State-Dependent Organization of Microscale Functional Circuitry in Visual Cortex

Rahul Biswas[1,2], Hasika Wickrama Senevirathne[3], Yujing Wang[3], Jiang Zhang[3], Somabha Mukherjee[3], and Reza Abbasi-Asl[1,2,4,*]

[1]Department of Neurology, University of California, San Francisco, CA, USA
[2]UCSF Weill Institute for Neurosciences, San Francisco, CA, USA
[3]Department of Statistics and Data Science, National University of Singapore, Singapore
[4]Department of Bioengineering and Therapeutic Sciences, University of California, San Francisco, CA, USA

Brain state modulates sensory processing across visual cortex, yet how it relates to the organization of functional circuitry at the level of individual neurons and cell types remains largely unknown. To address this, we constructed one of the largest microscale directed functional circuit maps in mouse visual cortex from calcium imaging of more than 57,000 neurons across four visual areas and five cortical layers. Using a time-aware causal inference framework, we found that intra-areal connections dominate across arousal states, consistent with experimental findings on the local bias of cortical anatomy. Among intra-areal connections, anterolateral area (AL) had the highest density, and among inter-areal connections, the AL↔rostrolateral area (RL) axis formed the strongest pathway. Laminar circuit organization was dominated by layer 6 recurrence within-layer, while the most prominent between-layer pathway was layer 5-to-layer 6 in low arousal and layer 4-to-layer 5 in high arousal. Spatial extent was selectively greater for excitatory-to-inhibitory connections in high arousal, but not for excitatory-to-excitatory connections. Across 6,597 electron-microscopy reconstructions of neuron pairs, synapse count predicted functional connection strength in both arousal states, but structure-function coupling was weaker in high arousal. In stimulus-driven response prediction, neuron pairs with stronger functional connections exhibited more similar predictive performance in both states, with performance varying by layer and cell type. Overall, our findings map, at single-neuron resolution, the multi-scale organization of directed functional circuitry in mouse visual cortex across brain states.

1 Introduction

Microscale *functional circuitry*, defined here as directed, statistically causal functional connections between individual neurons inferred from their activity, provides a window into the brain's circuit-level organization [1, 2]. Simultaneous structural and functional measurements have shown that

*Corresponding author: Reza.AbbasiAsl@ucsf.edu

functional circuitry is related to the anatomical wiring [3], however, the correlations are weak. More importantly, there is a large body of theoretical and computational studies demonstrating that functional circuitry captures the dynamic routing of information through the network [4–8]. Traditional functional connectivity methods rely on undirected associations and cannot capture directional interactions [9]. Recent methods have addressed this by applying statistical causal inference frameworks to neural activity recordings, enabling the estimation of directed functional connections and advancing the field from associative connectivity toward microscale functional circuitry [9–11]. Yet despite this progress, it remains unknown how such directed functional circuitry is organized across different brain states, particularly at the resolution of individual neurons.

Arousal state is a particularly important dimension along which to examine this question. Neural responses during sensory processing, attention, and behavioral tasks are known to vary with arousal, ranging from drowsiness to alertness [12–14]. Prior studies that examined functional connectivity across brain states relied on local field potential (LFP) array signals [15] or combined multi-unit spiking and LFP recordings [16]. While these revealed state-dependent shifts in population-level coordination, neither approach could resolve how circuit organization changes at the level of individual neurons. Therefore, how directed functional circuitry is shaped across arousal states, at the level of individual neurons, cell types, laminar origin, and multi-area hierarchy, remains an open question.

To answer these questions, we leveraged the MICrONS dataset [3], a unique resource combining large-scale calcium imaging of 146,388 neurons, of which 15,434 were coregistered with a dense electron microscopy reconstruction. This allowed us to map the functional circuitry of the mouse visual cortex with single-neuron resolution, estimating state-dependent connections across four distinct visual areas (V1, LM, AL, RL) and linking them to their anatomical positions, structural scaffold, and cell types [17]. To infer directed functional connections from the neural activity recordings, we employed the Time-Aware PC (TPC) algorithm [10], a causal discovery method that uses conditional independence tests on time-lagged neural activity. We mapped the directed *statistically-causal functional circuitry* (CFC) within each arousal state and identified state-specific differences in spatial extent, layer composition, and structure-function coupling.

We found that intra-areal connections were denser and stronger than inter-areal connections in both arousal states, with area AL showing the highest within-area density and the AL↔RL axis as the prominent inter-area pathway. Laminar organization within areas revealed state-specific motifs. Layer 6 (L6) recurrence was the dominant within-layer pathway in both states. The strongest between-layer pathway was layer 5 (L5)→L6 in low arousal and layer 4 (L4)→L5 in high arousal. Spatial extent showed cell-type selective expansion in high arousal. Excitatory-to-inhibitory ($E \rightarrow I$) connections extended further with inter-neuron distance while excitatory-to-excitatory connections did not. Structurally, synapse count predicted functional weight in both states, but this structure-function coupling was weaker in high arousal. Neuron pairs with stronger structural support showed smaller state-dependent changes. For stimulus-driven response prediction, neuron pairs with stronger functional connections exhibited more similar predictive performance in both states, with performance varying by cortical layer and cell type. Additionally, in high arousal, predictive correlation was redistributed: lower-performing neurons improved while higher-performing ones declined.

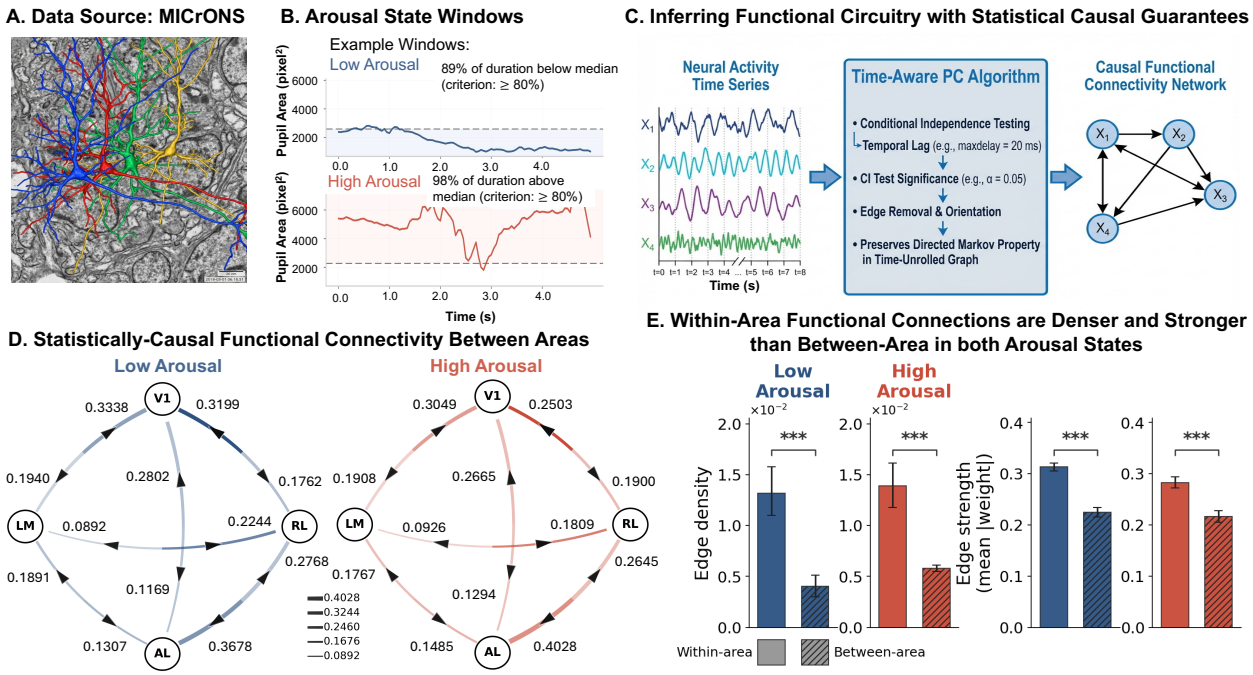

Figure 1: Large-scale multimodal recording, arousal-state definition, and directed functional circuit inference in mouse visual cortex. (A) Data source: The MICrONS dataset combines large-scale calcium imaging with dense EM reconstruction [3]. **(B)** Arousal state windows defined by median split of pupil area (computed per session and scan) for more than 80% of trial duration [12, 13]. **(C)** The TPC algorithm infers directed functional connections with statistical causal guarantees [10]. **(D)** Statistically-causal functional circuitry (CFC) adjacency matrices and area-level graph for a representative trial, in Low and High arousal states. **(E)** Within-area functional connections are denser and stronger than between-area connections in both arousal states (bootstrap mean ± 95% CI, $n = 101$ fields, 5000 iterations). Both edge density and edge strength are shown. Asterisks mark bootstrap significance (two-sided, FDR-corrected): * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

2 Results

2.1 Intra-areal connections dominate a heterogeneous functional circuit

The architecture of cortical functional circuitry, which pathways are active and how strongly, shapes how sensory information is routed within and across visual areas [17, 18]. Structural anatomy strongly predicts local dominance: most cortical synapses are intrinsic, and long-range projections represent a minority of total synaptic input [19]. Whether this local bias is preserved in *directed* functional circuitry, and how it is structured across areas and arousal states, is less well characterized at single-neuron resolution. We used the MICrONS dataset to establish the baseline architecture of directed functional circuitry in each arousal state, and to ask how this architecture varies by area and whether it is consistent across arousal states. Arousal state was continuously monitored via pupil area [12, 13, 20] (**Figure 1A,B**). Trials were classified as High Arousal if ≥ 80% of the trial duration had pupil area above the session-specific median, and as Low Arousal if ≥ 80% was below the median; trials not meeting this criterion were excluded.

To infer directed functional circuitry from the calcium traces, we employed the Time-Aware

PC (TPC) algorithm [10], a causal discovery method for time-series data that extends the PC algorithm with time-lagged conditional independence tests (**Figure 1C**); full algorithmic details, assumptions, and implementation parameters are given in Methods. We chose TPC because it has been shown to significantly outperform standard correlation and Granger causality in reconstructing neural circuits [10,11,21]. Therefore, the inferred functional circuitry could serve as a more reliable recommendation system generating robust hypotheses for future experiments.

The calcium imaging recordings cover 101 imaging fields (13 session-scans) comprising 146,388 units. Units showing detectable activity events in fewer than 12% of frames during all stimulus types were not included in TPC estimation, as silent neurons contribute no detectable functional connections; this yielded 57,686 active units across the imaging fields. Field-level CFC estimates were then compared between arousal states.

TPC was run on each trial independently, yielding a per-trial directed functional connectivity matrix over all neuron pairs within each imaging field (**Figure 1D**). We summarized this matrix at the area level by computing *edge density* for each directed area pair: the fraction of all directed neuron pairs within that area pair for which TPC inferred a nonzero connection (**Figure 1D–E**). To aggregate per-trial estimates robustly across arousal states, we used a paired bootstrap: for each of 5000 iterations, one trial was sampled randomly from the low-arousal pool and one from the high-arousal pool for each field, and metrics were averaged across fields ($n = 101$ fields). This pairing avoids averaging over unequal per-state trial counts and provides bootstrap mean ± 95% CI throughout.

Consistent with the local bias of cortical anatomy, intra-areal functional connections were denser and stronger than inter-areal connections in both arousal states (**Figure 1E**). In high arousal, within-area edge density (bootstrap mean = 0.014, 95% CI: [0.012, 0.016]) substantially exceeds between-area density (mean = 0.006; bootstrap $p < 0.001$). The same pattern holds in low arousal (within: 0.013, between: 0.004; $p < 0.001$). Within-area edge strength likewise exceeds between-area strength in both states (low: 0.313 vs 0.225; high: 0.283 vs 0.217; both $p < 0.001$). These values characterize the organization of each arousal state independently. Between-state differences in overall edge density and strength were not statistically significant (Supplementary Table 1).

2.1.1 Functional circuit organization varies across visual areas

The mouse visual cortex comprises multiple areas (V1, LM, AL, RL) with distinct functional specializations and hierarchical positions [17,18]. We characterized the organization of functional circuitry across all 16 directed area pairs separately for each arousal state, revealing heterogeneous structure within the intra-areal dominant baseline (**Figure 2A**).

In both arousal states, functional circuitry is structured: within-area connections are denser and stronger than between-area connections (bootstrap $p < 0.001$; **Figure 2A**). Within-area density varies substantially across areas: AL exhibits the highest within-area density in both states (low: 0.0235, high: 0.0271), while V1 is the lowest (low: 0.0049, high: 0.0053), with LM and RL intermediate (∼0.011–0.013). Between-area connections are markedly sparser (0.002–0.021), with the AL↔RL axis forming the densest inter-area pathway in both states (AL→RL: low = 0.014, high = 0.021; RL→AL: low = 0.011, high = 0.017), substantially above all other between-area pairs.

