


Interactions across hemispheres in prefrontal cortex reflect global cognitive processing

Received: 3 July 2025

Accepted: 25 March 2026

Published online: 11 April 2026

Megan E. McDonnell^{1,2,7}, Akash Umakantha^{1,2,3,7}, Ryan C. Williamson^{1,2,3,4,7},
Matthew A. Smith^{1,2,5,8} ✉ & Byron M. Yu^{1,2,5,6,8} ✉

Brain functions involve processing in local networks as well as modulation from brainwide signals, such as arousal. Dissecting the contributions of populations of neurons to these functions requires knowledge of interactions between brain areas. We investigated these interactions using dual hemisphere recordings of prefrontal cortex in monkeys performing a spatial memory task. To tease apart global processing from local interactions, we applied a novel statistical approach called pCCA-FA (a combination of probabilistic canonical correlation analysis and factor analysis) to analyze trial-to-trial variability in neuronal activity. We found substantial shared variability among neurons within each population, much of which was actually shared across populations and linked to an arousal process. Our work presents a path by which we can leverage multi-area recordings to reveal aspects of brain functions that are hidden in single-area recordings.

Brain functions such as attention^{1,2}, decision making³, and working memory^{4,5} rely on coordination among neurons within brain areas as well as across brain areas. Thus, the spiking activity of individual neurons can be influenced by neurons from the same brain area, neurons in other brain areas, global modulatory signals, and more⁶. Teasing apart the impact of each of these influences on neural activity would enable a deeper understanding of how neurons in different brain areas interact to give rise to such brain functions.

To dissect the various influences on neural activity, we can consider how the activity of one neuron covaries with that of other nearby neurons. Many studies have examined how the activity of pairs of neurons in a single brain area covaries from trial-to-trial by computing spike count correlation (r_{sc})^{7,8}. These pairwise correlations have been widely used as a window to investigate a range of phenomena from cognitive states to circuit architectures^{9,10}. To go beyond pairs of neurons, a related body of work has identified the population-wide covariability structure in single brain areas. These population-level

approaches have been used to study the trial-to-trial variability underlying sensory encoding^{11,12}, attention^{13,14}, motor control^{15–17}, learning^{18–20}, decision-making^{21–23}, network structure^{24–26}, and more²⁷. Taken together, these studies inform our understanding of how neural activity is shared among neurons within a brain area.

Given that most brain functions rely on interactions across brain areas, it is likely that some component of the activity of each brain area is shaped by interactions with other brain areas. Developments in neural recording technology have enabled simultaneous recordings of populations of neurons in multiple brain areas^{28–30}, which give rise to the possibility of distinguishing interactions across brain areas from interactions solely within a single brain area^{31–34}. Simultaneous recordings from multiple brain areas have been leveraged to study widespread brain functions such as sensory processing^{35–40}, motor control^{41–44}, learning^{45–47}, decision-making^{48–50}, working memory^{51,52}, arousal^{53,54}, and attention^{55–60}. However, most statistical methods for analyzing populations of neurons are not designed to distinguish

¹Neuroscience Institute, Carnegie Mellon University, Pittsburgh, PA, USA. ²Center for the Neural Basis of Cognition, Carnegie Mellon University & University of Pittsburgh, Pittsburgh, PA, USA. ³Machine Learning Department, Carnegie Mellon University, Pittsburgh, PA, USA. ⁴School of Medicine, University of Pittsburgh, Pittsburgh, PA, USA. ⁵Department of Biomedical Engineering, Carnegie Mellon University, Pittsburgh, PA, USA. ⁶Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, USA. ⁷These authors contributed equally: Megan E. McDonnell, Akash Umakantha, Ryan C. Williamson. ⁸These authors jointly supervised this work: Matthew A. Smith, Byron M. Yu. ✉e-mail: mattsmith@cmu.edu; byronyu@cmu.edu

activity shared across areas from activity shared solely among neurons within each area.

Here, we seek to identify components of neural activity that involve interactions across two areas (i.e., an across-area component), as well as interactions solely within each area (i.e., a within-area component). We simultaneously recorded spiking activity from each of the two hemispheres (here referred to as two areas) of prefrontal cortex (PFC) in macaque monkeys while they performed a spatial memory task. Utilizing a statistical method we developed called pCCA-FA (a combination of probabilistic canonical correlation analysis and factor analysis), we found substantial shared variance among neurons in PFC, much of which involved neurons across both areas. In fact, the across-area component of PFC activity was larger than the within-area component. We then uncovered a link between the across-area component and a brainwide arousal process. Our work motivates a reinterpretation of shared variance among neurons within a brain area as arising from multiple sources, both within that area as well as via interactions across the brain.

Results

To support perception, cognition, and action, each neuron's activity can be shared with other neurons in the same brain area, as well as with neurons in other brain areas. Consider the spiking activity of two distinct populations of neurons, recorded simultaneously from two brain areas (area 1 and area 2, Fig. 1a). For each neuron, we would like to identify which aspects of its activity are shared with neurons in the other area (pink arrow) and which aspects are shared exclusively with neurons in the same area (blue arrow). To understand how the activity of neuron 1 in area 1 is shared with other neurons (i.e., to quantify the relationships described by the pink and blue arrows in Fig. 1a), we would like to partition its trial-to-trial spike count variance into an across-area component (termed across-area variance; Fig. 1b, pink), a within-area component (termed within-area variance; Fig. 1b, blue), and a component independent to this neuron (termed independent variance; Fig. 1b, gray).




Fig. 1 | Understanding neuronal variability using spike count variance. **a** Given two populations of neurons, we seek to understand what aspects of each neuron's activity (e.g., area 1, neuron 1, shaded triangle; box indicates spike count bin) are shared with neurons in the other area (pink arrow) versus only with other neurons in the same area (blue arrow). **b** By considering how a neuron's activity covaries with other neurons in the same area and the other area, we can partition its spike count variance into a component that is only shared with neurons in the same area (within-area variance, blue partition), a component that is shared with neurons in the other area (across-area variance, pink partition), and a component that is not shared with any other recorded neuron (independent variance, gray partition).

Mean spike count correlation does not adequately partition shared variance

We simultaneously recorded spiking activity from neurons in both hemispheres of dorsolateral PFC to understand how neural activity was shared either across or within areas (Fig. 2a). We began by characterizing the pairwise interactions among neurons within each area using r_{sc} . The spike counts used to compute r_{sc} were taken from a one-second window during the delay period of a memory-guided saccade task (Fig. 2b). Throughout this work, we combined trials across conditions after removing condition-specific means, and removed slow-timescale trends to avoid spurious correlations⁶¹ (see Methods, Supplementary Fig. 1).

The mean r_{sc} among pairs of recorded neurons in a brain area is commonly taken to indicate the strength of shared activity among neurons. We computed r_{sc} for each pair of neurons within each area (Fig. 2c, left and middle, example session), then averaged over the pairs to obtain a within-area mean r_{sc} per session (right PFC within-area mean $r_{sc} = 0.029$, left PFC within-area mean $r_{sc} = 0.033$, example session). In the majority of sessions, we found that within-area mean r_{sc} values were positive (Fig. 2d, top and middle).

To understand how activity is shared across areas, one approach is to compute r_{sc} between pairs of neurons across areas (i.e., one neuron in each area). Mean r_{sc} decreases for pairs of neurons as the physical distance between the neurons increases^{62–64}, and has been shown to be smaller for pairs of neurons across versus within brain areas⁶⁵. We observed an across-area mean r_{sc} that was close to zero (Fig. 2c, right; d, bottom). This might be taken to suggest that neurons within areas interact more strongly than across areas (Fig. 2d; $p = 1.0e - 13$, pooled across animals and brain areas; one-tailed t-test), and that there are little to no interactions across the hemispheres of PFC.

However, there are several limitations to mean r_{sc} measurements. First, by averaging over all pairs of neurons (either across or within areas), we ignore the breadth of r_{sc} distributions (Fig. 2c). For example, if there are pairs of neurons with strong positive correlations and also pairs with strong negative correlations, the mean r_{sc} might be close to zero even though many neuron pairs exhibit strong correlations. One way to identify neuron pairs with strong positive and negative correlations is to consider the relative tuning preferences of the neurons^{66–70}. Indeed, when we computed mean r_{sc} among subsets of across-area neuron pairs with similar tuning (Fig. 2e, top) and opposite tuning (Fig. 2e, bottom), we found that similarly tuned pairs of neurons exhibited positive r_{sc} (mean $r_{sc} = 0.019$; $p = 1.6e - 10$, pooled across animals; one-tailed t-test), whereas oppositely tuned pairs of neurons exhibited negative r_{sc} (mean $r_{sc} = -0.013$; $p = 1.6e - 7$, pooled across animals; one-tailed t-test). Second, the within-area mean r_{sc} does not reveal whether the activity shared among neurons in an area is also shared with other areas. Rather, all co-fluctuations among neurons in a brain area are collapsed into the within-area mean r_{sc} . As such, within-area mean r_{sc} could reflect interactions with other areas and/or global neuromodulation. Third, mean r_{sc} (either across- or within-area) is not designed to identify the many possible population-wide co-fluctuations among neurons. Several distinct changes in population-wide co-fluctuations can result in the same change in mean r_{sc} ^{71,72}. By collapsing all interactions in one metric, we are unable to separately characterize the many influences on neural activity. Here we desire to leverage the simultaneity of our dual PFC recordings to dissect the multiple population-wide influences on across- and within-area activity (as in Fig. 1b).

pCCA-FA partitions across- and within-area shared variance

To explicitly partition across- and within-area variance, we can leverage statistical methods that take into account all simultaneously-recorded neurons together. In a single brain area, we can identify population-wide co-fluctuation patterns using factor analysis (FA) or




Fig. 2 | Spike count correlations during a memory-guided saccade task. **a** Raster plots showing spiking activity from each hemisphere of PFC recorded with multi-electrode arrays during an example session (Monkey Sa). Our goal was to partition neuronal variability in PFC into across- and within-area components. **b** Monkeys performed a memory-guided saccade task in which a brief target cue was followed by a variable-length delay period (see Methods). After the delay period, they reported the remembered location of the cue with a saccade. Spike counts recorded during the delay period were used for all analyses. **c** r_{sc} distributions for pairs of neurons within right PFC (left), within left PFC (middle), and across areas (right) from an example session (Monkey Pe). Within-area mean $r_{sc} > 0$ ($p = 1.5e - 158$, pooled across brain areas, $n = 6,484$ neuron pairs; one-tailed t-test against zero). Across-area mean r_{sc} was not > 0 ($p = 0.55$, $n = 6,557$ neuron pairs; one-tailed t-test against zero). The mean of each distribution is marked with a triangle. **d** Aggregated across sessions, within-area mean r_{sc} (top: right PFC, middle: left PFC) was greater than across-area mean r_{sc} (bottom) ($p = 1.0e - 13$, pooled across animals and brain areas, $n = 84$ brain areas, 2 per session; Monkey Pe: $p = 8.6e - 21$, $n = 32$; Monkey Wa:

$p = 0.039$, $n = 20$; Monkey Sa: $p = 1.1e - 6$, $n = 32$; one-tailed t-tests on the differences in mean r_{sc} for each session). Each box plot covers the first to third quartiles of the values of mean r_{sc} in all sessions (across- or within-area), with the mean across sessions marked with a dashed line. Each whisker covers an additional $1.5 \times$ the inter-quartile range. Values for individual sessions are shown above each box plot, with the example session in **c** marked with stars. **e** Similarly tuned (tuned to target location) pairs of neurons in the two areas exhibited positive r_{sc} (top; mean $r_{sc} = 0.019$; $p = 1.6e - 10$, pooled across animals, $n = 42$ sessions; one-tailed t-test against zero), whereas oppositely tuned pairs of neurons in the two areas exhibited negative r_{sc} (bottom; mean $r_{sc} = -0.013$; $p = 1.6e - 7$, pooled across animals, $n = 42$ sessions; one-tailed t-test against zero). The bar plots indicate the mean across sessions, and the error bars indicate the standard error of the mean computed across the $n = 42$ sessions. Values for individual sessions are shown above each bar plot, with the example session in **c** marked with stars. * Source data are provided as a Source Data file.

related methods⁷³. While this provides a more detailed view of population-wide covariability than mean r_{sc} ⁷², these methods do not reveal whether activity fluctuations are shared with other brain areas. When recording from two brain areas simultaneously, we can identify variance shared among neurons across areas by employing canonical correlation analysis (CCA) or related methods³³. However, neither CCA nor its probabilistic variant, pCCA⁷⁴, are designed to also identify

variance shared only among neurons within each area (see Methods). Given that none of the aforementioned methods are designed to partition variance into both across-area and within-area components, a new approach is needed.

