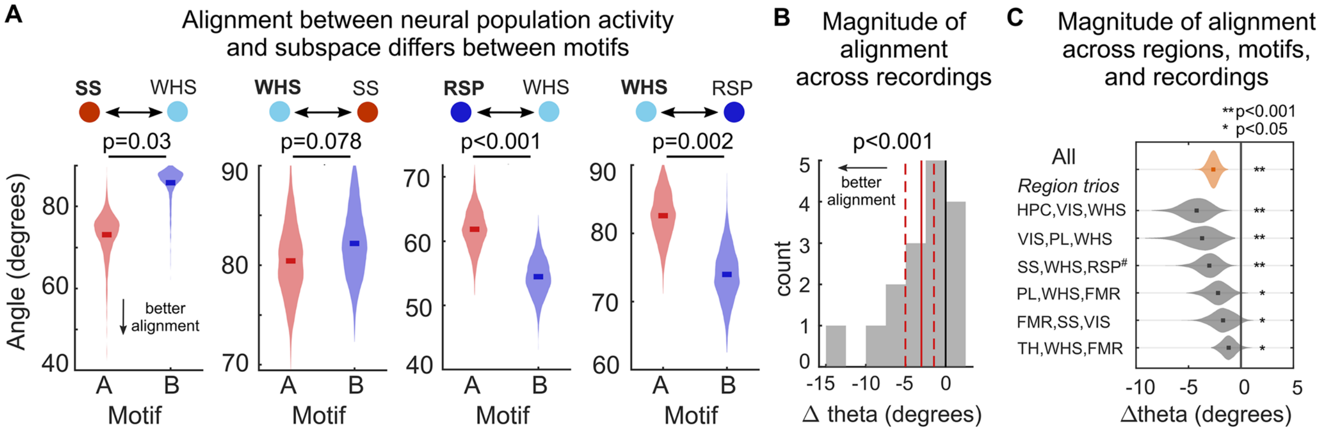

Fig. 8 | Neural responses were aligned to shared subspaces across regions, motifs, and recordings. A Bootstrapped distribution of the average alignment between the neural response in SS, WHS, or RSP (denoted in bold) during motif A or B and the subspace between that region and another region (indicated by bidirectional arrows along top). Lower angles (y-axis) reflect stronger alignment. $N = 1000$ paired bootstraps. Panels reflect data from the same example recording as Fig. 7. **B** Distribution of average difference in alignment angle between motif A and B across recordings for SS, WHS, and RSP ($n = 16$ comparisons). Difference in angle was computed relative to the motif with greater engagement of a target brain region, i.e., motif A minus B for comparisons involving SS, but motif B minus A for comparisons involving RSP. Thus, negative values indicate that the motif with greater engagement exhibited better alignment between local neural activity and subspace activity. Vertical red lines show mean (solid) and 95% CI of mean (dashed). **C** Bootstrapped distribution of difference in alignment from multiple pairs of motifs and trios of brain regions across recordings (see Fig. S14). Orange violin shows the overall alignment across all comparisons (120 alignment comparisons, see Methods). Overall, the local neural representation and the associated subspace dimension were more aligned when the stronger motif was engaged ($p < 0.001$). Gray violins show alignment distribution for example pairs of motifs/trios of regions ($n = 32–48$ comparisons each, see Methods). Individually, neural activity during all compared motifs/regions exhibited significant alignment to the associated subspace ($p < 0.001$; $p = 0.009$; $p < 0.001$; $p = 0.01$; $p = 0.035$; $p = 0.048$, for plots top-to-bottom, respectively). # indicates the example motif pair/region trio used in **A**, **B** and Fig. 7. For all panels p-values estimated with one-sided bootstrap test, $n = 1000$ bootstraps. See also Fig. S14.

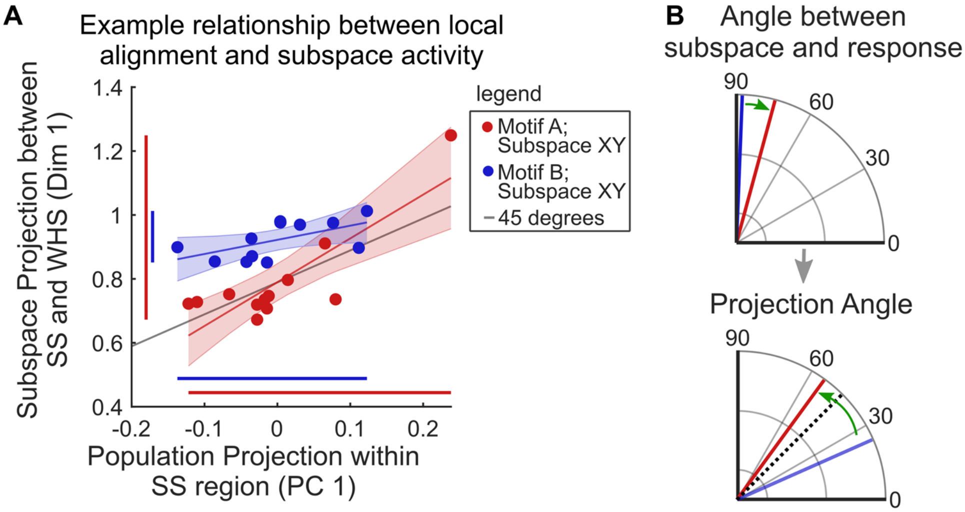

Fig. 9 | Decreasing the angle between the neural response and subspace increases projection of neural response onto subspace. A Projection of neural activity during Motif A and Motif B along both the first PC of activity in SS (x-axis) and the first SS-WHS subspace dimension (y-axis; as in Fig. 7). Greater alignment between the local neural response and subspace during Motif A was associated with a larger neural response along subspace. The domain and range of neural response during both motifs are shown by red and blue lines along axes. Points show timepoints of the neural activity during each motif (as in 7B). Line and shaded region show least squares fit and 95% confidence bounds, respectively. **B** Modest rotations in the neural representation can lead to large changes in the projection of the neural response along the subspace. Upper plot shows a relatively modest change in the alignment between the local representation of activity in SS (taken as the first PC) and the SS-WHS subspace dimension for Motif A (red) and Motif B (blue). Lower plot shows this translates into a large change in the response when projecting onto the SS-WHS subspace. The efficacy of the neural response is measured as the angle between the projection along the SS-WHS subspace and the first PC of the local response within SS. From example dataset shown in Fig. 7. Dotted line in **B** (lower) indicates 45 degrees (i.e., perfect alignment). See Fig. S15.

projecting to different networks of brain regions. Multiple subspace networks existed within each region; some engaged broad, cortex-wide, networks while others engaged more specific networks of regions (Figs. 4–6). These networks were often overlapping; each cortical area was involved in approximately one-third of the subspace networks for any given brain region (Fig. 6).

These results suggest each brain region can be functionally connected with multiple networks at the same time. For example, a part of the neural representation in frontal motor cortex might be functionally connected with a broad network spanning both anterior and posterior cortex (e.g., the neural response along dimension 1 in Fig. 6) while other parts of the neural representation are functionally connected

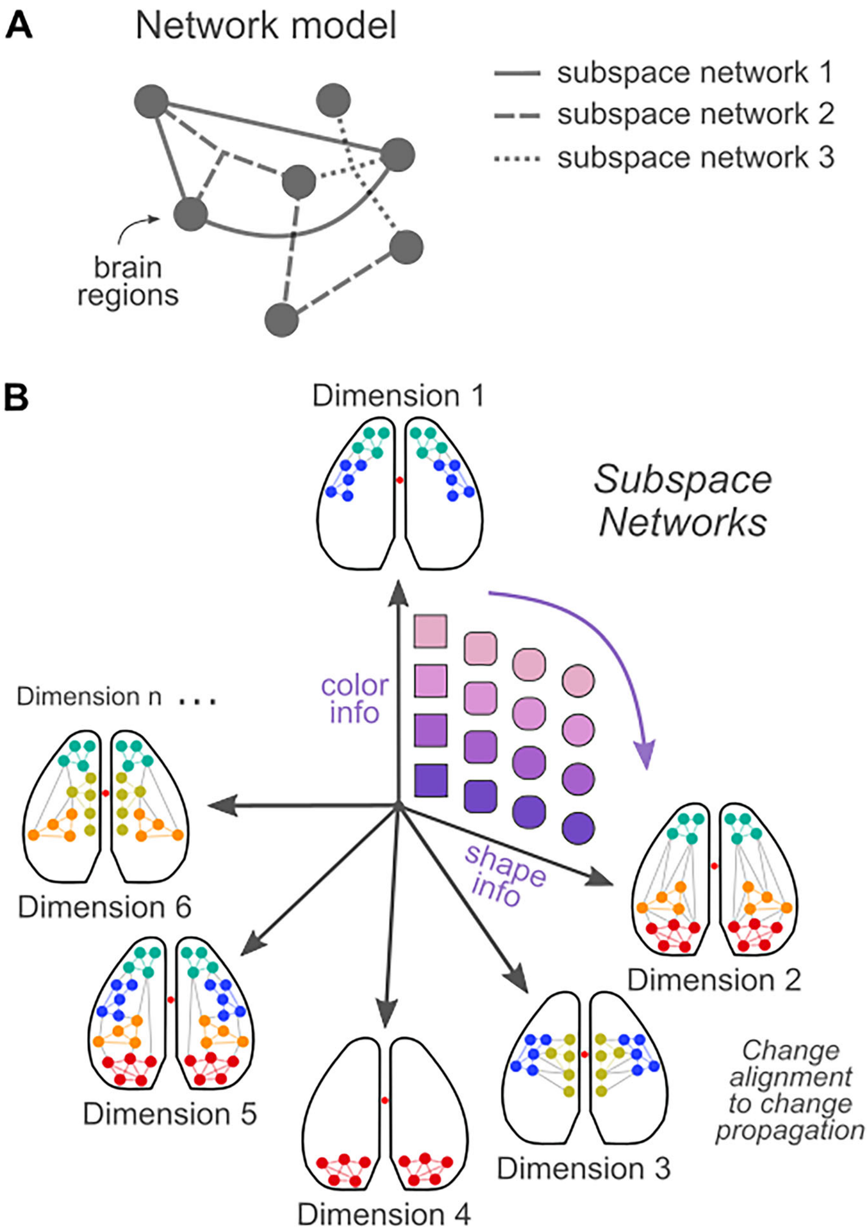

Fig. 10 | Schematic modeling how changes in local representations can propagate activity to multiplexed subspace networks. A Individual brain regions are connected to a network of independent but overlapping regions (i.e., multiplexed networks). Different style lines show the network of regions that are functionally connected through a shared subspace network. **B** Different aspects of the neural representation within a brain region can engage different subspace networks when aligned with different subspace dimensions. This is schematized as purple-to-pink and square-to-circle gradients being differentially aligned with different subspace dimensions, with gray arrows indicating the alignment of neural activity within a region. For example, aligning a region's neural activity with subspace dimension 1 propagates 'color' information to the associated network. Conversely, alignment to subspace dimension 2 propagates 'shape' information to a different network.

with only part of the cortex (e.g., dimensions 2, 6, and 9 were shared with anterior regions alone and dimension 4 was shared with posterior regions alone in Fig. 6). Importantly, the orthogonal nature of these shared subspaces means that this selective sharing of information can occur simultaneously. This supports a model of brain dynamics with parallel, distributed, processing such that different aspects of the representation in a region can be shared simultaneously with different networks of regions (schematized in Fig. 10B).

