

The Thesis Committee for Mohamed Salim M'jahad
certifies that this is the approved version of the following thesis:

**Neural Coding of Orientation in Macaque Primary Visual
Cortex Revealed by Two-Photon Calcium Imaging**

SUPERVISING COMMITTEE:

Eyal Seidemann, Supervisor

Wilson S. Geisler

**Neural Coding of Orientation in Macaque Primary Visual
Cortex Revealed by Two-Photon Calcium Imaging**

by
Mohamed Salim M'jahad

Thesis

Presented to the Faculty of the Graduate School of
The University of Texas at Austin
in Partial Fulfillment
of the Requirements
for the Degree of

Master of Science

**The University of Texas at Austin
July 2025**

Acknowledgments

I thank my advisor, Eyal Seidemann, for his guidance and for creating the space in which this project could come to life. I also thank the members of the lab for their support throughout this work. I believe the lab is quietly doing important work, and I'm grateful to have spent formative years within it.

I am especially grateful to Yuzhi Chen, whose expertise, persistence, and foundational work made this dataset possible. Without him, this project—and the insights within it—would not exist.

I also want to acknowledge the nonhuman primates with whom I had the privilege of working. It is not lost on me what it means to collaborate with other sentient beings in the name of science and human understanding. I feel a deep responsibility for their care and contribution.

Although this technique failed with most animals, I want to thank the monkeys who carried this work forward: Shushu, Jupiter, Salt, and especially Earth, whom I trained from scratch as a naïve monkey and who, four years after the previous subject (Loki), gave us the majority of the data in this study. Their cooperation and presence made this research possible.

I was supported in part by a fellowship from UM6P (Mohammed VI Polytechnic University), and I'm thankful for their support.

To my parents: thank you for your love, your strength, and for giving me life.

To my sisters: thank you for your constant encouragement.

To my friends—every one of whom is close, and all of whom know who they are—thank you for your companionship, your ideas, and your presence.

Abstract

Neural Coding of Orientation in Macaque Primary Visual Cortex Revealed by Two-Photon Calcium Imaging

Mohamed Salim M'jahad, M.S.
The University of Texas at Austin, 2025

SUPERVISOR: Eyal Seidemann

Abstract

Neurons in primary visual cortex (V1) exhibit selective responses to oriented stimuli, but how their population-level correlation structure is modulated by sensory input remains incompletely understood. This thesis presents a characterization of orientation coding and noise correlations in macaque V1 using two-photon calcium imaging of layer 2/3 neurons in awake animals.

Our central finding is that visual stimulation imposes a modest but detectable tuning-dependent pattern on shared variability: noise correlations show a gradual decline from similarly tuned pairs (0.15) to orthogonal pairs (0.08). In contrast, spontaneous activity exhibits flat correlation structure, lacking any tuning dependence. However, this functional organization is largely abolished when controlling for spatial distance, suggesting it primarily reflects local circuit connectivity rather than long-range functional relationships.

We recorded 2,667 neurons across 16 imaging sessions in two macaques (M22 and M36), identifying 1,340 (50%) as orientation-selective based on circular Gaussian fits to tuning curves. These neurons exhibited narrow bandwidths (mean = 17.3° , median = 16.6°), high response amplitudes (mean = 1.93 z-score units), and strong normalized tuning (mean NTA = 83.2%).

Correlation analyses using 64,168 neuron pairs revealed modest but significant increases in noise correlations during evoked compared to spontaneous conditions (mean $r_{\text{evoked}} = 0.097$, $r_{\text{spont}} = 0.072$, $p = 8.8 \times 10^{-113}$). However, the key difference was structural: evoked correlations showed modest orientation dependence, while spontaneous correlations remained flat.

These results suggest that visual input not only increases activity but may also reshape the correlation architecture of cortical populations. Given our current findings, the tuning-dependent reorganization of shared variability could potentially

influence coding efficiency, though the functional implications of this organization require further investigation.

Table of Contents

Abstract	5
List of Figures	9
Chapter 1: Introduction	19
1.1 Visual Processing in Primary Visual Cortex	19
1.2 Population Coding and Noise Correlations	19
1.3 Technical Advances in Population Recording	20
1.4 Research Objectives	21
Chapter 2: Results	23
2.1 Dataset Overview	23
2.2 Representative Neural Responses	27
2.3 Single-Neuron Orientation Tuning Properties	30
2.3.1 Tuning Parameter Distributions	30
2.3.2 Cross-Parameter Relationships	34
2.4 Population-Level Orientation Organization	40
2.4.1 Orientation Preference Distributions	40
2.5 Spatial Organization of Orientation Preferences	40
2.5.1 Orientation Preference Maps	40
2.5.2 Quantitative Spatial Organization Analysis	42
2.5.3 Signal Quality and Functional Classification	46
2.6 Stimulus-Dependent Modulation of Response Variability	47
2.6.1 Individual Neuron Variability Profiles	47
2.6.2 Population-Level Variability Analysis	49
2.6.3 Correlations with Orientation Tuning Properties	56
2.7 Noise Correlations and Population Dynamics	58
2.7.1 Global Noise Correlation Statistics	58
2.7.2 Orientation-Dependent Correlation Structure	61
2.7.3 Distance-Dependent Correlation Structure	63
2.8 Session-wise Variability and Consistency	66
2.9 Population Decoding of Orientation Information	68
2.9.1 Overall Decoding Performance	68
2.9.2 Fine versus Coarse Orientation Discrimination	68
2.9.3 Impact of Noise Correlations on Population Decoding	73

Chapter 3: Methods	76
3.1 Experimental Subjects and Surgical Procedures	76
3.2 Viral Expression of GCaMP6f	76
3.3 Two-Photon Calcium Imaging	76
3.3.1 Imaging System	76
3.3.2 Motion Correction and Image Processing	77
3.4 Visual Stimulation	77
3.4.1 Stimulus Generation and Presentation	77
3.4.2 Behavioral Control	78
3.5 Data Analysis	78
3.5.1 Fluorescence Signal Processing	78
3.5.2 Orientation Tuning Analysis	79
3.5.3 Functional Classification	80
3.5.4 Noise Correlation Analysis	80
3.5.5 Response Variability Analysis	81
3.5.6 Statistical Analysis	83
3.5.7 Orientation Dependence Analysis	83
3.5.8 Spatial Organization Analysis	84
3.5.9 Population Decoding Analysis	86
3.5.10 Cross-Session and Cross-Animal Analysis	90
Chapter 4: Discussion	91
4.1 Summary of Main Findings	91
4.2 Comparison with Previous Studies	92
4.2.1 Single-Neuron Tuning Properties	92
4.2.2 Noise Correlations and Functional Structure	92
4.3 Implications for Visual Processing and Coding Fidelity	92
4.3.1 Noise Correlations and Population Coding Constraints	92
4.3.2 Coding Regime Transitions	94
4.3.3 Implications for Sensory Processing	94
4.4 Technical Considerations and Limitations	95
4.5 Future Directions	95
4.6 Conclusions	96
Chapter 5: Conclusion	98
Works Cited	100
Vita	103