Within-area edge strength (mean |weight| over nonzero edges) also exceeds between-area strength in both states. The AL↔RL pathway has a relatively strong between-area weights (RL→AL: low = 0.368, high = 0.403), while RL→LM is consistently the weakest between-area pathway (low = 0.089, high = 0.093). These patterns reveal that functional circuitry in both arousal states is intra-areal dominant, with a hierarchically structured inter-area backbone dominated by specific

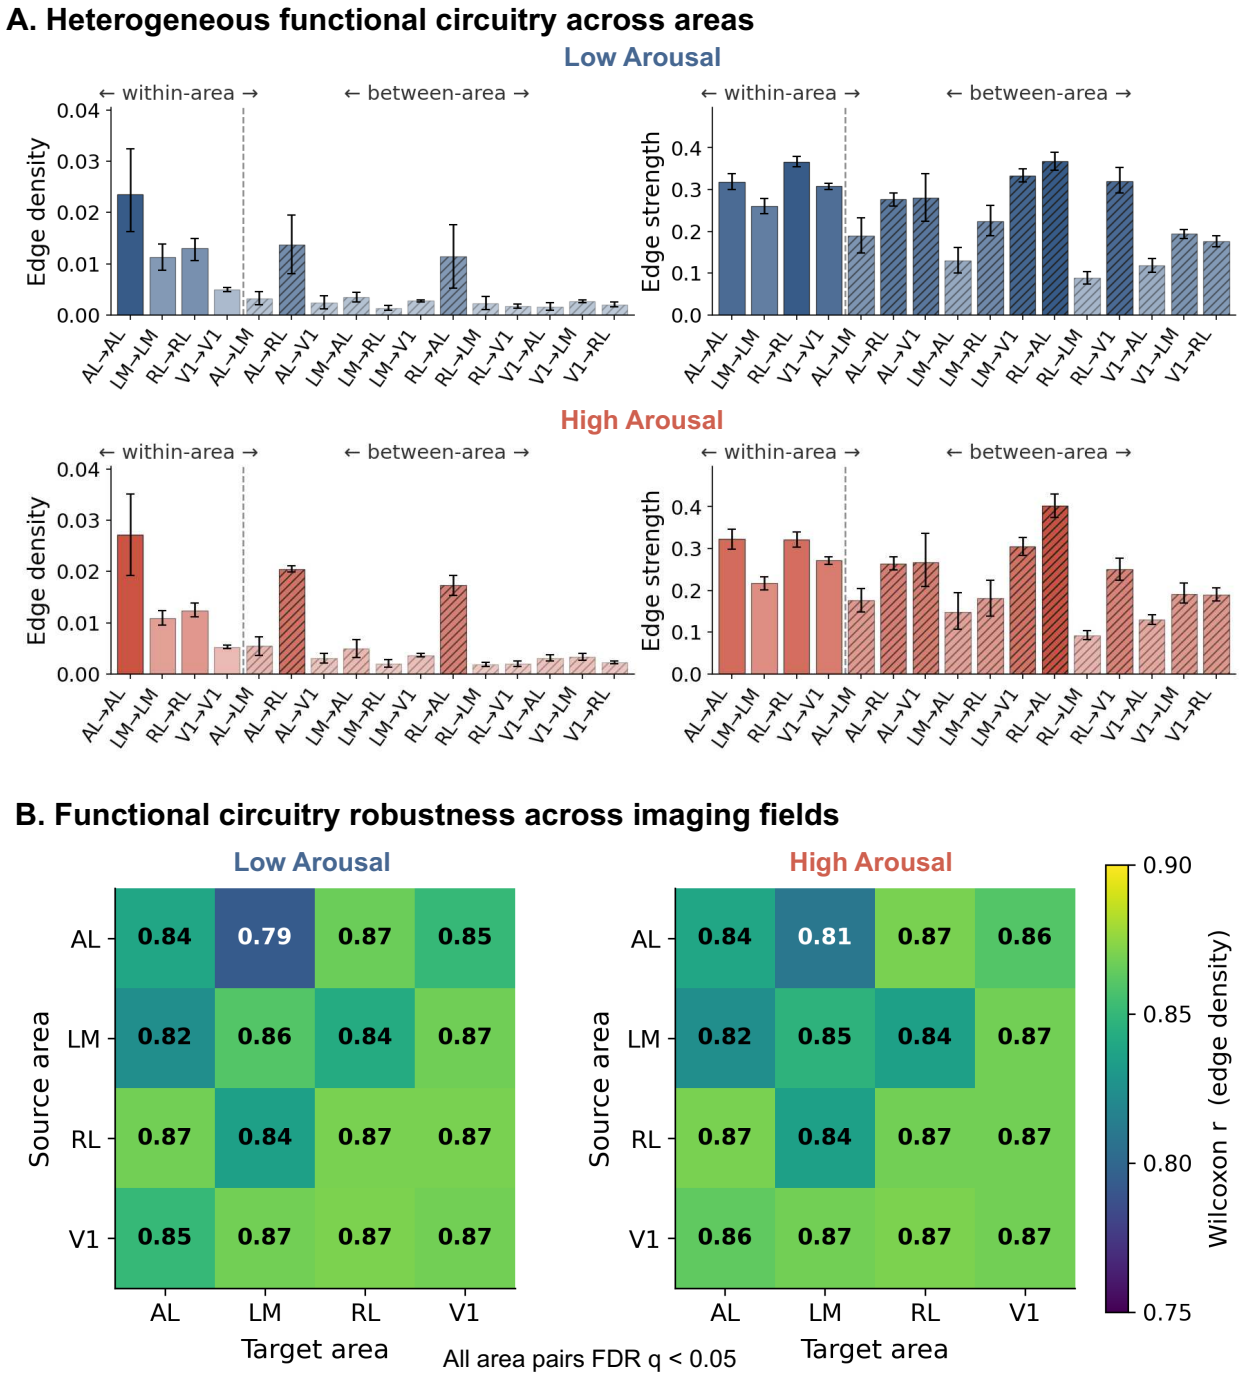

Figure 2: Areal functional circuitry is heterogeneous across directed pairs in both arousal states. (A) Heterogeneous functional circuitry across areas: edge density and edge strength for all 16 directed area pairs, shown separately for Low arousal (top row) and High arousal (bottom row). Bootstrap mean ± 95% CI, $n = 101$ fields. **(B)** Functional circuitry robustness across imaging fields: Wilcoxon effect size ($r = Z/\sqrt{N}$) per directed area pair, measuring the consistency of within-field functional circuitry across imaging fields. Higher r indicates greater consistency across fields.

pathways.

Among within-area pairs, AL exhibits the highest edge density in both arousal states, and the AL↔RL pathway is the dominant inter-area connection in both states. Within-area density in AL is approximately five-fold higher than in V1 (0.0235 vs. 0.0049 in low arousal), consistent with a higher-order area with stronger local recurrence [22–24].

Robustness of functional circuitry across imaging fields, measured as the Wilcoxon effect size ($r = Z/\sqrt{N}$) per area pair (**Figure 2B**), is consistently high across all within-area pairs (Wilcoxon r range 0.839–0.870), and the AL↔RL pathway exhibits similarly high robustness, indicating that both local and this dominant inter-area pathway are stable across imaging fields.

2.2 Functional circuitry differs across arousal states by cell type, layer, and spatial extent

Functional circuitry in cortex is organized along multiple dimensions: cell type, laminar position, and spatial extent of interactions [17,25,26]. Having characterized areal organization in each arousal state, we next asked how cell type and layer structure functional circuitry, and whether the spatial extent of functional connections differs between states.

2.2.1 Excitatory subtypes display distinct functional circuit motifs in each arousal state

Excitatory neurons in visual cortex comprise molecularly and morphologically distinct subtypes across layers, each with characteristic projection targets and functional roles [17]. We constructed aggregate cell-type functional circuitry networks separately for each arousal state using a rigorous bootstrap approach ($n = 43$ fields with EM coregistration, 5000 iterations; Methods) and examined the organization of cell-type functional circuitry in each state (**Figure 3A**).

Each arousal state is characterized by a different dominant circuit motif. The two states differ in circuit topology, not merely in the magnitude of the same connections.

In **low arousal**, the dominant functional circuit motif is deep-layer recurrence: the strongest functional link is L6tall-a→L6short-b (bootstrap mean normalized weight 0.235), followed by the reciprocal L6short-b→L6tall-a (0.141). L5→L6 projections (e.g., L5ET→L6short-a, mean 0.013) are present but substantially weaker than the dominant L6 recurrence. Cell-type identities follow the transcriptomic taxonomy of mouse visual cortex [27–29]. This architecture suggests that in low arousal, layer 6 sustains the dominant recurrent functional dynamics, with L5-to-L6 feedforward contributions.

In **high arousal**, the dominant functional circuit motif centers on L6short-a: the strongest functional links are L6short-a→L6tall-a (0.079) and L6short-a→L5ET (0.030), with L5a→L5b connections (mean 0.003) present at substantially lower weights. This architecture implicates L6short-a as the key hub, with projections targeting both deep-layer recurrent partners and corticospinal/corticotectal output neurons (L5ET).

These patterns are descriptive, from the bootstrap mean network without formal edge-level statistical tests. Layer-level corroboration is provided by the statistical analysis in Section 2.2.2. The per-state graphs reveal distinct functional architectures across states rather than a simple scaling of connectivity. Individual fields contain too few neurons of each specific excitatory subtype to support reliable per-field estimates. Formal hypothesis testing of subtype-pair differences is therefore not possible; these patterns remain hypothesis-generating.

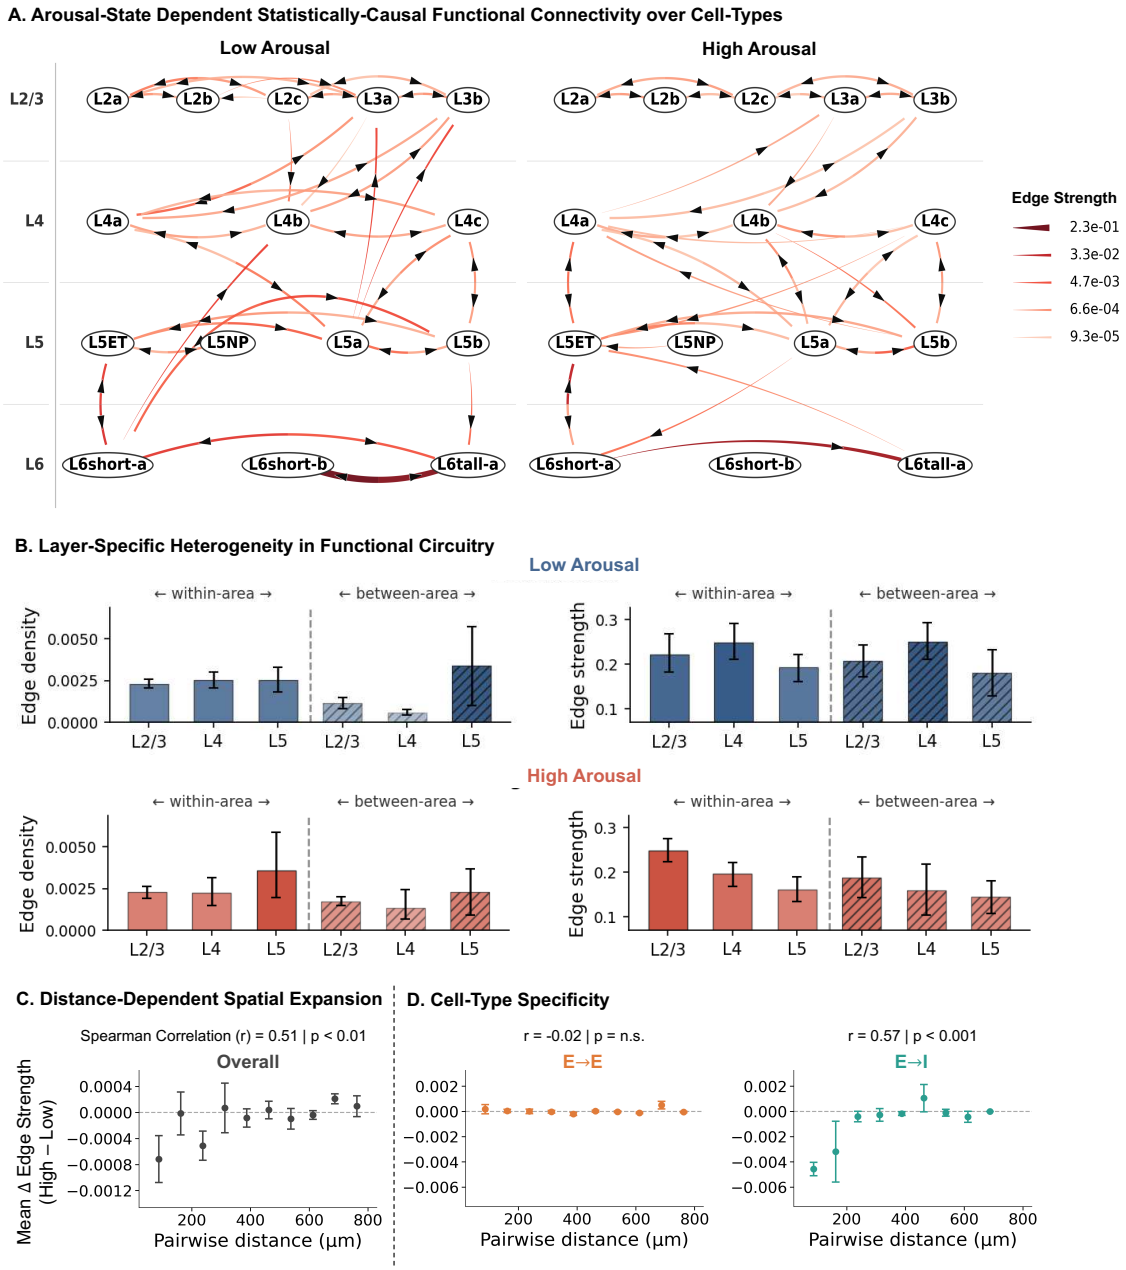

Figure 3: Cell-type, layer-specific, and spatial organization of functional circuitry. (A) Excitatory cell-type functional circuitry networks for low (left) and high (right) arousal. Bootstrap mean, $n = 43$ fields with EM coregistration, 5000 iterations. **(B)** Within- and between-layer edge density and strength per arousal state, showing layer-specific heterogeneity. **(C)** Functional circuitry change (high − low) vs. inter-neuron distance, showing distance-dependent spatial expansion (bootstrap mean Spearman r, p-value). **(D)** Cell-type specificity of spatial expansion. E→E connections (left) exhibit no significant distance dependence (bootstrap mean Spearman $r = -0.017$, FDR $p = 0.884$). E→I connections (right) had a significant positive correlation with distance (bootstrap mean Spearman $r = 0.570$, 95% CI $[0.206, 0.850]$, FDR $p < 0.001$). Spatial expansion is specific to excitatory-to-inhibitory functional connections.

2.2.2 Layer 6 hosts the densest functional connections

Cortical layers differ in their connectivity patterns and computational roles: layer 4 receives thalamic input, layers 2/3 perform local and inter-areal integration, layer 5 provides cortical output, and layer 6 maintains corticothalamic feedback projections [30–32]. We characterized the layer organization of functional circuitry across excitatory subtypes in each arousal state using the per-state bootstrap (**Figure 3B–D**).

Within each arousal state, layer 6 neurons sustain by far the strongest within-layer functional connections. In the bootstrap-mean cell-type network (**Figure 3A**), the mean within-L6 functional connection weight in low arousal (0.065) exceeds that of L5 (0.0006), L4 (0.0004), and L2/3 (0.0003) by approximately two orders of magnitude; in high arousal, within-L6 weight (0.013) remains the dominant within-layer signal. Between-layer connections are overall much weaker than within-L6. In low arousal, the L5→L6 pathway is the strongest between-layer route (mean 0.0015). In high arousal, the most prominent between-layer connections are L4→L5 (mean 0.0005).

2.2.3 Spatial extent of functional circuitry is greater in high than low arousal

Functional circuitry in cortex is distance-dependent, with nearby neurons sharing more functional interactions than distant ones [1,2]. We found that spatial extent of functional circuitry is expanded in high arousal: the change in connectivity strength was significantly positively correlated with inter-neuron distance (bootstrap mean Spearman $r = 0.513$, 95% CI $[0.091, 0.879]$, $p = 0.008$, one-sided bootstrap), meaning long-range functional connections are preferentially enhanced in high arousal relative to short-range ones. The absolute changes are small in magnitude (on the order of 10^{-3}) but consistent across distance bins. This pattern is consistent with enhanced long-range integration alongside preserved local precision (**Figure 3C**).