A geometric model of cognitive control

Cognitive control is thought to control how information flows between task-relevant brain regions, allowing the brain to select task-appropriate cognitive computations. Several mechanisms have been proposed for controlling how neural activity propagates between brain regions. Increasing the gain of a response may ensure a representation is successfully propagated by 'winning' the competition with other representations in downstream brain regions[14]. Alternatively, synchronizing neural activity within or between brain regions may facilitate the flow of information[16,42]. Our results support a geometric mechanism for cognitive control: when neural responses were aligned with (i.e., projected along) different subspace dimensions, it predicted how neural activity propagated between brain regions (Figs. 7–9).

Previous work has shown neural representations within a region can be dynamically transformed from one subspace to another[43,44]. These transformations depend on the task, suggesting they are under cognitive control[19]. In this way, cognitive control can transform neural representations within a region such that they become aligned to different dimensions of the shared subspace and, thus, engage different subspace networks and flexibly change how the encoded information is routed to other neural circuits (Fig. 7C). This may allow cognitive control to determine what, how, and when information is computed on. One advantage of this model is that it only requires a local change in the neural activity (avoiding the need to coordinate across regions). This may allow for flexible, dynamic changes in the alignment of neural responses and subspace dimensions. Indeed, we found alignment changed for different motifs, suggesting it can change on the timescale of hundreds of milliseconds, possibly supporting the flexible routing of neural activity across different brain networks.

Previous work has found large changes in neural representations, such that they moved into and out of orthogonal subspaces[19,45–47]. Orthogonality may be important for selective routing of information to different circuits – representations that are orthogonal to a subspace dimension are 'invisible' to the associated subspace network. However, we also found small changes in the angle of a neural representation could lead to considerable changes in how it was projected onto an interregional subspace (Figs. 9 and S15). This may allow for graded modulation of information flow between regions. In addition, oblique representations may allow a representation to be aligned with multiple dimensions simultaneously and, therefore, propagate to multiple subspace networks.

Indeed, the multiplexed nature of subspace networks may allow for nuanced control of how information represented in one region is propagated to downstream regions. By leveraging subspace networks, brain regions can represent multiple pieces of information simultaneously, encoding each in a different dimension within the local neural space[19,44,48]. When these dimensions are aligned with different dimensions of the shared subspace, then the represented information could propagate to different subspace networks (schematized in Fig. 10B). For example, a coarse action signal represented in frontal motor cortex could be broadcast widely by aligning it to one subspace network while, at the same time, a detailed, efferent copy, of the motor representation could be shared exclusively with somatosensory regions by aligning it to another subspace network.

Limitations and future directions

Our approach has several limitations that motivate future research. While we leveraged both widefield imaging and multiregional electrophysiology to record neural activity across the brain, our recordings remain limited to a subset of brain regions and a subset of neurons within each region. This likely constrained estimates of dimensionality within and across regions. For example, what appeared to be a 'local' dimension within a region (Fig. 1) may, in actuality, be a dimension of activity that is functionally connected to unrecorded regions. It is also possible that subspace dimensions capturing what seem like exclusive interactions between brain regions may include connections to unrecorded regions. Similarly, our widefield imaging approach captured the activity of the excitatory neurons thought to form the majority of connections between regions (e.g., Thy1-GCaMP mice). However recent work has shown that cortex-wide dynamics vary across different neural populations[49,50]. Future research is needed to understand inter-areal interactions across different neuron subtypes.

Brain state and behavioral state change the dynamics of interactions between neural populations[19,27,51–53]. For example, we have previously found that the relative frequency with which different cortical motifs are expressed varies with different behavioral contexts[28] and between individuals[51]. Our findings were robust across datasets, however, technical limitations constrained the results in this work to six total recording sessions from three individual mice. Future work is needed to test how different task environments and brain states affect the functional connectivity of brain regions.

It is also important to note that directionality is difficult to ascertain with our approach. So, we do not know whether correlations reflected directed flow of information between regions or bidirectional engagement. Traditionally, directed flow has been thought to be important for serial processing, but bi-directional recurrent connections may allow multiple regions to engage in a shared computation. Future advances in our capacity for recording large-scale neural activity[54] and improved analytical approaches for dissecting multiregional interactions[21,55] will begin to address these limitations.

Subspace networks were fairly low-dimensional: 10 subspace dimensions explained ~90% of the variance within a brain region (Fig. 1). This may reflect the nature of our recording approaches. Yet, we also note that previous work has found neural representations in the brain are often low-dimensional[7,56–58]. Cognition is similarly limited - attention can only be allocated to a few items at once[59], only a few items can be held in working memory at once[60], and decisions can only be made between a few items at once[61,62]. Our results may provide an explanation for such limits on cognition: they reflected the limited dimensionality of interactions between brain regions (Fig. 1). Future work is needed to test this hypothesis.

It is also important to note that electrophysiological recordings were restricted to a part of each region. For topographically organized regions, such as somatosensory, visual, or motor cortex, recordings were likely restricted to a region of stimulus/motor space. This may further explain the low-dimensional nature of the subspace networks. Future work is needed to understand how the subspace networks within a part of a brain region generalize to the rest of the region. For example, whether engagement of a subspace network is correlated or independent between different parts of a brain area.

Finally, future work is needed to understand the mechanisms by which neural populations align their activity to different subspace dimensions. Previous work has begun to identify potential mechanisms[44], but uncovering these mechanisms will provide a more complete understanding of how neural activity flows across neural populations in support of cognition.

Methods

Experimental model and subject details

All experiments and procedures were approved by the Animal Care and Use Committee (IACUC) of Princeton University and were carried out in accordance with the standards of the National Institutes of Health. Animals were kept on a 12 hr dark/light reverse light cycle throughout experiments.

Experiments were performed on three adult (> 8 weeks old) male ($N = 2$) and female ($N = 1$) mice expressing GCaMP6f in cortical excitatory neurons (Thy1-GCaMP6f line[63]). Each mouse was recorded twice for a total of $n = 6$ recordings. Previous studies in animals expressing GCaMP under the Thy1 promoter have not shown epileptiform activity[10,64]. Consistent with this, we screened all recordings for potential epileptiform activity and no events were observed.

Surgical procedures

Surgical procedures were performed in two stages. First, the skull was made optically transparent and a headplate was installed to allow for awake, in vivo widefield imaging. Second, 6-10 days later, four craniotomies were made to allow for acute electrophysiology.

Headplate implantation and preparation of dorsal cortex for widefield imaging followed previous work[7,10]. In brief, mice were anesthetized and given analgesics (Buprenorphine, 0.1 mg/kg; Meloxicam, 1 mg/kg) and sterile saline. The skin over the dorsal cranium was shaved, disinfected (betadine and 70% isopropanol), and resected. To make the skull optically accessible, the periosteum was removed, and a thin layer of clear dental acrylic was applied to the skull surface (C&B Metabond Quick Cement System). After drying, the acrylic was polished with a rubber rotary tool tip (Shofu, part #0321; Dremel, Series 7700) and coated with clear nail polish (Electron Microscopy Sciences, part #72180). To allow for head fixation, a titanium headplate with 11 mm trapezoidal window was cemented to the skull. Mice recovered, single-housed, in a clean home cage with post-op analgesia (Meloxicam, 1 mg/kg 24 hours post-surgery).

In the second surgery, mice were anesthetized and given analgesics (Buprenorphine, 0.1 mg/kg; Meloxicam, 1 mg/kg; Dexamethasone, 3 mg/kg) and sterile saline, and ~1 mm diameter craniotomies were drilled through the acrylic and bone of the dorsal cortex. No visible bleeding was observed during craniotomies. Exposed tissue was kept submerged in sterile saline until coated with Duragel (Dowsil, part #4680) and sealed with Kwiksil (World Precision Instruments). Mice recovered for multiple days prior to recording, during which they were single-housed in a clean home cage with post-op analgesia (Meloxicam, 1 mg/kg or Buprenorphine, 0.1 mg/kg; 48 hours post-surgery).

To determine craniotomy location, we used a combination of custom code and brainrender[65] to model electrode insertion paths. This allowed us to design insertion trajectories that targeted our desired brain areas while 1) maintaining insertion angles that would not obscure the imaging field of view and 2) fitting within the confined space of a single hemisphere. We then cross-referenced these potential trajectories with the location of large superficial blood vessels in each mouse—visualized by a brief imaging session prior to the second surgery—to find four craniotomy sites that were consistently situated above frontal motor (FMR), visual (VIS), somatosensory (SS), and retrosplenial (RSP) regions. These sites were centered approximately [−2mm, 1 mm], [3 mm, 2 mm], [2.6um, 0um], [1.2um, 0.9um] along the anterior-posterior, medial-lateral axis relative to bregma, respectively (positive values reflect posterior and lateral to bregma). Insertion angles for each electrode were approximately [32°, 54°, 42°, and 52°] degrees, respectively. This allowed us to target multiple regions with each electrode: 1) FMR and PL; 2) VIS, HPC, and TH; 3) SS and WHS; and 4) RSP.

Habituation and behavior

Animals were habituated to head fixation and running on a horizontal treadmill under the microscope between the first and second surgeries. Over four days the animals were acclimated for 30, 60, 120, and 180 minutes.

As we were interested in understanding how the brain flexibly engages different cognitive computations (and the associated networks of regions), the animals were free to perform a variety of different behaviors. Once habituated, animals spontaneously switched behaviors, including periods of running on the treadmill, grooming, whisking, and relative quiescence.

Widefield imaging

Widefield imaging was performed using an Optimos CMOS Camera (Photometrics) through back-to-back 50 mm objective lens (Leica, 0.63x and 1x magnification), separated by a 495 nm dichroic mirror (Semrock Inc, FF495-Di03-50×70). Excitation light (470 nm, 0.4 mW/mm2) was delivered through the objective lens from an LED (Luxeon, 470 nm Rebel LED, part #SP-03-B4) with a 470/22 clean-up bandpass filter (Semrock, FF01-470/22-25). Fluorescence was captured at 30 frames per second (FPS; 33.3 ms exposure) using Micro-Manager software (V1.4) at 980×540 resolution (~ 34um/pixel) for 90 minutes. Neuropixel recordings and imaging were aligned to the exposure out signal of the CMOS camera (captured on a NIDAQ PXIe-8381 in SpikeGLX; aligned with Tprime v1.6 available from https://billkarsh.github.io/SpikeGLX/#tprime).