List of Figures

2.1	Example recording session showing cellular spatial organization. This figure illustrates data from a representative recording session (M22D20210302R2TS, 88 total neurons). (Left) Mean fluorescence image of the field of view showing the spatial distribution and morphology of recorded neurons. (Center) Segmented cell masks colored by functional classification: orientation-selective neurons (green), non-orientation-selective neurons (orange), and non-responsive neurons (blue). (Right) The same cell masks colored by preferred orientation for orientation-selective neurons only, using an HSV color scheme where hue represents preferred orientation (0° - 180°). This visualization demonstrates the cellular resolution of our two-photon imaging approach and illustrates the spatial distribution of functionally characterized neurons across the cortical field of view. Scale bar and field dimensions are provided in the original figure.	24
2.2	Detailed breakdown of a representative recording session. This figure illustrates the per-session analysis pipeline using data from a single, representative recording with a total of 234 neurons. The top row shows the functional classification results, presented as a pie chart of percentages (left) and a bar chart of absolute counts (right). In this session, 119 neurons (50.9%) were classified as orientation-selective, 26 (11.1%) as non-orientation-selective, and 89 (38.0%) as non-responsive. The bottom-left panel shows the distribution of tuning bandwidths specifically for the 119 orientation-selective neurons, which exhibit a mean bandwidth of 18.0° . The text summary (bottom right) confirms these statistics and notes that this session was programmatically selected for its balanced representation of all functional cell types. . .	25
2.3	Summary of neuron classification across all 16 recording sessions. The top panel presents a stacked bar chart showing the percentage of neurons within each session classified as orientation-selective (green), non-orientation-selective (orange), or non-responsive (blue). This visualization highlights the variability in the proportion of functional cell types across different recordings. The bottom panel complements this by showing the absolute number of orientation-selective neurons identified in each corresponding session. Across all sessions, a total of 1,340 orientation-selective neurons were analyzed.	26

2.4	Comparison of neuron classification statistics between animals M22 and M36.	The top-left panel shows that on average, sessions from animal M36 yielded a higher percentage of orientation-selective neurons compared to M22. The top-right panel confirms this by showing a higher average count of orientation-selective neurons per session for M36. The bottom-left panel provides context by showing the number of recording sessions for each animal (12 for M36 and 4 for M22). Finally, the scatter plot on the bottom-right displays the distribution of each individual session based on its percentage of orientation-selective and non-orientation-selective cells, with points colored by animal.	27
2.5	Time course of neural activity across example neurons (z-scored).	Time series of z-scored $\Delta F/F$ responses (normalized by blank trial standard deviation) for four example neurons, selected to represent neurons with different orientation selectivity levels. Each panel shows the responses to preferred (red), null (blue), and blank (gray) stimulus conditions. Solid lines show trial-averaged responses with SEM shading. Yellow shaded regions mark averaged windows used for analysis. The stimulus was presented for 200 ms and we averaged 300ms. Z-scoring provides standardized units that facilitate comparison across neurons with different baseline variability.	28
2.6	Orientation tuning curves for example neurons (z-scored).	Tuning curves for the same four neurons shown in Figure 2.5, showing z-scored $\Delta F/F$ responses (normalized by blank trial standard deviation). Each panel displays individual trial responses (gray dots) and trial-averaged responses with error bars (red circles) to different stimulus orientations (0° – 180°), along with fitted circular Gaussian functions (blue curves). Gray shaded regions show blank trial variability (mean \pm SD). Z-scoring enables direct comparison of response magnitudes across neurons and provides appropriate units for variance-mean scaling analysis.	29
2.7	Distributions of orientation-tuned neuron properties across all sessions (N = 1,340).	Top row: Tuning bandwidths (σ_θ) were narrowly distributed with a median around 18° ; baseline responses were centered around 0.3 z-score units; and amplitude responses showed a right-skewed distribution centered around 1.9 z-score units. Bottom row: Normalized Tuning Amplitude (NTA) values clustered tightly around 83%, indicating robust tuning strength; and Circular Variance (CV) showed a mean of 0.51 (median 0.57), consistent with moderate orientation selectivity. Together, these distributions summarize the diversity and typical range of tuning properties among the 1,340 neurons classified as orientation-selective.	32

2.8	Session-wise averages of orientation-tuned neuron properties.	Each bar represents the mean value for a given session across the 1,340 orientation-selective neurons in total. (Top row) Tuning bandwidth (σ_θ), baseline response (z-score), and amplitude response (z-score) all exhibit meaningful variability across sessions, reflecting differences in signal quality and tuning specificity. (Bottom row) Normalized tuning amplitude (NTA) and circular variance (CV) were generally stable across sessions, with only minor fluctuations. These summaries align with the quantitative metrics reported in the session-wise summary table (e.g., mean $\sigma_\theta = 18.05^\circ$, mean CV = 0.562).	33
2.9	Comparison of orientation-tuned neuron properties between animals M22 and M36.	Each subplot summarizes the across-session mean and standard error of key tuning metrics for M22 and M36. (Top-left) Tuning bandwidths (σ_θ) were significantly broader in M36 compared to M22. (Top-right) Baseline response levels (z-score) were also consistently higher in M36, suggesting possible differences in imaging quality or viral expression. (Bottom-left) Amplitude responses (z-score) showed higher average responses in M36, albeit with more variability across sessions. (Bottom-right) Normalized tuning amplitude (NTA) was slightly higher in M22, but both animals showed strong tuning strength overall.	34
2.10	Session-wise correlation matrix of tuning parameters.	Heatmap displaying Pearson correlation coefficients between key orientation tuning parameters across all 16 recording sessions. Each row corresponds to a session, and each column to a unique pair of parameters: tuning bandwidth (σ_θ), baseline response (z-score), response amplitude (z-score), normalized tuning amplitude (NTA), and circular variance (CV). Color intensity represents the strength and sign of the correlation: red for positive, blue for negative, and white for near-zero. Annotated values within each cell show the exact Pearson correlation coefficient. This matrix reveals consistent trends such as strong negative correlations between CV and amplitude, and positive correlations between baseline and CV, suggesting structured interdependencies in the tuning landscape of V1 neurons.	36