To address this need, we developed a dimensionality reduction method for two simultaneously recorded brain areas called pCCA-FA (a combination of probabilistic canonical correlation analysis and




Fig. 3 | pCCA-FA partitions across- and within-area shared variance. a Simulated neural population activity from two brain areas (illustrated here using 20 neurons in each brain area) share an across-area latent variable (z , pink trace, same for both areas). Each area has its own within-area latent variable (z_1 and z_2 , light and dark blue trace for areas 1 and 2, respectively). An example time bin of interest (trial t) is highlighted by the gray shading. **b** The across-area latent variable is mapped to the activity of each neuron via a pair of across-area co-fluctuation patterns, one for each area (top and bottom, pink). Each within-area latent variable is mapped to the activity of each neuron in the corresponding area via a within-area co-fluctuation pattern (area 1: top, light blue; area 2: bottom, dark blue). **c** The co-fluctuation patterns can be represented as vectors in a population activity space for each area (area 1: top; area 2: bottom), where each coordinate axis represents the activity of one neuron. Each pair of points (one in area 1 and one in area 2) represents the spike

counts in a particular time bin of interest (e.g., trial t in **a**). The across-area co-fluctuation patterns (pink arrows) are paired across areas. There is a separate within-area co-fluctuation pattern (area 1: top, light blue arrow; area 2: bottom, dark blue arrow) for each area. **d** pCCA-FA partitions each neuron's variance into three components as in Fig. 1b. **e** pCCA-FA accurately recovered ground truth values of across-area (top) and within-area (bottom) %sv with sufficient trials (horizontal axis) in simulated data. Error bars indicate the mean $\pm$ one standard deviation across $n = 30$ simulated datasets. Dark lines indicate %sv estimates from pCCA-FA, and light lines indicate %sv estimates from pCCA. **f** pCCA-FA accurately recovered ground truth values of across-area (top) and within-area (bottom) d_{shared} with sufficient trials in simulated data. Same conventions as in **e**, but for d_{shared} . Note that the same pCCA-FA and pCCA models are analyzed in both **e**, **f**. * Source data are provided as a Source Data file.

factor analysis), which partitions the variance of each area's activity into an across-area component, a within-area component, and a component independent to each neuron. pCCA-FA can be applied to any two populations of simultaneously recorded neurons, and here we focus on the case where each population is recorded in a distinct brain area (Fig. 3a, area 1 and area 2). We seek to identify latent variables that capture co-fluctuations among neurons across areas (Fig. 3a, z , pink),

as well as latent variables specific to each area that capture co-fluctuations exclusively among neurons within each area (Fig. 3a, area 1: z_1 , light blue; area 2: z_2 , dark blue). In total, pCCA-FA identifies three sets of latent variables: one across-area set, and a separate within-area set for each area. Because the within-area latent variables are specific to each area, the number of within-area latent variables for each area need not be the same.




Fig. 4 | Across-area interactions were substantial and often greater than within-area interactions. **a** pCCA-FA partitions the variance of dual-hemisphere PFC recordings (same example session as Fig. 2c, Monkey Pe). Same conventions as in Fig. 3d. Across-area %sv was 14.97% and 17.47%, and within-area %sv was 5.65% and 7.47%, for right and left PFC, respectively. For visual clarity, only 14/79 neurons in right PFC and 14/83 neurons in left PFC are shown. Note, though the across-area latent variables are common to both areas, the across-area %sv can have a unique value, depending on the individual variances of each neuron. **b** Aggregated across sessions, across-area %sv (mean 11.50%) was larger than within-area %sv (mean 7.66%; $p = 2.6 \times 10^{-7}$, pooled across animals and brain areas, $n = 84$ brain areas, 2 per session; Monkey Pe: $p = 0.00048$, $n = 32$; Monkey Wa: $p = 0.40$, $n = 20$; Monkey Sa: $p = 4.2 \times 10^{-6}$, $n = 32$; paired one-tailed t-tests). Each session contributes two points, one each for right and left PFC. The example session in **a** is marked in light

and dark blue for right and left PFC, respectively. The gray histogram shows the difference between across- and within-area %sv for each session ($\Delta = \text{within-area \%sv} - \text{across-area \%sv}$), with the mean difference indicated by a triangle. **c** Aggregated across sessions, across-area d_{shared} (mean 5.41) was larger than within-area d_{shared} (mean 4.06; $p = 3.1 \times 10^{-6}$, pooled across animals and brain areas, $n = 84$ brain areas, 2 per session; Monkey Pe: $p = 0.0013$, $n = 32$; Monkey Wa: $p = 0.32$, $n = 20$; Monkey Sa: $p = 9.0 \times 10^{-5}$, $n = 32$; paired one-tailed t-tests). Same conventions as in **b**. Each session contributes two values, one for each area. Because d_{shared} is integer-valued, the values from multiple sessions overlap. The size of each symbol indicates the number of areas across sessions with the particular values of d_{shared} (larger symbol indicates more sessions). Symbols for each animal are jittered slightly for visual clarity. * Source data are provided as a Source Data file.

pCCA-FA approximates neural activity as a weighted sum of the across-area latent variables (Fig. 3b, pink) and the within-area latent variables (Fig. 3b, light and dark blue for areas 1 and 2, respectively). The co-fluctuation patterns comprise the weights, which are estimated from the neural activity and fixed across all trials. The values of the latent variables can vary from trial to trial. Each co-fluctuation pattern can be represented as a vector in a population activity space, where each coordinate axis represents the activity of one neuron (Fig. 3c). Whereas the across-area co-fluctuation patterns are paired across areas (Fig. 3c, pink), the within-area co-fluctuation patterns are separate for each area (i.e., not paired across areas) (Fig. 3c, light and dark blue for areas 1 and 2, respectively).

For each neuron, we can quantify the proportion of its spike count variance explained by each set of latent variables. pCCA-FA partitions each neuron's variance into a component that is described by the across-area latent variables (Fig. 3d, pink) and a component that is described by that area's within-area latent variables (Fig. 3d, light and dark blue for areas 1 and 2, respectively). The residual variance not accounted for by across- or within-area latent variables is independent to each neuron (Fig. 3d, gray). Thus, pCCA-FA explicitly partitions the spike count variance of each neuron into across-area, within-area, and independent components (as in Fig. 1b).

To summarize the outputs of pCCA-FA, we define the following two metrics. First, to understand the extent to which a particular set of latent variables explains a neuron's spike count variance, we define the across- and within-area percentage of variance shared among neurons (%sv). The values of this metric reflect the relative sizes of the bars in Fig. 3d. Across-area %sv is defined as the proportion of spike count variance explained by the across-area latent variables, averaged across neurons in each brain area (see Eq. (8)). Within-area %sv is defined analogously using the within-area latent variables (see Eq. (9)). Thus, each brain area has both an across-area %sv and a within-area %sv. Second, we quantify the complexity of the interactions described by each set of latent variables by computing its shared dimensionality

(d_{shared})^{25,75}. A larger d_{shared} would indicate more complex interactions among neurons, whereas a smaller d_{shared} would indicate less complex interactions among neurons. Across- and within-area d_{shared} are defined as the number of latent variables required to explain the across- or within-area variance, respectively (see Methods).

Before applying pCCA-FA to neural recordings, we performed extensive model validation. We verified that pCCA-FA accurately identified the ground truth %sv and d_{shared} using simulated data with properties chosen to be similar to our neural recordings (Supplementary Figs. 2–3). We also found that pCCA-FA required fewer trials to recover the ground truth %sv and d_{shared} than pCCA (Fig. 3e, f) because pCCA-FA has fewer parameters than pCCA (Supplementary Fig. 4). Further, pCCA-FA outperformed other dimensionality reduction techniques (pCCA and FA) on neural recordings (Supplementary Fig. 4). Thus, the pCCA-FA framework successfully captures how neurons in two brain areas covary by partitioning variance into across- and within-area components.

Across-area interactions were substantial and often greater than within-area interactions

To identify across- and within-area variance in our PFC recordings, we fit pCCA-FA to the same neural activity (with the same preprocessing) as in the r_{sc} analyses (Fig. 2). For each session, pCCA-FA estimated a single set of across-area latent variables common to both areas and a separate set of within-area latent variables for each area. From the mean r_{sc} values in Fig. 2, one might expect to uncover a large within-area component and small across-area component using pCCA-FA.

To determine if this was the case, we computed across- and within-area %sv and d_{shared} . For the same example session as in Fig. 2c, we found that a substantial portion of variance was shared across areas (14.97% and 17.47%, for right and left PFC, respectively), compared to the variance shared exclusively within each area (5.65% and 7.47%, for right and left PFC, respectively) (Fig. 4a). This trend was evident across sessions, with the across-area %sv being larger than the within-area %sv




Fig. 5 | Across- and within-area components involved different co-fluctuation patterns. **a** The across- (pink arrow) and within-area (blue arrow) co-fluctuation patterns identified by pCCA-FA can be represented as vectors in each population activity space. To determine how similar the top across- and within-area co-fluctuation patterns were, we computed the angle between them ($\theta_{\text{estimated}}$). **b** Weights of the top across- (pink) and within-area (blue) co-fluctuation patterns from an example session (Monkey Sa, right PFC). The height of each bar indicates the weight of one neuron for that co-fluctuation pattern. For this example session, there was a high diversity of weights, many with a positive sign and many with a negative sign. The angle between the across- and within-area co-fluctuation patterns ($\theta_{\text{estimated}}$) was 72° . **c** $\theta_{\text{estimated}}$ for all sessions and all animals (mean $\theta_{\text{estimated}} = 68.2^\circ$ marked with a triangle). Each session contributed one angle per brain area (i.e., two angles per session, $n = 84$ angles). **d** Simulations, in which the ground truth angle between the top across- and within-area co-fluctuation patterns was specified (θ_{sim} , see Methods), were performed to generate chance

distributions. Same conventions as in **c**. Top, three simulations were performed in which θ_{sim} was specified as a constant value across all simulated datasets ($\theta_{\text{sim}} \in [30, 60, 90^\circ]$). The $\theta_{\text{estimated}}$ from neural recordings (**c**) was greater than the $\theta_{\text{estimated}}$ from simulated data for $\theta_{\text{sim}} = 30^\circ$ ($p = 9.0 \times 10^{-28}$, pooled across animals and brain areas, $n = 84$ brain areas, 2 per dataset; paired one-tailed t-test) and for $\theta_{\text{sim}} = 60^\circ$ ($p = 0.00092$, pooled across animals and brain areas, $n = 84$ brain areas, 2 per dataset; paired one-tailed t-test), but not for $\theta_{\text{sim}} = 90^\circ$ ($p = 1.0$, pooled across animals and brain areas, $n = 84$ brain areas, 2 per dataset; paired one-tailed t-test). Bottom, an additional simulation was performed in which the value of θ_{sim} was drawn uniformly and independently for each dataset between 45° and 90° . The $\theta_{\text{estimated}}$ from neural recordings (**c**) was not statistically different from the $\theta_{\text{estimated}}$ from these simulated data ($p = 0.40$, pooled across animals and brain areas, $n = 84$ brain areas, 2 per dataset; paired two-tailed t-test). * Source data are provided as a Source Data file.

(Fig. 4b; $p = 2.6 \times 10^{-7}$, pooled across animals and brain areas; paired one-tailed t-test; see Supplementary Fig. 5 for a shuffle control). Thus, substantial across-area interactions were present in the PFC activity and identified by pCCA-FA (see also Supplementary Fig. 1e for the across-area correlation values computed using pCCA-FA), contrary to the intuition from measuring pairwise across-area r_{sc} . This discrepancy can be understood by considering the sign of pairwise correlations among neurons (Fig. 2e, see Math Note A of ref. 72). Mean r_{sc} averages over positive and negative correlations found in individual pairs of neurons, whereas %sv accounts for both positive and negative correlations among neurons. Note that the values of across- and within-area %sv (like mean r_{sc}) can depend on neuronal properties, including their physical distance and tuning preferences.

Having found that neurons tended to more strongly covary across areas than within each area, we next considered the complexity of these interactions. We found that the across-area d_{shared} was also larger than the within-area d_{shared} (Fig. 4c; $p = 3.1 \times 10^{-6}$, pooled across animals and brain areas; paired one-tailed t-test). This indicated that, in addition to being stronger, the interactions we identified across areas were more complex than the interactions within each area. We further note that these findings are not a consequence of statistical bias in measuring across- versus within-area %sv and d_{shared} (Supplementary Fig. 2).