This aggregate distance-dependent expansion is, however, not uniform across connection types, as we elaborate next.

2.2.4 Greater spatial extent in high arousal is specific to excitatory-to-inhibitory functional connections

Neuromodulatory systems associated with arousal act preferentially on specific cell types, particularly through disinhibitory circuits involving excitatory-to-inhibitory connections [25, 26, 33]. To identify whether specific functional connection types account for the spatial expansion, we analyzed distance-dependent connectivity changes separately for different cell-type connection pairs. The analysis included 312,970 E→E, 1,467 E→I, 1,443 I→E, and 60 I→I functional connections. I→E functional connections showed no significant distance dependence (not shown). Excitatory-to-Excitatory (E→E) functional connections showed essentially no distance dependence (bootstrap mean Spearman $r = -0.017$, FDR-corrected $p = 0.884$, not significant). In contrast, Excitatory-to-Inhibitory (E→I) functional connections showed a statistically significant positive correlation with distance (bootstrap mean Spearman $r = 0.570$, 95% CI $[0.206, 0.850]$, FDR-corrected $p < 0.001$), despite the much smaller number of E→I pairs relative to E→E. Thus, within the limits of our data, the distance-dependent expansion appears to be specific to E→I functional connections and subtle in magnitude rather than a large, visually dramatic effect (**Figure 3D**). The small number of inhibitory neurons (105 of 7,815, 1.3%) and correspondingly sparse I→I functional connections (60 total) precluded a meaningful analysis of inhibitory spatial circuitry patterns.

2.3 Structure-function coupling is weaker in high than low arousal

Having characterized the within-state architecture of functional circuitry at the area, cell-type, and spatial level, we next ask how the two states relate to the underlying anatomical scaffold and to sensory encoding quality.

Electron microscopy (EM) provides synapse-resolution structural connectivity maps [1], enabling a direct test of whether state-dependent functional changes are constrained by anatomical wiring. EM reconstruction is sparse, covering $n = 43$ fields and $n = 6{,}597$ connection-level observations with structural data (mean 1,176 EM-reconstructed synapses per field, range 18–3,365; pooled across fields; see Methods). We applied two complementary approaches: field-level Spearman correlations and a connection-level mixed-effects model. At the field level, synapse count positively predicted TPC-inferred functional weight in both arousal states (mean Spearman $r = 0.141$ in low, $r = 0.098$ in high arousal; one-sample Wilcoxon against zero, $n = 43$ fields; FDR-corrected $p = 3.41 \times 10^{-13}$ and $p = 2.39 \times 10^{-12}$), with the association present in essentially every field. Pairs for which TPC inferred no direct edge were assigned functional weight zero. This reflects a genuine absence of a detectable functional connection, not missing data. The field-level Spearman correlation therefore estimates the degree to which synapse count predicts functional weight within the EM-coregistered subset of synaptically connected pairs ($n = 6{,}597$). This structure-function coupling was weaker during high arousal, and synapse count was negatively associated with arousal-dependent functional changes (mean Spearman $r = -0.029$, FDR-corrected $p = 1.6 \times 10^{-3}$), indicating that pairs with stronger structural support exhibit smaller state-dependent change (**Figure 4A**).

The connection-level analysis fitted the model:

$$\mathrm{FC} = \beta_0 + \beta_{SC} \cdot \mathrm{SC} + \beta_{state} \cdot \mathrm{state} + \beta_{SC\times state} \cdot (\mathrm{SC} \times \mathrm{state}) + \alpha_f \tag{1}$$

where state = 0 (low) or 1 (high arousal) and α_f is a random field effect. Synapse count showed a significant positive association with TPC functional weight in both states: the structural coefficient at baseline (low arousal) was $\beta_{SC} = 0.0087$, and the effective slope in high arousal ($\beta_{SC} + \beta_{SC\times state}$) was 0.0054. The interaction $\beta_{SC\times state} = -0.0033$ ($p = 1.20 \times 10^{-17}$) was significantly negative, confirming that structure-function coupling is weaker in high arousal and that functional circuitry deviates further from the structural scaffold in high arousal (**Figure 4B**, left).

Equivalently, fitting a separate model with ΔFC directly as the outcome yields the same relationship:

$$\Delta\mathrm{FC} = \beta_0 + \beta_{SC} \cdot \mathrm{SC} + \alpha_f \tag{2}$$

yielding $\beta_{SC} = -0.0034$ ($p = 1.20 \times 10^{-17}$), indicating that neuron pairs with more structural synapses exhibit smaller functional changes (**Figure 4B**, right). Strong structural connections have smaller state-dependent changes, while weaker structural connections have larger changes. Both field-level and connection-level analyses point to the same conclusion: structure provides the scaffold for functional circuitry, but arousal-dependent changes are larger at neuron pairs with fewer EM-reconstructed synapses. These results remained significant after including inter-neuron distance, brain area, and cell type as covariates (Supplementary **Table 5**).

2.4 Stimulus-driven predictive correlation covaries with functional circuitry strength

A key question in systems neuroscience is whether circuit-level organization is associated with measurable difference in the brain's ability to represent sensory stimuli [34]. To address this, we computed predictive correlation (how well a stimulus-driven model predicts each neuron's response)

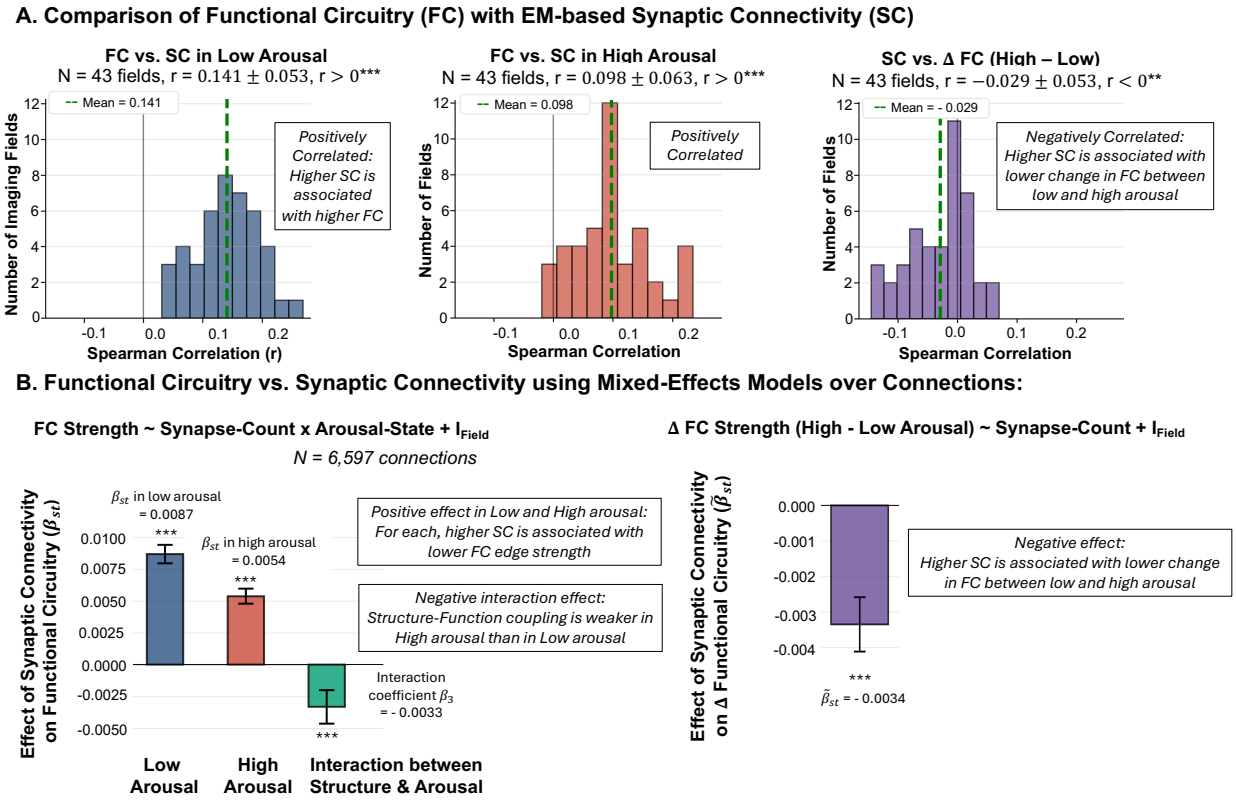

Figure 4: Structure-function relationships and state-dependent deviation from anatomy. (A) Field-level Spearman correlations between EM synapse count and TPC functional weight (undetected pairs set to zero) across $n = 43$ fields. Synapse count positively predicts TPC-inferred functional weight in both arousal states; synapse count is negatively associated with arousal-dependent functional change (ΔFC). Significance assessed by one-sample Wilcoxon signed-rank tests ($n = 43$ field-level correlations vs. zero), with FDR correction across the three tests. Asterisks mark FDR-corrected significance (key as in Figure 1). **(B)** Confirmatory connection-level mixed-effects models ($n = 6{,}597$ SC-positive observations across 43 fields; field as random effect) exhibit a positive association between synapse count and functional weight in both states, weaker in high than low arousal (negative interaction), and a negative association between synapse count and ΔFC. Asterisks mark FDR-corrected significance for model coefficients.

as a proxy for sensory encoding quality [35], separately for each arousal state. We used the winning architecture from the Sensorium 2023 NeurIPS competition [36], trained on MICrONS data with 7-fold cross-validation and evaluated on held-out folds (**Figure 5A**; see Methods). Arousal was not a training target, so state differences in predictive correlation are not an artifact of the training procedure. All findings were replicated on an independent held-out test set ($n = 752$ trials). The analysis covers 119,913 units across 13 session-scans.

Neurons with stronger functional connections exhibit more correlated predictive performance across trials: TPC functional circuitry weight was positively correlated with predictive correlation profile similarity (the across-trial Spearman correlation between source and target neurons' predictive correlation profiles) in both arousal states (**Figure 5B**). This relationship was significant for all within-area pairs and several between-area pairs (Wilcoxon one-sample vs. 0, FDR-corrected).

Predictive correlation varied by cortical layer and cell type in both arousal states (**Figure 5C**).

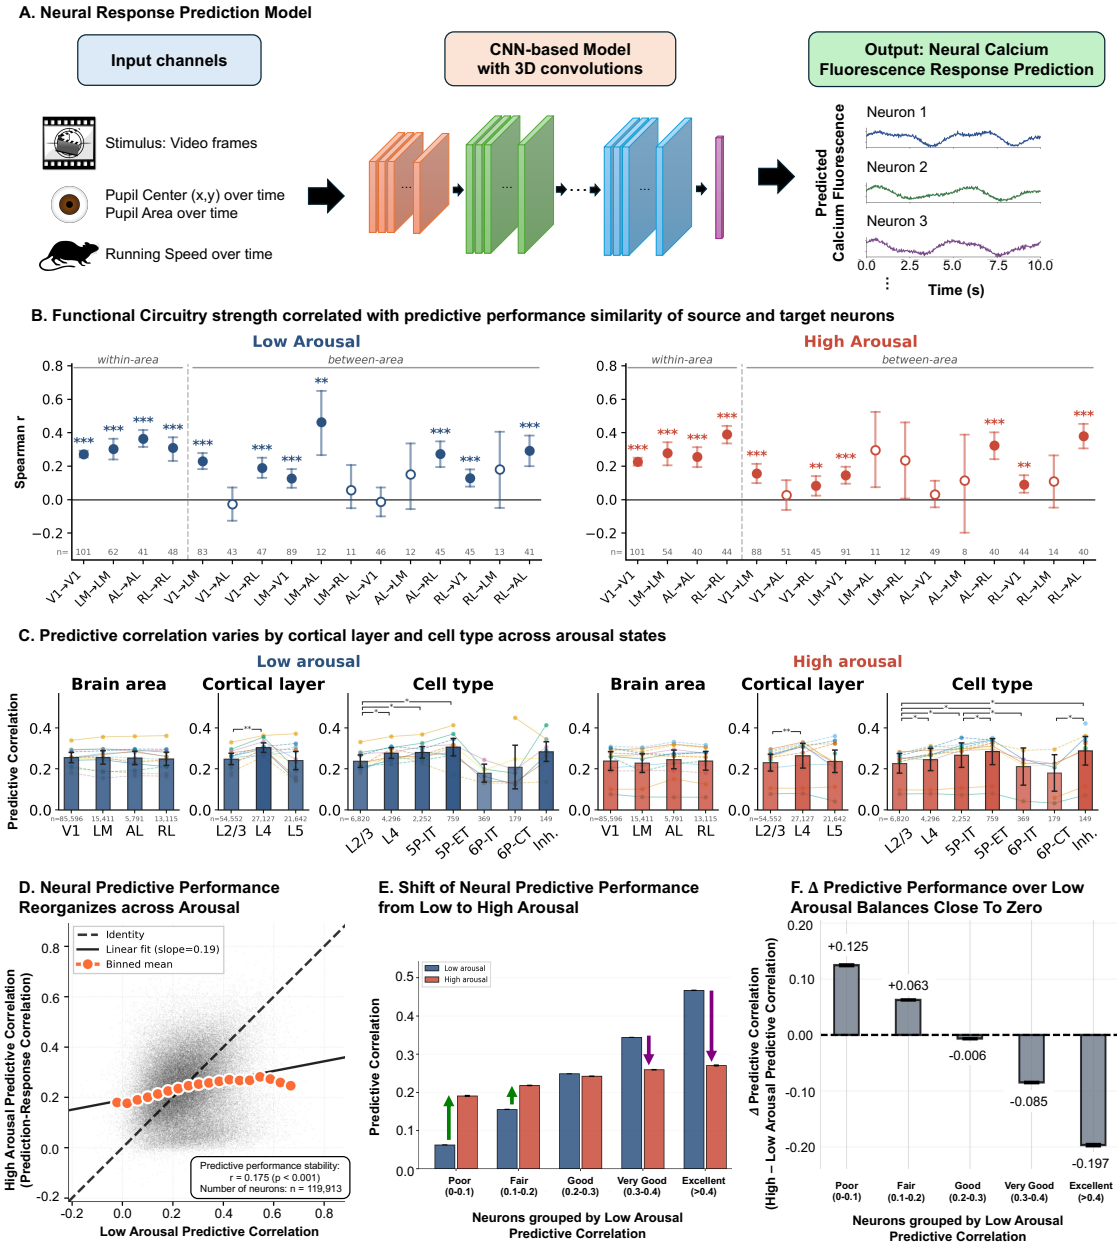

Figure 5: Neural sensory-response predictive performance varies with functional circuit organization and arousal state. **(A)** Schematic of the neural encoding model used to compute predictive correlation. **(B)** Functional circuitry weight correlates with encoding similarity: bootstrap mean Spearman correlation between TPC weight and predictive correlation profile similarity, per directed area pair and arousal state (mean ± 95% CI). **(C)** Predictive correlation by cortical layer and cell type across arousal states; asterisks indicate FDR-corrected pairwise comparisons (Supplementary Table 4). **(D)** Scatter of low vs. high arousal predictive correlation across neurons (n = 119,913); linear fit shown. **(E)** Mean predictive correlation by low-arousal performance bin across arousal states. **(F)** Change in predictive correlation (high − low) by low-arousal performance bin, showing redistribution from high- to low-performing neurons.