- 2.11 **Scatter plots of key tuning parameter relationships.** Direct visualization of the relationships between key orientation tuning parameters across all 1,340 orientation-selective neurons from 16 sessions. Each panel shows individual neurons as points colored by animal (blue = M22, orange = M36), with black dashed trend lines indicating the overall relationship. (Left) Tuning bandwidth (σ_θ) vs. baseline response (z-score) shows a negative correlation, suggesting that sharper tuning co-occurs with lower baseline activity. (Center) Tuning bandwidth vs. amplitude response (z-score) reveals a negative relationship, indicating that neurons with narrower tuning exhibit stronger peak responses. (Right) Amplitude vs. baseline response shows a positive correlation, suggesting that neurons with higher baseline activity also tend to have larger stimulus-driven responses. Correlation coefficients, p-values, and sample sizes are displayed for each relationship. These scatter plots complement the session-wise correlation analyses by showing the underlying data distributions and the consistency of relationships across animals. 37
- 2.12 **Comparison of tuning parameter correlations between animals.** Violin plots summarize the distribution of Pearson correlation coefficients across sessions for seven key parameter pairs, separated by animal (M22 vs. M36). Each violin represents the empirical distribution of session-level correlation values for one animal, with red dots indicating the mean correlation across sessions. The parameter pairs include: (top row) tuning bandwidth (σ_θ) vs. response amplitude (z-score), baseline response (z-score), and normalized tuning amplitude (NTA); (middle row) tuning bandwidth vs. circular variance (CV), and amplitude (z-score) vs. CV; (bottom row) baseline response (z-score) vs. CV, and NTA vs. CV. The gray dashed line at $r = 0$ provides a reference for the absence of correlation. These second-order statistics highlight robust and animal-consistent dependencies such as the strong negative correlation between NTA and CV, and positive correlations between baseline and CV, while also revealing more variable or animal-specific trends in relationships involving tuning bandwidth. . . 38
- 2.13 **Session-wise consistency of tuning parameter correlations.** Each panel shows the Pearson correlation coefficient (r) for one of seven key parameter pairs, computed independently for each of the 16 recording sessions. Points are colored by significance level: red indicates $p < 0.05$, blue indicates $p \geq 0.05$; point size is proportional to the number of orientation-selective neurons contributing to that session's correlation estimate. Green lines mark the cross-session mean correlation for each parameter pair. Titles include the number of sessions with statistically significant correlations and the mean r value across sessions. These scatter plots reveal which relationships are robust (e.g., NTA vs. CV, Baseline vs. CV) and which vary more substantially across recordings (e.g., Sigma vs. Amplitude). 39

- 2.14 **Spatial organization of orientation preferences across all recording sessions.** Grid of orientation preference maps for 16 imaging sessions showing HSV-colored cell masks overlaid on mean fluorescence images. Each panel represents one recording session, with orientation-selective neurons colored by their preferred orientation using the HSV color scheme (legend shows 0°-180° mapping). Non-orientation-selective and non-responsive neurons are not colored. The maps reveal heterogeneous spatial organization across sessions, with some showing clear clustering of similar orientations (indicative of orientation columns) while others show more distributed patterns. Session names and neuron counts are provided for each field of view. Scale bars and field dimensions vary across sessions but are indicated in the individual panels. 41
- 2.15 **Quantitative analysis of spatial organization.** Combined spatial analysis of all neuron pairs from 1,340 orientation-selective neurons across 16 sessions. **Top-left:** Hexagonal density plot showing the relationship between spatial distance (μm) and orientation difference ($^\circ$). The density distribution demonstrates that nearby neurons tend to have more similar orientations than distant pairs. **Top-right:** Spatial autocorrelation function showing orientation similarity as a function of distance, with SEM error bars. The blue line with error bars shows empirical data; the green line shows the two-Gaussian model fit ($R^2 = 0.995$). The oscillatory pattern reveals periodic organization with hypercolumn spacing of 688 μm . Parameter box displays the complete two-Gaussian fit parameters. **Bottom-left:** Distribution of orientation differences across all pairs, showing mean = 42.4° (close to random expectation of 45°). **Bottom-right:** Distance-binned analysis showing mean orientation difference as a function of spatial distance. The increasing trend confirms systematic spatial organization, with orientation similarity decreasing with distance. 45
- 2.16 **Individual neuron variance-mean relationships with linear fits.** Scatter plots showing the relationship between response variance and mean for four representative orientation-selective neurons from session M22D20210127R0TS. Each point represents a different stimulus orientation, colored by the specific orientation presented. The red line shows the linear fit $\sigma^2 = \text{slope} \cdot \mu + \sigma_0^2$, with fitted parameters displayed in the legend. The black star indicates the reference point ($\mu = 0, \sigma^2 = 1$) for illustration. These examples illustrate the range of variance-mean scaling indices observed across individual neurons, where both the slope and σ_0^2 parameters are fitted from stimulus conditions. 48

- 2.17 **Combined variance-mean relationship across all orientation-selective neurons.** Hexagonal density plot showing the relationship between response variance and mean across all 1,340 orientation-selective neurons and stimulus conditions. Data points with extreme values (mean responses > 10 z-score units or variances > 100) were excluded as outliers to focus on physiologically relevant response ranges. All responses are z-scored using σ_{blank} (blank trial standard deviation) for normalization. Color intensity represents the density of data points at each location. The red line shows the population-level linear fit (slope = 0.983), while the black dashed line indicates perfect linear scaling (slope = 1.0) for comparison. The analysis includes 9,507 variance-mean data points from diverse neurons and stimulus conditions. 51
- 2.18 **Distribution of variance-mean scaling indices for orientation-selective neurons.** Histogram showing the distribution of variance-mean scaling indices computed using the two-parameter fitting approach for all 1,340 orientation-selective neurons. Both slope and σ_0^2 parameters were fitted from stimulus conditions using least-squares regression. The distribution is concentrated at values below 1.0 (mean = 0.540 ± 0.715), indicating that response variance increases less than proportionally with mean response for most neurons. 53
- 2.19 **Distribution of fitted baseline variance (σ_0^2) across orientation-selective neurons.** Histogram of fitted baseline variance values (σ_0^2) computed from the two-parameter linear model for all 1,340 orientation-selective neurons. The distribution shows that most neurons have relatively modest fitted baseline variance levels (mean $\sigma_0 = 0.577 \pm 0.116$ z-score units), indicating that the intercept of the variance-mean relationship is generally smaller than stimulus-driven variance changes for many neurons. Note that σ_0^2 represents the fitted baseline variance from stimulus conditions and is distinct from σ_{blank} used for z-score normalization. 54
- 2.20 **Relationship between variance-mean scaling index and fitted baseline variance.** Scatter plot of fitted baseline variance (σ_0^2) versus variance-mean scaling index (slope) for all 1,340 orientation-selective neurons. The red dashed line marks slope = 1. Neurons with large baseline variance ($\sigma_0^2 > 0.5$) are almost exclusively found at slopes < 1 , whereas neurons with slopes > 1 overwhelmingly have σ_0^2 near zero. Neurons with very low baseline variance span the full range of scaling indices. 55