Across- and within-area components involved different co-fluctuation patterns

We next sought to understand the population-wide co-fluctuation patterns represented by each set of latent variables, as they provide insight into the nature of the co-fluctuations among neurons (Fig. 5a). There are many possible population-wide co-fluctuations. Perhaps the simplest example is all neurons sharing the same-signed loading onto a latent variable, such that the population would all increase their activity together or decrease their activity together. We examined whether this was the case in our PFC recordings. For this analysis, we focused on the first co-fluctuation pattern of each ordered set, as it reflected a large proportion of the across- or within-area variance (Supplementary Fig. 6). We found that all neurons increasing their activity together or decreasing their activity together was not consistent with the across- or within-area co-fluctuation patterns we identified (Fig. 5b, Supplementary Fig. 7).

We next asked whether across- and within-area co-fluctuation patterns were similar to each other by computing the angle between them in population activity space ($\theta_{\text{estimated}}$; Fig. 5a). They might be similar if two different processes (one global and one local) had a similar effect among the neurons. Alternatively, each process might engage distinct circuit mechanisms, yielding different co-fluctuation




Fig. 6 | Across-area latent variables were related to pupil diameter. **a** We examined whether across- (pink) or within-area (light and dark blue for right and left PFC, respectively) latent variables explained variance in the pupil diameter. **b** Fluctuations in pupil diameter (black, one normalized value per trial, see Methods) for an example session (Monkey Sa). For visual clarity, only 40/1258 trials are shown. Regression onto the pupil diameter using across-area latent variables (pink) or within-area latent variables identified by pCCA-FA (light and dark blue for right and left PFC, respectively) are overlaid. For this session, across-area latent variables explained more trial-to-trial variance in the pupil diameter ($r^2 = 0.35$) than within-area latent variables (right PFC: $r^2 = 0.013$, left PFC: $r^2 = 0.052$). **c** Aggregated across sessions, r^2 values for trial-by-trial pupil diameter were greater for across-area latent

variables (mean $r^2 = 0.25$) than for within-area latent variables (right PFC: mean $r^2 = 0.038$; left PFC: mean $r^2 = 0.030$) ($p = 1.3 \times 10^{-30}$; pooled across animals and brain areas, $n = 82$ brain areas, 2 per session; Monkey Pe: $p = 6.7 \times 10^{-12}$, $n = 32$; Monkey Wa: $p = 1.3 \times 10^{-8}$, $n = 18$; Monkey Sa: $p = 2.7 \times 10^{-21}$, $n = 32$; one-tailed t-tests on the differences in r^2 for each session). Error bars indicate the mean $\pm$ the standard error of the mean across sessions for each animal (Monkey Pe: $n = 16$, Monkey Wa: $n = 9$, Monkey Sa: $n = 16$). Gray bars indicate 95% confidence intervals of the shuffle distribution, obtained by regressing onto the trial-by-trial pupil diameter on session j using the latent variables on session i , where $i \neq j$. * Source data are provided as a Source Data file.

patterns. In an example session, we found that the across- and within-area co-fluctuation patterns appeared visually distinct, and produced a $\theta_{\text{estimated}}$ of 72° (Fig. 5b). In all sessions, we found that the across- and within-area co-fluctuation patterns had large $\theta_{\text{estimated}}$ (Fig. 5c), often close to 90° , which would indicate distinct co-fluctuation patterns.

To further assess the extent to which the measured $\theta_{\text{estimated}}$ reflected true differences in the across- and within-area co-fluctuation patterns, we created a simulation in which we defined the ground truth angle between across- and within-area co-fluctuation patterns (θ_{sim} , see Methods, Supplementary Fig. 3). We simulated activity in each area with a specified value of θ_{sim} that was otherwise matched to the neural recordings in terms of the number of neurons, latent variables, and trials (see Methods). We found that, for any given θ_{sim} , the resulting distributions of $\theta_{\text{estimated}}$ (Fig. 5d, top) were narrower than what we observed in the neural recordings (Fig. 5c). Instead, if we sampled a different value of θ_{sim} for each simulated dataset uniformly from 45° to 90° , we recovered a distribution of $\theta_{\text{estimated}}$ (Fig. 5d, bottom) that was similar to that observed in the neural recordings ($p = 0.40$, pooled across animals and brain areas; paired two-tailed t-test). We confirmed that the observed variation in $\theta_{\text{estimated}}$ from session to session was not related to the number of neurons or trials recorded (Supplementary Fig. 8). Furthermore, we determined that the variation in $\theta_{\text{estimated}}$ did not arise solely from sampling different subsets of neurons from session to session, although it still likely contributed (Supplementary Fig. 9). Together, these results indicate that across- and within-area

components involved different co-fluctuation patterns that varied in their similarity from session to session.

Across-area latent variables were related to pupil diameter

Given that we found substantial across-area variance (Fig. 4) involving different co-fluctuation patterns than the within-area component (Fig. 5), we considered whether the across-area variance might be related to a brainwide process. Arousal is known to modulate activity of neurons throughout the brain^{53,54,76–78}. In addition, arousal state has been linked to changes in both the tonic diameter of the pupil^{79,80} and event-related changes in pupil diameter^{81,82}. Changes in pupil diameter have also been linked to arousal-related changes in neural activity in both cortical^{83–86} and subcortical^{79,87–89} regions. Given that arousal has been found to be a brainwide phenomenon, we wondered if it was related to the across-area variance we observed. Recall that the across-area component is defined as variance that is shared with one or more neurons in the other area (regardless of whether that variance is also shared with neurons in the same area). We therefore asked how well the across- and within-area latent variables identified by pCCA-FA explained variance in the pupil diameter (Fig. 6a).

To address this question, we measured the trial-by-trial pupil diameter during the delay period and adjusted the pupil diameter values to remove slow fluctuations in an analogous manner to the neural activity, yielding one pupil diameter measure per trial (see Methods, Supplementary Fig. 10a, b). We used the across- or within-area latent variables identified by pCCA-FA (i.e., either d across-area

latent variables or d_m within-area latent variables for area m ; see Methods) to compute a trial-by-trial regression of the pupil diameter for each set of latent variables (Fig. 6b). Across sessions, we found that the across-area latent variables explained substantially more variance in the pupil diameter (mean $r^2 = 0.25$) than the within-area latent variables (right PFC, mean $r^2 = 0.038$; left PFC, mean $r^2 = 0.030$) (Fig. 6c; $p = 1.3e - 30$, pooled across animals and brain areas; one-tailed t-test). Furthermore, we found that the across-area latent variables explained variance in the event-related changes in pupil diameter observed just after target onset (Supplementary Fig. 10; $p = 1.1e - 5$, pooled across animals and brain areas; one-tailed t-test). Thus, the across-area latent variables may be related to multiple aspects of pupil-linked arousal. Taken together, these results are consistent with a link between the coordinated activity across brain areas and global cognitive processes such as arousal.

Discussion

In this work, we asked to what extent neural activity is shared across versus within hemispheres of PFC by considering how the spike count variance of the two brain areas covaried from trial-to-trial. Applying a novel statistical approach called pCCA-FA, we found that of the substantial shared trial-to-trial variability among PFC neurons, a large proportion involved neurons in both areas. The component of neural activity shared among neurons in both areas (across-area variance) was larger in magnitude (%sv) and complexity (d_{shared}) than the component shared only among neurons in one area (within-area variance), in contrast to what was suggested by mean r_{sc} analyses. Additionally, we found that the across-area component was related to pupil diameter, a common readout of arousal state. Our approach thus leveraged multi-area recordings in PFC to uncover signatures of global cognitive processing that were otherwise hidden in pairwise and single-area analyses.

Across-area variance may arise from a variety of circuit motifs. First, it may come from circuits that provide common inputs to both hemispheres of PFC. One such pathway involves the encoding of brainwide cognitive states, such as arousal. Arousal state modulation is thought to propagate through brainwide signaling that originates in the locus coeruleus^{90,91}. The locus coeruleus projects to both hemispheres of PFC^{92,93} and influences pupil diameter⁹⁴, a common readout of arousal state. Correlations between neural activity and pupil diameter have been observed across the brain^{36,43,54,95}, supporting the notion that the across-area variance we observed in PFC may be in part due to common inputs originating in the locus coeruleus. Common input to both hemispheres of PFC may also arise through circuit motifs involved in sensory processing and motor control. Each hemisphere of PFC represents both visual fields¹⁷, and common motor plans have been shown to impact neural activity across the brain^{36,41,43}. Second, variance shared across hemispheres may arise from direct anatomical projections⁹⁶, which have been shown to be spatially patterned⁹⁷. Interactions across the hemispheres are central to various brain functions such as working memory^{51,52}, motor control^{141,42}, sensory perception⁴⁰, and spatial attention^{56,57}.

Although we emphasized the variance shared across areas, we also identified a substantial amount of variance that was shared only among the neurons within each area. We labeled this variance as within-area, but we note that it could be shared with other non-recorded brain areas. The within-area variance we observed may originate through a variety of circuit motifs. First, representations of visual fields in PFC are topographically arranged with a contralateral bias, resulting in each hemisphere of PFC encoding different, although partially overlapping, information^{98–100}. Second, internal states like spatial attention tie neurons together with nearby receptive fields, especially those occupying the same hemisphere of visual space. For example, previous studies of visual area V4 during spatial attention tasks have identified attention-related components of neural activity

that are hemisphere-specific^{56,57}. This type of selective modulation may account for some of the within-area variance we report here. Third, within-area variance may result from constraints on patterns of neural activity imposed by the cortical circuitry (e.g., horizontal connections) in each hemisphere^{101,102}. In network models, spatial clustering architecture that may underlie local circuits within a hemisphere can have a profound influence on trial-to-trial variability^{24–26,103,104}.

Underlying the across- and within-area circuit motifs are various cellular and circuit mechanisms, which have been implicated in the shaping of correlations among neurons both across and within brain areas. First, the structure of inhibition is known to shape correlations among neurons, where the influence of inhibition depends on its spatial scale^{26,105} and timing relative to excitation^{106–109}. Second, changes at the level of the synapse, including synaptic depression¹¹⁰ and connection strength¹¹¹, can also regulate correlated variability. Third, neuronal correlations are influenced by neuromodulators such as acetylcholine^{112–114} and norepinephrine⁹¹, which together have wide-ranging modulatory effects including altering the strength of thalamocortical inputs¹¹⁵ and regulating global brain state¹¹⁶. Finally, the strength of excitatory synaptic connectivity has been linked to the strength of correlation¹¹⁷. Anatomically, long-range horizontal connections within a brain area^{118–120}, across brain areas⁹⁷, and across the corpus callosum⁹⁶, are all weaker than local recurrent connectivity. However, the strength of the detected correlations need not be directly linked to the strength of the anatomical connection between hemispheres or among distant neurons¹²¹. For example, weak neuromodulatory inputs could drive strong correlations among neurons by influencing the much stronger local circuit mechanisms that exist in local circuitry. Our analyses using pCCA-FA, like any other assessment based solely on spiking activity, do not explicitly distinguish between these cellular and circuit mechanisms. Instead, we focused on the spiking variability yielded by these mechanisms and dissected the components of spiking variability that are shared across versus within brain areas.

With the identification of within-area variance, pCCA-FA allowed us to isolate components of neural activity that were not shared between the two hemispheres of PFC. Previous work has described a gating of signals across areas via a communication subspace¹²². pCCA-FA identified what would be termed a communication subspace between the two hemispheres of PFC: we found distinct across- and within-area co-fluctuation patterns, indicating that not all signals were shared between areas. In other words, the existence of the within-area latent variables in our PFC recordings indicated that a communication subspace exists between the hemispheres.

pCCA-FA improves upon CCA and its probabilistic variant (pCCA) in identifying across-area interactions. First, CCA and pCCA are designed only to identify across-area variance and do not partition the remaining variance into a component that is shared among neurons within each area from variance that is independent to each neuron. Thus, pCCA-FA provides a more thorough dissection of the population activity in two brain areas. Second, even in cases where only across-area variance is of interest, pCCA-FA produces a more accurate estimate of across-area variance than pCCA (or CCA), especially in the limited data regime (Fig. 3e, f). This is due to the pCCA-FA model having a smaller number of parameters than pCCA (see Eqs. (2) and (4), and Supplementary Fig. 4). The tradeoff is the increased computation time of pCCA-FA compared with pCCA for jointly identifying the across- and within-area dimensionalities in pCCA-FA compared to identifying only the across-area dimensionality in pCCA (or CCA). This can be partially mitigated by employing efficient hyperparameter optimization techniques^{123,124}. For these reasons, we advocate for the use of pCCA-FA in the place of pCCA or CCA for studying interactions between brain areas.

pCCA-FA fills an unmet need in the set of exploratory tools tailored for studying the trial-to-trial variability of neural activity. When

studying neurons from a single brain area, FA is among the simplest approaches for separating variance that is shared among neurons from variance that is independent to each neuron. For two brain areas, pCCA-FA is among the simplest methods for separating variance that is shared across areas from variance that is only shared among neurons within each area. To incorporate the concept of time, one can consider using Delayed Latents Across Groups (DLAG)¹²⁴ or related methods^{32,125–130} to separate activity time courses that are shared across areas from those that are shared only among neurons within each area. pCCA-FA can be extended in the following ways. First, one may seek to identify latent variables that are explicitly related to task stimuli or the animal's behavior by using a targeted dimensionality reduction approach^{95,131–134} or by applying pCCA-FA to raw spike counts. Second, one may seek to identify non-linear interactions between brain areas^{135–138}. As multi-area population recordings continue to increase in prevalence^{28–30}, pCCA-FA and related methods will become increasingly essential for dissecting the rich structure of across- and within-area interactions.