Preliminary experiments revealed that the strobed excitation lights often used during widefield imaging for alternating GCaMP6f excitation (blue) and hemodynamic correction (violet or green) produced a large, saturating electrophysiological artifact. Potential mitigating measures, such as excitation intensity ramping[66] did not sufficiently remove this artifact. As our analyses relied on spiking data, we prioritized eliminating electrical artifacts. To this end, we opted to not strobe LEDs for excitation of GCaMP and hemodynamic correction. Instead, excitation was maintained throughout the recording session. To mitigate hemodynamic effects of large vasculature, we masked non-neural pixels around blood vessels[7,10]. Slow drift in the imaging signal was corrected by normalizing fluorescence to a rolling baseline variance (described below). These steps, combined with subsequent deconvolution (described below), minimized hemodynamic contributions, as evidenced by the fact that our observed strength of correlation between deconvolved fluorescence and ground truth spiking activity matched previous work that used hemodynamic correction[66]. After normalization, we spatially binned our imaging signal to 68×68 pixels (~ 135um/pixel) for subsequent analysis (as in previous work[7,51]). This resolution was chosen because preliminary analyses revealed that higher resolutions produced very similar results; adjacent pixels at finer resolutions were highly correlated and spatial binning to 68×68 pixels preserved the structure of cortex wide patterns in neural activity while increasing signal to noise and limiting computational constraints of the multiplicative update algorithm in CNMF[67] (Fig. S16).

Widefield recordings were registered within mice across days using rigid registration on user-drawn fiducials ($n > 10$ points) of notable vasculature and craniotomy edges. Recordings across mice were registered using rigid registration on user-labeled skull landmarks.

Estimating the neural signal underlying widefield imaging

One-photon widefield imaging is a useful measure of large-scale neural dynamics. However, the precise relationship between the recorded fluorescence signal and the spiking activity of underlying neural populations remains unclear[29,68]. Furthermore, the temporal relationship between fluorescence and spiking is lagged (due to indicator kinetics), which needed to be corrected in order to accurately map between spiking and cortical activity. Finally, differences between the signal recorded from craniotomies (covered by Duragel) and through bone can change the variability of the recorded signal[69]. Together, these effects could bias analyses. To avoid this, we performed an extensive set of supplemental analyses aimed at identifying the normalization procedure and deconvolution method that best estimated the neural population firing activity from a widefield imaging signal (Figs. S8–10).

As shown in Fig. S8, we found imaging signals were most comparable across regions when the fluorescence of each pixel was normalized by its rolling standard deviation (i.e., $\Delta F/\sigma_F$, where σ_F is the standard deviation during a 120 s rolling window). By definition, this corrected for differences in signal variance across the cranium as well as between exposed tissue (craniotomies) and intact skull. Variance normalization also improved subsequent estimations of firing rate in comparison to alternative approaches of normalizing by dividing by a baseline fluorescence value (i.e., $\Delta F/F_0$).

We then estimated relative firing rate by performing non-linear deconvolution on the normalized fluorescence signal. Specifically, on each recording, we trained a shallow feedforward neural network to predict spiking activity at time point t_0 using the one second of fluorescence signal on either side of that timepoint (i.e., t_{-30} to t_{+30}). This method resolved the temporal offset of the neural signal, better generalized to withheld data, and better recapitulated the log-normal statistics of neural spiking than alternative approaches[66,68,70] (Figs. S9 and S10). For this analysis alone, spiking data was binned to 30FPS to match the imaging data and normalized to standard deviation across the recording for deconvolution. As expected, widefield imaging predominantly reflected population-level activity of superficial layers of cortex (~ 200-600um; Fig. S10).

Identifying spatiotemporal motifs of cortex-wide neural activity

Our goal was to understand how neural activity flows between brain regions. This requires quantifying these dynamics, allowing us to categorize the cortex-wide flow of neural activity at each moment in time. This allows us to 1) estimate the functional connectivity (i.e., co-variations in activity) between regions (Figs. 1–6) and 2) understand how changes in neural representations influence the flow of neural activity between regions (Figs. 7–9).

To quantify the "flow" of neural activity across cortex, we used a convolutional factorization approach[7,37] to identify recurring patterns in the activity of cortical regions in our widefield calcium signal. Each of these'motifs' captured a unique cortex-wide spatiotemporal patterns of neural activity that lasts for ~1 s (examples are shown in Fig. S1F). Different motifs have been associated with different cortical processes. Some motifs capture bursts of activity in sensory cortex thought to reflect the processing of sensory stimuli. Other motifs capture bursts of activity in motor regions thought to reflect the preparation and execution of motor movements. Still other motifs capture traveling cortical waves thought to reflect cortex-wide integrative processes[38]. By identifying when each motif occurred, we were able to parse the stream of spontaneous activity into distinct 'trials', each with a different pattern of information flow. Importantly, our approach does this in an unbiased, data-driven manner.

Motif discovery followed previous work[10] and is schematized in Fig. S1A-E. We used convolutional non-negative factorization (CNMF)[37] with spatiotemporal penalty terms to identify unique 1-second recurring spatiotemporal sequences in our widefield signal. Specifically, CNMF approximates the full (pixel by time) image sequence X_{pt} as a sum of serial convolutions between a spatiotemporal tensor W and its temporal weighting matrix H:

$$X_{pt} \approx \tilde{X}_{pt} = \sum_{k=1}^{K} \sum_{\ell=0}^{L-1} W_{pk\ell} H_{\mathrm{k(t-\ell)}} \equiv (W \circledast H)_{pt} \tag{1}$$

where K and L are the maximum number of motifs and the maximum length of each motif, respectively and $\circledast$ indicates the convolutional operator. The maximum length (L) of motifs was set to 1-second. This duration was informed by previous work[7,28] and chosen to match the duration of dynamics in neural activity observed in our recordings while minimizing free parameters and computational demands. Supplementary analyses revealed that most neural activity within motifs occurred during a ~500 ms window (Fig. S17), well within the 1-second duration used. As in previous work[7], using motifs of longer length (i.e., 2-seconds) explained similar amount of variance in neural activity while dramatically increasing computational demands.

Motif discovery was performed on deconvolved imaging data binned to 15FPS and then split into one-minute chunks. Alternating chunks were used for motif discovery and withheld for cross-validation. The resulting n = 4311 cross-validated motifs were clustered across all animals using a graph-based nearest-neighbor clustering of the spatiotemporal correlation between all pairs of motifs (Phenograph[71]; with k = 12 neighbors used for construction of the graph). This process identified 14 clusters of motifs, each with a unique spatiotemporal pattern. A 15th motif was identified and excluded from subsequent analysis as it reflected an obvious artifact of hair entering the dorsal field of view in one animal on one recording. The average motif from each cluster were refit to the full length of all recordings, providing a measure of the activity of each motif at each moment in time. This refitting revealed that the 14 motifs captured the majority (66% CI: 65–67%) of variance in our cortical signal across all across all animals (consistent with our previous work[7,10]).

To determine when a motif was active, we applied a threshold to the motif's temporal weightings. Thresholding allowed us to capture the moments when cortical dynamics were well described by the motifs and avoid small transients in motif activity that may reflect noise. To determine the optimal threshold, we swept through 50 linearly spaced threshold values for each motif and identified the threshold at which the threshold-triggered pattern in the deconvolved mesoscale activity around the time of the threshold had the strongest spatiotemporal correlation with the original motif. After identifying the best threshold for each motif, we were able to determine the onset times when each motif occurred. On average, each motif occurred 299 times during a recording session (range 226–451). In other words, each motif occurred once every 11–24 seconds.

Electrophysiological recordings

Electrophysiological recordings used four Neuropixels[20] 1.0 probes (phase 3B2), inserted simultaneously (see above for trajectories). Electrodes were inserted under 2x magnification with micromanipulators (Siskiyou, part# MX-1131). Probes were coated with DiI (ThermoFisher Scientific, item #V22885) prior to insertion to allow for post-hoc histological reconstruction. Probes were lowered to desired depths (1.5-5 mm, pre-determined from modeling described above) and allowed to settle for >30 minutes before starting recording.

Recordings were 90 minutes in length. Data was acquired using SpikeGLX (v3.0 available from https://billkarsh.github.io/SpikeGLX/), with tip reference mode and high-pass filtered at 300 Hz. After collection, data was re-referenced by subtracting the global average across all channels using CatGT (v2), and spike-sorted using Kilosort[54] (v2.5). Automatically identified units were then manually curated with Phy[72] into well-isolated single units and 'multiunits' that may have aggregated the activity of multiple neurons.

Histology and reconstruction of electrode placement

Approximately 2 weeks after recording was completed, animals were transcardially perfused and their brains fixed (24-48hrs in 4%PFA) and cryopreserved (30% Sucrose). Brains were coronally sectioned with a cryostat at 60uM thickness. DiI electrode tracks were imaged with a NanoZoomer (Hamamatsu Photonics). Probe trajectories were reconstructed, following previous work[66] (http://github.com/petersaj/AP_histology; with minor adaptations). Trajectories were cross referenced with the A/P and M/L positioning of insertion sites derived from widefield images, to determine their final positioning. Correlograms of spiking activity showed strong concordance between estimated anatomical boundaries and functional boundaries in spiking activity (Fig. 1D). In one animal, probe insertion location had to be shifted for two electrodes to avoid large blood vessels: resulting in two recordings missing retrosplenial cortex, and one recording missing whisker somatosensory cortex.

Recorded neurons were grouped by anatomical location, as labeled in the Allen Brain Atlas Common Coordinates Framework[73] (CCF v3). *Prelimbic* (PL; $n = 1527$) included neurons from CCF parent regions Prelimbic (PL), Infralimbic (ILA), Dorsal Anterior Cingulate Area (ACAd). *Frontal Motor* (FMR; $n = 1257$) included neurons from Secondary Motor Area (Mos). *Visual* (VIS; $n = 716$) included neurons from Posteromedial (VISpm), Anterior (VISa), and Anteromedial (VISam) visual areas. *Somatosensory* (SS; $n = 833$) included neurons from nose (SSp-n) mouth (SSp-m) and unassigned (SSp-un) primary somatosensory areas. *Whisker* (WHS; $n = 805$) included neurons from Primary Somatosensory Barrel Field area (SSp-bfd). *Retrosplenial* (RSP; 640) included neurons from Dorsal and Lateral Agranular Retrosplenial areas (RSPd and RSPagl, respectively). *Hippocampus* (HPC; $n = 353$) included neurons from Dentate Gyrus (DG) and Ammon's horn (CA). *Thalamus* (TH, $n = 389$) included neurons from the Lateral Group (LAT), Medial Group (MED), and Intralaminar nuclei (ILM) of the dorsal Thalamus and Epithalamus (EPI).