- 2.21 **Summary of correlations between response variability and tuning properties.** Horizontal bar chart showing Pearson correlation coefficients between variance-mean scaling index, baseline noise (σ_0), and orientation tuning parameters across all 1,340 orientation-selective neurons. Red bars indicate statistically significant correlations ($p \leq 0.05$), while gray bars show non-significant relationships. Correlation values are displayed next to each bar. The analysis reveals several significant relationships, including correlations between tuning bandwidth and baseline activity, amplitude and baseline responses, and various tuning parameter interdependencies. Individual scatter plots for significant relationships provide detailed visualization of these pairwise associations. 57
- 2.22 **Summary of noise correlation distributions across all sessions.** (Top-left) Overlaid histograms of true noise correlations for spontaneous (blue) and evoked (red) conditions across 64,168 neuron pairs. Evoked correlations are higher on average with a notably narrower distribution. (Top-right) Scatter plot comparing each pair's spontaneous vs. evoked correlation. The black dashed line indicates unity ($y = x$); points above this line reflect increased correlation during stimulus presentation. The overall Pearson correlation between spontaneous and evoked values was $r = 0.205$, $p < 10^{-16}$. (Bottom-left) Boxplots of correlation distributions by animal and condition. Both M22 and M36 show higher medians for evoked relative to spontaneous activity. (Bottom-right) Histogram of correlation differences (evoked - spontaneous), showing a clear rightward shift. The mean difference was 0.025, and a paired t -test confirmed a significant increase in evoked correlations ($p = 8.8 \times 10^{-113}$), with a meaningful effect size (Cohen's $d = 0.089$). Together, these results demonstrate that visual stimulation systematically enhances noise correlations while constraining their variability. 60
- 2.23 **Orientation dependence of noise correlations: original vs distance-controlled analysis.** (Top row) Individual session traces showing correlations as a function of orientation preference difference ($\Delta\theta_{\text{pref}}$), binned into 15° intervals. Left panel shows spontaneous correlations, which exhibit no systematic dependence on orientation difference. Center panel shows evoked correlations with modest orientation dependence, declining from 0.15 to 0.08. Right panel shows distance-controlled evoked correlations (neuron pairs $\geq 300 \mu\text{m}$ apart), revealing minimal orientation dependence when spatial confounds are removed. (Bottom row) Animal-averaged traces with SEM shading for each condition. The distance-controlled analysis demonstrates that the observed functional organization primarily reflects local circuit connectivity rather than long-range functional relationships across orientation columns. Text annotations show the percentage of neuron pairs retained after distance filtering. Lines are colored by animal (M22: gray, M36: red). 62

2.24	Spatial distance dependence of noise correlations: original vs orientation-controlled analysis. (Top row) Individual session traces showing correlations as a function of spatial distance (μm). Left panel shows spontaneous correlations with gradual distance-dependent decline. Center panel shows evoked correlations with steeper distance dependence. Right panel shows orientation-controlled evoked correlations (neuron pairs with $30^\circ \pm 10^\circ$ orientation differences), suggesting that distance dependence may persist even among functionally dissimilar neuron pairs. (Bottom row) Animal-averaged traces with SEM shading for each condition. The orientation-controlled analysis suggests that spatial proximity may contribute to correlated variability independent of orientation preference similarity. While correlation magnitudes are reduced after orientation filtering (consistent with intermediate functional similarity), the systematic decline with distance remains, which could support the role of local circuit connectivity. Text annotations show the percentage of neuron pairs retained after orientation filtering. Lines are colored by animal (M22: gray, M36: red).	65
2.25	Overview of session-level metrics for noise correlation analysis. (Top-left) Number of orientation-selective neurons per session, colored by animal (M22 in pink, M36 in gold). (Top-center) Number of trials per session for each condition (blank vs. stimulus), indicating strong trial coverage overall. (Top-right) Number of neuron pairs analyzed per session, proportional to the square of neuron count. (Bottom-left) Boxplots of mean noise correlations per session, separated by condition and animal. While variability exists, evoked correlations tend to be slightly higher than spontaneous. (Bottom-center) Bar plot of the correlation difference (evoked – spontaneous) for each session, showing modest but often positive changes in stimulus-evoked synchrony. (Bottom-right) Standard deviation of correlation distributions, which tended to decrease under stimulus presentation. This figure provides a quality and consistency check across the dataset, confirming sufficient coverage, stable correlation metrics, and a general pattern of modest stimulus-induced decorrelation.	67
2.26	Example session orientation discrimination analysis. Discriminability performance curves for fine and coarse orientation discrimination in a representative session (M36D20250327R1TS) across four decoding methods. Top row shows fine discrimination (2° , 4° , 8° differences) limited to small population sizes ($N \leq 8$), while bottom row shows coarse discrimination (15° - 90° differences) across the full population range. Each panel displays mean discriminability as a function of orientation difference for different population sizes, with color intensity representing population size. This example illustrates the typical scaling patterns: fine discrimination shows gradual performance improvements with modest asymptotic values, while coarse discrimination exhibits steeper scaling and higher peak performance levels.	70