Among layers, L4 showed the highest mean predictive correlation in both low arousal (mean = 0.304, 95% CI [0.281, 0.327]) and high arousal (mean = 0.264, 95% CI [0.204, 0.325]), significantly exceeding L2/3 in both states (Wilcoxon, FDR-corrected: $p < 0.01$). Among cell types, 5P-ET showed the highest predictive correlation in both states (low: mean = 0.306; high: mean = 0.285), while L6 excitatory subtypes showed the lowest (6P-IT low: 0.179; 6P-CT low: 0.209). More cell-type pairs were significantly different in high arousal than in low arousal (FDR-corrected; Supplementary Table 4). Brain areas did not differ significantly in predictive correlation in either state.

The scatter of low versus high arousal predictive correlation across all 119,913 neurons (**Figure 5D**) reveals systematic redistribution rather than uniform change. The population mean showed a net decrease ($\Delta = -0.015$, $p < 0.001$, paired t-test, high − low), but this average masks substantial heterogeneity. When neurons were stratified by their low-arousal predictive correlation, a clear gradient emerged: low-performing neurons improved while high-performing neurons declined (**Figure 5E,F**). Neurons in the lowest performance bin (correlation < 0.1) showed a mean improvement of $\Delta = +0.125$ ($p < 0.001$); neurons in the highest performance bin (correlation > 0.4) showed a mean decrease of $\Delta = -0.197$ ($p < 0.001$). At the individual-neuron level, the correlation between low-arousal predictive correlation and the arousal-associated change was $r = -0.598$ ($p < 0.001$; 119,913 units). A permutation test confirmed the monotonic gradient ($p = 0.019$). To verify the gradient is not an artifact of regression to the mean, we performed split-half cross-validation: bin assignment used one half of low-arousal trials and the gradient was evaluated on held-out trials. The gradient was fully preserved (Spearman $r = -1.0$, $p = 0.019$). All bins remained significant after FDR correction. This gradient held consistently across all visual areas (V1: $r = -0.600$, LM: $r = -0.587$, AL: $r = -0.589$, RL: $r = -0.604$; all $p < 0.001$).

Neurons with high, low, or changing predictive correlation in high arousal were distributed broadly across the imaged cortical volume rather than clustered in specific subregions (Supplementary **Figure S1**). We also quantified the *correlation of predictive performance* (the Pearson r between low- and high-arousal predictive correlations across neurons), which measures how stably neurons maintain their relative predictability across arousal states. This varied by area, layer, and cell type (Supplementary **Figure S2**; full pairwise statistics in Supplementary **Table 4**), and was highest in LM among areas, L6 among layers, and L2b among cell types.

3 Discussion

This study characterizes microscale functional circuitry in mouse visual cortex across arousal states at single-neuron resolution. The characterization spans area-level topology, cell-type and laminar motifs, spatial extent, structure-function coupling, and sensory encoding.

3.1 Intra-areal dominance and hierarchical asymmetry across visual areas

The finding that within-area connections dominate functional circuitry in both arousal states is consistent with the cortical column as a fundamental unit of information processing [17,37]. Dense recurrent functional connectivity within areas supports local computation. Sparse between-area functional connections carry the output of this local processing across the visual hierarchy [18, 38]. The observation that AL exhibits the highest within-area connection density in both states may reflect its role as a higher-order area specialized for motion processing with stronger local recurrence [22, 23]. All within-area pairs have consistently high cross-field robustness (Wilcoxon r range 0.839–0.870). The AL↔RL pathway exhibits similarly high robustness. The dominance of the AL↔RL inter-area axis is consistent with the lateral, non-hierarchical connectivity between AL

and RL identified in anatomical studies [39]. Unlike feedforward V1→LM projections, the AL↔RL connection links areas at similar hierarchical levels. This is consistent with lateral integration within the dorsal motion-processing stream rather than hierarchical feedforward transmission.

3.2 L6 recurrence and L5ET output as state-dependent functional circuit motifs

L6 within-layer functional connection strength exceeds that of all other layers by approximately 100-fold. This indicates that L6 is the primary substrate for within-area functional recurrence. This likely reflects L6 corticocortical neurons, which receive approximately 39% of presynaptic inputs from within L6 [40]. This circuit-level finding is compatible with single-neuron studies showing that arousal-related gain modulation is weakest in L5/6 [41]. Those results describe firing rate changes between states. The L6 recurrence finding describes connection structure within each arousal state. In high arousal, L6→L5 output pathways dominate in place of L6 recurrence, with E/I balance implications. Feedforward-like projections from deep layers to L5 may recruit stronger PV-mediated inhibition than the recurrent L6 motif. This is consistent with the known asymmetry in inhibitory recruitment between feedforward and feedback pathways [42].

At the cell-type level, the per-state analysis is based on grand-averaged networks across all fields; no per-field statistics were computed for cell-type motifs. These observations are therefore hypothesis-generating rather than confirmed findings. In low arousal, L6 recurrence between morphologically distinct subtypes (L6tall-a↔L6short-b) dominates. In high arousal, this deep-layer recurrence is absent. Instead, L6short-a-mediated projections to L5ET dominate, consistent with selective engagement of the corticospinal/corticotectal output pathway [17, 43]. These observations are consistent with the hypothesis that arousal state is associated with selective engagement of different deep-layer circuits for output routing rather than uniform modulation of all cell types.

3.3 Structure-function coupling is selectively weakened in high arousal

Previous studies have documented that high-arousal states are associated with elevated single-neuron firing rates and gain [14]. Prior work has also shown that arousal is associated with reduced pairwise noise correlations between neurons [44], capturing undirected pairwise structure. The microscale functional circuitry changes we observe represent a distinct level of analysis [45]: directed, network-level interactions across thousands of simultaneously recorded neurons. These state-dependent changes are not accessible through pairwise or single-neuron measures alone.

An additional dimension not previously examined at the level of individual EM-reconstructed synapses is the relationship between state-dependent functional circuitry and underlying structural anatomy. EM synapse count positively predicts TPC functional weight in both arousal states. Pairs with more synapses are more likely to carry a functional edge and tend to have stronger functional weights. However, structure-function coupling is weaker in high arousal. State-dependent changes in functional circuitry are greatest at neuron pairs with fewer EM-reconstructed synapses. Functional circuitry thus departs more from structural anatomy in high arousal. One interpretation is that strongly-wired neuron pairs operate near a functional ceiling imposed by their synaptic weights. This leaves less room for state-dependent modulation. Weakly-wired pairs may represent a more flexible substrate whose functional weight varies more across arousal states.

3.4 Functional circuitry weight correlates with sensory encoding similarity

The positive correlation between TPC functional circuitry weight and predictive correlation profile similarity (**Figure 5B**) demonstrates that stronger functional connections link neurons with more similar sensory encoding performance. This relationship holds across both arousal states and is

significant for all within-area directed pairs. It validates that TPC connections are functionally meaningful: they capture not only directionality but also the degree to which connected neurons share a common sensory encoding profile across trials.

The variation in predictive correlation across layers and cell types (**Figure 5C**) is consistent with known hierarchical specialization of visual cortex [17, 18]. L4 carries the highest predictive correlation among all layers in both states. This is expected: L4 is the primary recipient of thalamocortical input and is therefore most tightly coupled to the sensory stimulus. Among cell types, 5P-ET neurons show the highest predictive correlation. Layer 5 extratelencephalic neurons are a principal output class of cortex; their strong sensory responses are consistent with their role in broadcasting visual information to subcortical targets [17]. At the other extreme, L6 excitatory subtypes (6P-IT, 6P-CT) show the lowest predictive correlation. The asymmetry between states is also informative: more cell-type pairs are significantly different in high arousal than in low arousal. This suggests that arousal state is associated with a broader differentiation of encoding quality at the cell-type level, beyond what is captured by mean firing-rate changes. The stability of relative encoding quality across states is measured as the Pearson r between low- and high-arousal predictive correlations across neurons. This stability is highest in LM among areas, L6 among layers, and L2b among cell types. LM's high stability is consistent with its strong reciprocal connectivity with V1, which may preserve representational structure across states [38, 46]. The high stability of L6 is notable given that L6 neurons show the lowest mean predictive correlation; neurons with weak absolute encoding can nonetheless maintain a stable relative rank across arousal states.

The redistribution of predictive correlation across neurons represents a further dimension not accessible through pairwise measures alone [34]. The population mean shows a net decrease in high arousal ($\Delta = -0.015$), but this average conceals a pronounced gradient (**Figure 5D–F**). Neurons with low baseline encoding quality improve substantially in high arousal (lowest bin: $\Delta = +0.125$), while neurons with high baseline encoding quality decline (highest bin: $\Delta = -0.197$). This gradient ($r = -0.598$) is not an artifact of regression to the mean: split-half cross-validation fully preserves it. The gradient holds consistently across all visual areas. This pattern is consistent with a selective rebalancing of sensory representation across the population in high arousal, rather than uniform suppression or amplification.

3.5 E-to-I spatial expansion is consistent with a disinhibitory mechanism

One candidate mechanism linking the circuit-level and coding-level findings involves the E→I spatial expansion demonstrated here. In high arousal, expanded long-range E→I connections may recruit weakly-driven neurons into the functional circuit by providing broader excitatory input. This could explain why low-baseline neurons improve while high-baseline neurons decline. This account is speculative. Causal perturbation experiments are required to test it directly.

The spatial expansion pattern is not consistent with simple gain modulation. E→I connections selectively expand their spatial extent in high arousal; E→E connections do not. Uniform gain scaling would affect all functional interactions equally [47,48]. This selectivity suggests that specific circuit mechanisms are associated with reconfiguration of inhibitory pathway reach in high arousal [49].

The E→I-driven spatial expansion is consistent with disinhibitory circuit mechanisms [26,49,50]. During arousal, neuromodulatory systems (e.g., acetylcholine, noradrenaline) selectively suppress parvalbumin (PV) interneurons, which provide dense proximal inhibition to local excitatory neurons [25]. Suppression of proximal PV interneurons releases excitatory neurons from local inhibitory control. These disinhibited excitatory neurons can then recruit more distant inhibitory targets, specifically somatostatin (Sst) interneurons with long-range axonal projections [51]. Different Sst

MET-types have distinct synaptic connectivity patterns, with some subtypes projecting at distances that could account for the observed expansion. This two-step mechanism (proximal PV suppression → distal Sst recruitment) would selectively expand E→I spatial extent without proportionally affecting E→E connections. That is the pattern observed here.

3.6 Limitations and Future Directions

This study has several methodological boundaries that motivate future work. First, pupil area is a well-established proxy for arousal [12, 13, 20], but does not capture the full complexity of arousal-related neuromodulation. Pupil area correlates with locomotion speed, and this analysis does not distinguish their independent contributions. Future work could separate them by regressing out running speed or restricting to stationary epochs. Future studies combining optogenetic manipulation of neuromodulatory systems with functional imaging would provide more direct causal evidence for the mechanisms proposed here. Relevant targets include locus coeruleus and basal forebrain. Second, our analysis focused on visual cortex. Extending these findings to other sensory and association areas would clarify whether state-dependent circuit reconfiguration is a general property of neocortex or specific to sensory hierarchies [52, 53]. Third, causal sufficiency is a standard assumption in statistical causal inference that is challenging to satisfy in vivo: all common causes of recorded neurons must be observed [54], yet unobserved inputs such as thalamic drive and neuromodulatory fluctuations cannot be fully excluded. TPC partially addresses this by testing, for each connection, whether it holds after controlling for subsets of other neurons' activity in the lag window. When a recorded neuron is driven by the same unobserved input as the pair being tested, controlling for it absorbs that input's influence and can remove the spurious edge [54]. Perturbation experiments targeting specific cell types would provide direct validation of the inferred connections. Fourth, the cell-type classification used here aggregates many subtypes. Future work with more granular cell-type identification (e.g., specific interneuron subtypes) could reveal more detailed mechanisms [25, 26, 33]. Finally, this analysis was limited to awake, head-fixed mice viewing natural movies. How these findings extend to other behavioral states (e.g., sleep, active locomotion) remains an open question [55, 56].

4 Methods

4.1 Data Source

We used the MICrONS dataset (Minnie65) [3], containing electron microscopy reconstruction and calcium imaging from the mouse visual cortex [57, 58]. The dataset includes detrended calcium fluorescence traces from 146,388 units across 16 session-scans, of which 15,434 were coregistered with a dense EM volume. For the predictive correlation analysis, the unit of observation is a session-scan-unit combination (119,913 units in total across 13 session-scans after excluding 3 sessions with corrupted pupil recordings; mean 9,224 units per session-scan, range 8,395–12,675). Stratified analyses (e.g., by brain area or cell type) use subsets of these units with available annotations, and report the corresponding sample size in each case.

4.2 Arousal State Classification

Arousal states were defined based on pupil area, a well-established proxy for arousal [12, 13, 20, 59]. For each session and scan, we computed the median pupil area ($\pi\times$ major radius $\times$ minor radius) across all trials in that session. This median split was computed separately for each session and

scan to account for session-specific variations in pupil size. To ensure robust state classification, we applied a consistency criterion. Trials were classified as High Arousal only if $\geq 80\%$ of trial duration had pupil area above the session-specific median. Trials were classified as Low Arousal only if $\geq 80\%$ of trial duration had pupil area below the median. We excluded trials that did not meet this consistency criterion. This approach ensures that each trial represents a relatively stable arousal state and improves the reliability of state-dependent comparisons.