Methods

Animal model

Three male rhesus macaque monkeys (*Macaca mulatta*) were used in this study. One animal was sourced from Tulane National Primate Research Center (Monkey Wa) and two animals were from Covance (Monkey Pe and Monkey Sa). They were 10 (Monkey Pe), 7 (Monkey Wa), and 7 (Monkey Sa) years old at the time of data collection. All experimental procedures were approved by the Institutional Animal Care and Use Committees of Carnegie Mellon University and the University of Pittsburgh, and were in accordance with the United States National Research Council's Guide for the Care and Use of Laboratory Animals.

Surgical preparation

Each animal was surgically implanted with a titanium headpost, which was fixed onto the skull using titanium screws. This was necessary to limit head movement such that the position of the eyes could be tracked and neural activity recorded. Animals were then each implanted with two 96-electrode "Utah" arrays (Blackrock Neurotech, Salt Lake City, UT). Electrode arrays were placed on the prearcuate gyrus just anterior to the arcuate sulcus, one in each hemisphere. Surgeries were performed in sterile conditions under general anesthesia using isoflurane.

Electrophysiological methods

Animals were positioned 36 cm from a 21" cathode ray tube monitor with 1024 × 768-pixel resolution and a refresh rate of 100 Hz. Task stimuli were generated using custom-designed software in Matlab (Mathworks, Natick, MA) that utilized the Psychophysics toolbox extensions^{139–141}. Signals from the implanted electrodes were band-pass filtered (0.3–7500 Hz) and then digitized at 30,000 Hz. For each electrode, spiking waveforms were defined as a 52-sample (1.73 ms) window of the filtered voltage signal triggered by the signal crossing a predefined threshold, and stored for offline analysis. The threshold was defined as a multiple of the root-mean-square voltage of a brief epoch of the raw signal on each electrode collected at the beginning of each session. All behavioral data and neural activity were recorded using a Grapevine recording system (Ripple, Salt Lake City, UT) for further offline processing and analysis.

Behavioral task

Animals were trained to perform a standard memory-guided saccade task¹⁴². At the beginning of each trial, a 0.5 degree diameter blue fixation circle appeared at the center of a gray screen. The animal initiated fixation within an invisible window (Monkey Pe: 2.4, Monkey Wa: 1.8, Monkey Sa: 2.9 degrees radius) centered on the fixation circle

and then 200 ms later a white target flashed in the animal's periphery (Monkey Pe: 11.8 or 17.5 degrees from fixation depending on the session, Monkey Wa: 11.6 degrees from fixation, Monkey Sa: a single value in the range of 8.0–15.6 degrees from fixation depending on the session, with one session containing multiple randomly interleaved distances in this range) at one of 4, 8, or 16 locations depending on the session. The white target remained on the screen for either 100, 200, or 400 ms (depending on the session) after which it was removed from the screen. The animal maintained fixation for a delay period drawn randomly from a uniform distribution between 1.5 and 3 seconds. After the delay, the fixation circle disappeared, which served as the go cue for the animal to initiate a saccade toward the remembered location where the white target flash had appeared. The animal had 500 ms (Monkey Sa: 400 ms) to initiate a saccade, and an additional 200 ms to reach the target location. Successful completion of a saccade was defined by the animal maintaining fixation within an invisible window (Monkey Pe: 4.1, Monkey Wa: 5.9 or 7.0, Monkey Sa: 4.1 degrees radius) centered at the target location for at least 100 ms. A liquid reward was delivered for a saccade to the correct location. For a subset of sessions, a dim white circle was flashed at the target location after saccade initiation to assist the animal in target acquisition. Trials were pseudo-randomized in blocks during which the animal was required to correctly complete all target directions before beginning a new block. In total, we collected 42 sessions of data (16 in Monkey Pe, 10 in Monkey Wa, and 16 in Monkey Sa). For each session, we considered only trials in which the animal successfully acquired the saccade target (i.e., correct trials). The number of correct trials per session for Monkey Pe was 729.8 ± 173.4 (mean ± standard deviation), for Monkey Wa was 463.5 ± 170.1, and for Monkey Sa was 850.1 ± 284.2.

Preprocessing of neural activity

We used a previously trained and published neural network to classify waveforms as spikes or noise¹⁴³. Briefly, the network was trained on human-sorted waveforms to distinguish between waveforms putatively of neural origin and waveforms not of neural origin. Using this network, we removed threshold crossings that were unlikely to be of neural origin, and refer to the remaining waveforms as spikes. The spikes from each electrode on the Utah array were treated as multi-unit activity, and no further differentiation of single units was performed. Thus, each electrode was treated as one neural unit. Using multi-unit activity tends to yield higher correlation values than single unit activity⁸, and this applies equally to the across-area and within-area measurements.

We removed electrodes that had very low firing rates or very high trial-to-trial variability. To do this, we first binned neural activity by counting spikes during the target presentation and delay period between target onset and saccade initiation cue. We then removed electrodes with mean spike count lower than 2 spikes/second or Fano factor greater than 10 during this time window. This served to eliminate electrodes contaminated by recording artifacts, or units that were not present for the entirety of a recording session. We also removed electrodes affected by artifactual cross-talk due to electrical coupling. For each pair of electrodes, we flagged spikes as coincident if they occurred within 100 μs of each other. If either electrode in the pair had 20% or more of its spikes flagged as coincident, we flagged that pair as having artifactual crosstalk. This procedure provided an index for each electrode of how many other electrodes it was coincident with. We then iteratively removed the electrode coincident with the highest number of other electrodes until no crosstalk remained, which ensured that the fewest possible number of electrodes were removed. After this process the number of remaining electrodes in Monkey Pe was 78.3 ± 7.8 and 79.5 ± 8.5 (mean ± standard deviation), in Monkey Wa was 85.3 ± 8.1 and 25.2 ± 4.4, and in Monkey Sa was 62.6 ± 9.8 and 75.3 ± 13.2 for right and left PFC respectively.

After the above preprocessing steps, the spiking activity was binned into spike counts. The bin was defined as the one second of the delay period immediately preceding the earliest possible saccade initiation cue in the session, such that each trial contributed one bin. For example, if the earliest possible saccade initiation cue in a given session was 1.5 s after the target disappeared, the bin would be defined as 0.5–1.5 s after the target disappeared for every trial in that session. This one-second window aggregated the interactions across and within brain areas observed during the delay period on the trial-to-trial timescale.

Computing signal correlation

In Fig. 2e, we investigated if the distribution of across-area r_{sc} , despite having a mean near zero, contained pairs of neurons with positive r_{sc} and other pairs of neurons with negative r_{sc} , in a way that was systematically related to the neurons' task tuning. To do this, we identified each neuron's tuning across all of the saccade target locations. We determined whether each pair of neurons had similar or opposite tuning by computing their signal correlation. This was defined as the Pearson correlation between the two neurons' raw trial-averaged responses to each of the possible targets (4, 8, or 16 depending on the session) during the one-second window described above. To assess significance for each neuron pair, we generated a chance distribution of correlation values by using a permutation test (shuffling the target location labels). We then labeled a neuron pair as having significant signal correlation (positive for similarly tuned and negative for oppositely tuned pairs) if their signal correlation was more extreme than the 99th percentile of the chance distribution. Neuron pairs not meeting this criterion were excluded from this analysis.

Combining trials across conditions

To increase the number of observations for model fitting, we combined trials across conditions (saccade target direction). To focus our analyses on trial-to-trial variability, for each neuron, we first subtracted the condition mean from the binned spike counts in each condition and then combined residual spike counts across conditions. Because the spike count variance can differ across conditions, the residual spike counts can still contain a small amount of condition-specific signal. An alternative approach is to z-score the spike counts for each neuron and condition before combining across conditions⁶³. We repeated the key analyses related to the identification of shared variance among neurons in Figs. 2 and 4 using condition-specific z-scoring and found qualitatively similar results. Note, the computation of signal correlation (above) was performed before the subtraction of condition-specific means.

Separation of slow and fast components

In this work, we were interested in trial-to-trial variability of neural activity on a timescale of seconds. As we investigated interactions between left and right PFC, it soon became apparent that both areas also contained a component that varied slowly on a timescale of minutes to hours. While slow-timescale variability presents its own set of scientific questions (e.g., ref. 54), these fluctuations complicated our analyses by introducing dependencies between one trial and the next, violating the basic independence assumption of correlation analysis (including regression, Pearson correlation, pCCA, and pCCA-FA). Therefore, we removed the slow component from the neural activity (see Supplementary Fig. 1), and focused all analyses in this work on faster-timescale trial-to-trial variability, except where noted in Supplementary Fig. 1 and for computing signal correlation. We computed the slow component for each session using a centered boxcar filter of length 25 trials, computed for each remaining electrode after the preprocessing described above. We then subtracted this component from the residual spike counts to remove slow-timescale trends that could have induced spurious correlations⁶¹. We performed this

preprocessing procedure independently for each neuron. All analyses in this study were performed on the residual faster-timescale component, except where noted.

Preprocessing of pupil diameter

Eye position and pupil diameter were monitored using monocular infrared tracking at a 1000 Hz sample rate (EyeLink 1000, RS Research, Mississauga, Canada). On each trial, we measured the pupil diameter during the delay period, from target offset until the saccade initiation cue. Each trial was smoothed using a median filter with a 50 ms window (i.e., the median value in a 50 ms centered window was subtracted from each sample) to remove high frequency noise.

For the analyses in Fig. 6, we computed mean smoothed pupil diameter (arbitrary units) over the full time window described above to obtain the most robust estimate of the pupil diameter on each trial. This yielded one average pupil diameter value per trial. Like the preprocessing of neural activity, we removed slow-timescale fluctuations in pupil diameter using a centered boxcar filter of length 25 trials. This yielded what we termed the trial-by-trial pupil diameter (Supplementary Fig. 10a, b).

To ensure pupil data quality, we computed the variance across trials of the trial-by-trial pupil diameter. We then removed outlier sessions that exceeded a variance threshold of two times the interquartile range of the variance values across sessions. This was performed for each animal individually. This procedure excluded one session from Monkey Wa from analyses of the pupil (Fig. 6, Supplementary Fig. 10), resulting in $n = 9$ sessions for analysis. To ensure pupil diameter was on a similar scale across sessions, we z-scored the trial-by-trial pupil diameter for each session by subtracting the mean and dividing by the standard deviation across trials in the session.

Probabilistic canonical correlation analysis (pCCA)

A commonly used approach for identifying interactions between brain areas is canonical correlation analysis (CCA)³³. Briefly, CCA attempts to find a co-fluctuation pattern in area 1 and a co-fluctuation pattern in area 2, such that neural activity along these co-fluctuation patterns is maximally correlated across areas. Further co-fluctuation patterns can be found subject to the constraint that they are uncorrelated with previous co-fluctuation patterns. The correlations along these co-fluctuation patterns are known as canonical correlations (ρ). Note that CCA identifies activity shared across areas, but is not designed to partition activity into across- and within-area components.

The probabilistic interpretation of CCA (pCCA⁷⁴) is defined by the following model. pCCA defines a linear-Gaussian relationship between the spike count vectors $\mathbf{x}_1 \in \mathbb{R}^{n_1}$ from area 1 and $\mathbf{x}_2 \in \mathbb{R}^{n_2}$ from area 2 and the across-area latent variables $\mathbf{z} \in \mathbb{R}^d$ as:

$$\begin{aligned}\mathbf{z} &\sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d) \\ \mathbf{x}_1 | \mathbf{z} &\sim \mathcal{N}(\boldsymbol{\mu}_1 + \mathbf{W}_1 \mathbf{z}, \boldsymbol{\Psi}_1) \\ \mathbf{x}_2 | \mathbf{z} &\sim \mathcal{N}(\boldsymbol{\mu}_2 + \mathbf{W}_2 \mathbf{z}, \boldsymbol{\Psi}_2),\end{aligned}\quad (1)$$

where n_1 and n_2 are the number of neurons recorded in areas 1 and 2, respectively, and d is the number of across-area latent variables. The identity matrix $\mathbf{I}_d \in \mathbb{R}^{d \times d}$ describes prior covariance of the latent variables. The vectors $\boldsymbol{\mu}_1 \in \mathbb{R}^{n_1}$ and $\boldsymbol{\mu}_2 \in \mathbb{R}^{n_2}$ represent the mean spike counts of areas 1 and 2, respectively. The loading matrices $\mathbf{W}_1 \in \mathbb{R}^{n_1 \times d}$ and $\mathbf{W}_2 \in \mathbb{R}^{n_2 \times d}$ map the latent variables onto neural activity in areas 1 and 2 respectively, and $\boldsymbol{\Psi}_1 \in \mathbb{R}^{n_1 \times n_1}$ and $\boldsymbol{\Psi}_2 \in \mathbb{R}^{n_2 \times n_2}$ are full covariance matrices.