2.27	<p>Comparison of fine versus coarse orientation discrimination performance with error bars. Direct comparison of discriminability for fine (2°, 4°, 8° differences) and coarse (15°-90° differences) orientation discrimination across the 10 sessions with both discrimination types. Left panels show Logistic Regression decoder performance; right panels show SVM decoder performance. Top row presents fine discrimination (limited to $N \leq 8$ neurons, red color scale); bottom row shows coarse discrimination (full population range, blue color scale). Color intensity varies with population size N, ranging from light ($N=1$) to dark ($N=32$ for coarse, $N=8$ for fine). Error bars represent standard error of the mean (SEM) across sessions for each population size and orientation difference. Fine discrimination shows more gradual scaling with population size, while coarse discrimination exhibits steeper performance gains and higher asymptotic values. Both analyses demonstrate that orientation information scales systematically with population size, with discrimination difficulty appropriately reflected in performance curves.</p>	71
2.28	<p>Fine orientation discrimination performance across decoders with error bars. Performance curves for fine orientation discrimination (2°, 4°, 8° differences) across four decoding methods. Each subplot shows mean and median d' values as a function of population size ($N=1, 2, 4, 8$ neurons) for a single decoder, with error bars representing standard error of the mean (SEM) for mean values. Different colors represent different orientation differences within the fine discrimination range. Performance scales systematically with population size, while orientation differences within the fine range show minimal and non-significant variation, with performance appearing to decrease slightly from 2° to 8° differences despite the counterintuitive nature of this pattern. Fine discrimination was assessable in 10 sessions (all from M36), representing conditions with dense orientation sampling. Peak performance levels ($d' \approx 4-6$) demonstrate that even small orientation differences can be reliably discriminated by appropriately sized V1 populations.</p>	72
2.29	<p>Coarse orientation discrimination performance across decoders with error bars. Performance curves for coarse orientation discrimination (15°-90° differences) across four decoding methods. Each subplot shows mean and median d' values as a function of population size ($N=1, 2, 4, 8, 16, 32, 64, 128$ neurons) for a single decoder, with error bars representing standard error of the mean (SEM) for mean values. Different colors represent different orientation differences within the coarse discrimination range. Performance shows strong scaling with both population size and orientation difference, with 90° differences (orthogonal orientations) achieving the highest discriminability values. Coarse discrimination was assessable across all 16 sessions with 7,062 total pairs. High peak performance levels ($d' > 10$) demonstrate the substantial orientation information available in V1 populations for behaviorally relevant discrimination tasks.</p>	73

2.30 **Impact of noise correlations on population decoding performance.** Comparison between empirical decoding performance and shuffle controls for 45° and 90° orientation differences. Each subplot shows one decoder (Nearest Centroid, Diagonal Covariance, Logistic Regression, SVM). Solid lines represent empirical data; dashed lines represent shuffled controls where trial correlations are eliminated. Green indicates 45° differences; purple indicates 90° differences. Circles show empirical data points; squares show shuffled data points. Text boxes display saturating exponential fit parameters for empirical curves: $d'(N) = a[1 - \exp(-bN)] + c$. Shuffled performance consistently exceeds empirical performance for $N \geq 4$, suggesting that noise correlations may influence population decoding efficiency. The performance gap increases with both population size and orientation difference, which could suggest that correlations have the greatest effect on the most discriminable stimulus pairs. 75

Chapter 1: Introduction

1.1 Visual Processing in Primary Visual Cortex

The primary visual cortex (V1) serves as the principal cortical gateway for visual information processing in mammals, transforming retinal inputs into increasingly complex feature representations. One of the most fundamental and well-studied properties of V1 neurons is their selective response to oriented visual stimuli (Hubel and Wiesel). This orientation selectivity underpins more complex processing in downstream cortical areas and has long served as a model for studying cortical computation.

Orientation-selective neurons in V1 respond maximally to edges or gratings of a particular orientation, with response magnitude decreasing systematically as the stimulus orientation deviates from the neuron's preferred orientation. This tuning is thought to arise from convergent input of lateral geniculate nucleus (LGN) cells with spatially offset receptive fields (Hubel and Wiesel), although the precise biophysical mechanisms continue to be refined (Priebe).

1.2 Population Coding and Noise Correlations

While single-neuron orientation tuning has been extensively characterized, understanding how populations of V1 neurons collectively encode orientation requires considering both individual tuning properties and the correlations between neurons. Noise correlations—trial-to-trial covariability in neural responses after removing stimulus-driven components—are known to shape the information content of population codes (Averbeck et al.).

The impact of noise correlations depends critically on their structure relative to signal correlations (i.e., similarity in tuning). When noise correlations align with signal correlations, they can limit information transmission by inducing redundant fluctuations among similarly tuned neurons (Abbott and Dayan). Conversely, when

orthogonal, they may have minimal effect or even enhance decoding performance (Moreno-Bote et al.). However, the stimulus dependence of these correlation structures, particularly in large populations of V1 neurons in awake primates, remains poorly understood.

1.3 Technical Advances in Population Recording

Electrophysiological recordings provide excellent temporal resolution and, through cortical penetrations, access to laminar organization across cortical layers. However, electrode-based approaches are fundamentally limited in their ability to simultaneously sample large populations within the same cortical plane, typically recording from only a few dozen neurons distributed across different depths and cortical locations.

Two-photon calcium imaging with genetically encoded calcium indicators (GECIs) provides a complementary approach that excels at population-level recording within specific cortical layers (Chen et al., a). This technique utilizes the nonlinear optical properties of two-photon excitation, where fluorophore excitation occurs only at the focal point where photon density is highest. This localized excitation, combined with the deeper tissue penetration of near-infrared light, enables optical sectioning at specific cortical depths (typically 150-300 μm below the surface) while minimizing out-of-focus fluorescence. The result is planar access to hundreds of individual neurons within a single cortical layer at cellular resolution—a capability that is difficult to achieve with traditional electrophysiological methods.

The development of fast, sensitive indicators like GCaMP6f has made this approach viable for detecting spiking-related activity in awake, behaving animals, enabling simultaneous functional characterization of large neuronal populations within the tangential plane of cortex. This planar access is particularly advantageous for studying retinotopically organized areas like V1, where sampling across the cortical surface within a single processing layer provides access to neurons representing

different regions of visual space at the same stage of cortical computation. Spatial organization principles including the clustering of similarly tuned neurons and periodic hypercolumn structure provide the anatomical framework within which population coding operates. However, a detailed characterization of how sensory input shapes both spatial organization and noise correlation structure in these populations—especially in V1 layer 2/3—remains lacking.