Statistics and reproducibility

All analyses were performed in MATLAB (Mathworks Versions 2019-2023). Unless otherwise noted, we combined across $n = 14$ motifs and

$n = 6$ recordings (from three mice) when performing statistical analyses (this follows previous work[18]). Thus, throughout the main text and methods a *dataset* refers to the collection of neural activity during all instances of one type of motif (out of 14) during one recording (out of 6) in each of the 8 brain regions (6-8 regions per recording); for a total of 640 datasets. When applicable, tests were corrected for multiple comparisons. Permutation tests and bootstrap statistics were performed using 1000 shuffles/resamples. Unless otherwise noted, statistical tests were one-tailed tests of specific predictions.

Choice of representative examples

Our recordings spanned multiple brain regions. To limit bias in data presentation when illustrating general results (e.g., Figs. 1H, and 2B-D), all examples use data from visual cortex. In general, similar results were seen in all regions. The choice to use visual cortex was made prior to final analysis.

Identifying functional connectivity between neural populations

To quantify the functional connectivity between brain regions, we measured the degree to which neural activity in one region could predict the activity in another region. To avoid large evoked potentials in neural activity from obscuring underlying interactions between brain regions[74], we fit all models to the moment-to-moment variability in spontaneous spiking activity.

Specifically, we split spiking activity into discrete 'trials' based on the 14 recurring cortex-wide motifs observed in the widefield calcium signal (see above). We considered a 'trial' to be the 1333 ms (10 timebins) following the onset of a motif. To get the change in activity during a trial relative to baseline, we normalized spiking activity to the 400 ms baseline period prior to that trial: $FR_{\text{trial}} = FR_{\text{trial}}/(FR_{\text{baseline}}+1)$, as in previous work[5]. Equivalent results were found when normalizing by subtracting the baseline. Unless otherwise specified, spiking activity was binned with 133 ms time bins (7.5FPS) for all analyses. Binning with 66.7 ms bins (15FPS) gave comparable results.

After identifying the trials, we subtracted the mean activity during all trials of a single motif from both the imaging signal and spiking activity (Fig. S1). The residual activity captured the moment-to-moment variability in the neural population[18]. If two regions interact, then the moment-to-moment fluctuations in neural activity in one region should predict the fluctuations in the other region regardless of the broader patterns in brainwide activity[74]. Subtracting the trail-average response is consistent with previous work[18] and allowed us to quantify how two regions were interacting without the variability between motifs. Nevertheless, similar experimental results were observed when analyzing the full response (Fig. S18).

For analyses supporting Figs. 1–6, we identified interactions by fitting regression models to predict the spiking variability of each 'target' brain area from the spiking variability of all other 'source' brains regions. A total of 630 models were fit: one for each brain region ($n = 6$-8 regions per recording) during each motif ($n = 14$) and during each recording ($n = 6$). To fit each model, we concatenated neural activity across all instances of a motif, resulting in a two-dimensional matrix of neural activity per brain area (*[instances x timepoints] x neurons*). In analyses supporting Figs. 7 and 9, models were fit in a pairwise manner across brain areas, with only one region as the source and another region as the target. This resulted in a total of 4592 models.

All models were fit using reduced rank regression[18,75] (RRR) in MATLAB. Model fitting followed previous work[18]. RRR performs simultaneous regression and dimensionality reduction by identifying a m-dimensional set of predictive dimensions in a source neural population that best predicts the trial-to-trial variability in activity of a target neural population according to a linear model:

$$\hat{Y} = XBV^T \tag{2}$$

where $\hat{Y}$ is a (N_{trials} x $N_{\text{timepoints}}$) by $N_{\text{target neurons}}$ matrix of predicted activity in the target brain areas. X is a (N_{trials} x $N_{\text{timepoints}}$) by $N_{source\ \text{neurons}}$ matrix of activity in source brain areas (i.e., the independent variable). B is a $N_{source\ \text{neurons}}$ x m matrix and V is a $N_{\text{target neurons}}$ x m matrix that contain the m-dimensional subspace beta weights (β) for the source and target neurons, respectively (i.e., $\beta_{i,j}$ captures the contribution of neuron i to subspace dimension j). Unless otherwise noted, the first 10 subspace dimensions ($m = 10$) were used in all analyses, which captured an average of 90% of the total predictive performance of the fitted models (i.e., "all subspace dimensions" in the main text means the top 10 dimensions).

When fitting β, ridge (L2) regularization was used to avoid overfitting noise in the data. This was particularly important at higher values of m, allowing us to better estimate the contributions of higher dimensions. In addition, regularization was important for estimating the total explainable variance captured by the full dimensional model, as shown in Fig. 1F (i.e., where $m = N_{\text{target neurons}}$). The strength of L2 regularization was determined through 10-fold cross-validation[18]; the value for regularization was taken as the greatest value for which the mean cross-fold performance (r^2) was within a standard error of the mean of the peak performance across all regularization values.

Results of model performance (Fig. 1F, H, main text statistics) show cross validated performance of 10 repeats of 10% held-out data. Fitted models were stable across cross-validation folds, i.e., β were highly similar with an average rho between folds of 0.83 CI: 0.82-0.84, for the first 10 subspace dimensions). So, subsequent analyses used models fit to full datasets.

For a given model fit, the positive/negative direction of βs are arbitrary (since the final direction is the product of B and V). So, in order to facilitate the comparison of βs across areas and models (e.g., Figs. 2B-D, 3A-B, S5, and S7), or when projecting activity along β (e.g., Fig. 4B), we used the convention that β weights were oriented such that the majority of neurons had positive weights.

To confirm our results were robust to the analysis approach, we tested two other methods for quantifying interactions. First, as detailed below, we repeated our analyses but reversed the direction of fitting of the reduced-rank-regression models. Instead of predicting the activity of one region from all other regions, we used the activity in one region to predict activity in all other regions. Second, we repeated our initial analyses using Canonical Correlation Analysis[21,76] (CCA). CCA is a different, but related, technique for uncovering interregional relationships. CCA found highly similar results to RRR. For example, CCA also found interareal relationships existed in a 'subspace' of neural activity, with an average of 5 significant subspace dimensions (significance determined with permutation testing as in previous work[23]).

Estimating the dimensionality of subspaces

To understand the proportion of neural activity that was shared with other regions, we compared the dimensionality of the 'shared subspace' within each region to the 'local' dimensionality of population activity within that same region (Fig. 1I-J). Of note, our goal was not to estimate the 'true' dimensionality of neural activity, as this is inherently difficult to estimate, especially given that we only recorded a subset of neurons in each brain area. Instead, we wanted to compare the relative local and subspace dimensionality of neural activity and compare this ratio across regions.

Thus, to estimate dimensionality, we used cross-validated reduced rank regression, following a similar procedure described above. A region's neural population was randomly split into equal halves: e.g., into "source" and "target" populations *within* a region. To compute local dimensionality, we predicted the target population activity from this source population activity. Dimensionality was taken as the total number of dimensions in the source population that reliably predicted variance in the target population (e.g., predicted variability in withheld trials). This was repeated 100 times for each prediction to assure reliability in our random sampling of the population. The mean dimensionality from these 100 estimates was taken as the final value for that neural population.

Similarly, to estimate the shared subspace dimensionality, we again randomly split a region's neural population into equal halves. We then predicted the target population activity within a region using a random subset of neurons from all other brain regions. Subsets of neurons across other regions matched the firing rate distribution and number of neurons of the local source population. Again, this process was repeated 100 times for every prediction to assure adequate sampling across all brain regions and the number of dimensions was averaged across repetitions. The same 100 random splits of the local population into halves were used for computing both local and shared subspace dimensionality, allowing for direct comparison between the two estimates. To control for differences in the number of neurons recorded in each brain area, we subsampled to $n = 25$ neurons for target and source populations (i.e. in Fig. 1I-J).

Importantly, our finding that the 'shared subspace' within each region was significantly smaller than the local space of that region's population activity was robust to different estimates of dimensionality including, shared variance components analysis[8], factor analysis[18] and principal components analysis.

Estimating the strength of interactions between brain regions

The contribution of each neuron to predicting the neural activity of a target region is estimated by the magnitude of its β weight (for schematic, see Fig. 2B). By examining the βs of neurons in different regions, we can quantify the relative strength of interactions between each source region and the target region. For example, Fig. 2C shows the distribution of βs of neurons predicting the first dimension of activity of the visual region, sorted by decreasing magnitude. The fraction of neurons within a source region that contributed βs of each magnitude is shown in Fig. 2C, right (smoothed with a 50-neuron gaussian kernel). Figure 2D shows the cumulative distribution for each region, averaged over all 84 datasets. To quantify the contribution of a source region to predicting the activity of a target region, we looked for a greater concentration of high β weights in one region, relative to what would be expected by randomly distributing β weights across regions.

To do so, we calculated the area under the curve (AUC) of the source region's cumulative distribution function. A higher AUC indicates the source region contributed stronger βs, relative to other source regions (see inset of Figs. 2D and S5A for examples). In other words, a region with a higher AUC had stronger interactions with the target region (note that this approach does assume that the fundamental unit of communication between brain regions is the action potential). In contrast, to estimate the number of regions with a relatively stronger influence on the target region, we counted the number of source regions with an AUC greater than 0.5. In contrast, if neurons were randomly distributed across regions, then our analysis would lead to AUCs in all areas being equal (at 0.5; this is not what was observed). As detailed in the main text, and shown in Fig. S5B, earlier subspace dimensions tended to integrate activity from more regions.

To quantify which regions contributed the most to another region, we averaged the AUCs across all subspace dimensions, and all datasets. This allowed us to generate a graph of the strongest interactions between regions (Figs. 3 and S7).

Testing whether subspace dimensions are functionally connected with multiple brain regions

As detailed in the main text, our results suggest each subspace dimension was functionally connected with a network of regions. One concern might be that this is an artifact of fitting a single model that predicts the activity in one target region from all other source regions (note: regularization discourages this). Therefore, we tested whether similar results were seen when fitting a set of independent models that estimated the 'projection vectors' from one source region to each target region (Fig. S6A). If the same source representation was functionally connected to a network of target regions, then the projection vectors for different target regions should be correlated (despite the models being fit independently). Indeed, we found strong correlations in many source regions. For example, the vector of hippocampal activity influencing the first dimension of somatosensory cortex was strongly correlated with the vector influencing the first dimension of prelimbic cortex (Fig. S6B; $r = 0.65$, $p < 0.001$ bootstrap test versus zero correlation). Similarly, the projection vectors from frontal motor to the third dimension of the visual and somatosensory subspaces were significantly correlated (Fig. S6C; $r = 0.26$, $p < 0.001$ bootstrap test). Overall, 30.4% of vectors across all regions were significantly correlated, suggesting that many source regions shared their projections across target regions.