From Eq. (1), the marginal distributions for $\mathbf{x}_1$ and $\mathbf{x}_2$ are:

$$\begin{aligned}\mathbf{x}_1 &\sim \mathcal{N}(\boldsymbol{\mu}_1, \mathbf{W}_1 \mathbf{W}_1^T + \boldsymbol{\Psi}_1) \\ \mathbf{x}_2 &\sim \mathcal{N}(\boldsymbol{\mu}_2, \mathbf{W}_2 \mathbf{W}_2^T + \boldsymbol{\Psi}_2)\end{aligned}\quad (2)$$

We can see that pCCA decomposes the covariance of neural activity from each area into a low-rank component shared across areas (e.g., $W_1 W_1^T$) and a full noise covariance matrix (e.g., Ψ_1). Note that pCCA treats all variance which cannot be explained by across-area latent variables ($\mathbf{z}$), including within-area interactions and variance independent to each neuron, as noise.

It can be shown that pCCA (Eq. (1)) finds the same co-fluctuation patterns (described by W_1 and W_2) as CCA that maximize correlation across areas⁷⁴. CCA and pCCA therefore identify the same canonical correlations.

Probabilistic canonical correlation analysis - factor analysis (pCCA-FA)

We developed a model called pCCA-FA (a combination of probabilistic canonical correlation analysis and factor analysis (FA¹⁴⁴)), which explicitly partitions covariance of neural activity into an across-area component, a within-area component, and a component independent to each neuron. Though we term the components across-area and within-area, the pCCA-FA framework can be applied to any two simultaneously-recorded populations of neurons. pCCA-FA is related to extensions of FA that have been proposed in other contexts, such as refs. 145 and 146.

pCCA-FA defines the linear-Gaussian relationship between the spike count vectors $\mathbf{x}_1 \in \mathbb{R}^{n_1}$ from area 1 and $\mathbf{x}_2 \in \mathbb{R}^{n_2}$ from area 2 and across-area latent variables $\mathbf{z} \in \mathbb{R}^d$ and within-area latent variables $\mathbf{z}_1 \in \mathbb{R}^{d_1}$ and $\mathbf{z}_2 \in \mathbb{R}^{d_2}$ as:

$$\begin{aligned} \mathbf{z} &\sim N(\mathbf{0}, I_d) \quad \mathbf{z}_1 \sim N(\mathbf{0}, I_{d_1}) \quad \mathbf{z}_2 \sim N(\mathbf{0}, I_{d_2}) \\ \mathbf{x}_1 | \mathbf{z}, \mathbf{z}_1 &\sim N(\boldsymbol{\mu}_1 + W_1 \mathbf{z} + L_1 \mathbf{z}_1, \Psi_1) \\ \mathbf{x}_2 | \mathbf{z}, \mathbf{z}_2 &\sim N(\boldsymbol{\mu}_2 + W_2 \mathbf{z} + L_2 \mathbf{z}_2, \Psi_2) \end{aligned} \quad (3)$$

where d is the number of across-area latent variables, and d_1 and d_2 are the number of within-area latent variables for areas 1 and 2, respectively. The identity matrices $I_d \in \mathbb{R}^{d \times d}$, $I_{d_1} \in \mathbb{R}^{d_1 \times d_1}$, and $I_{d_2} \in \mathbb{R}^{d_2 \times d_2}$ describe the prior covariance of each set of latent variables. Given n_1 neurons from area 1, $W_1 \in \mathbb{R}^{n_1 \times d}$ is the loading matrix for the across-area component in area 1, and $L_1 \in \mathbb{R}^{n_1 \times d_1}$ is the loading matrix for the within-area component in area 1. The matrix $\Psi_1 \in \mathbb{R}^{n_1 \times n_1}$ is a diagonal matrix containing the independent variance of each neuron in area 1, and $\boldsymbol{\mu}_1 \in \mathbb{R}^{n_1}$ represents the mean spike counts of area 1. The parameters for area 2 are defined analogously.

From Eq. (3), the marginal distributions for $\mathbf{x}_1$ and $\mathbf{x}_2$ are:

$$\begin{aligned} \mathbf{x}_1 &\sim N(\boldsymbol{\mu}_1, W_1 W_1^T + L_1 L_1^T + \Psi_1) \\ \mathbf{x}_2 &\sim N(\boldsymbol{\mu}_2, W_2 W_2^T + L_2 L_2^T + \Psi_2) \end{aligned} \quad (4)$$

We can observe that pCCA-FA decomposes the covariance of neural activity in area 1 as the sum of an across-area component ($W_1 W_1^T$), a within-area component ($L_1 L_1^T$), and an independent neuron component (Ψ_1). The covariance of neural activity in area 2 is decomposed into analogous across-area ($W_2 W_2^T$), within-area ($L_2 L_2^T$), and independent neuron (Ψ_2) components. In the pCCA-FA model, variance shared with neurons in the other area is assigned to the across-area component (appearing in the pCCA-FA model as W_m for each area m , and corresponding to latent variable $\mathbf{z}$, which involves both areas), regardless of whether that variance is also shared with neurons in the same area. Consequently, variance is assigned to the within-area component only when it is not shared with any neuron in the other area (appearing in the pCCA-FA model as L_m , and corresponding to latent variable $\mathbf{z}_m$, which only involves area m).

pCCA-FA can be viewed as a combination of two existing latent variable methods, namely pCCA, which finds co-fluctuation patterns that maximize correlation between two populations, and FA, which maximizes covariance between neurons within a single population. Importantly, pCCA-FA explicitly identifies both types of covariance in a

single probabilistic framework. The key difference between pCCA-FA and pCCA is that pCCA-FA constrains the full covariance matrices identified by pCCA (e.g., Ψ_1 in Eq. (2)) as low-rank plus diagonal (e.g., $L_1 L_1^T + \Psi_1$ in Eq. (4), see Supplementary Fig. 4). It follows that pCCA-FA has fewer parameters than pCCA, and is thus more robust in the regime of low trial counts (i.e., able to recover ground truth with fewer trials; Fig. 3e, f). With sufficient trials, pCCA-FA and pCCA identify the same across-area co-fluctuation patterns.

Fitting pCCA-FA

To fit pCCA-FA to neural activity, we define a joint vector of neural activity in both areas ($\tilde{\mathbf{x}}$) and a joint vector of across- and within-area latent variables ($\tilde{\mathbf{z}}$):

$$\tilde{\mathbf{x}} = \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix} \quad \tilde{\mathbf{z}} = \begin{bmatrix} \mathbf{z} \\ \mathbf{z}_1 \\ \mathbf{z}_2 \end{bmatrix}$$

The probabilistic model of pCCA-FA from Eq. (3) can be rewritten as:

$$\begin{aligned} \tilde{\mathbf{z}} &\sim N(\mathbf{0}, I_{d+d_1+d_2}) \\ \tilde{\mathbf{x}} | \tilde{\mathbf{z}} &\sim N(\tilde{\boldsymbol{\mu}} + \tilde{L} \tilde{\mathbf{z}}, \tilde{\Psi}) \end{aligned} \quad (5)$$

where

$$\tilde{\boldsymbol{\mu}} = \begin{bmatrix} \boldsymbol{\mu}_1 \\ \boldsymbol{\mu}_2 \end{bmatrix} \quad \tilde{L} = \begin{bmatrix} W_1 & L_1 & \mathbf{0} \\ W_2 & \mathbf{0} & L_2 \end{bmatrix} \quad \tilde{\Psi} = \begin{bmatrix} \Psi_1 & \mathbf{0} \\ \mathbf{0} & \Psi_2 \end{bmatrix}$$

and $\boldsymbol{\mu}_1$, $\boldsymbol{\mu}_2$, W_1 , W_2 , L_1 , L_2 , Ψ_1 , and Ψ_2 are defined as in Eq. (3). The formulation in Eq. (5) makes clear how we can think of pCCA-FA as a constrained and structured FA model. The zeros in the loading matrix $\tilde{L}$ ensure that the within-area latent variables ($\mathbf{z}_1$ and $\mathbf{z}_2$) are only used to explain activity in their respective areas, whereas across-area latent variables ($\mathbf{z}$) are used to explain activity in both areas. Based on Eq. (5), the marginal distribution of the activity in both areas is:

$$\tilde{\mathbf{x}} \sim N(\tilde{\boldsymbol{\mu}}, \tilde{L} \tilde{L}^T + \tilde{\Psi}) \quad (6)$$

We fit pCCA-FA model parameters to neural activity using an exact expectation-maximization (EM) algorithm¹⁴⁷. The EM algorithm for pCCA-FA resembles that used for FA. Specifically, the M-step update equations for the parameters $\tilde{L}$ and $\tilde{\Psi}$ mirror those derived in FA for the loading matrix and noise covariance matrix. Each M-step parameter update for pCCA-FA (i.e., for parameters W_1 , L_1 , etc.) effectively applies a zeroing mask to the contents of $\tilde{L}$, ensuring that the blocked structure of the loading matrix is maintained. Additionally, there is no constraint on the relationship between the across- and within-area co-fluctuation patterns (e.g., they are not constrained to be orthogonal).

The columns of each loading matrix (W_m and L_m for each area m) as estimated using the EM algorithm are not ordered. As a post-processing step for all analyses, we applied the singular value decomposition to transform each loading matrix, such that the first column contained the co-fluctuation pattern that explained the most shared variance, the second column explained the second-most shared variance, etc. For across-area co-fluctuation patterns, it may also be desirable to order the columns according to shared correlation (as in CCA, see below).

For each session, we jointly chose the dimensionalities for across- (d) and within-area (d_1 and d_2) latent variables using 10-fold cross-validation and selected the dimensionalities which maximized the cross-validated data likelihood $P(\tilde{\mathbf{x}})$ (see Eq. (6)). We tested integer dimensionality values between 1 and 15, inclusive, for both across- and within-area latent variables for each cross-validation run. After across- and within-area dimensionalities were selected, a

final pCCA-FA model was fit with the specified dimensionalities using all trials. The parameters of this final model were used in all analyses.

Estimating latent variables in pCCA-FA

To estimate the across- and within-area latent variables in pCCA-FA (e.g., for use in predicting pupil diameter in Fig. 6), we computed the posterior mean of latent variables:

$$\mathbb{E}[\tilde{\mathbf{z}}|\tilde{\mathbf{x}}] = \tilde{\mathbf{L}}^T (\tilde{\mathbf{L}}\tilde{\mathbf{L}}^T + \tilde{\Psi})^{-1} (\tilde{\mathbf{x}} - \tilde{\boldsymbol{\mu}}) \quad (7)$$

where the first d entries of $\mathbb{E}[\tilde{\mathbf{z}}|\tilde{\mathbf{x}}]$ are $\mathbb{E}[\mathbf{z}|\mathbf{x}_1, \mathbf{x}_2]$ (across-area latent variables), the next d_1 entries are $\mathbb{E}[\mathbf{z}_1|\mathbf{x}_1]$ (within-area latent variables for area 1), and the final d_2 entries are $\mathbb{E}[\mathbf{z}_2|\mathbf{x}_2]$ (within-area latent variables for area 2).