4.3 Functional Circuitry Analysis

Functional circuitry was estimated using the Time-Aware PC (TPC) algorithm, a causal discovery method for time-series data [10]. TPC extends the PC (Peter-Clark) algorithm [54] to time-lagged data. Standard correlation cannot distinguish a direct connection between two neurons from an indirect one mediated by a third neuron. Pairwise Granger causality [60] addresses directionality but does not account for network-wide dependencies. Multivariate extensions do, but assume linear interactions between neurons and condition on all other neurons simultaneously [61]. TPC makes no such assumption and its performance advantage over these methods on simulated neural circuits is established in prior work [10, 11]. TPC tests whether each neuron pair remains statistically dependent after controlling for progressively larger subsets of other neurons' activity in a lag window, retaining and edge between i and j only if no subset renders them conditionally independent. This removes indirect dependencies and yields a directed, weighted connectivity matrix per trial.

Algorithm details: TPC operates on a time-expanded representation in which each neuron at each time lag is treated as a distinct variable, enabling conditional independence tests that respect temporal ordering and capture inter-temporal causal relationships. For neurons i and j with activity $X_i(t)$ and $X_j(t)$, TPC tests whether $X_j(t)$ is conditionally independent of $X_i(t-\tau)$ given subsets of activity of all other neurons in the lag window, where τ is the time lag. Edges that survive this test are retained with a signed weight reflecting the direction and magnitude of inferred influence. Positive weights indicate excitatory-like coupling; negative weights indicate inhibitory-like coupling. We used a maximum lag of one frame (approximately 159 ms at 6.3 Hz [3]) and a significance threshold of $\alpha = 0.05$. The one-frame lag is appropriate for calcium imaging, where temporal resolution is limited by indicator kinetics rather than neural spiking.

Statistical causal inference guarantees: TPC recovers the correct directed functional connections consistently as the number of observed time points grows [11], under two standard assumptions for causal discovery methods [54]. First, the directed Markov property: the graph structure matches the conditional independence patterns in the data. Second, faithfulness: every direct connection must produce a statistical dependency that survives conditioning on all other observed neurons, with no exact cancelations. Faithfulness is applied to temporally lagged variables, so inter-temporal causal relationships that contemporaneous methods miss are captured. Faithfulness violations are theoretically possible but unlikely to be systematic across the large number of neuron pairs analyzed here. Performance on simulated neural circuits is described in Benchmarking below. Limitations on causal sufficiency are discussed in the Limitations subsection above. These connections constitute the statistically causal functional circuitry analyzed throughout.

Unit selection: TPC was applied separately within each of 101 imaging fields across 13 session-scans. A unit was retained if its deconvolved spike trace exceeded a value of one in more than 12% of frames during at least one stimulus type (Clip, Monet, or Trippy; union criterion). Different stimulus types selectively engage different neurons; the union captures all neurons with detectable stimulus-evoked activity across any part of the recording. Neurons below this threshold in all three

stimulus types contribute no detectable functional connections. Of 146,388 total units, this yielded 57,686 active units (mean 465 per field) used for causal functional circuitry (CFC) inference.

Benchmarking: TPC has been benchmarked against correlation, Granger causality, and the standard PC algorithm on simulated neural time series. It outperformed these methods in recovering known connectivity patterns [11, 21]. These results support TPC's suitability for inferring directed functional circuitry from neural time-series data.

We applied TPC to calcium imaging traces separately for trials in each arousal state to infer the directed functional circuitry for each state.

4.4 Data Processing Pipeline

CFC Matrix Computation: For each field, session, and scan, TPC was applied to calcium imaging traces (as provided in the MICrONS processed dataset) separately for trials in each arousal state. This generated causal functional circuitry matrices (CFC matrices) where each entry (i, j) represents the inferred connectivity strength from neuron i to neuron j. CFC values can be positive (excitatory-like) or negative (inhibitory-like), reflecting the direction of inferred influence.

Data Aggregation: Before aggregation, absolute values of all CFC values were taken to ensure that connectivity strength is represented as a magnitude, regardless of whether the coupling is positive or negative. This approach was chosen because the sign of CFC values may reflect measurement artifacts or indirect effects rather than true excitatory vs. inhibitory connectivity. For each session and scan, we computed connectivity ratios (sum CFC / sum possible pairs) for each brain area pair and cell-type pair, pooling CFC values and possible pairs across all fields within that session and scan. These ratios were then averaged across all sessions and scans to obtain the final connectivity estimates. This unweighted aggregation method ensures that each session and scan contributes equally to the final connectivity estimate, regardless of the number of neuron pairs observed in that session and scan. We chose unweighted aggregation to avoid biasing results toward session-scans with more neurons.

4.5 Statistical Analysis

4.5.1 Data Structure and Unit of Analysis

Each recording session comprised multiple scans, each of which contained up to four imaging fields (spatially distinct two-photon fields of view, each capturing a population of neurons at a specific cortical location; referred to as "fields" throughout, following MICrONS dataset conventions). For all connectivity analyses, the unit of analysis was the individual trial-level CFC matrix per field. Averaging CFC matrices across trials before computing metrics would conflate trial-to-trial variability with state estimates. We therefore used a paired bootstrap procedure. For each of 5000 iterations, one trial was drawn at random from the low-arousal pool and one from the high-arousal pool for each field. Connectivity metrics were computed separately for each drawn trial, then averaged across fields. This yields a bootstrap distribution of field-averaged metrics for each state and their difference, from which we report bootstrap mean $\pm$ 95% confidence intervals (2.5th–97.5th percentile) and two-sided bootstrap p-values. Two-sided bootstrap p-values were computed as twice the proportion of bootstrap iterations in which the sign of the difference (high minus low arousal) was opposite to the observed mean difference. A fixed random seed (seed = 42) was used throughout to facilitate reproducibility given the deposited code. The number of fields used was $n = 101$ for area-level analyses and $n = 43$ for cell-type analyses (fields with EM coregistration).

Independence of Observations: Fields within the same session and scan may share some neurons, potentially violating independence assumptions. However, we treated fields as independent observations for the following reasons: (1) each field represents a distinct spatial location with largely non-overlapping neuron populations, (2) the number of shared neurons between fields is minimal (estimated $< 5\%$ based on unit ID overlap across fields within each session-scan), and (3) our primary analyses compare within-field changes (low vs. high arousal) rather than across-field comparisons, reducing concerns about pseudoreplication. For analyses comparing across fields (e.g., spatial correlation), we report both field-level statistics and note that results are robust to session-level aggregation. The total number of fields analyzed varied by hypothesis (ranging from $n = 43$ for structure-function analysis to $n > 100$ for connectivity comparisons), with exact sample sizes reported for each test.

4.5.2 Hypothesis Testing Procedures

Area-Specific Decomposition: For each field, within-area connectivity metrics (edge density, edge strength) were computed separately for each area (V1, LM, AL, RL) and between-area metrics separately for each of the 12 directed between-area pairs. Within-area metrics used submatrices containing only neurons from the same area; between-area metrics used cross-area submatrices. The primary analysis characterizes within-state organization using the paired bootstrap procedure described above: bootstrap distributions of per-state means are computed for each area pair, and within-state contrasts (within-area vs. between-area) are assessed by comparing bootstrap distributions. Between-state comparisons (low vs. high arousal for each area pair) were performed with Wilcoxon signed-rank tests and are reported in Supplementary Table 1; FDR correction (Benjamini-Hochberg) was applied across all such tests within one family.

Structure-Function (field-level): For structure-function analysis, we obtained unit_id $\rightarrow$ pt_root_id mappings from the coregistration data for each field. Structural connectivity matrices were constructed by querying direct synapse counts from the MICrONS EM reconstruction for all neuron pairs within each field. Synapse counts were aggregated by grouping pre- and post-synaptic root_ids, with multiple direct synapses between the same neuron pair summed. This measure represents the total number of direct synapses (not indirect/polysynaptic paths) between neuron pairs.

Functional circuitry matrices were computed from CFC values averaged across all trials within each arousal state per field. Absolute values of CFC were taken before computing correlations (consistent with aggregation method). The excitatory-to-inhibitory (E$\rightarrow$I) versus excitatory-to-excitatory (E$\rightarrow$E) distinction throughout is based on the EM-coregistered cell-type label of the source neuron, not on the sign of the CFC entry. TPC-undetected entries (NaN) were set to zero, reflecting the absence of a direct inferred functional connection for those pairs. For each field, both structural (synapse counts) and functional circuitry (absolute CFC values, undetected pairs set to zero) matrices were flattened (off-diagonal entries only (upper and lower triangles combined)) and correlated using Spearman correlation.

This analysis was performed for three comparisons: (1) structural connectivity vs. functional circuitry in low arousal, (2) structural connectivity vs. functional circuitry in high arousal, and (3) structural connectivity vs. functional circuitry changes (high arousal $-$ low arousal). Fields were included if at least one neuron pair had EM-reconstructed synaptic data and the Spearman correlation could be computed ($n = 43$ fields; mean 1,176 EM-reconstructed synapses per field, range 18–3,365; 0.02–3.39% of possible neuron pairs per field). The connection-level mixed-effects model used all neuron pairs with structural data across these fields ($n = 6{,}597$ pairs). Generalizability to the full network is therefore uncertain, though the observed relationships are robust within the

sampled connections. Statistical significance was assessed using the one-sample Wilcoxon signed-rank test (median correlation across fields vs. zero); FDR correction (Benjamini-Hochberg) was applied across the three field-level structure-function comparisons.

Within-Area/Between-Area: Connections were classified as within-area (same area) or between-area (different areas). Mean connectivity per category was computed per field and arousal state. For paired comparisons (low vs. high arousal within the same field), Wilcoxon signed-rank tests were used. For unpaired comparisons (within-area vs. between-area within the same state), we used Mann-Whitney U tests.

Spatial Extent: Euclidean distance between neuron pairs was computed from 3D coordinates. Connectivity change (high − low) was averaged within each distance bin (75 μm bins, spanning 50–800 μm; neuron pairs closer than 50 μm were excluded). For each bootstrap iteration, Spearman correlation between bin centers and the field-averaged bin means was computed. Bootstrap mean r and 95% CI across 5000 iterations are reported.

Layer-Specific: Neurons were assigned to layers by z-coordinate (L2/3: 150-350 μm, L4: 350-450 μm, L5: 450-650 μm, L6: >650 μm). Mean within-area and between-area connectivity per layer (or layer pair) was computed per field. For paired comparisons (low vs. high arousal within the same field and layer or layer pair), Wilcoxon signed-rank tests were used.

Cell-Type Network Graph: To construct the aggregate cell-type functional circuitry network (**Figure 3A**), neurons were assigned to 17 cell types (DTC, Inhibitory, L2a, L2b, L2c, L3a, L3b, L4a, L4b, L4c, L5ET, L5NP, L5a, L5b, L6short-a, L6short-b, L6tall-a) based on coregistration with morphological classification from the EM reconstruction. For each bootstrap iteration ($n = 5000$; one low-arousal and one high-arousal trial drawn per field), the normalized connection weight for each directed cell-type pair was computed as the sum of absolute CFC values divided by the total number of possible neuron pairs for that cell-type pair, then averaged across fields ($n = 43$ fields with EM coregistration). The bootstrap mean across iterations constitutes the reported network weight for each cell-type pair and arousal state. Note that all analyses use 13 session-scans (three excluded due to corrupted pupil recordings).

Cell-Type Spatial: Distance-dependent connectivity changes were computed separately for E→E and E→I pairs using the same field-averaged bootstrap approach as the overall spatial analysis (ten 75 μm-wide distance bins spanning 50–800 μm; neuron pairs closer than 50 μm were excluded). Spearman correlation between bin centers and field-averaged connectivity change was computed per bootstrap iteration; bootstrap mean r and 95% CI across 5000 iterations are reported. The analysis included 312,970 E→E, 1,467 E→I, and 1,443 I→E functional connections. I→E functional connections showed no significant distance dependence and are not shown. I→I connections (60 total) were excluded due to insufficient sample size.

Structure-Function Mixed-Effects Models: To test structure-function relationships while accounting for nested structure (connections within fields), we used mixed-effects models with field as a random intercept (random slopes were not included due to model convergence constraints). Both the mixed-effects model and the field-level analysis use the same $n = 43$ fields (the mixed-effects model uses $n = 6{,}597$ connections across these fields). The structural connectivity covariate is the total number of direct synapses between neuron pairs, obtained from the MICrONS EM synapse reconstruction. The reference model tested whether structural synapses predict functional circuitry, with an interaction term to test whether the relationship differs between arousal states. The changes model tested whether structural synapses predict functional circuitry changes. To

demonstrate robustness, we also ran models controlling for potential confounds including inter-neuron distance, same brain area, and same cell type (Supplementary **Table 5**). Same cortical layer was excluded as a covariate because all synaptically connected neuron pairs in the EM data are within the same cortical layer, leaving no variation in this variable.

Predictive Correlation: Pearson correlation between model predictions and actual responses was computed per neuron and trial. The architecture was chosen based on the winning model at the Sensorium 2023 competition [36]; model predictions were obtained using this architecture, which employs an ensemble of 3D convolutional networks with depthwise separable convolutions and temporal attention mechanisms (see Supplementary Methods for architectural and training details). We trained this model architecture on MICrONS data using 7-fold cross-validation (out-of-fold predictions), so that each trial's prediction comes from a model trained on the other six folds; evaluation is on these held-out trials only. Model inputs include video frames, pupil center, pupil area, and running speed; the model was not trained to predict arousal state, and fold assignment did not use arousal labels, so differences in predictive correlation between low- and high-arousal trials are not an artifact of the training procedure. Mean predictive correlation was compared between arousal states using a paired t-test (high $-$ low; 119,913 units, 13 session-scans). Neurons were stratified by their low-arousal predictive correlation into fixed-threshold performance bins: Poor (< 0.1), Fair ($0.1 - 0.2$), Good ($0.2 - 0.3$), Very Good ($0.3 - 0.4$), and Excellent (> 0.4). The gradient across performance bins (Spearman correlation between per-bin mean baseline predictive correlation and mean Δ) was assessed by permutation test (10,000 shuffles of the five per-bin mean Δ values; one-sided). FDR correction (Benjamini-Hochberg) was applied to all bin-level comparisons. Brain area-specific correlations were computed separately for each visual area (V1, LM, AL, RL) with SEM error bars (Fisher z approximation).

All connectivity comparisons used non-parametric tests (Wilcoxon signed-rank, Mann-Whitney U, Spearman) due to non-normal distributions. For paired comparisons (e.g., low vs. high arousal within the same field), we used Wilcoxon signed-rank tests. For unpaired comparisons (e.g., within-area vs. between-area within the same state), we used Mann-Whitney U tests. Predictive correlation related hypothesis tests used paired t-tests given large sample size and paired design. All tests were two-sided except the following, whose one-sided direction was pre-specified: (1) within-area vs. between-area Jaccard similarity (one-sided Mann-Whitney U; local bias of cortical anatomy predicts intra $>$ inter); (2) overall spatial reach (one-sided bootstrap Spearman; prior literature predicts expanded spatial reach in high arousal [12]).