However, it is important to note that not all projection vectors were correlated (Fig. S6D). This reflects the fact that there were multiple dimensions within a source region that influenced different networks of regions. To quantify this, we measured the effective dimensionality of the projection vectors:

$$effective\ dimensionality = \frac{1}{\sum_i^N p_i^2} \tag{3}$$

where p_i is the percent of variance explained by each of the N principal components of the projection vectors (Fig. S6E). The null hypothesis is that the projection vectors are random and, thus, uncorrelated across regions. Therefore, to generate a null distribution, we randomly permuted the β weights across projection vectors and recalculated the effective dimensionality of each permutation (Fig. S6E). Consistent with partially-shared projections, the observed projection vectors had a dimensionality that was 45.4% of what was expected by random, uncorrelated, projections (Fig. S6E; CI: 44.6–46.2%, $p < 0.001$). Interestingly, although all regions were less than random, there was variability in the dimensionality of the projection space across regions (Fig. S6F). Frontal regions, including frontal motor and prelimbic regions, exhibited the greatest degree of sharing, spanning only 40% of the space (39.2% CI: 38.0–40.4% and 42.2% CI: 40.5–43.9%, respectively). This suggests representations in frontal regions are projected widely. In contrast, hippocampal and retrosplenial regions exhibited the least sharing, spanning 50% of the space (52.3% CI: 50.1–54.6% and 50.3% CI: 47.7–52.8%, respectively; both less than frontal motor and prelimbic at $p < 0.001$, bootstrap test). This suggests HPC and RSP, while having broad impacts on other regions, tend to have more independent projections to other regions.

Testing whether different subspace dimensions are functionally connected with shared or independent populations of neurons

Our results show there is a subspace of neural activity within each brain region that is functionally connected with other regions (Fig. 1). Different subspace dimensions within a target region are connected with different networks of source regions (Figs. 2–6). Within the target region, each dimension of the subspace was supported by a population of neurons. We were interested in understanding how the network of neurons is associated with each subspace dimension related to other subspace dimensions.

Our goal was to discriminate between two hypotheses. First, one might expect that independent sub-populations of neurons support different subspace dimensions (Fig. S3, upper). Consistent with this, anatomical tracing has found individual neurons send axonal collaterals to a constellation of brain regions, with different sub-populations of neurons projecting to different target regions[73,77]. Similarly, sub-populations of neurons have been found to be correlated with different cortex-wide networks of regions[78]. Different populations of neurons may even carry different types of information; for example, neurons in primary somatosensory cortex (S1) projecting to either secondary somatosensory cortex (S2) or primary motor cortex (M1) are biased to carry different types of somatosensory information[79].

Alternatively, the same population of neurons could support multiple subspace dimensions (Fig. S3A, lower; either within or between brain regions). Importantly, if the representations associated with each subspace dimension exist along orthogonal dimensions within the local population, then this would still allow different information to be projected to different subspace dimensions. Consistent with this hypothesis, previous work has shown neural representation of sensory stimuli[80], short-term memories[44], motor movements[6], and task representations[81] are distributed across the entire neural population.

To discriminate these hypotheses, we quantified the distribution of β weights associated with pairs of subspace dimensions (i.e., the vectors β_i^n and β_j^n for all neurons n in the recorded population of N neurons for subspace dimensions i and j). The two hypotheses make different predictions on how the β weights should be distributed across neurons. If each subspace dimension was influenced by an independent sub-population of neurons, then we would expect neurons to contribute to one subspace dimension but not the other. In other words, β weights should be large for dimension i and small for dimension j. In a two-dimensional plot of the absolute magnitude of β weights, these neurons would have β weights lie along the x-axis or the y-axis (see Fig. S3B, left, for schematic). Alternatively, if both subspace dimensions rely on the same population of neurons, then the β weights should be evenly distributed across the x-y plane (Fig. S3B, middle). Finally, if the same neurons contribute to both subspace dimensions, then the β weights should cluster around the diagonal (Fig. S3B, right). We used two statistics to quantify the distribution of β weights and, thus, discriminate these hypotheses.

First, we examined the angular distribution of β weights to each pair of subspace dimensions. The angle of the β weights was calculated for each neuron and for each pair of subspaces (taken as $\theta = \mathrm{angle}(|\beta_n^i| + \sqrt{-1}|\beta_j^n|)$, using the *angle* function in MATLAB). To ensure the β weight distribution was not influenced by the number of neurons or the average firing rate of a region, the β weights within each subspace dimension were normalized by their standard deviation. The distribution of angles was measured across all simultaneously recorded neurons (binned into steps of $\frac{\pi}{20}$). A distribution was estimated for each pair of subspace dimensions, during each motif, and for each recording session. Figure S3C shows the angular histogram for each region; averaged across all subspace dimensions, motifs, and recordings sessions (similar results were seen for individual motifs and recording sessions). Note, all pairs are considered, including mirror opposites (i.e., (i,j) and (j,i)) and so this distribution is guaranteed to be symmetrical. This does not affect the above logic.

As seen in Fig. S3C, all regions showed a relatively uniform distribution of β weight angles. To estimate the distribution that would be expected given a random distribution of β weights, we randomly shuffled the β weights for each subspace dimension and re-estimated the distribution (Fig. S3C, red lines). In all eight brain regions, the observed distribution of β weights was more clustered along the diagonal than expected by random chance. This was not due to the fact that some pairs of subspace dimensions had correlated β weights; the distribution of β weights for pairs of subspace dimensions that were not significantly correlated (i.e., $pval(\rho(\beta_i^n, \beta_j^n)) > 0.05$, where ρ was Pearson's linear correlation coefficient) were also uniformly distributed (Fig. S3D, example shown for HPC).

To quantify the distribution of angles, we estimated the curvature of the histogram for each pair of subspace regions. As schematized in Fig. S3E, the three models make three different predictions for the curvature of the histogram. If the β weights are from independent subpopulations, then there should be a concentration of β weights along the two axes, resulting in a histogram that is convex. If the β weights are randomly distributed across the population, then we expect a uniform histogram. Finally, if neurons involved in one subspace dimension also tend to be involved in the other subspace dimension, then we expect a concave angle distribution, reflecting the clustering of β weights along the diagonal. The curvature of the histogram was estimated for each pair of subspace dimensions by taking the second derivative of a quadratic equation fit to the histogram: a positive second derivative reflects a convex shape, while a negative second derivative reflects a concave shape (similar results were seen when using the second derivative of the raw data). To test whether the observed curvature was significantly different from chance, we compared the observed curvature to the distribution of curvatures across permuted β weight distributions. This comparison was done for each pair of subspace dimensions, allowing us to calculate a z-score of the observed curvature for each pair of subspace dimensions. Figure S3F shows the distribution of curvatures for each region, across all subspace dimensions. In support of the hypothesis that the same population of neurons influenced both subspace dimensions, the majority of curvatures were below zero for all eight regions (Fig. S3F). A significant percentage of pairs of subspace dimensions were more concave than expected by chance (FMR: 38.3%, PL: 39.0%, WHS: 30.3%, SS: 23.2%, RSP: 21.1%, VIS: 18.6%, HPC: 11.7%, TH: 8.3% of distributions were significantly more concave than chance, $p \le 0.01$ by permutation test; counts in all regions were greater than expected by chance, all $p < 10^{-26}$). In contrast, few pairs of subspace dimensions had significantly more convex distributions than expected by chance (FMR: 0.0%, PL: 0.0%, WHS: 0.0%, SS: 0.0%, RSP: 0.0%, VIS: 0.0%, HPC: 0.1%, TH: 0.1% of distributions were significantly more convex than chance, $p \le 0.01$ by permutation test; counts in all regions were less than expected by chance, all $p < 10^{-26}$). Similar results were seen when taking only those subspace dimensions with unsignificant correlations in β weights.

Second, we used a receiver operating characteristic (ROC) approach to quantify the distribution of beta-weights across neurons. The β weights were sorted according to their contribution to the first (i) subspace dimension: $\beta_i^{n=1} \ge \beta_i^{n=2} \ge \ldots \ge \beta_i^{n=N}$ and the ROC was taken as the proportion of the sorted β weights for the j dimension that were above a specific percentile:

$$\mathrm{ROC(p)} = \sum_{n=1}^{pN} \beta_j^n \ge P_p\left(\beta_j\right) \tag{4}$$

where $P_p(x)$ is the value of the p$^{\text{th}}$ percentile of distribution x. To quantify the ROC, we calculated the area under the ROC curve: $AUC = \sum_{p=0}^{1} ROC(p)$. In this way, the AUC quantifies the distribution of β weights. If independent populations of neurons interact with the two subspace dimensions, then the associated AUC will be ¼. In contrast, if the same representation projected to two subspace dimensions (i.e., the β weights are perfectly correlated) then the AUC will be ½. In general, a larger AUC indicates greater overlap in the population of neurons influencing the two subspace dimensions.

The average AUC, across all pairs of subspace dimensions, was ~0.35 in all eight regions. A majority of pairs within each region had higher AUCs than expected by chance (Fig. S3G), reflecting overlap in the populations of neurons interacting with both subspace dimensions. AUCs greater than chance reflect pairs of subspace dimensions with significantly more overlap in the neural populations than expected by chance ($p \le 0.01$, permutation test). This was the case for a significant number of pairs of subspace dimensions in all eight regions (FMR: 50.3%, PL: 48.6%, WHS: 41.3%, SS: 29.1%, RSP: 29.4%, VIS: 24.9%, HPC: 15.7%, TH: 10.6%; counts in all regions were greater than expected by chance, all $p < 10^{-26}$). In contrast, an AUC that is less than chance ($p \le 0.01$, permutation test) would indicate independent populations of neurons. This was the case for significantly fewer pairs of subspace dimensions than expected by chance (FMR: 0.0%, PL: 0.0%, WHS: 0.0%, SS: 0.0%, RSP: 0.0%, VIS: 0.0%, HPC: 0.0%, TH: 0.1% of distributions were significantly more convex than chance; counts in all regions were less than expected by chance, all $p < 10^{-26}$)

Interestingly, there was slightly more overlap in the populations of neurons that influenced a pair of subspace dimensions when the two dimensions were in the same target region as supposed to being in two different regions. This was reflected in a significant difference in both the curvature of the β weight histogram and the AUC. The β weight histograms were more concave when comparing two subspace dimensions within the same target region compared to between two different target regions (quantified using a contrast measure $\left(curve_{between} - curve_{within}\right)/\left(curve_{between} + curve_{within}\right)$; FMR: −0.085, PL: −0.091, WHS: −0.062, SS: −0.112, RSP: −0.084, VIS: −0.104, HPC: −0.111, and TH: −0.104). The distributions of curvature were significantly different in all eight regions (FMR: t(174753) = 43.64, PL: t(157848) = 36.02, WHS: t(155188) = 24.47, SS: t(177168) = 39.25, RSP: t(135238) = 28.70, VIS: t(177168) = 37.95, HPC: t(177168) = 32.57, and TH: t(177168) = 25.86; all $p < 10^{-146}$ by t-test). Similarly, the AUCs were significantly higher when subspace dimensions were within the same target region than between regions (contrast: FMR: −0.079, PL: −0.085, WHS: −0.056, SS: −0.098, RSP: −0.075, VIS: −0.094, HPC: −0.098, TH: −0.098; t-test: FMR: t(174753) = −46.68, PL: t(157848) = −36.48; WHS: t(155188) = −25.28, SS: t(177168) = −38.23, RSP: t(135238) = −29.77, VIS: t(177168) = −39.43, HPC: t(177168) = −33.92, TH: t(177168) = −28.66, all $p < 10^{-139}$ by t-test).