Computation of canonical variables

The pCCA-FA parameters allow for the computation of the canonical variables analogous to those found for CCA (Supplementary Fig. 6a). Following ref. 74, the canonical directions and correlations can be obtained applying the singular value decomposition on the cross-correlation matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$:

$$\mathbf{K} = \Sigma_{11}^{-1/2} \Sigma_{12} \Sigma_{22}^{-1/2}$$

where n is the total number of neurons (i.e., $n_1 + n_2$),

$$\begin{aligned} \Sigma_{11} &= \mathbf{W}_1 \mathbf{W}_1^T + \mathbf{L}_1 \mathbf{L}_1^T + \Psi_1 \\ \Sigma_{22} &= \mathbf{W}_2 \mathbf{W}_2^T + \mathbf{L}_2 \mathbf{L}_2^T + \Psi_2 \\ \Sigma_{12} &= \mathbf{W}_1 \mathbf{W}_2^T \end{aligned}$$

are the blockwise entries of $\tilde{\mathbf{L}}\tilde{\mathbf{L}}^T + \tilde{\Psi}$. The canonical pairs of directions ($\mathbf{u}_{1i} \in \mathbb{R}^{n_1}, \mathbf{u}_{2i} \in \mathbb{R}^{n_2}$) for $i = 1, 2, \dots, d$, are defined as:

$$(\mathbf{u}_{1i}, \mathbf{u}_{2i}) = (\Sigma_{11}^{-1/2} \mathbf{v}_{1i}, \Sigma_{22}^{-1/2} \mathbf{v}_{2i}),$$

where ($\mathbf{v}_{1i} \in \mathbb{R}^{n_1}, \mathbf{v}_{2i} \in \mathbb{R}^{n_2}$) are a pair of left-right singular vectors of $\mathbf{K}$. The corresponding singular value ρ_i is the canonical correlation.

Note that the Σ matrices are defined in terms of the estimated pCCA-FA parameters rather than the sample covariance matrices (which is what CCA and pCCA use). This is the regularization that pCCA-FA provides that enables it to perform well in a limited data regime (Fig. 3e, f).

Measuring %sv and d_{shared}

We defined two metrics to summarize the outputs of pCCA-FA: the percentage of variance shared among neurons (%sv) and shared dimensionality (d_{shared}), which were originally introduced for FA^{25,72} and adapted here for pCCA-FA.

%sv. We used %sv to measure the prominence of (i.e., the proportion of variance explained by) the across- or within-area components. For a given pCCA-FA model, there are four instances of %sv: an across-area and a within-area instance, for each of the two areas. The across-area %sv quantifies the percentage of each neuron's variance that was explained by across-area latent variables (i.e., the percentage of each neuron's activity that was shared with at least one other neuron in the other area). More specifically:

$$\text{Across-area \%sv for neuron } k \text{ in area } m = \frac{\mathbf{w}_{mk} \mathbf{w}_{mk}^T}{\mathbf{w}_{mk} \mathbf{w}_{mk}^T + \mathbf{l}_{mk} \mathbf{l}_{mk}^T + \psi_{mk}} \quad (8)$$

where $\mathbf{w}_{mk}$ is the k^{th} row of the across-area loading matrix for area m , $\mathbf{l}_{mk}$ is the k^{th} row of the within-area loading matrix for area m , and ψ_{mk} is the independent variance for the k^{th} neuron in area m (i.e., the k^{th} entry of the diagonal of Ψ_m). The numerator represents the across-area variance for neuron k in area m , and the denominator represents the total variance for that neuron. In our analyses of %sv, we averaged over neurons in a brain area to obtain a single across-area %sv metric for each area.

Similarly, the within-area %sv quantifies the percentage of each neuron's variance that was explained by within-area latent variables (i.e., the percentage of each neuron's activity that was shared with at least one other neuron in the same area, but was not shared with any neurons in the other area). More specifically:

$$\text{Within-area \%sv for neuron } k \text{ in area } m = \frac{\mathbf{l}_{mk} \mathbf{l}_{mk}^T}{\mathbf{w}_{mk} \mathbf{w}_{mk}^T + \mathbf{l}_{mk} \mathbf{l}_{mk}^T + \psi_{mk}} \quad (9)$$

d_{shared} . We used the metric, d_{shared} , to measure the complexity of across- and within-area interactions. For a given pCCA-FA model, there are four instances of d_{shared} , an across-area and a within-area instance, for each of the two areas. Note that d_{shared} measures the dimensionality of interactions shared among neurons and excludes activity related to each neuron's independent variance. We measured d_{shared} in an analogous way to ref. 25, but adapted for across- and within-area dimensionalities.

Briefly, the first step in measuring across-area d_{shared} is to fit the loading matrices $\mathbf{W}_m$ with d^* across-area latent variables, where d^* is the number of latent variables that maximizes the cross-validated data likelihood. Across-area d_{shared} is defined for each area m as the minimum number of latent variables needed to explain at least 95% of the across-area shared covariance in $\mathbf{W}_m \mathbf{W}_m^T$, where $d_{\text{shared}} \leq d^*$. A 95% threshold was used in order to focus on latent variables that explain a large amount of shared variance. Note that although there is a single set of d^* across-area latent variables, it is possible to obtain a different value of across-area d_{shared} for each area. We analogously defined within-area d_{shared} for each area using the procedure described above by replacing $\mathbf{W}_m \mathbf{W}_m^T$ with $\mathbf{L}_m \mathbf{L}_m^T$.

Simulated data generation

To validate the pCCA-FA model (Fig. 3e, f), as well as generate chance distributions (Fig. 5d), we generated simulated data $\mathbf{x}_m$ for each area m from the pCCA-FA model according to Eq. (4). This involved defining ground truth model parameters $\boldsymbol{\mu}_m, \mathbf{W}_m, \mathbf{L}_m$, and Ψ_m for each area m .

For model validation in Fig. 3e, f, we defined these parameters in the following way. We generated random loading matrices $\mathbf{W}_m$ and $\mathbf{L}_m$ for each area m using independent draws from a standard normal distribution. We scaled the %sv in each loading matrix by adjusting the singular values of the across- and within-area loading matrices ($\mathbf{W}_m$ and $\mathbf{L}_m$) relative to each other to achieve the desired ratios ($\approx 15\%$ across-area and $\approx 10\%$ within-area). The independent variance values (i.e., entries of Ψ_m) were then scaled to achieve a level of 75%. Each simulated dataset had 30 neurons per area, 5 across-area latent variables ($d = 5$), and 3 within-area latent variables per area ($d_m = 3$ for each area m). For each trial count configuration, we generated 30 sets of ground truth parameters and a single simulated dataset per set of parameters.

For Fig. 5d, we defined the ground truth model parameters to be copies of those that were estimated from each of the neural recordings, with one manipulation, described below. For the chance distributions in Fig. 5d and the simulated data in Supplementary Fig. 3, we specified the angle between across- and within-area co-fluctuation patterns (θ_{sim}). To achieve this, we held the across-area loading matrix ($\mathbf{W}_m$) constant and set the first column of the within-area loading matrix ($\mathbf{L}_m$) to be a vector that was θ_{sim} degrees apart from the first column of $\mathbf{W}_m$. The remaining columns of the within-area loading

matrix were randomly generated from a standard normal distribution, and then orthogonalized relative to the first column (i.e., the first column remained fixed and the remaining columns were set to be mutually orthogonal to it).

Fitting pCCA-FA and pCCA to simulated data

To fit pCCA-FA to simulated data for the analysis in Fig. 3e, f, we used the EM algorithm described above. Across- and within-area dimensionalities were estimated in a cross-validated manner. Candidate dimensionalities for across- and within-area ranged from 1 to 10 (inclusive), for both across- and within-area latent variables. After the across- and within-area dimensionalities were selected, a final pCCA-FA model was fit with the specified dimensionalities using all simulated trials. The parameters of this final model were used to compute %sv and d_{shared} .

To fit pCCA, we used an analogous EM algorithm, presented in ref. 74. We fit pCCA to the same sets of simulated data as pCCA-FA. The across-area dimensionality was estimated in a cross-validated manner. Candidate dimensionalities for the across-area latent variables ranged from 1 to 10 (inclusive). After the across-area dimensionality was selected, a final pCCA model was fit with the specified dimensionality using all simulated trials. The parameters of this final model were used to compute across-area %sv and d_{shared} .

Pupil diameter regression

To explain pupil-related variance using the latent variables estimated by pCCA-FA, the across-area latent variables $\mathbb{E}[\mathbf{z}|\mathbf{x}_1, \mathbf{x}_2]$ and the within-area latent variables $\mathbb{E}[\mathbf{z}_1|\mathbf{x}_1]$ and $\mathbb{E}[\mathbf{z}_2|\mathbf{x}_2]$ were estimated using Eq. (7). We regressed either the across-area or the within-area (right or left PFC) latent variables onto the trial-by-trial pupil diameter using a linear regression model for each set of latent variables. We then assessed r^2 , the proportion of variance in the trial-by-trial pupil diameter explained by the regression model.

To compute chance distributions, we repeated the procedure above, except that we performed the regression using latent variables estimated from the neural activity from a different session than the one in which we measured pupil diameter (i.e., the latent variables from session i were used to explain variance in the pupil diameter from session j , where $i \neq j$). Trials were truncated in the session with more trials to ensure equal trial numbers. Chance distributions contained 240 samples for Monkeys Pe and Sa with 16 sessions, and 72 samples for Monkey Wa with 9 sessions. We report the 95% confidence intervals of this chance distribution (Fig. 6c, gray bars, termed shuffle). An alternative method to generate the chance distribution involves performing a within-session shuffle, in which the latent variables from session i are used to explain variance in the pupil diameter from session i but with permuted trial labels. This procedure yielded a similar, but slightly more lenient, chance distribution.

To control for differences in the number of across- and within-area latent variables, we repeated the entire procedure using only one predictor (i.e., the top across- or within-area latent variable) rather than the full set of predictors (i.e., d across- or d_m within-area latent variables). This yielded similar results ($p = 4.3e - 22$; pooled across animals and brain areas, $n = 82$ brain areas, 2 per session; one-tailed t-test on the differences in r^2 for each session).

Statistical tests and reproducibility

Data preprocessing was performed using custom code in MATLAB (Mathworks, Natick, MA). Criteria for inclusion of units are described in Methods section: Preprocessing of neural data. All analyses were performed in Python (version 3.8.18). Statistical tests (t-test) were performed using the SciPy (version 1.10.0, <https://scipy.org>) statistics module. Sample sizes and p -values are directly reported in the corresponding figure caption. Averaged data shown in figures are averaged across sessions ($n = 42$) or brain areas ($n = 84$, 2 per session), noted in

each caption. For data shown in histograms, means are marked with triangles. Detailed descriptions of plotted data are provided in each caption. Code to reproduce the figures and statistics is publicly available on GitHub at <https://github.com/meganmcd13/figures-dual-pfc>.

Reporting summary

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

Data availability

The preprocessed electrophysiological and pupillometry data that support the findings of this study have been deposited in the KiltHub database with the identifier <https://doi.org/10.1184/R1/31431415>¹⁴⁸. Source data are provided with this paper.

Code availability

Python code for fitting and using the pCCA-FA model is publicly available on GitHub at https://github.com/SmithLabNeuro/pcca_fa and archived on Zenodo at <https://doi.org/10.5281/zenodo.18913386>¹⁴⁹. Matlab and Python code for figure generation is publicly available on GitHub at <https://github.com/meganmcd13/figures-dual-pfc> and archived on Zenodo at <https://doi.org/10.5281/zenodo.18929016>¹⁵⁰. Matlab behavioral control software for running the animal experiments used in this study is publicly available at <https://github.com/SmithLabNeuro/Ex>.

References

- Petersen, S. E. & Posner, M. I. The attention system of the human brain: 20 years after. *Annu. Rev. Neurosci.* **35**, 73–89 (2012).
- Maunsell, J. H. R. Neuronal mechanisms of visual attention. *Annu. Rev. Vis. Sci.* **1**, 373–391 (2015).
- Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
- D'Esposito, M. & Postle, B. R. The cognitive neuroscience of working memory. *Annu. Rev. Psychol.* **66**, 115–142 (2015).
- Miller, E. K., Lundqvist, M. & Bastos, A. M. Working memory 2.0. *Neuron* **100**, 463–475 (2018).
- Truccolo, W., Hochberg, L. R. & Donoghue, J. P. Collective dynamics in human and monkey sensorimotor cortex: predicting single neuron spikes. *Nat. Neurosci.* **13**, 105–111 (2010).
- Averbeck, B. B., Latham, P. E. & Pouget, A. Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* **7**, nrr1888 (2006).
- Cohen, M. R. & Kohn, A. Measuring and interpreting neuronal correlations. *Nat. Neurosci.* **14**, 811–819 (2011).
- Doiron, B., Litwin-Kumar, A., Rosenbaum, R., Ocker, G. K. & Josić, K. The mechanics of state-dependent neural correlations. *Nat. Neurosci.* **19**, 383–393 (2016).
- Ben Hadj Hassen, S. & Ben Hamed, S. Functional and behavioural correlates of shared neuronal noise variability in vision and visual cognition. *Curr. Opin. Physiol.* **16**, 85–97 (2020).
- Mazor, O. & Laurent, G. Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron* **48**, 661–673 (2005).
- 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).
- Cohen, M. R. & Maunsell, J. H. R. A neuronal population measure of attention predicts behavioral performance on individual trials. *J. Neurosci.* **30**, 15241–15253 (2010).
- Snyder, A. C., Yu, B. M. & Smith, M. A. Distinct population codes for attention in the absence and presence of visual stimulation. *Nat. Commun.* **9**, 1–14 (2018).
- Afshar, A. et al. Single-trial neural correlates of arm movement preparation. *Neuron* **71**, 555–564 (2011).