Multiple Comparisons Correction: We performed multiple statistical tests across different hypotheses. To control the false discovery rate (FDR), we applied the Benjamini-Hochberg procedure [62] within each hypothesis family. Families were defined as follows: (1) **Connectivity decomposition:** within-area vs. between-area density and edge strength (aggregate and per area pair), in one family. (2) **Layer-specific:** all layer-pair comparisons in one family. (3) **Cell-type spatial:** E→E and E→I distance-dependence tests in one family. (4) **Structure-function field-level:** the three field-level comparisons (FC vs. SC in low arousal, FC vs. SC in high arousal, SC vs. ΔFC) in one family. (5) **Predictive correlation:** bin-level and area-specific comparisons in one family. This family-wise approach is appropriate because each family addresses a distinct scientific question. For all hypotheses with multiple tests, we report FDR-corrected p-values in the main text. For single-test hypotheses (e.g., overall spatial correlation), no correction is needed and we report uncorrected p-values. Tests with FDR-corrected $p < 0.05$ (or $p < 0.05$ for single tests) were considered significant.

Effect Sizes: For correlation analyses, we report Spearman or Pearson correlation coefficients

as appropriate. For differences between conditions, we report raw differences (Δ) along with the relevant test statistic and p-value (e.g., Wilcoxon or Mann-Whitney for connectivity; paired t-test for predictive correlation comparisons where used).

All analyses were implemented in Python.

Data and Code Availability

All analysis code is publicly available at `https://github.com/abbasilab/mfc`. Raw calcium imaging data and electron microscopy reconstructions are available through the MICrONS project (`https://www.microns-explorer.org/`). The Time-aware PC (TPC) algorithm used to compute functional circuitry matrices is described in [10] and publicly available at `https://github.com/shlizee/TimeAwarePC`.

Acknowledgments

We thank members of the Abbasi Lab for their insightful feedback on this work. This study was partially supported by the National Institute of Mental Health of the National Institutes of Health under award number RF1MH128672 (model development and initial validation) and by the National Library of Medicine of the National Institutes of Health under award number R01LM014619 (celltype and layer specific analysis). The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

References

[1] Olaf Sporns et al. The human connectome: a structural description of the human brain. *PLoS computational biology*, 1(4):e42, 2005.

[2] Ed Bullmore and Olaf Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature reviews neuroscience*, 10(3):186–198, 2009.

[3] The MICrONS Consortium. Functional connectomics spanning multiple areas of mouse visual cortex. *Nature*, 640(8058):435–447, 2025.

[4] Wei-Chung Allen Lee and R Clay Reid. Specificity and randomness: structure–function relationships in neural circuits. *Current opinion in neurobiology*, 21(5):801–807, 2011.

[5] Nancy J Kopell, Howard J Gritton, Miles A Whittington, and Mark A Kramer. Beyond the connectome: the dynome. *Neuron*, 83(6):1319–1328, 2014.

[6] Rachel M Rahn, Lindsey M Brier, Annie R Bice, Matthew D Reisman, Joseph D Dougherty, and Joseph P Culver. Functional connectivity of the developing mouse cortex. *Cerebral Cortex*, 32(8):1755–1768, 2022.

[7] Viktor K Jirsa. Connectivity and dynamics of neural information processing. *Neuroinformatics*, 2(2):183–204, 2004.

[8] Emily S Finn, Xilin Shen, Dustin Scheinost, Monica D Rosenberg, Jessica Huang, Marvin M Chun, Xenophon Papademetris, and R Todd Constable. Functional connectome fingerprinting:

identifying individuals using patterns of brain connectivity. *Nature neuroscience*, 18(11):1664–1671, 2015.

[9] Andrew T Reid, Drew B Headley, Ravi D Mill, Ruben Sanchez-Romero, Lucina Q Uddin, Daniele Marinazzo, Daniel J Lurie, Pedro A Valdés-Sosa, Stephen José Hanson, Bharat B Biswal, et al. Advancing functional connectivity research from association to causation. *Nature neuroscience*, 22(11):1751–1760, 2019.

[10] Rahul Biswas and Eli Shlizerman. Statistical perspective on functional and causal neural connectomics: The time-aware pc algorithm. *PLOS Computational Biology*, 18(11):e1010653, 2022.

[11] Rahul Biswas and Somabha Mukherjee. Consistent causal inference from time series with pc algorithm and its time-aware extension. *Statistics and Computing*, 34(1):14, 2024.

[12] Matthew J McGinley et al. Cortical membrane potential signature of optimal states for sensory signal detection. *Neuron*, 87(1):179–192, 2015.

[13] Jacob Reimer et al. Pupil fluctuations track fast switching of cortical states during quiet wakefulness. *Neuron*, 84(2):355–362, 2014.

[14] Pierre-Olivier Polack et al. Cellular mechanisms of brain state–dependent gain modulation in visual cortex. *Nature neuroscience*, 16(9):1331–1339, 2013.

[15] Iain Stitt, Karl J Hollensteiner, Edgar Galindo-Leon, Florian Pieper, Eva Fiedler, Thomas Stieglitz, Gerhard Engler, Guido Nolte, and Andreas K Engel. Dynamic reconfiguration of cortical functional connectivity across brain states. *Scientific reports*, 7(1):8797, 2017.

[16] Iain Stitt et al. Arousal dependent modulation of thalamo-cortical functional interaction. *Nature Communications*, 9(1):2455, 2018.

[17] Julie A Harris et al. Hierarchical organization of cortical and thalamic connectivity. *Nature*, 575(7781):195–202, 2019.

[18] Daniel J Felleman and David C Van Essen. Distributed hierarchical processing in the primate cerebral cortex. *Cerebral cortex (New York, NY: 1991)*, 1(1):1–47, 1991.

[19] Yong-Jun Liu, Markus U Ehrengruber, Moritz Negwer, Han-Juan Shao, Ali H Cetin, and David C Lyon. Tracing inputs to inhibitory or excitatory neurons of mouse and cat visual cortex with a targeted rabies virus. *Current Biology*, 23(18):1746–1755, 2013.

[20] Siddhartha Joshi et al. Relationships between pupil diameter and neuronal activity in the locus coeruleus, colliculi, and cingulate cortex. *Neuron*, 89(1):221–234, 2016.

[21] Rahul Biswas and Eli Shlizerman. Statistical perspective on functional and causal neural connectomics: A comparative study. *Frontiers in Systems Neuroscience*, 16:817962, 2022.

[22] James H Marshel, Marina E Garrett, Ian Nauhaus, and Edward M Callaway. Functional specialization of seven mouse visual cortical areas. *Neuron*, 72(6):1040–1054, 2011.

[23] Mark L Andermann, Aaron M Kerlin, Demetris K Roumis, Lindsey L Glickfeld, and R Clay Reid. Functional specialization of mouse higher visual cortical areas. *Neuron*, 72(6):1025–1039, 2011.

[24] Quanxin Wang and Andreas Burkhalter. Area map of mouse visual cortex. *Journal of Comparative Neurology*, 502(3):339–357, 2007.

[25] Bassam V Atallah et al. Parvalbumin-expressing interneurons linearly transform cortical responses to visual stimuli. *Neuron*, 73(1):159–170, 2012.

[26] Hyun-Jae Pi et al. Cortical interneurons that specialize in disinhibitory control. *Nature*, 503(7477):521–524, 2013.

[27] Bosiljka Tasic, Vilas Menon, Thuc Nghi Nguyen, Tae Kyung Kim, Tim Jarsky, Zizhen Yao, Boaz Levi, Lucas T Gray, Staci A Sorensen, Tim Dolbeare, et al. Adult mouse cortical cell taxonomy revealed by single cell transcriptomics. *Nature neuroscience*, 19(2):335–346, 2016.

[28] Bosiljka Tasic, Zizhen Yao, Lucas T Graybuck, Kimberly A Smith, Thuc Nghi Nguyen, Darren Bertagnolli, Jeff Goldy, Emma Garren, Michael N Economo, Sarada Viswanathan, et al. Shared and distinct transcriptomic cell types across neocortical areas. *Nature*, 563(7729):72–78, 2018.

[29] Zizhen Yao, Cindy TJ Van Velthoven, Thuc Nghi Nguyen, Jeff Goldy, Adriana E Sedeno-Cortes, Fahimeh Baftizadeh, Darren Bertagnolli, Tamara Casper, Megan Chiang, Kirsten Crichton, et al. A taxonomy of transcriptomic cell types across the isocortex and hippocampal formation. *Cell*, 184(12):3222–3241, 2021.

[30] Wenzhi Sun, Zhongchao Tan, Brett D Mensh, and Na Ji. Thalamus provides layer 4 of primary visual cortex with orientation-and direction-tuned inputs. *Nature neuroscience*, 19(2):308–315, 2016.

[31] Matthew T Colonnese and Marnie A Phillips. Thalamocortical function in developing sensory circuits. *Current opinion in neurobiology*, 52:72–79, 2018.

[32] Farran Briggs and W Martin Usrey. Corticogeniculate feedback and visual processing in the primate. *The Journal of physiology*, 589(1):33–40, 2011.

[33] Seung-Hee Lee et al. Activation of specific interneurons improves v1 feature selectivity and visual perception. *Nature*, 488(7411):379–383, 2012.

[34] Stefano Panzeri, Jakob H Macke, Joachim Gross, and Christoph Kayser. Neural population coding: combining insights from microscopic and mass signals. *Trends in cognitive sciences*, 19(3):162–172, 2015.

[35] Reza Abbasi-Asl, Yuansi Chen, Adam Bloniarz, Michael Oliver, Ben DB Willmore, Jack L Gallant, and Bin Yu. The deeptune framework for modeling and characterizing neurons in visual cortex area v4. *bioRxiv*, page 465534, 2018.

[36] Polina Turishcheva, Paul Fahey, Michaela Vystrčilová, Laura Hansel, Rachel Froebe, Kayla Ponder, Yongrong Qiu, Konstantin Willeke, Mohammad Bashiri, Ruslan Baikulov, et al. Retrospective for the dynamic sensorium competition for predicting large-scale mouse primary visual cortex activity from videos. *Advances in Neural Information Processing Systems*, 37:118907–118929, 2024.

[37] Vernon B Mountcastle. The columnar organization of the neocortex. *Brain: a journal of neurology*, 120(4):701–722, 1997.

[38] Quanxin Wang, Olaf Sporns, and Andreas Burkhalter. Network analysis of corticocortical connections reveals ventral and dorsal processing streams in mouse visual cortex. *Journal of Neuroscience*, 32(13):4386–4399, 2012.

[39] Rinaldo D D'Souza, Quanxin Wang, Weiqing Ji, Andrew M Meier, Henry Kennedy, Kenneth Knoblauch, and Andreas Burkhalter. Hierarchical and nonhierarchical features of the mouse visual cortical network. *Nature communications*, 13(1):503, 2022.

[40] Mateo Vélez-Fort, Charly V Rousseau, Christian J Niedworok, Ian R Wickersham, Ede A Rancz, Alexander PY Brown, Molly Strom, and Troy W Margrie. The stimulus selectivity and connectivity of layer six principal cells reveals cortical microcircuits underlying visual processing. *Neuron*, 83(6):1431–1443, 2014.

[41] Cristopher M Niell and Michael P Stryker. Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron*, 65(4):472–479, 2010.

[42] Rinaldo David D'Souza, Andrew Max Meier, Pawan Bista, Quanxin Wang, and Andreas Burkhalter. Recruitment of inhibition and excitation across mouse visual cortex depends on the hierarchy of interconnecting areas. *elife*, 5:e19332, 2016.

[43] Gyorgy Lur, Martin A Vinck, Lan Tang, Jessica A Cardin, and Michael J Higley. Projection-specific visual feature encoding by layer 5 cortical subnetworks. *Cell reports*, 14(11):2538–2545, 2016.

[44] Alexander S Ecker, Philipp Berens, R James Cotton, Manivannan Subramaniyan, George H Denfield, Cathryn R Cadwell, Stelios M Smirnakis, Matthias Bethge, and Andreas S Tolias. State dependence of noise correlations in macaque primary visual cortex. *Neuron*, 82(1):235–248, 2014.

[45] Danielle S Bassett and Olaf Sporns. Network neuroscience. *Nature neuroscience*, 20(3):353–364, 2017.

[46] Lindsey L Glickfeld, Mark L Andermann, Vincent Bonin, and R Clay Reid. Cortico-cortical projections in mouse visual cortex are functionally target specific. *Nature neuroscience*, 16(2):219–226, 2013.

[47] Frances S Chance et al. Gain modulation from background synaptic input. *Neuron*, 35(4):773–782, 2002.

[48] Emilio Salinas and Peter Thier. Gain modulation: a major computational principle of the central nervous system. *Neuron*, 27(1):15–21, 2000.

[49] Yu Fu et al. A cortical circuit for gain control by behavioral state. *Cell*, 156(6):1139–1152, 2014.

[50] Johannes J Letzkus et al. A disinhibitory microcircuit for associative fear learning in the auditory cortex. *Nature*, 480(7377):331–335, 2011.

[51] Casey M Schneider-Mizell, Agnes L Bodor, Derrick Brittain, JoAnn Buchanan, Daniel J Bumbarger, Leila Elabbady, Clare Gamlin, Daniel Kapner, Sam Kinn, Gayathri Mahalingam, et al. Inhibitory specificity from a connectomic census of mouse visual cortex. *Nature*, 640(8058):448–458, 2025.

[52] M-Marsel Mesulam. From sensation to cognition. *Brain: a journal of neurology*, 121(6):1013–1052, 1998.

[53] Yuri B Saalmann et al. The pulvinar regulates information transmission between cortical areas based on attention demands. *Science*, 337(6095):753–756, 2012.

[54] Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. MIT Press, 2nd edition, 2000.

[55] Garrett T Neske. The slow oscillation in cortical and thalamic networks: mechanisms and functions. *Frontiers in neural circuits*, 9:88, 2016.

[56] Carsen Stringer et al. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364(6437):eaav7893, 2019.

[57] Shinsuke Shibata, Yuji Komaki, Fumiko Seki, Michiko O Inouye, Toshihiro Nagai, and Hideyuki Okano. Connectomics: comprehensive approaches for whole-brain mapping. *Microscopy*, 64(1):57–67, 2015.

[58] Tsai-Wen Chen, Trevor J Wardill, Yi Sun, Stefan R Pulver, Sabine L Renninger, Amy Baohan, Eric R Schreiter, Rex A Kerr, Michael B Orger, Vivek Jayaraman, et al. Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature*, 499(7458):295–300, 2013.