Altogether, these results support the hypothesis that the same population of neurons is functionally connected with multiple subspace dimensions in other brain regions. We found this was true even when the representations influencing two different subspace dimensions were uncorrelated. Our results are consistent with anatomical studies showing axons from a single neuron can branch, sending collaterals to several brain regions. Furthermore, it is important to note that measures of functional connectivity may include multi-synaptic influences, and so a shared pool of neurons may project to other brain regions by acting locally on different sub-populations of neurons. This would allow the same information to be routed to multiple regions simultaneously, even if the connections between brain regions were anatomically separate.

Identifying subspace networks in imaging data

To understand how the neural population within a region was functionally connected with regions across the cortex, we correlated the activity projected along each subspace dimension with the deconvolved widefield fluorescence signal. Activity along a subspace dimension was defined as the dot product of the activity of a region and the β of a subspace dimension (resulting in a $N_{\text{trials}} * N_{\text{timepoints}}$ vector). This vector was then correlated with the ($N_{\text{trials}} * N_{\text{timepoints}}$) *by* N_{pixels} matrix of cortical fluorescence activity, yielding a correlation value for each pixel.

To visualize significant cortical networks (Fig. 4B), the observed correlation maps were compared with maps generated from trial shuffled data. We generated an FDR-corrected null distribution of correlation values by taking the strongest correlation across all pixels for each of 1000 trial-permuted maps. Pixels with positive correlations greater than the 95% of this null distribution were taken as significant. The resulting map defined the "subspace networks".

The similarity of two subspace networks, shown in Fig. 4C, was computed as the percentage of pixels that were significant in both subspace network *A* and subspace network *B*:

$$\% \, overlap = \frac{2 * (A \cap B)}{A + B} \tag{5}$$

To validate our imaging results, we tested whether the order of the strongest-to-weakest interactions between brain areas observed in our electrical recordings was correlated with the order of the strength of correlation between brain areas in our imaging data. Specifically, we computed the spearman's correlation coefficient between the vector of AUCs (e.g., as in Fig. 2D) across cortical regions FMR, RSP, WHS, SS and VIS, and the strength of correlation between brain regions in our imaging data (e.g., in Fig. 4B). See main text for statistics.

Clustering subspace networks

As shown in Fig. 4, each dimension of neural activity within a brain region was correlated with the activity in a cortex-wide, distributed network of regions. Many of these 'subspace networks' were distinct from one another. For example, dimension 1, 2, 4, and 6 of area FMR were associated with different subspace networks (Fig. 4B). However, other dimensions were associated with similar subspace networks (e.g., FMR-6, FMR-9, TH-5, RSP-5, and RSP-10, Fig. 4B). To quantify the distribution of subspace networks, we clustered the subspace networks associated with the first 10 dimensions in each brain region, across all datasets. To minimize the impact of noise, only subspace networks that included significant correlations with more than 5% of the pixels were used. After exclusions, a total of 6,279 subspace networks were clustered. Clustering was done using the Phenograph algorithm[71]. Distance between data points was defined as the inverted correlation of the correlation vectors for area a_i, dimension d_i and area a_j, dimension d_j: $\text{dist} = 1 - corr\left(\rho_{area_i, dim_i}, \rho_{area_j, dim_j}\right)$, where ρ_{area_i, dim_i} is a vectorized form of the map of correlation between the electrophysiologically recorded neural activity in area a_i and dimension d_i and the cortex-wide neural activity as imaged with widefield calcium imaging (i.e., the non-thresholded subspace network). Similar results were seen with other metrics of the distance between subspace networks, including measuring the correlation between the z-transformed correlation maps or measuring the city-block distance between the cortex-wide correlation maps, threshold by significance.

We chose to use the Phenograph algorithm as it does not require defining the number of clusters a priori; instead, depending only on the parameter *k*, which defines the number of edges connecting each data point to other similar data points. Phenograph is robust to changes in *k*, with a wide range of values providing good clustering results[71], although, generally, lower *k* values lead to more clusters and higher *k* values lead to fewer clusters. Our null hypothesis is that each subspace network is unique. Therefore, we chose $k = 16$ to be conservative in our estimation of the number of clusters. However, our results did not depend on the value of *k*; similar results were seen across several values of *k*. Furthermore, alternative clustering algorithms, such as identifying the *n* exemplars that minimized distance of each data point to an exemplar, yielded similar results.

After Phenograph clustering, a self-organizing map approach was used to assign individual subspace networks to the most similar cluster. On each iteration, the median subspace network was calculated for each cluster. Each individual subspace network was then assigned to the cluster with which it had the highest correlation to the median. This process was repeated until no subspace networks changed clusters. Combining the Phenograph and self-organizing map algorithms improved the quality of the clustering by increasing the within cluster similarity.

Clustering identified 13 clusters of subspace networks. Figure 5A shows the similarity matrix of all 6279 subspace networks. Subspace networks were more similar within a cluster than between clusters. Each cluster captured a different network of regions—Fig. S5B shows the median subspace network correlation map for each cluster.

As noted in the main text, many dimensions, both within and between regions, had similar subspace networks, quantified by the fact that they were members of the same cluster. All regions had subspace dimensions that were correlated with subspace networks belonging to each cluster (Fig. S13B). These results suggest the same subspace network was functionally connected with several brain regions, which could allow a subspace network to integrate information across multiple regions. Note, the distribution of subspace networks across areas was not uniform—the entropy of the distribution was significantly less than expected by chance ($p \le 0.0002$, permutation test). Therefore, while each cluster of subspace networks was engaged by each brain region, they were engaged to different extents by different areas. This is consistent with differences in anatomical connections between regions.

Similarly, the same subspace networks were engaged during different motifs (Fig. S13C). This is consistent with the observation that motifs did not change the underlying structure of interactions, as detailed below.

Correlation with motor activity

Recent work has found representations of movements are widely distributed across cortex[6,8]. Motivated by this, we tested whether activity in each subspace network was correlated with motor activity. To this end, we used cameras (PS3 EYE webcam; 640×480 resolution) to track movements of the animal's nose, whisker-pad, and front shoulder as indicators or generalized motor activity/arousal[7,8]. Regions of interest (ROIs) were manually defined for each region (Fig. 2E). Motion energy was calculated as the average absolute temporal derivative of pixels within all three ROIs (similar results were seen for each ROI independently). Motion energy was then correlated with neural activity projected along each subspace dimension (as detailed above). To account for delays between the signals, the strength of the correlation was taken as the maximum cross-correlation value between the subspace activity vector and behavioral activity vector, up to a 533 ms offset between the signals.

The first subspace dimension—and to a lesser extent the second dimension—of most regions was significantly correlated with movements (Fig. 2F). These networks tend to be broad, interacting with many different cortical regions (Figs. 4B and S11). This is consistent with the idea that broad subspace networks may broadcast broadly relevant information, such as motor actions, across cortex.

Consistency of subspace networks

If subspace networks are capturing the interactions between brain regions, then one would expect them to be reliable across time. To quantify the stability of subspace networks between motifs we tested whether the subspace βs fit on a subset of instances of one motif could generalize to withheld instances of the same motif. Specifically, we tested the cross-generalization of the reduced rank regression:

$$Y = X_A B_A V_A^T \tag{6}$$

$$\hat{Y}' = X_{A'} B_A V_A^T \tag{7}$$

where B_A and V_A are the βs of the source and target neurons defined during trials of motif A (trained on a subset of trials). X_A and Y are the activity of neurons in the source and target regions during motif A, respectively. Withheld trials are indicated as A'.

For each subspace dimensions, the similarity between motifs was computed as the percent of variance in Y' explained by $\hat{Y}'$. The result was a 1xn vector (S) of percent similarity between the subspaces for each dimension in motif A. Total percent generalization was then computed by weighting S by the amount of explainable variance of the subspace captured by each dimension (1xn vector, P):

$$\%\,\text{generalization} = \sum_{i=1}^{10} S_i P_i \tag{8}$$

Thus, the total percent generalization reflects the total amount of variance of activity in a subspace that can be explained when generalizing across motifs.

Models were trained and tested on 4-folds of trials of a motif. The average percent explained variance on withheld data was computed for all datasets (69.4% CI: 68.7–70.0%). To compare this to chance, we shuffled the trials of a motif when fitting the models, allowing the neural activity to remain the same but breaking the association between target and source region activity. These shuffled models explained no variance in the withheld data (−0.00020% CI: −0.0031% −0.0034%). This was significantly lower than the observed generalization ($p = 0.001$).

Consistency of subspace networks across motifs

As detailed in the main text, one hypothesis is that neural activity engages different regions during different moments in time because the subspace networks themselves change. If true, then we might expect different subspace networks to occur during different motifs. This is not what we found. Instead, we found subspace networks were relatively stable across motifs, explaining the majority of possible variance.

Similar to above, we quantified the stability of subspace networks between motifs by testing whether the subspace βs fit on one motif could generalize to another motif. Specifically, we tested the cross-generalization of the reduced rank regression:

$$Y = X_A B_A V_A^T \tag{9}$$

$$\hat{Y} = X_A B_B V_B^T \tag{10}$$

where $B_{A,B}$ and $V_{A,B}$ are the βs of the source and target neurons defined during trials of motif A and B. X_A and Y are the activity of neurons in the source and target regions during motif A, respectively.