16. Li, N., Daie, K., Svoboda, K. & Druckmann, S. Robust neuronal dynamics in premotor cortex during motor planning. *Nature* **532**, 459–464 (2016).
17. Khanna, S. B., Snyder, A. C. & Smith, M. A. Distinct sources of variability affect eye movement preparation. *J. Neurosci.* **39**, 4511–4526 (2019).
18. Athalye, V. R., Ganguly, K., Costa, R. M. & Carmena, J. M. Emergence of coordinated neural dynamics underlies neuroprosthetic learning and skillful control. *Neuron* **93**, 955–970.e5 (2017).
19. Ni, A. M., Ruff, D. A., Alberts, J. J., Symmonds, J. & Cohen, M. R. Learning and attention reveal a general relationship between population activity and behavior. *Science* **359**, 463–465 (2018).
20. Hennig, J. A. et al. Learning is shaped by abrupt changes in neural engagement. *Nat. Neurosci.* **24**, 727–736 (2021).
21. Harvey, C. D., Coen, P. & Tank, D. W. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* **484**, 62–68 (2012).
22. Kiani, R., Cueva, C. J., Reppas, J. B. & Newsome, W. T. Dynamics of neural population responses in prefrontal cortex indicate changes of mind on single trials. *Curr. Biol.* **24**, 1542–1547 (2014).
23. Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex. *eLife* **4**, e04677 (2015).
24. Mazzucato, L., Fontanini, A. & La Camera, G. Stimuli reduce the dimensionality of cortical activity. *Front. Syst. Neurosci.* **10**, 11 (2016).
25. Williamson, R. C. et al. Scaling properties of dimensionality reduction for neural populations and network models. *PLOS Comput. Biol.* **12**, e1005141 (2016).
26. Huang, C. et al. Circuit models of low-dimensional shared variability in cortical networks. *Neuron* **101**, 337–348 (2019).
27. Churchland, M. M. et al. Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* **13**, 369–378 (2010).
28. Ahrens, M. B., Orger, M. B., Robson, D. N., Li, J. M. & Keller, P. J. Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nat. Methods* **10**, 413–420 (2013).
29. Jun, J. J. et al. Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**, 232–236 (2017).
30. Yang, W. & Yuste, R. In vivo imaging of neural activity. *Nat. Methods* **14**, 349–359 (2017).
31. Kang, B. & Druckmann, S. Approaches to inferring multi-regional interactions from simultaneous population recordings: Inferring multi-regional interactions from simultaneous population recordings. *Curr. Opin. Neurobiol.* **65**, 108–119 (2020).
32. Keeley, S. L., Zoltowski, D. M., Aoi, M. C. & Pillow, J. W. Modeling statistical dependencies in multi-region spike train data. *Curr. Opin. Neurobiol.* **65**, 194–202 (2020).
33. Smedo, 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).
34. Machado, T. A., Kauvar, I. V. & Deisseroth, K. Multiregion neuronal activity: the forest and the trees. *Nat. Rev. Neurosci.* **23**, 683–704 (2022).
35. Bastos, A. M. et al. Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron* **85**, 390–401 (2015).
36. Stringer, C. et al. Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, eaav7893 (2019).
37. Javadzadeh, M. & Hofer, S. B. Dynamic causal communication channels between neocortical areas. *Neuron* **110**, 2470–2483 (2022).
38. Jia, X. et al. Multi-regional module-based signal transmission in mouse visual cortex. *Neuron* **110**, 1585–1598 (2022).
39. Smedo, J. D. et al. Feedforward and feedback interactions between visual cortical areas use different population activity patterns. *Nat. Commun.* **13**, 1099 (2022).
40. Park, H., Keri, H. V. S., Yoo, C., Bi, C. & Pluta, S. R. Bilateral integration in somatosensory cortex is controlled by behavioral relevance. *Nat. Neurosci.* **28**, 1300–1310 (2025).
41. Cohen, J. Y. et al. Cooperation and competition among frontal eye field neurons during visual target selection. *J. Neurosci.* **30**, 3227–3238 (2010).
42. Ames, K. C. & Churchland, M. M. Motor cortex signals for each arm are mixed across hemispheres and neurons yet partitioned within the population response. *eLife* **8**, e46159 (2019).
43. 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).
44. Hagan, M. A. & Pesaran, B. Modulation of inhibitory communication coordinates looking and reaching. *Nature* **604**, 708–713 (2022).
45. Perich, M. G., Gallego, J. A. & Miller, L. E. A neural population mechanism for rapid learning. *Neuron* **100**, 964–976 (2018).
46. 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).
47. Han, S. & Helmchen, F. Coordinated multi-level adaptations across neocortical areas during task learning. *Nat. Commun.* **16**, 7719 (2025).
48. Tang, H., Bartolo, R. & Averbach, B. B. Reward-related choices determine information timing and flow across macaque lateral prefrontal cortex. *Nat. Commun.* **12**, 894 (2021).
49. Tseng, S., Chettih, S. N., Arlt, C., Barroso-Luque, R. & Harvey, C. D. Shared and specialized coding across posterior cortical areas for dynamic navigation decisions. *Cell* **184**, 3717–3730 (2021).
50. Stine, G. M., Trautmann, E. M., Jeurissen, D. & Shadlen, M. N. A neural mechanism for terminating decisions. *Neuron* **111**, 2601–2613 (2023).
51. Brincat, S. L. et al. Interhemispheric transfer of working memories. *Neuron* **109**, 1055–1066.e4 (2021).
52. Chen, G., Kang, B., Lindsey, J., Druckmann, S. & Li, N. Modularity and robustness of frontal cortical networks. *Cell* **184**, 3717–3730 (2021).
53. Steinmetz, N. A., Zarka-Haas, P., Carandini, M. & Harris, K. D. Distributed coding of choice, action and engagement across the mouse brain. *Nature* **576**, 266–273 (2019).
54. Cowley, B. R. et al. Slow drift of neural activity as a signature of impulsivity in macaque visual and prefrontal cortex. *Neuron* **108**, 551–567 (2020).
55. Saalmann, Y. B., Pinsk, M. A., Wang, L., Li, X. & Kastner, S. The pulvinar regulates information transmission between cortical areas based on attention demands. *Science* **337**, 753–756 (2012).
56. Rabinowitz, N. C., Goris, R. L., Cohen, M. & Simoncelli, E. P. Attention stabilizes the shared gain of V4 populations. *eLife* **4**, e08998 (2015).
57. Mayo, J. P. & Maunsell, J. H. R. Graded neuronal modulations related to visual spatial attention. *J. Neurosci.* **36**, 5353–5361 (2016).
58. Herman, J. P., Katz, L. N. & Krauzlis, R. J. Midbrain activity can explain perceptual decisions during an attention task. *Nat. Neurosci.* **21**, 1651–1655 (2018).
59. Snyder, A. C., Yu, B. M. & Smith, M. A. A stable population code for attention in prefrontal cortex leads a dynamic attention code in visual cortex. *J. Neurosci.* **41**, 9163–9176 (2021).
60. Srinath, R., Ruff, D. A. & Cohen, M. R. Attention improves information flow between neuronal populations without changing the communication subspace. *Curr. Biol.* **31**, 5299–5313 (2021).

61. Harris, K. D. Nonsense correlations in neuroscience Preprint at <https://www.biorxiv.org/content/10.1101/2020.11.29.402719v3> (2021).
62. Lee, D., Port, N. L., Kruse, W. & Georgopoulos, A. P. Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J. Neurosci.* **18**, 1161–1170 (1998).
63. Smith, M. A. & Kohn, A. Spatial and temporal scales of neuronal correlation in primary visual cortex. *J. Neurosci.* **28**, 12591–12603 (2008).
64. Solomon, S. S., Chen, S. C., Morley, J. W. & Solomon, S. G. Local and global correlations between neurons in the middle temporal area of primate visual cortex. *Cereb. Cortex* **25**, 3182–3196 (2015).
65. Ruff, D. A. & Cohen, M. R. Stimulus dependence of correlated variability across cortical areas. *J. Neurosci.* **36**, 7546–7556 (2016).
66. Zohary, E., Shadlen, M. N. & Newsome, W. T. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* **370**, 140–143 (1994).
67. Romo, R., Hernández, A., Zainos, A. & Salinas, E. Correlated neuronal discharges that increase coding efficiency during perceptual discrimination. *Neuron* **38**, 649–657 (2003).
68. Kohn, A. & Smith, M. A. Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci.* **25**, 3661–3673 (2005).
69. Huang, X. & Lisberger, S. G. Noise correlations in cortical area MT and their potential impact on trial-by-trial variation in the direction and speed of smooth-pursuit eye movements. *J. Neurophysiol.* **101**, 3012–3030 (2009).
70. Leavitt, M. L., Pieper, F., Sachs, A., Joobar, R. & Martinez-Trujillo, J. C. Structure of spike count correlations reveals functional interactions between neurons in dorsolateral prefrontal cortex area 8a of behaving primates. *PLOS ONE* **8**, e61503 (2013).
71. Goris, R. L. T., Movshon, J. A. & Simoncelli, E. P. Partitioning neuronal variability. *Nat. Neurosci.* **17**, 858–865 (2014).
72. Umakantha, A. et al. Bridging neuronal correlations and dimensionality reduction. *Neuron* **109**, 2740–2754 (2021).
73. Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* **17**, 1500–1509 (2014).
74. Bach, F. R. & Jordan, M. I. A Probabilistic Interpretation of Canonical Correlation Analysis. Technical Report 688, University of California, Berkeley, Department of Statistics (2005).
75. Bittner, S. R. et al. Population activity structure of excitatory and inhibitory neurons. *PLOS ONE* **12**, e0181773 (2017).
76. Niell, C. M. & Stryker, M. P. Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron* **65**, 472–479 (2010).
77. Allen, W. E. et al. Thirst regulates motivated behavior through modulation of brainwide neural population dynamics. *Science* **364**, eaav3932 (2019).
78. Fujimoto, A., Murray, E. A. & Rudebeck, P. H. Interaction between decision-making and interoceptive representations of bodily arousal in frontal cortex. *Proc. Natl. Acad. Sci. USA* **118**, e2014781118 (2021).
79. Strauch, C., Wang, C.-A., Einhäuser, W., Stigchel, S. V. D. & Naber, M. Pupillometry as an integrated readout of distinct attentional networks. *Trends Neurosci.* **45**, 635–647 (2022).
80. Burlingham, C. S., Mirbagheri, S. & Heeger, D. J. A unified model of the task-evoked pupil response. *Sci. Adv.* **8**, eabi9979 (2022).
81. Mathôt, S. & Van der Stigchel, S. New light on the mind's eye: The pupillary light response as active vision. *Curr. Direct. Psychol. Sci.* **24**, 374–378 (2015).
82. Ebitz, R. B. & Moore, T. Both a gauge and a filter: cognitive modulations of pupil size. *Front. Neurol.* **9**, 2018 (2019).
83. Reimer, J. et al. Pupil fluctuations track fast switching of cortical states during quiet wakefulness. *Neuron* **84**, 355–362 (2014).
84. McGinley, M. J., David, S. V. & McCormick, D. A. Cortical membrane potential signature of optimal states for sensory signal detection. *Neuron* **87**, 179–192 (2015).
85. van Kempen, J. et al. Behavioural and neural signatures of perceptual decision-making are modulated by pupil-linked arousal. *eLife* **8**, e42541 (2019).
86. Vinck, M., Batista-Brito, R., Knoblich, U. & Cardin, J. A. Arousal and locomotion make distinct contributions to cortical activity patterns and visual encoding. *Neuron* **86**, 740–754 (2015).
87. Wang, C.-A. & Munoz, D. P. A circuit for pupil orienting responses: implications for cognitive modulation of pupil size. *Curr. Opin. Neurobiol.* **33**, 134–140 (2015).
88. Joshi, S., Li, Y., Kalwani, R. & Gold, J. I. Relationships between pupil diameter and neuronal activity in the locus coeruleus, colliculi, and cingulate cortex. *Neuron* **89**, 221–234 (2016).
89. Johnston, R. & Smith, M. A. Brain-wide arousal signals are segregated from movement planning in the superior colliculus. *eLife* **13**, RP99278 (2024).
90. Moore, R. Y. & Bloom, F. E. Central catecholamine neuron systems: Anatomy and physiology of the norepinephrine and epinephrine systems. *Annu. Rev. Neurosci.* **2**, 113–168 (1979).
91. Sara, S. J. The locus coeruleus and noradrenergic modulation of cognition. *Nat. Rev. Neurosci.* **10**, 211–223 (2009).
92. Gatter, K. C. & Powell, T. P. S. The projection of the locus coeruleus upon the neocortex in the macaque monkey. *Neuroscience* **2**, 441–445 (1977).
93. Porrino, L. J. & Goldman-Rakic, P. S. Brainstem innervation of prefrontal and anterior cingulate cortex in the rhesus monkey revealed by retrograde transport of HRP. *J. Comp. Neurol.* **205**, 63–76 (1982).
94. Joshi, S. & Gold, J. I. Pupil size as a window on neural substrates of cognition. *Trends Cognit. Sci.* **24**, 466–480 (2020).
95. Raut, R. V. et al. Arousal as a universal embedding for spatio-temporal brain dynamics. *Nature* **647**, 454–461 (2025).
96. Barbas, H., Hilgetag, C. C., Saha, S., Dermon, C. R. & Suski, J. L. Parallel organization of contralateral and ipsilateral prefrontal cortical projections in the rhesus monkey. *BMC Neurosci.* **6**, 32 (2005).
97. Schwartz, M. L. & Goldman-Rakic, P. S. Callosal and intrahemispheric connectivity of the prefrontal association cortex in rhesus monkey: Relation between intraparietal and principal sulcal cortex. *J. Comp. Neurol.* **226**, 403–420 (1984).
98. Boch, R. A. & Goldberg, M. E. Participation of prefrontal neurons in the preparation of visually guided eye movements in the rhesus monkey. *J. Neurophysiol.* **61**, 1064–1084 (1989).
99. Funahashi, S., Bruce, C. J. & Goldman-Rakic, P. S. Visuospatial coding in primate prefrontal neurons revealed by oculomotor paradigms. *J. Neurophysiol.* **63**, 814–831 (1990).
100. Khanna, S. B., Scott, J. A. & Smith, M. A. Dynamic shifts of visual and saccadic signals in prefrontal cortical regions 8Ar and FEF. *J. Neurophysiol.* **124**, 1774–1791 (2020).
101. Kritzer, M. F. & Goldman-Rakic, P. S. Intrinsic circuit organization of the major layers and sublayers of the dorsolateral prefrontal cortex in the rhesus monkey. *J. Comp. Neurol.* **359**, 131–143 (1995).
102. Pucak, M. L., Levitt, J. B., Lund, J. S. & Lewis, D. A. Patterns of intrinsic and associational circuitry in monkey prefrontal cortex. *J. Comp. Neurol.* **376**, 614–630 (1996).
103. Litwin-Kumar, A. & Doiron, B. Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nat. Neurosci.* **15**, 1498–1505 (2012).
104. Driscoll, L. N., Shenoy, K. & Sussillo, D. Flexible multitask computation in recurrent networks utilizes shared dynamical motifs. *Nat. Neurosci.* **27**, 1349–1363 (2024).