[59] Chin-An Wang and Douglas P Munoz. A circuit for pupil orienting responses: implications for cognitive modulation of pupil size. *Current opinion in neurobiology*, 33:134–140, 2015.

[60] Clive WJ Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, pages 424–438, 1969.

[61] Mingzhou Ding, Yonghong Chen, and Steven L Bressler. Granger causality: basic theory and application to neuroscience. *Handbook of time series analysis: recent theoretical developments and applications*, pages 437–460, 2006.

[62] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, 57(1):289–300, 1995.

[63] Roman Solovyev. Sensorium 2023 NeurIPS competition - 1st place solution. `https://github.com/lRomul/sensorium`, 2023. Accessed: 2026-02-21.

[64] Menglong Zhu and Andrey Zhmoginov Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4510–4520, 2018.

[65] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

Supplementary Methods

Model Architecture and Training

For predictive correlation analysis, the architecture was chosen based on the winning model at the Sensorium 2023 NeurIPS competition [36, 63], which achieved first place in predicting large-scale mouse visual cortex activity from videos. The model employs an ensemble approach combining multiple architectural innovations.

Architecture Components: The base architecture consists of 3D convolutional networks with depthwise separable convolutions [64], enabling efficient spatiotemporal feature extraction from video inputs. The model incorporates temporal attention mechanisms [65] to capture long-range dependencies across time. The ensemble combines predictions from multiple models trained with different configurations, improving robustness and generalization.

Training on MICrONS Data: We trained this model architecture on the MICrONS dataset, using the same visual stimuli (natural movies and synthetic parametric videos) and behavioral signals (pupil center, pupil area, running speed) that were available during calcium imaging. Trials were divided into 7 folds for out-of-fold predictions, so that each trial's prediction comes from a model trained on the other six folds, ensuring that all predictions were evaluated on held-out data. The model was trained to predict neural responses (calcium fluorescence traces from the MICrONS processed dataset) from video frames and the behavioral inputs (pupil center, pupil area, running speed), learning a mapping from these inputs to neural activity patterns.
Relationship to Arousal: Model inputs include video frames, pupil center, pupil area, and running speed (i.e., the model can use moment-to-moment state when predicting responses). The train/test and fold splits were performed at the trial level without regard to arousal state, and the model was not trained to predict arousal state as a target. Trial-level arousal classification (median split of pupil area) was applied only after obtaining predictions, for comparison. Thus differences in predictive correlation between low- and high-arousal trials are not an artifact of training on arousal labels.

Performance: On the original Sensorium 2023 competition dataset, this architecture achieved single-trial correlations of 0.2913 on the main track and 0.2215 on the bonus track (final test set) [63]. On the MICrONS data used in this study, we computed the same single-trial correlation metric as in the Sensorium competition (without regard to arousal state): mean correlation between predicted and actual responses per trial, averaged over trials and units. The model achieved a single-trial correlation of 0.2878 on out-of-fold predictions and 0.3198 on the held-out test set (752 trials). The model's ability to capture stimulus-response relationships makes it well-suited for assessing predictive correlation as a measure of how well neural responses can be predicted from the visual input.

Predictive correlation: full pipeline. For transparency and reproducibility, the end-to-end pipeline is as follows.

(1) Data. MICrONS calcium imaging data: for each session and scan, we use processed calcium fluorescence traces (responses) and the corresponding inputs for every trial. Inputs to the model are video frames, pupil center, pupil area, and running speed; target is the trial-wise response matrix (units $\times$ time). The unit of analysis is session-scan-unit (146,388 units across 16 session-scans).

(2) Train/test split and out-of-fold predictions. A subset of 752 trials was held out from fold assignment and not used in training (separate test set). The remaining trials (6,672 across 16

session-scans) were split into $K = 7$ folds. For each fold, the model is trained on the other six folds (∼86% of those trials) and predicts on the held-out fold (∼14%). Thus each fold trial receives a prediction from a model that did not see that trial during training (out-of-fold predictions). The split is at the trial level, is identical for all neurons within a session-scan, and does not use arousal state or any trial metadata. Reported predictive correlations are computed from out-of-fold predictions. The same analysis was repeated on the held-out test set; the finding (higher predictive correlation in high arousal and the redistribution pattern) replicates on this set.

(3) Model and training. The model architecture was chosen based on the winning model at the Sensorium 2023 competition: a depthwise separable 3D CNN ("dwiseneuro") with temporal attention, with approximately 309 million trainable parameters, trained to predict neural responses from video and behavioral inputs (pupil center, pupil area, running speed). Training minimizes Poisson negative log-likelihood (NLL) between predicted and actual responses on the training trials of each fold. Optimizer: AdamW with base learning rate 3×10^{-4} (with warmup and cosine decay to 1% of base), weight decay 0.05, batch size 32. Training runs for 21 epochs (3-epoch linear warmup followed by 18 epochs with cosine learning-rate decay); the held-out fold is used only for monitoring validation loss and validation correlation and is never used for training or early stopping. Training was performed on two NVIDIA RTX 5000 Ada Generation GPUs; each fold took approximately 12 hours. Model inputs include video, pupil center, pupil area, and running speed; arousal state (median-split label) is not used for fold assignment, as a training target, or in model selection.

(4) Evaluation. For each trial, the trained model (or ensemble) produces a predicted response matrix of shape (units × time). Predictions are saved per trial so that each trial has a matching pair (predicted trace, actual trace) for every unit. The out-of-fold predictions are used in the arousal analysis reported in the main text.

(5) Arousal assignment and predictive correlation. After all predictions are obtained, each trial is classified as low or high arousal using the same criterion as in the main Methods: session-scan-specific median split of mean pupil area over the trial, with $\geq 80\%$ of trial duration above (high) or below (low) the median. For each session-scan and each unit, we collect all low-arousal trials and all high-arousal trials. Within each state, we concatenate the predicted and actual traces across trials (along the time dimension), yielding one long predicted trace and one long actual trace per unit per state. The predictive correlation in that state is the Pearson correlation between these concatenated predicted and actual traces. Thus each unit has exactly two values: $corr_{\text{low}}$ (low-arousal trials) and $corr_{\text{high}}$ (high-arousal trials). These are the quantities used in the main text (stratification by low-arousal performance bin, area- and cell-type analyses, and all statistical tests). Arousal state (median-split label) was not used for fold assignment or as a training target; it is applied only after predictions are obtained, to label trials and compute state-specific correlations.

Training dynamics and convergence. Training was run for 21 epochs with no early stopping across seven cross-validation folds. Loss decreased and validation correlation increased over training epochs, confirming convergence (Supplementary **Figure S3**). All predictive correlation results reported in the paper use out-of-fold predictions from this run across all 13 session-scans.

Supplementary Tables: Full Statistical Test Results

Notation: $p < 0.05$ = *, $p < 0.01$ = **, $p < 0.001$ = ***, n.s. = not significant. All FDR corrections use Benjamini–Hochberg (BH) procedure.

Table 1: Supplementary Table 1. Within-state area-level functional circuitry. Bootstrap mean ± 95% CI and Wilcoxon one-sample test (vs. 0, paired by field) per directed area pair, per arousal state. FDR-BH corrected across all 32 area-pair × state combinations per metric (density: 32 tests; strength: 32 tests). Density values are connection fractions (range 0–1); strength values are mean Spearman r among connected pairs. n = number of fields with both states present. Significance: *** $p < 0.001$.

Area pair	**Low [95% CI]**	**High [95% CI]**	$p_{\mathbf{fdr}}$ **low**	$p_{\mathbf{fdr}}$ **high**
Density (connection fraction)				
Within-area pairs (n fields: V1=101, LM=94, AL=68, RL=50)				
V1→V1	0.00492 [0.00451, 0.00535]	0.00525 [0.00493, 0.00559]	4.26e-17***	4.26e-17***
LM→LM	0.01130 [0.00874, 0.01385]	0.01091 [0.00951, 0.01230]	5.34e-15***	3.67e-14***
AL→AL	0.02351 [0.01627, 0.03246]	0.02714 [0.01927, 0.03518]	1.87e-10***	1.87e-10***
RL→RL	0.01306 [0.01065, 0.01496]	0.01242 [0.01112, 0.01380]	1.48e-09***	1.93e-09***
Between-area pairs (n fields per pair in parentheses)				
V1→LM (n=95)	0.00268 [0.00240, 0.00294]	0.00334 [0.00265, 0.00403]	2.44e-16***	2.44e-16***
V1→AL (n=82)	0.00164 [0.00088, 0.00240]	0.00308 [0.00247, 0.00373]	8.14e-13***	1.05e-13***
V1→RL (n=50)	0.00206 [0.00155, 0.00255]	0.00229 [0.00200, 0.00256]	6.32e-15***	6.32e-15***
LM→V1 (n=95)	0.00275 [0.00255, 0.00296]	0.00367 [0.00335, 0.00400]	2.97e-16***	2.44e-16***
LM→AL (n=76)	0.00345 [0.00256, 0.00437]	0.00498 [0.00326, 0.00668]	7.01e-11***	7.01e-11***
LM→RL (n=44)	0.00136 [0.00089, 0.00183]	0.00201 [0.00126, 0.00276]	1.85e-07***	2.48e-07***
AL→V1 (n=82)	0.00237 [0.00122, 0.00373]	0.00307 [0.00217, 0.00397]	8.14e-13***	2.06e-13***
AL→LM (n=76)	0.00322 [0.00199, 0.00451]	0.00546 [0.00362, 0.00730]	5.61e-10***	1.87e-10***
AL→RL (n=50)	0.01375 [0.00811, 0.01946]	0.02051 [0.01991, 0.02110]	2.74e-09***	1.48e-09***
RL→V1 (n=50)	0.00173 [0.00127, 0.00219]	0.00202 [0.00149, 0.00256]	1.48e-09***	1.93e-09***
RL→LM (n=44)	0.00232 [0.00101, 0.00367]	0.00187 [0.00148, 0.00226]	2.48e-07***	2.48e-07***
RL→AL (n=50)	0.01137 [0.00518, 0.01761]	0.01735 [0.01534, 0.01929]	1.93e-09***	1.48e-09***
Strength (mean Spearman r among connected pairs)				
Within-area pairs (n fields: V1=101, LM=94, AL=68, RL=50)				
V1→V1	0.308 [0.300, 0.315]	0.271 [0.262, 0.280]	4.26e-17***	4.26e-17***
LM→LM	0.260 [0.242, 0.279]	0.217 [0.202, 0.233]	5.34e-15***	2.16e-14***
AL→AL	0.318 [0.299, 0.338]	0.322 [0.299, 0.346]	1.14e-10***	1.14e-10***
RL→RL	0.366 [0.354, 0.379]	0.321 [0.304, 0.339]	8.75e-15***	1.42e-14***
Between-area pairs (n fields per pair in parentheses)				
V1→LM (n=95)	0.194 [0.183, 0.205]	0.191 [0.170, 0.217]	2.44e-16***	2.44e-16***
V1→AL (n=82)	0.119 [0.102, 0.135]	0.129 [0.118, 0.141]	5.18e-13***	6.07e-14***
V1→RL (n=50)	0.176 [0.163, 0.190]	0.190 [0.175, 0.205]	6.32e-15***	6.32e-15***
LM→V1 (n=95)	0.334 [0.317, 0.350]	0.305 [0.284, 0.327]	2.97e-16***	2.44e-16***
LM→AL (n=76)	0.131 [0.101, 0.161]	0.148 [0.107, 0.194]	4.49e-11***	4.49e-11***
LM→RL (n=44)	0.224 [0.189, 0.262]	0.181 [0.138, 0.224]	4.05e-11***	6.65e-11***
AL→V1 (n=82)	0.280 [0.224, 0.337]	0.266 [0.209, 0.336]	5.18e-13***	1.24e-13***
AL→LM (n=76)	0.189 [0.148, 0.232]	0.177 [0.148, 0.205]	3.50e-10***	1.14e-10***
AL→RL (n=50)	0.277 [0.261, 0.293]	0.264 [0.248, 0.281]	2.53e-14***	8.75e-15***
RL→V1 (n=50)	0.320 [0.292, 0.353]	0.250 [0.224, 0.277]	8.75e-15***	1.42e-14***
RL→LM (n=44)	0.089 [0.074, 0.104]	0.093 [0.081, 0.104]	6.65e-11***	6.65e-11***
RL→AL (n=50)	0.368 [0.346, 0.389]	0.403 [0.375, 0.431]	1.42e-14***	8.75e-15***

Continued on next page

Table 4 - continued

Comparison	Z	p (raw)	p (FDR)	Sig.

Table 4: Supplementary Table 4. Predictive performance pairwise comparisons (best group vs. others). Fisher z-test on neuron-level Spearman correlations. Brain area (best = LM), cortical layer (best = L6), cell type (best = L2b). FDR-BH corrected within each family. Test directions: best vs. other (one-tailed); for layer family, L4 vs. L5 (two-tailed) and L2/3 > L5 (one-tailed); for cell type, within-L5 directional (one-tailed), L4 and L2 pairs (two-tailed).

Comparison	Z	p (raw)	p (FDR)	Sig.
Brain area (best = LM, $r = 0.4876$, $n = 18{,}703$)				
LM > V1	3.410	0.0003	0.0005	***
LM > RL	3.262	0.0006	0.0006	***
LM > AL	4.934	0.0000	0.0000	***
Cortical layer (best = L6, $r = 0.5282$, $n = 680$)				
L6 > L2/3	2.974	0.0015	0.0018	**
L6 > L4	6.209	0.0000	0.0000	***
L6 > L5	6.052	0.0000	0.0000	***
L4 vs. L5 (two-tailed)	0.334	0.7385	0.7385	n.s.
L2/3 > L5	5.787	0.0000	0.0000	***
Cell type (best = L2b, $r = 0.5582$, $n = 627$)				
L2b > L3b	3.035	0.0012	0.0022	**
L2b > L3a	5.890	0.0000	0.0000	***
L2b > L4b	6.744	0.0000	0.0000	***
L2b > L4a	6.652	0.0000	0.0000	***
L2b > L2a	1.902	0.0286	0.0429	*
L2b > L2c	2.515	0.0060	0.0097	**
L2b > L5a	3.459	0.0003	0.0005	***
L2b > L4c	4.427	0.0000	0.0000	***
L2b > L5ET	6.543	0.0000	0.0000	***
L2b > L5b	6.373	0.0000	0.0000	***
L2b > L6short-a	1.101	0.1354	0.1625	n.s.
L5a > L5b	3.634	0.0001	0.0004	***
L5a > L5ET	3.462	0.0003	0.0005	***
L5ET > L5b	0.593	0.2765	0.3111	n.s.
L4a vs. L4b (two-tailed)	−0.392	0.6952	0.6952	n.s.
L4a vs. L4c (two-tailed)	−2.037	0.0417	0.0577	n.s.
L4b vs. L4c (two-tailed)	−1.834	0.0666	0.0857	n.s.
L2a vs. L2c (two-tailed)	0.788	0.4305	0.4558	n.s.