Consistent with the hypothesis that subspaces were consistent across time, we found that regression models fit during one motif generalized to capture over 45.2% of the variance in neural activity of the other motifs across target regions and datasets (Fig. S19 CI: 44.8–45.5%). While less than the stability of subspace networks within a motif (see above), the generalization was still significantly greater than the shuffled models which predicted no variance ($p = 0.001$, shuffled models described above). This supports the hypothesis that subspace networks were relatively stable over time, both within and across motifs. Indeed, the cortex-wide maps of subspace networks from the imaging data (e.g., Fig. 4B) were similar between motifs (Figs. S13C and S17C; average of 63.48% overlap between motifs, CI: 63.40–63.56%).

Comparing evoked responses in three-region neural space

As shown in Figs. 7A and S1, motifs reflected neural activity in multiple cortical regions. Different motifs were associated with different magnitudes of neural response in each brain region. To visualize the pattern of neural activity in a triad of brain regions, we projected the neural activity in all three regions into a three-dimensional space defined by the first principal component (PC) of each region (Fig. 7B). For these analyses, we were interested in comparing the evoked response across different motifs. For this reason, PCs were fit to the raw neural activity evoked during the compared motifs (a $N_{neurons}$ *by* $[N_{\text{trials}} * N_{\text{timepoints}}]$ matrix). This is in contrast to the analyses of subspaces, which were defined using RRR on the trial-to-trial variability in firing rate (detailed above).

PCs were fit to the neural activity during one of the two motifs (chosen at random). Similar results were seen when fitting PCs to neural activity from the alternative motif or from both motifs (the latter of which required dropping data to balance occurrences between motifs).

As seen in Fig. 7B, neural activity evolved differently during the two motifs. To quantify this, we fit planes to the evoked response within the three-dimensional co-response space. Planes of best fit were estimated from the first PC fit to the evoked response across all three regions (Fig. 7B, upper insert). The angle between planes was taken as the angle between the vectors normal to each plane (Fig. 7B, lower insert).

Calculating alignment between subspace activity and local neural activity

Finally, we tested the hypothesis that aligning the neural response with a particular subspace dimension would lead to neural activity in the associated network of brain regions (Figs. 5–9). If true, then the alignment of the evoked neural response and the subspace should be related to how neural activity flows across cortex. In other words, alignment should change as a function of the motifs.

To test this hypothesis, we measured the angle between 1) the vector of evoked activity in a region and 2) the subspace between two regions (both N-length vectors, where N is the number of neurons in a region). The evoked neural activity within each region was estimated by the first principal component of the neural response during each motif ($EV^{Motif\ X}_{Region1}$). As above, the subspace connecting two regions was estimated by reduced rank regression. A separate reduced rank regression model was fit to each pair of regions independently. This isolated the interactions between a pair of regions. To avoid any potential circularity in our analyses, models were fit to the trial-to-trial variability in firing rate during each motif (i.e., after removing the evoked response, as described above). This resulted in a subspace ($Sub^{Motif\ X}_{Region1\_Region2}$) for each region pair during each motif, defined as the vector β weights for the estimated subspace.

For example, the panels in Fig. 8A show the angles between the evoked response in each region and the subspace connecting two regions during each motif. The left panel shows the distribution of angles between the subspace connecting whisker and somatosensory cortex, defined using the trial-by-trial variability during motif A ($Sub^{Motif\ A}_{WHS\_SS}$), and the evoked response in somatosensory cortex during motif A (red; $EV^{Motif\ A}_{SS}$) or motif B (blue; $EV^{Motif\ B}_{SS}$). The angle between the evoked response during motif A and the subspace was less than the angle between the evoked response during motif B and the subspace. In other words, the evoked response during motif A was more 'aligned' with the subspace. This is consistent with the hypothesis: a stronger neural response was seen in WHS and SS during motif A, which may reflect better alignment of the evoked response and the subspace between WHS and SS.

Similar results were seen for the other pairs of regions in the triad. For example, the second from the left panel shows the angle between the subspace between whisker and somatosensory cortex, defined using the trial-by-trial variability during motif A ($Sub^{Motif\ A}_{WHS\_SS}$), and the evoked response in whisker cortex during motif A (red; $EV^{Motif\ A}_{WHS}$) or motif B (blue; $EV^{Motif\ B}_{WHS}$). As above, the angle between the evoked response and subspace was less during motif A, consistent with better alignment and a stronger response in the whisker and somatosensory network.

The third panel from the left shows the angle between the subspace between whisker and retrosplenial cortex, defined using the trial-by-trial variability during motif B ($Sub^{Motif\ B}_{WHS\_RSP}$), and the evoked response in whisker cortex during motif A (red; $EV^{Motif\ A}_{RSP}$) or motif B (blue; $EV^{Motif\ B}_{RSP}$). Finally, the rightmost panel shows the angle between the WHS-RSP subspace, defined using the trial-by-trial variability during motif B ($Sub^{Motif\ B}_{WHS\_RSP}$), and the evoked response in whisker cortex during motif A (red; $EV^{Motif\ A}_{WHS}$) or motif B (blue; $EV^{Motif\ B}_{WHS}$). Again, consistent with our hypothesis, the angle between the evoked response in both regions and the WHS-RSP subspace was lower during motif B, which was the motif that more strongly engaged RSP.

Note that these examples use subspaces defined on the trial-to-trial variability of the motif that more strongly engaged a pair of regions (i.e., Motif A for WHS-SS and Motif B for WHS-RSP). However, this was not critical—similar results were observed when defining the subspace using residual activity during the other motif (i.e., Motif A for WHS-RSP and Motif B for WHS-SS).

In the main text (Figs. 7–9) and supplement (Fig. S14), alignment was calculated for six example triads of regions and pairs of motifs. These were manually chosen based on the pattern of activity seen within the motifs. All pairings were screened to confirm that 1) the motifs engaged all three brain regions (i.e., neural activity was greater than baseline activity) and 2) the magnitude of the evoked response differed between motifs (i.e., each motif preferentially engaged at least one brain region; Figs. 7A and S14). In addition, these example pairings were selected to reflect a distributed sample of the overall data; collectively including all 8 regions and 8 different motifs. Choices of region triads and pairs of motifs were made while blinded to the alignment results and all pre-chosen comparisons were included in calculation of the final statistics (Fig. 8C).

Reporting summary

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

Data availability

Pre-processed, deconvolved, widefield images, and spike-sorted neural data are publicly available through DataDryad: https://doi.org/10.5061/dryad.gxd2547x8. Source data are provided with this paper.

Code availability

Code for figure generation, and figure data, are publicly available through a publicly available GitHub repository: https://github.com/buschman-lab/MultiplexedSubspaces. https://doi.org/10.5281/zenodo.14721950.

References

1. Han, Y. et al. The logic of single-cell projections from visual cortex. *Nature* **556**, 51–56 (2018).
2. Hutchison, R. M. et al. Dynamic functional connectivity: promise, issues, and interpretations. *Neuroimage* **80**, 360–378 (2013).
3. Siegel, M., Buschman, T. J. & Miller, E. K. Cortical information flow during flexible sensorimotor decisions. *Science* **348**, 1352–1355 (2015).
4. Pinto, L. et al. Task-Dependent Changes in the Large-Scale Dynamics and Necessity of Cortical Regions. *Neuron* **104**, 810–824.e9 (2019).
5. Steinmetz, N. A., Zatka-Haas, P., Carandini, M. & Harris, K. D. Distributed coding of choice, action and engagement across the mouse brain. *Nature* **576**, 266–273 (2019).
6. Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.* **22**, 1677–1686 (2019).
7. MacDowell, C. J. & Buschman, T. J. Low-Dimensional Spatiotemporal Dynamics Underlie Cortex-wide Neural Activity. *Curr. Biol.* **30**, 2665–2680.e8 (2020).
8. Stringer, C. et al. Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, 255 (2019).
9. Mohajerani, M. H. et al. Spontaneous cortical activity alternates between motifs defined by regional axonal projections. *Nat. Neurosci.* **16**, 1426–1435 (2013).
10. MacDowell, C. J., Briones, B. A., Lenzi, M. J., Gustison, M. L. & Buschman, T. J. *Differences in the Expression of Cortex-Wide Neural Dynamics Are Related to Behavioral Phenotype*. https://doi.org/10.1016/j.cub.2024.02.004 (2024).