105. Rosenbaum, R., Smith, M. A., Kohn, A., Rubin, J. E. & Doiron, B. The spatial structure of correlated neuronal variability. *Nat. Neurosci.* **20**, 107–114 (2017).
106. Van Vreeswijk, C. & Sompolinsky, H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* **274**, 1724–1726 (1996).
107. Amit, D. J. & Brunel, N. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex* **7**, 237–252 (1997).
108. Renart, A. et al. The asynchronous state in cortical circuits. *Science* **327**, 587–590 (2010).
109. Hennequin, G., Ahmadian, Y., Rubin, D. B., Lengyel, M. & Miller, K. D. The dynamical regime of sensory cortex: Stable dynamics around a single stimulus-tuned attractor account for patterns of noise variability. *Neuron* **98**, 846–860 (2018).
110. Rosenbaum, R., Rubin, J. E. & Doiron, B. Short-term synaptic depression and stochastic vesicle dynamics reduce and shape neuronal correlations. *J. Neurophysiol.* **109**, 475–484 (2013).
111. Ko, H. et al. Functional specificity of local synaptic connections in neocortical networks. *Nature* **473**, 87–91 (2011).
112. Herrero, J. L. et al. Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. *Nature* **454**, 1110–1114 (2008).
113. Goard, M. & Dan, Y. Basal forebrain activation enhances cortical coding of natural scenes. *Nat. Neurosci.* **12**, 1444–1449 (2009).
114. Mincses, V., Pinto, L., Dan, Y. & Chiba, A. A. Cholinergic shaping of neural correlations. *Proc. Natl. Acad. Sci. USA* **114**, 5725–5730 (2017).
115. McCormick, D. A. Cholinergic and noradrenergic modulation of thalamocortical processing. *Trends Neurosci.* **12**, 215–221 (1989).
116. Lee, S.-H. & Dan, Y. Neuromodulation of brain states. *Neuron* **76**, 209–222 (2012).
117. Ts'o, D. Y., Gilbert, C. D. & Wiesel, T. N. Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J. Neurosci.* **6**, 1160–1170 (1986).
118. Gilbert, C. D. & Wiesel, T. N. Morphology and intracortical projections of functionally characterised neurones in the cat visual cortex. *Nature* **280**, 120–125 (1979).
119. McGuire, B. A., Gilbert, C. D., Rivlin, P. K. & Wiesel, T. N. Targets of horizontal connections in macaque primary visual cortex. *J. Comp. Neurol.* **305**, 370–392 (1991).
120. Lund, J. S., Yoshioka, T. & Levitt, J. B. Comparison of intrinsic connectivity in different areas of macaque monkey cerebral cortex. *Cereb. Cortex* **3**, 148–162 (1993).
121. Goulas, A., Schaefer, A. & Margulies, D. S. The strength of weak connections in the macaque cortico-cortical network. *Brain Struct. Funct.* **220**, 2939–2951 (2015).
122. Semedo, J. D., Zandvakili, A., Machens, C. K., Yu, B. M. & Kohn, A. Cortical areas interact through a communication subspace. *Neuron* **102**, 249–259 (2019).
123. Bergstra, J., Bardenet, R., Bengio, Y. & Kégl, B. Algorithms for hyper-parameter optimization. *Adv. Neural Inf. Process. Syst.* **24**, 2546–2554 (2011).
124. Gokcen, E. et al. Disentangling the flow of signals between populations of neurons. *Nat. Comput. Sci.* **2**, 512–525 (2022).
125. Bong, H. et al. Latent dynamic factor analysis of high-dimensional neural recordings. *Adv. Neural Inf. Process. Syst.* **33**, 16446–16456 (2020).
126. Glaser, J., Whiteway, M., Cunningham, J. P., Paninski, L. & Linderman, S. Recurrent switching dynamical systems models for multiple interacting neural populations. *Adv. Neural Inf. Process. Syst.* **33**, 14867–14878 (2020).
127. Balzani, E., Noel, J. P., Herrero-Vidal, P., Angelaki, D. E. & Savin, C. A probabilistic framework for task-aligned intra- and inter-area neural manifold estimation. In *International Conference on Learning Representations* <https://openreview.net/pdf?id=kt-dcBQcSA> (ICLR, 2023).
128. Gokcen, E. et al. Uncovering motifs of concurrent signaling across multiple neuronal populations. *Adv. Neural Inf. Process. Syst.* **36**, 34711–34722 (2023).
129. Li, W., Li, C., Wang, Y. & Wu, A. *Multi-region Markovian Gaussian Process: An efficient method to discover directional communications across multiple brain regions*, 28112–28131 (PMLR, 2024).
130. Jha, T., Sani, O. G., Pesaran, B. & Shanechi, M. M. Prioritized learning of cross-population neural dynamics. *J. Neural Eng.* **22**, 046040 (2025).
131. Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).
132. Kobak, D. et al. Demixed principal component analysis of neural population data. *eLife* **5**, e10989 (2016).
133. Sani, O. G., Abbaspourazad, H., Wong, Y. T., Pesaran, B. & Shanechi, M. M. Modeling behaviorally relevant neural dynamics enabled by preferential subspace identification. *Nat. Neurosci.* **24**, 140–149 (2021).
134. Schneider, S., Lee, J. H. & Mathis, M. W. Learnable latent embeddings for joint behavioural and neural analysis. *Nature* **617**, 360–368 (2023).
135. Akaho, S. A kernel method for canonical correlation analysis. International Meeting of Psychometric Society, Osaka, Japan <https://arxiv.org/pdf/cs/0609071> (2001).
136. Cowley, B. et al. *Distance covariance analysis*, 242–251 (PMLR, 2017).
137. Perich, M. G. & Rajan, K. Rethinking brain-wide interactions through multi-region “network of networks” models. *Curr. Opin. Neurobiol.* **65**, 146–151 (2020).
138. Liu, B., Sacks, J. & Golub, M. D. *Accurate identification of communication between multiple interacting neural populations*, 39381–39404 (PMLR, 2025).
139. Brainard, D. H. The Psychophysics Toolbox. *Spat. Vis.* **10**, 433–436 (1997).
140. Pelli, D. G. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat. Vis.* **10**, 437–442 (1997).
141. Kleiner, M. et al. What’s new in psychtoolbox-3. *Perception* **36**, 1–16 (2007).
142. Hikosaka, O. & Wurtz, R. H. Visual and oculomotor functions of monkey substantia nigra pars reticulata. I. Relation of visual and auditory responses to saccades. *J. Neurophysiol.* **49**, 1230–1253 (1983).
143. Issar, D., Williamson, R. C., Khanna, S. B. & Smith, M. A. A neural network for online spike classification that improves decoding accuracy. *J. Neurophysiol.* **123**, 1472–1485 (2020).
144. Everitt, B. S. Factor analysis. In *An Introduction to Latent Variable Models* (Springer, 1984).
145. Tucker, L. R. An inter-battery method of factor analysis. *Psychometrika* **23**, 111–136 (1958).
146. Klami, A., Virtanen, S., Leppäaho, E. & Kaski, S. Group factor analysis. *IEEE Trans. Neural Netw. Learn. Syst.* **26**, 2136–2147 (2015).
147. Dempster, A. P., Laird, N. M. & Rubin, D. B. Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc.: Ser. B (Methodol.)* **39**, 1–22 (1977).
148. McDonnell, M. E., Umakantha, A., Williamson, R. C., Smith, M. A. & Yu, B. M. Data from “Interactions across hemispheres in prefrontal cortex reflect global cognitive processing”. KiltHub, <https://doi.org/10.1184/R1/31431415> (2026).
149. McDonnell, M., Umakantha, A., Williamson, R. C., Smith, M. A. & Yu, B. M. pCCA-FA (v1.1). Zenodo, <https://doi.org/10.5281/zenodo.18913385> (2026).

150. McDonnell, M., Umakantha, A., Williamson, R. C., Smith, M. A. & Yu, B. M. Code from “Interactions across hemispheres in pre-frontal cortex reflect global cognitive processing” (v1.0). Zenodo, <https://doi.org/10.5281/zenodo.18929016> (2026).

Acknowledgements

The authors would like to thank Samantha Nelson for help with animal training and data collection, our animal care staff, and Evren Gokcen, Emilio Salazar-Gatzimas, and Sam Snyder for helpful discussions. M.E.M. was supported by NIH T32 EB029365. B.M.Y. and M.A.S. were supported by NIH R01 MH118929, NSF NCS BCS 1734916/1954107, NSF NCS DRL 2124066/2123911, and NIH R01 EB026953. B.M.Y. was supported by NIH RF1 NS127107, NIH R01 EY035896, and Simons Foundation 543065 and NC-GB-CULM-00003241-05. M.A.S. was supported by NIH R01 EY029250 and NIH R01 MH128393.

Author contributions

Conceptualization, M.E.M., A.U., R.C.W., M.A.S., and B.M.Y.; Data curation: M.E.M., A.U., and R.C.W.; Formal analysis, M.E.M. and A.U.; Funding acquisition, M.A.S. and B.M.Y.; Investigation, M.E.M., A.U., and R.C.W.; Methodology, M.E.M., A.U., R.C.W., M.A.S., and B.M.Y.; Software, M.E.M. and A.U.; Resources, M.A.S. and B.M.Y.; Supervision, M.A.S. and B.M.Y.; Validation, M.E.M. and A.U.; Visualization, M.E.M., A.U., and R.C.W.; Writing—original draft, M.E.M., A.U., and R.C.W.; Writing—review & editing, M.E.M., A.U., R.C.W., M.A.S., and B.M.Y.

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/s41467-026-71725-0>.

Correspondence and requests for materials should be addressed to Matthew A. Smith or Byron M. Yu.

Peer review information *Nature Communications* thanks Simon Jacob 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>

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

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

© The Author(s) 2026