1.4 Research Objectives

This thesis aims to characterize orientation coding in macaque V1 using two-photon calcium imaging. Specifically, we address the following objectives:

1. **Single-neuron characterization:** Quantify the distribution of orientation tuning parameters (preferred orientation, tuning bandwidth, response amplitude, baseline activity) across a large population of V1 neurons.
2. **Cross-parameter relationships:** Investigate systematic relationships between tuning parameters to uncover organizational principles of cortical selectivity.
3. **Spatial and population-level organization:** Assess the distribution and potential bias in preferred orientations across cortical regions, and quantify the spatial clustering of functionally similar neurons and hypercolumn organization.
4. **Noise correlation structure:** Quantify pairwise noise correlations during both spontaneous and evoked conditions, and analyze their relationship to tuning similarity.
5. **Stimulus-dependent modulation:** Determine how visual stimulation modulates the magnitude and structure of shared variability, with implications for population coding efficiency.

Using a large dataset from two animals across 16 recording sessions, this work uncovers key principles of population-level coding in V1. We quantify the spatial organization of orientation preferences, revealing systematic clustering of functionally similar neurons and hypercolumn periodicity. Most notably, we show that visual stimulation dynamically reshapes the structure of noise correlations in a tuning-dependent manner—enhancing shared variability among similarly tuned neurons and suppressing it among dissimilar ones. These findings illuminate how cortical circuits regulate redundancy and sensitivity during sensory processing within their spatial organizational framework, and provide a foundation for future work linking population dynamics to behavior and perception.

Chapter 2: Results

2.1 Dataset Overview

We recorded from a total of 2,667 neurons across 16 imaging sessions in two head-fixed macaques (M22: 4 sessions, M36: 12 sessions) using two-photon calcium imaging of layer 2/3 neurons expressing GCaMP6f. Neurons were classified into three categories based on their responses to oriented grating stimuli: orientation-selective, responsive but orientation-invariant, and non-responsive.

Across all sessions, 1,340 neurons (50.2%) were classified as orientation-selective, 416 neurons (15.6%) were responsive but orientation-invariant, and 911 neurons (34.2%) were non-responsive. The proportion of orientation-selective neurons varied across sessions (range: 27.5%–69.4%, mean \pm SD: $48.8 \pm 12.2\%$), consistent with sampling across different cortical domains with varying orientation column organization.

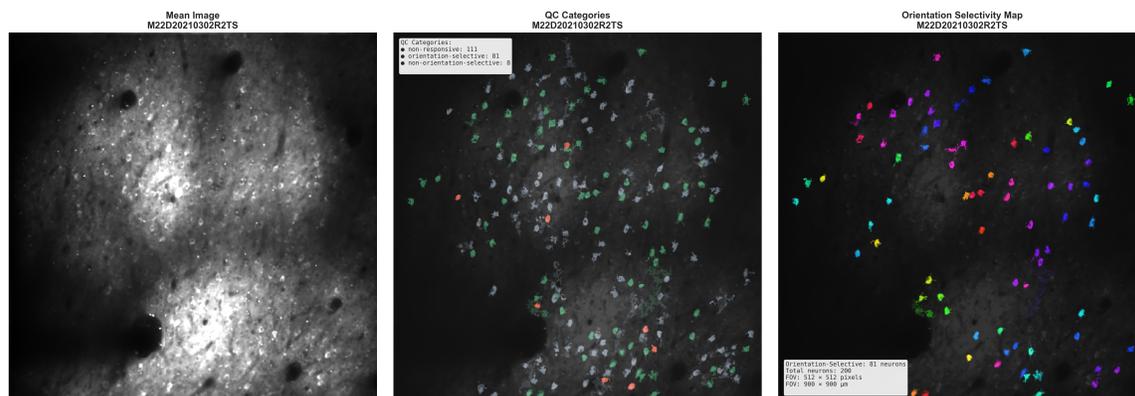


Figure 2.1: **Example recording session showing cellular spatial organization.** This figure illustrates data from a representative recording session (M22D20210302R2TS, 88 total neurons). (Left) Mean fluorescence image of the field of view showing the spatial distribution and morphology of recorded neurons. (Center) Segmented cell masks colored by functional classification: orientation-selective neurons (green), non-orientation-selective neurons (orange), and non-responsive neurons (blue). (Right) The same cell masks colored by preferred orientation for orientation-selective neurons only, using an HSV color scheme where hue represents preferred orientation (0°-180°). This visualization demonstrates the cellular resolution of our two-photon imaging approach and illustrates the spatial distribution of functionally characterized neurons across the cortical field of view. Scale bar and field dimensions are provided in the original figure.