Table 2: Supplementary Table 2. Within-state layer-specific functional circuitry. Bootstrap mean ± 95% CI per layer combination, per arousal state. Density values are connection fractions; strength values are mean Spearman r among connected pairs. p-values are from permutation bootstrap (two-sided test of low vs. high arousal difference); no within-state test vs. zero is reported here. FDR columns (str) indicate whether BH correction yielded significance for high > low and low > high directions, respectively.

		Density mean [95% CI]		Strength mean [95% CI]		High vs. Low (Δ)		
Layer pair	**mean presence**	**Low arousal**	**High arousal**	**Low arousal**	**High arousal**	**p (two-sided)**	**Δstr (FDR low>hi)**	**Δstr (FDR hi>lo)**
Within-layer (same source and target layer)								
Within-L2/3	low 0.00231, high 0.00268	0.00208 [– , –]†	0.00268 [– , –]†	0.241 [0.178, 0.310]	0.256 [0.231, 0.280]	0.1172 (density); 0.7204 (str)	No	No
Within-L4	low 0.00293, high 0.00212	0.00293 [– , –]†	0.00212 [– , –]†	0.233 [0.197, 0.272]	0.184 [0.158, 0.207]	0.3152 (density); 0.0324 (str)	Yes	No
Within-L5	low 0.00267, high 0.00383	0.00267 [– , –]†	0.00383 [– , –]†	0.210 [0.171, 0.249]	0.159 [0.133, 0.186]	0.3792 (density); 0.0308 (str)	Yes	No
Between-layer (source and target in different layers)								
Between-L2/3	low 0.00116, high 0.00172	0.00116 [– , –]†	0.00172 [– , –]†	0.223 [0.167, 0.277]	0.189 [0.145, 0.232]	0.0060 (density); 0.4096 (str)	No	No
Between-L4	low 0.00100, high 0.00151	0.00100 [– , –]†	0.00151 [– , –]†	0.245 [0.200, 0.290]	0.197 [0.145, 0.251]	0.4956 (density); 0.3580 (str)	No	No
Between-L5	low 0.00337, high 0.00231	0.00337 [– , –]†	0.00231 [– , –]†	0.171 [0.123, 0.219]	0.143 [0.108, 0.177]	0.5064 (density); 0.3984 (str)	No	No

† Density bootstrap CIs by layer were not computed; mean per-field density is shown instead. Strength CIs are from the bootstrap distribution.

Table 3: Supplementary Table 3. Structure-function field-level (Fig. 4A). One-sample Wilcoxon signed-rank test on per-field Spearman r values (median vs. zero). One FDR across three tests (FC vs. SC low; FC vs. SC high; SC vs. ΔFC). Field-level Wilcoxon on per-field Spearman r; $n = 43$ fields each.

Test	**n fields**	**mean r**	**p (raw)**	**p (FDR)**	**Sig.**
FC vs. SC (low arousal)	43	0.141	1.14e-13	3.41e-13	***
FC vs. SC (high arousal)	43	0.098	1.59e-12	2.39e-12	***
SC vs. ΔFC (change)	43	−0.029	1.65e-03	1.65e-03	***

Supplementary Tables: Additional Results

Table 5: Robustness of structure-function relationships to covariates. Mixed-effects model results comparing models with and without covariates (distance, same brain area, same cell type). The structure-function relationships remain significant after controlling for these potential confounds, demonstrating robustness of the main findings. Note: same cortical layer was excluded from models due to no variance (all connections were within the same layer). $n = 6,597$ connections across 43 fields for all models.

Model	Structural (Low)	Structural (High)	Interaction
Baseline Models (predicting functional circuitry)			
Without covariates	0.0087***	0.0054***	-
With covariates	0.0061***	0.0031***	-0.0030***

Model	Structural (Changes)
Changes Models (predicting functional circuitry changes)	
Without covariates	-0.0034***
With covariates	-0.0024***

***$p < 0.001$. All structural effects remain highly significant ($p = 1.00 \times 10^{-9}$) after controlling for covariates (distance, same area, same cell type). Interaction term (reference model) tests whether structure-function coupling differs between arousal states ($p = 3.49 \times 10^{-12}$). Reference models predict functional circuitry in low and high arousal states separately, with an interaction term. Change models predict functional circuitry changes (high − low arousal).

Table 6: Top 20 neurons by predictive correlation in high arousal. Neurons are ranked by predictive correlation (Pearson correlation between model-predicted and actual responses) in the high-arousal state. Columns show the same metric in low arousal, in high arousal, and the change (Delta = high − low). Among these top performers, some had moderate or low predictive correlation in low arousal; others were already high in low arousal.

Unit ID	Sess.	Region	Pred. corr. (low)	Pred. corr. (high)	Delta
2523	6	V1	0.859	0.836	−0.023
8183	7	V1	0.577	0.798	0.221
9024	7	V1	0.617	0.772	0.155
9018	7	AL	0.569	0.769	0.200
8045	6	V1	0.657	0.765	0.108
1491	7	RL	0.593	0.760	0.167
7374	5	AL	0.590	0.755	0.165
1162	7	LM	0.595	0.746	0.151
968	7	V1	0.510	0.743	0.233
9023	7	V1	0.441	0.740	0.299
8550	7	AL	0.441	0.738	0.297
8113	7	V1	0.621	0.734	0.113
5230	5	V1	0.104	0.734	0.630
3001	8	LM	0.549	0.733	0.184
1096	7	V1	0.588	0.733	0.145
9091	7	V1	0.441	0.733	0.292
4997	5	V1	0.571	0.729	0.158
7788	7	V1	0.431	0.726	0.295
8276	5	LM	0.617	0.725	0.108
567	6	V1	0.649	0.725	0.076

Supplementary Figures

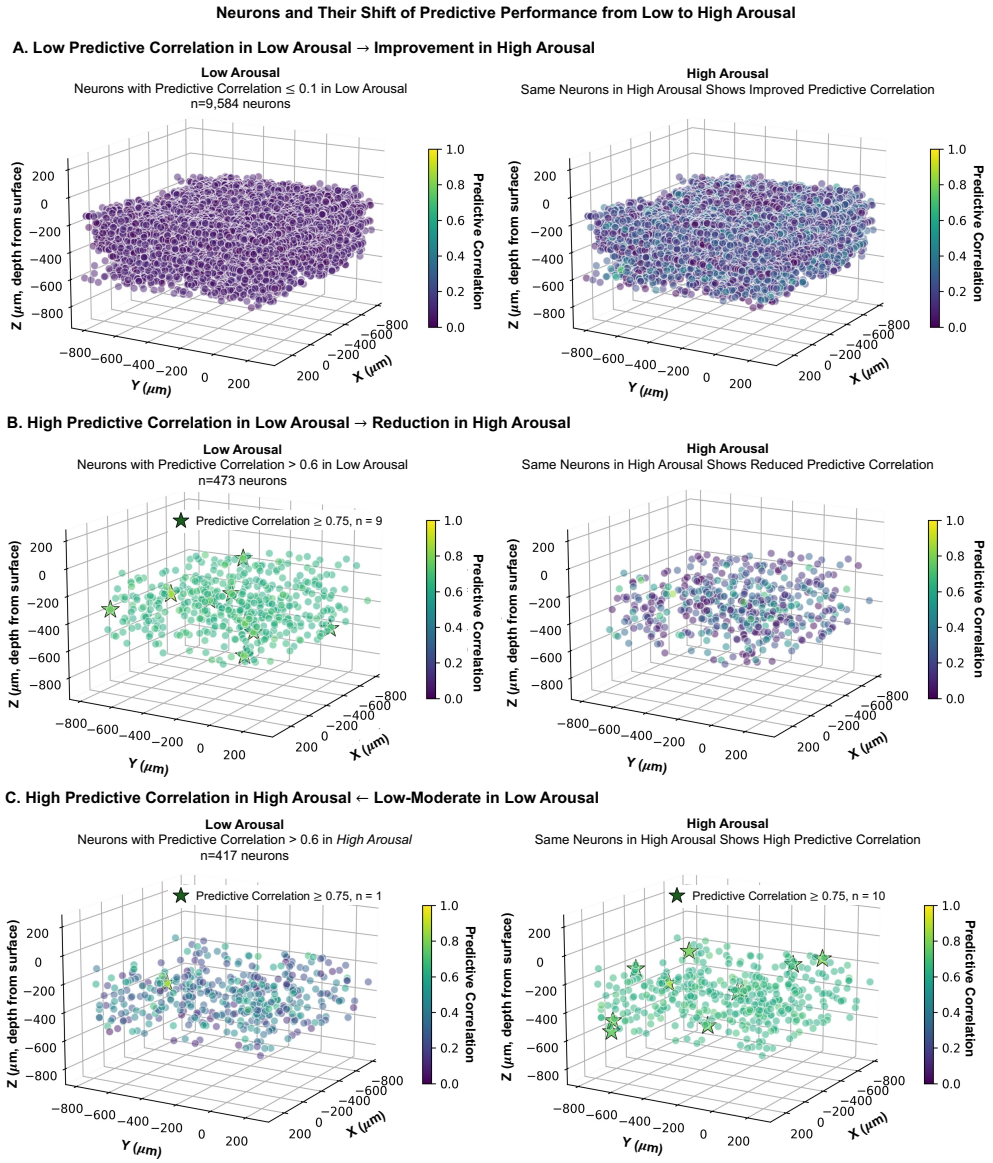

Figure S1: Spatial distribution of neurons by predictive-correlation change across arousal states. 3D scatter plots (inverted Z-axis: surface at top, depth downward) show cortical locations of neurons in each predictive-performance subgroup, colour-coded by predictive correlation. (A) Neurons with low predictive correlation in low arousal (correlation ≤ 0.1, $n = 9{,}584$); same neurons in high arousal has improved predictive correlation. (B) Neurons with high predictive correlation in low arousal (correlation > 0.6, $n = 473$; exceptional correlation ≥ 0.75, $n = 9$); same neurons in high arousal has reduced predictive correlation. (C) Neurons that reach high predictive correlation in high arousal (correlation > 0.6, $n = 417$) had low-moderate correlation in low arousal. Neurons are distributed broadly across the cortical volume rather than clustered in specific regions.

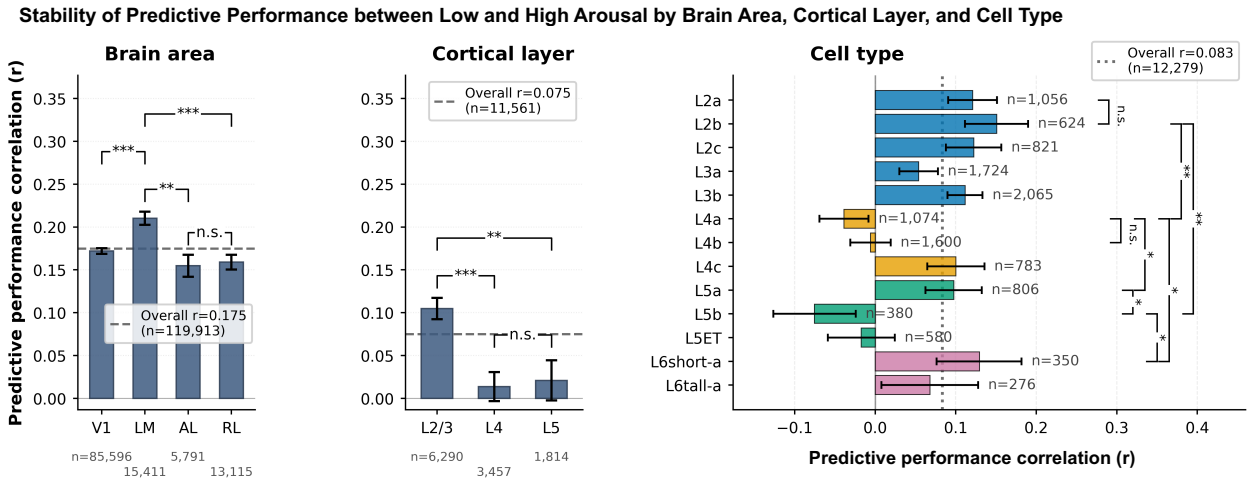

Figure S2: Correlation of predictive performance by brain area, cortical layer, and cell type. Pearson r between low- and high-arousal predictive correlations across neurons, per group. Asterisks mark FDR-corrected pairwise comparisons (Fisher z-test; key as in Figure 1). Full pairwise statistics in Supplementary Table 4.

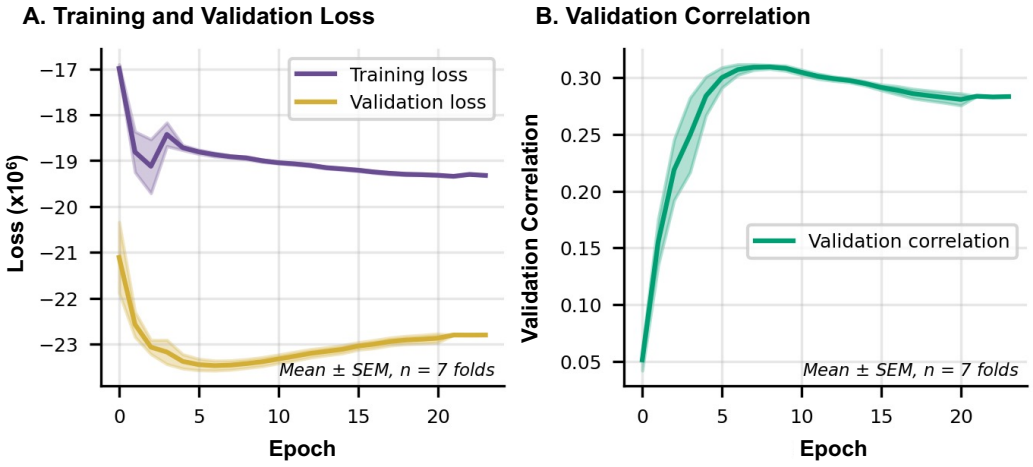

Figure S3: Predictive model training dynamics. (A) Training and validation loss (Poisson negative log-likelihood (NLL), scaled; lower is better) vs. epoch. (B) Validation correlation vs. epoch: mean prediction-response (Pearson) correlation across units on held-out validation trials; solid line is mean across 7 folds, shaded band is mean ± SEM. Results shown for all 13 session-scans. The model was trained for 21 epochs (3 warmup + 18 main) with no early stopping; reported results are from out-of-fold predictions.