11. Spadone, S. et al. Dynamic reorganization of human resting-state networks during visuospatial attention. *PNAS* **112**, 8112–8117 (2015).
12. Makino, H. et al. Transformation of Cortex-wide Emergent Properties during Motor Learning. *Neuron* **94**, 880–890.e8 (2017).
13. Morishima, Y. et al. Task-specific signal transmission from prefrontal cortex in visual selective attention. *Nat. Neurosci.* **12**, 85–91 (2009).
14. Miller, E. K. & Cohen, J. D. An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* **24**, 167–202 (2001).
15. Kohn, A. et al. Principles of Corticocortical Communication: Proposed Schemes and Design Considerations. *Trends Neurosci.* **43**, 725–737 (2020).
16. Fries, P. Rhythms for Cognition: Communication through Coherence. *Neuron* **88**, 220–235 (2015).
17. Palmigiano, A., Geisel, T., Wolf, F. & Battaglia, D. Flexible information routing by transient synchrony. *Nat. Neurosci.* **20**, 1014–1022 (2017).
18. Semedo, J. D., Zandvakili, A., Machens, C. K., Yu, B. M. & Kohn, A. Cortical Areas Interact through a Communication Subspace. *Neuron* **102**, 249–259.e4 (2019).
19. Panichello, M. F. & Buschman, T. J. Shared mechanisms underlie the control of working memory and attention. *Nature* **592**, 601–605 (2021).
20. Jun, J. J. et al. Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**, 232–236 (2017).
21. Semedo, J. D., Gokcen, E., Machens, C. K., Kohn, A. & Yu, B. M. Statistical methods for dissecting interactions between brain areas. *Curr. Opin. Neurobiol.* **65**, 59–69 (2020).
22. Stringer, C., Pachitariu, M., Steinmetz, N., Carandini, M. & Harris, K. D. High-dimensional geometry of population responses in visual cortex. *Nature* **1** https://doi.org/10.1038/s41586-019-1346-5 (2019).
23. Veuthey, T. L., Derosier, K., Kondapavulur, S. & Ganguly, K. Single-trial cross-area neural population dynamics during long-term skill learning. *Nat. Commun.* **11**, 4057 (2020).
24. Semedo, J. D. et al. Feedforward and feedback interactions between visual cortical areas use different population activity patterns. *Nat. Commun.* **13**, 1099 (2022).
25. Muñoz-Castañeda, R. et al. Cellular anatomy of the mouse primary motor cortex. *Nature* **598**, 159–166 (2021).
26. Aransay, A., Rodríguez-López, C., García-Amado, M., Clascá, F. & Prensa, L. Long-range projection neurons of the mouse ventral tegmental area: a single-cell axon tracing analysis. *Front Neuroanat.* **9**, 59 (2015).
27. Benisty, H. et al. Rapid fluctuations in functional connectivity of cortical networks encode spontaneous behavior. *Nat. Neurosci.* **27**, 148–158 (2024).
28. MacDowell, C. J., Briones, B. A., Lenzi, M. J., Gustison, M. L. & Buschman, T. J. Differences in the expression of cortex-wide neural dynamics are related to behavioral phenotype. *Curr. Biol.* https://doi.org/10.1016/j.cub.2024.02.004.
29. Cardin, J. A., Crair, M. C. & Higley, M. J. Mesoscopic Imaging: Shining a Wide Light on Large-Scale Neural Dynamics. *Neuron* **108**, 33–43 (2020).
30. Vanni, M. P., Chan, A. W., Balbi, M., Silasi, G. & Murphy, T. H. Mesoscale Mapping of Mouse Cortex Reveals Frequency-Dependent Cycling between Distinct Macroscale Functional Modules. *J. Neurosci.* **37**, 7513–7533 (2017).
31. Silasi, G., Xiao, D., Vanni, M. P., Chen, A. C. N. & Murphy, T. H. Intact skull chronic windows for mesoscopic wide-field imaging in awake mice. *J. Neurosci. Methods* **267**, 141–149 (2016).
32. Mandino, F. et al. A triple-network organization for the mouse brain. *Mol. Psychiatry* **27**, 865–872 (2022).
33. Grandjean, J. et al. Common functional networks in the mouse brain revealed by multi-centre resting-state fMRI analysis. *Neuroimage* **205**, 116278 (2020).
34. Vafaii, H. et al. Multimodal measures of spontaneous brain activity reveal both common and divergent patterns of cortical functional organization. *Nat. Commun.* **15**, 229 (2024).
35. Gozzi, A. & Schwarz, A. J. Large-scale functional connectivity networks in the rodent brain. *NeuroImage* **127**, 496–509 (2016).
36. Xu, N. et al. Functional Connectivity of the Brain Across Rodents and Humans. *Front Neurosci.* **16**, 816331 (2022).
37. Mackevicius, E. L. et al. Unsupervised discovery of temporal sequences in high-dimensional datasets, with applications to neuroscience. *Elife* **8**, e38471 (2019).
38. MacDowell, C. J., Tafazoli, S. & Buschman, T. J. A Goldilocks theory of cognitive control: Balancing precision and efficiency with low-dimensional control states. *Curr. Opin. Neurobiol.* **76**, 102606 (2022).
39. Park, H.-J. & Friston, K. Structural and functional brain networks: from connections to cognition. *Science* **342**, 1238411 (2013).
40. Giarrocco, F. & Averbeck, B. B. Neural Circuits: Organization of parietoprefrontal and temporoprefrontal networks in the macaque. *J. Neurophysiol.* **126**, 1289 (2021).
41. Moeller, S., Nallasamy, N., Tsao, D. Y. & Freiwald, W. A. Functional Connectivity of the Macaque Brain across Stimulus and Arousal States. *J. Neurosci.* **29**, 5897–5909 (2009).
42. Buschman, T. J. & Miller, E. K. Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science* **315**, 1860–1862 (2007).
43. Churchland, M. M. et al. Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).
44. Libby, A. & Buschman, T. J. Rotational dynamics reduce interference between sensory and memory representations. *Nat. Neurosci.* **24**, 715–726 (2021).
45. Jiang, X., Saggar, H., Ryu, S. I., Shenoy, K. V. & Kao, J. C. Structure in Neural Activity during Observed and Executed Movements Is Shared at the Neural Population Level, Not in Single Neurons. *Cell Rep.* **32**, 108–148 (2020).
46. Sun, X. et al. Cortical preparatory activity indexes learned motor memories. *Nature* **602**, 274–279 (2022).
47. Elsayed, G. F., Lara, A. H., Kaufman, M. T., Churchland, M. M. & Cunningham, J. P. Reorganization between preparatory and movement population responses in motor cortex. *Nat. Commun.* **7**, 13239 (2016).
48. Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Cortical activity in the null space: permitting preparation without movement. *Nat. Neurosci.* **17**, 440–448 (2014).
49. Musall, S. et al. *Pyramidal Cell Types Drive Functionally Distinct Cortical Activity Patterns during Decision-Making*. 2021.09.27.461599 https://www.biorxiv.org/content/10.1101/2021.09.27.461599v2 (2021) https://doi.org/10.1101/2021.09.27.461599.
50. O'Connor, D. et al. Functional network properties derived from wide-field calcium imaging differ with wakefulness and across cell type. *NeuroImage* **264**, 119735 (2022).
51. MacDowell, C. J., Briones, B. A., Lenzi, M. J., Gustison, M. L. & Buschman, T. J. Differences in the expression of cortex-wide neural dynamics are related to behavioral phenotype. *Curr. Biol.* **34**, 1333–1340 (2024).
52. Cardin, J. A. Functional flexibility in cortical circuits. *Curr. Opin. Neurobiol.* **58**, 175–180 (2019).
53. Lohani, S. et al. Spatiotemporally heterogeneous coordination of cholinergic and neocortical activity. *Nat. Neurosci.* **25**, 1706–1713 (2022).
54. Steinmetz, N. A. et al. Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science* **372**, eabf4588 (2021).
55. Gokcen, E. et al. Disentangling the flow of signals between populations of neurons. *Nat. Computational Sci.* **2**, 512–525 (2022).

56. Briggman, K. L., Abarbanel, H. D. I. & Kristan, W. B. From crawling to cognition: analyzing the dynamical interactions among populations of neurons. *Curr. Opin. Neurobiol.* **16**, 135–144 (2006).
57. Ganguli, S. et al. One-Dimensional Dynamics of Attention and Decision Making in LIP. *Neuron* **58**, 15–25 (2008).
58. Sadtler, P. T. et al. Neural constraints on learning. *Nature* **512**, 423–426 (2014).
59. Buschman, T. J. & Kastner, S. From Behavior to Neural Dynamics: An Integrated Theory of Attention. *Neuron* **88**, 127–144 (2015).
60. Cowan, N. The Magical Mystery Four: How is Working Memory Capacity Limited, and Why? *Curr. Dir. Psychol. Sci.* **19**, 51–57 (2010).
61. Iyengar, S. S. & Lepper, M. R. When choice is demotivating: can one desire too much of a good thing? *J. Pers. Soc. Psychol.* **79**, 995–1006 (2000).
62. Timmermans, D. The impact of task complexity on information use in multi-attribute decision making. *J. Behav. Decis. Mak.* **6**, 95–111 (1993).
63. Dana, H. et al. Thy1-GCaMP6 Transgenic Mice for Neuronal Population Imaging In Vivo. *PLOS ONE* **9**, e108697 (2014).
64. Steinmetz, N. A. et al. Aberrant Cortical Activity in Multiple GCaMP6-Expressing Transgenic Mouse Lines. *eNeuro* **4**, ENEURO.0207-17.2017 (2017).
65. Claudi, F. et al. Visualizing anatomically registered data with brainrender. *eLife* **10**, e65751 (2021).
66. Peters, A. J., Fabre, J. M. J., Steinmetz, N. A., Harris, K. D. & Carandini, M. Striatal activity topographically reflects cortical activity. *Nature* **591**, 420–425 (2021).
67. Degleris, A., Antin, B., Ganguli, S. & Williams, A. H. Fast Convolutive Nonnegative Matrix Factorization Through Coordinate and Block Coordinate Updates. *arXiv:1907.00139 [cs, stat]* (2019).
68. Stern, M., Shea-Brown, E. & Witten, D. Inferring the Spiking Rate of a Population of Neurons from Wide-Field Calcium Imaging. *bioRxiv* 2020.02.01.930040 https://doi.org/10.1101/2020.02.01.930040 (2020).
69. Waters, J. Sources of widefield fluorescence from the brain. *eLife* **9**, e59841 (2020).
70. Wekselblatt, J. B., Flister, E. D., Piscopo, D. M. & Niell, C. M. Large-scale imaging of cortical dynamics during sensory perception and behavior. *J. Neurophysiol.* **115**, 2852–2866 (2016).
71. Levine, J. H. et al. Data-Driven Phenotypic Dissection of AML Reveals Progenitor-like. *Cells that Correl. Prognosis. Cell* **162**, 184–197 (2015).
72. Rossant, C. et al. Spike sorting for large, dense electrode arrays. *Nat. Neurosci.* **19**, 634–641 (2016).
73. Wang, Q. et al. The Allen Mouse Brain Common Coordinate Framework: A 3D Reference Atlas. *Cell* **181**, 936–953.e20 (2020).
74. Brody, C. D. Correlations Without Synchrony. *Neural Comput.* **11**, 1537–1551 (1999).
75. Izenman, A. J. Reduced-rank regression for the multivariate linear model. *J. Multivar. Anal.* **5**, 248–264 (1975).
76. Hotelling, H. Relations between two sets of variates. *Biometrika* **28**, 321–377 (1936).
77. Gao, L. et al. Single-neuron projectome of mouse prefrontal cortex. *Nat. Neurosci.* **25**, 515–529 (2022).
78. Barson, D. et al. Simultaneous mesoscopic and two-photon imaging of neuronal activity in cortical circuits. *Nat. Methods* **17**, 107–113 (2020).
79. Chen, J. L., Carta, S., Soldado-Magraner, J., Schneider, B. L. & Helmchen, F. Behaviour-dependent recruitment of long-range projection neurons in somatosensory cortex. *Nature* **499**, 336–340 (2013).
80. Stringer, C., Pachitariu, M., Steinmetz, N., Carandini, M. & Harris, K. D. High-dimensional geometry of population responses in visual cortex. *Nature* **571**, 361–365 (2019).
81. Bernardi, S. et al. The Geometry of Abstraction in the Hippocampus and Prefrontal Cortex. *Cell* **183**, 954–967.e21 (2020).

Acknowledgements

We thank the Buschman lab for their feedback throughout this project and the Princeton Laboratory Animal Resources staff for their support. This work was funded by a grant from SFARI 670183 (T.J.B.), NIH NJCATS Award TL1TR003019 (C.J.M), and R01MH126022 (T.J.B).

Author contributions

Conceptualization: C.J.M., T.J.B., A.L., C.I.J., S.T.; Experimental Design: C.J.M., T.J.B.; Surgeries and Data Collection: C.J.M.; Analysis and Visualization: C.J.M., T.J.B., A.A.; Writing: C.J.M. and T.J.B.; Editing: T.J.B., C.J.M., A.L., C.I.J., S.T., A.A.; Funding Acquisition: T.J.B., C.J.M.; Supervision: T.J.B.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41467-025-58698-2.

Correspondence and requests for materials should be addressed to Timothy J. Buschman.

Peer review information *Nature Communications* thanks Giuliano Iurilli, Anna Letizia Allegra Mascaro and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. A peer review file is available.

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