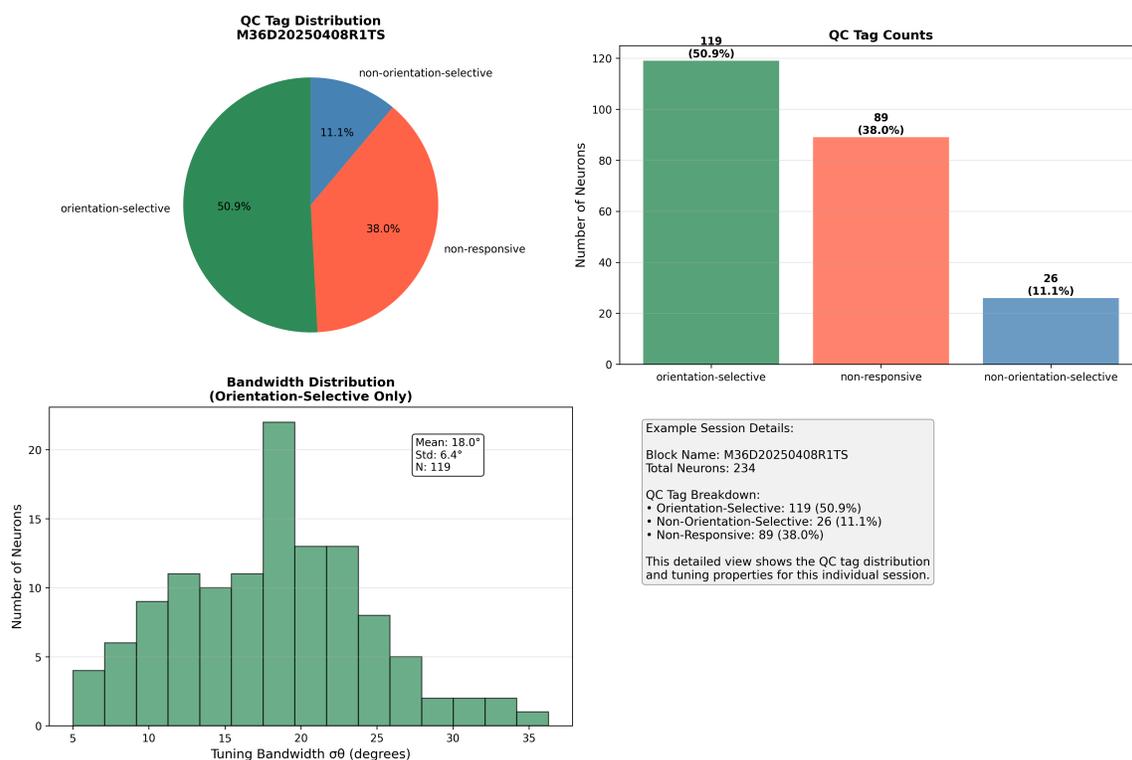


Figure 2.2: **Detailed breakdown of a representative recording session.** This figure illustrates the per-session analysis pipeline using data from a single, representative recording with a total of 234 neurons. The top row shows the functional classification results, presented as a pie chart of percentages (left) and a bar chart of absolute counts (right). In this session, 119 neurons (50.9%) were classified as orientation-selective, 26 (11.1%) as non-orientation-selective, and 89 (38.0%) as non-responsive. The bottom-left panel shows the distribution of tuning bandwidths specifically for the 119 orientation-selective neurons, which exhibit a mean bandwidth of 18.0°. The text summary (bottom right) confirms these statistics and notes that this session was programmatically selected for its balanced representation of all functional cell types.

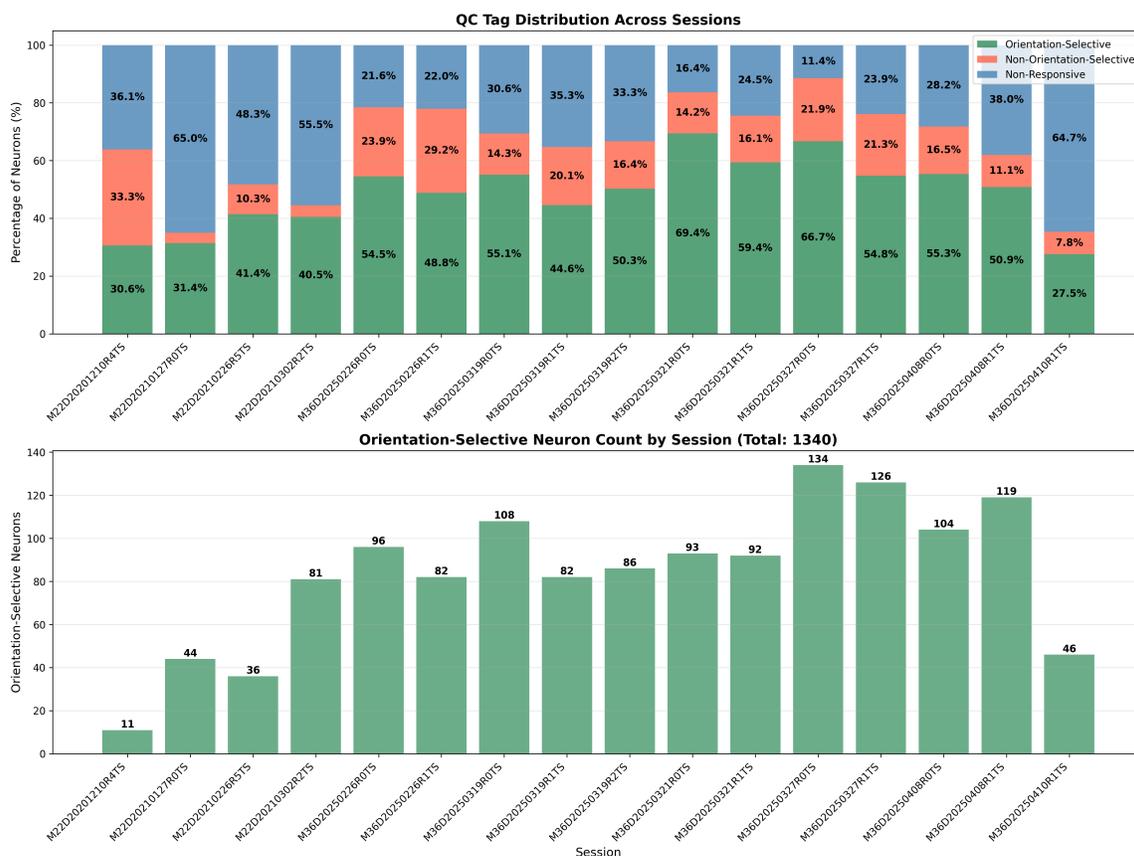


Figure 2.3: **Summary of neuron classification across all 16 recording sessions.** The top panel presents a stacked bar chart showing the percentage of neurons within each session classified as orientation-selective (green), non-orientation-selective (orange), or non-responsive (blue). This visualization highlights the variability in the proportion of functional cell types across different recordings. The bottom panel complements this by showing the absolute number of orientation-selective neurons identified in each corresponding session. Across all sessions, a total of 1,340 orientation-selective neurons were analyzed.

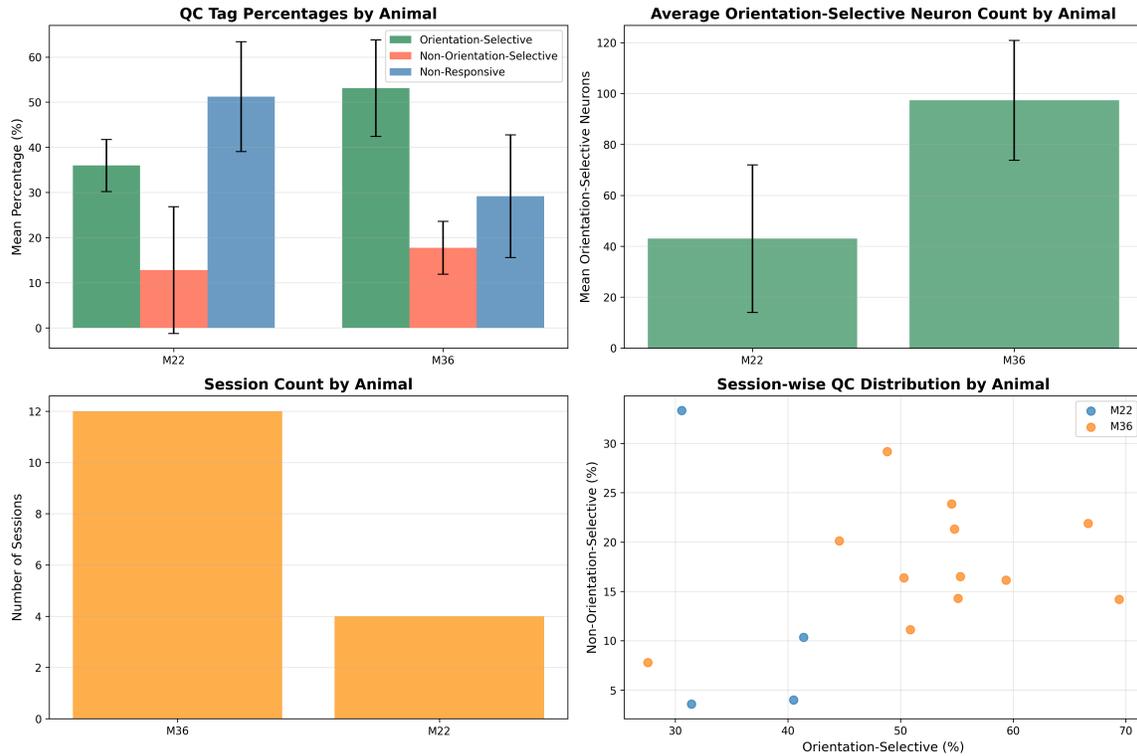


Figure 2.4: **Comparison of neuron classification statistics between animals M22 and M36.** The top-left panel shows that on average, sessions from animal M36 yielded a higher percentage of orientation-selective neurons compared to M22. The top-right panel confirms this by showing a higher average count of orientation-selective neurons per session for M36. The bottom-left panel provides context by showing the number of recording sessions for each animal (12 for M36 and 4 for M22). Finally, the scatter plot on the bottom-right displays the distribution of each individual session based on its percentage of orientation-selective and non-orientation-selective cells, with points colored by animal.

2.2 Representative Neural Responses

To illustrate the diversity of response profiles in our dataset, we examined representative time courses of z -scored $\Delta F/F$ responses from four neurons with different levels of orientation selectivity, showing how these cells respond to preferred, null, and blank stimulus conditions.

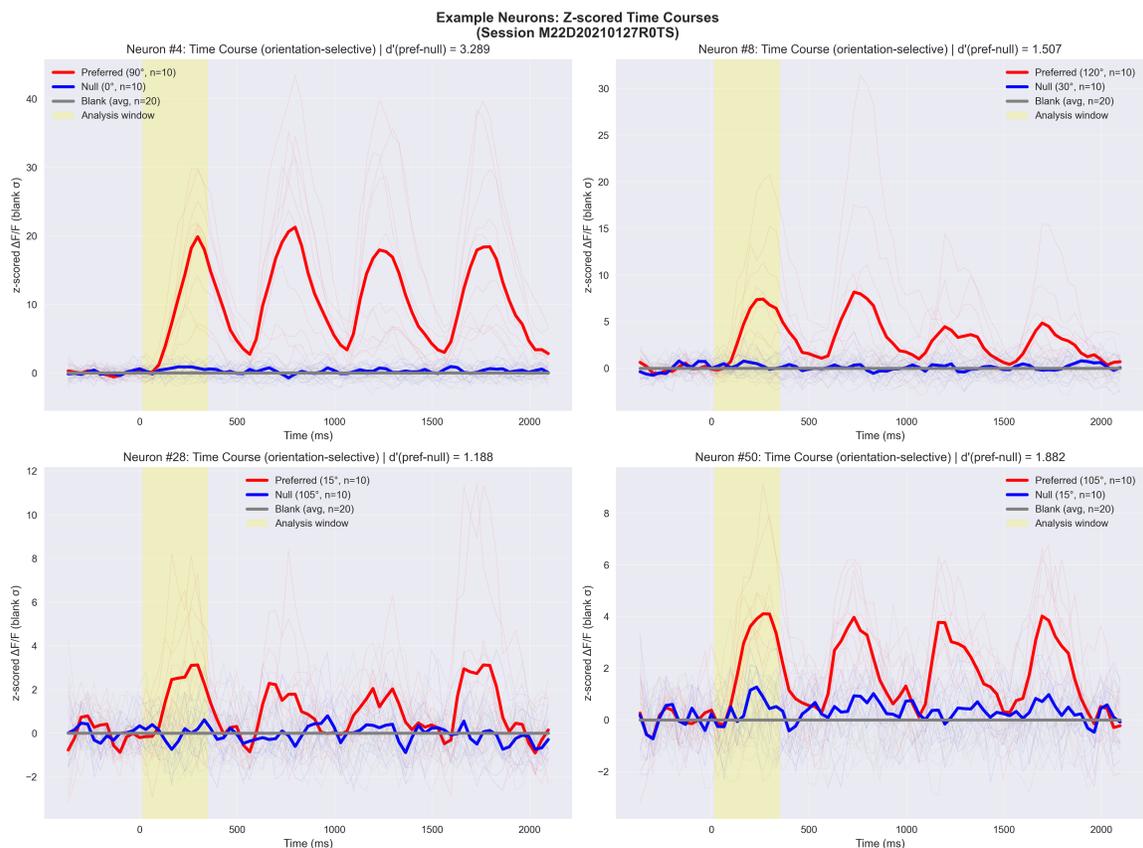


Figure 2.5: **Time course of neural activity across example neurons (z-scored).** Time series of z-scored $\Delta F/F$ responses (normalized by blank trial standard deviation) for four example neurons, selected to represent neurons with different orientation selectivity levels. Each panel shows the responses to preferred (red), null (blue), and blank (gray) stimulus conditions. Solid lines show trial-averaged responses with SEM shading. Yellow shaded regions mark averaged windows used for analysis. The stimulus was presented for 200 ms and we averaged 300ms. Z-scoring provides standardized units that facilitate comparison across neurons with different baseline variability.

These example neurons demonstrate the diversity of response patterns observed in our population. The neurons show varying degrees of orientation selectivity, from strong differential responses to preferred versus null orientations to more modest selectivity. All neurons show responses that are time-locked to stimulus presentation, with peak responses occurring during the stimulus period.

For each orientation-selective neuron, we fitted a circular Gaussian function to characterize its tuning properties. The tuning curves for the same four example neurons are shown in Figure 2.6, demonstrating the quality of fits achieved with our model.

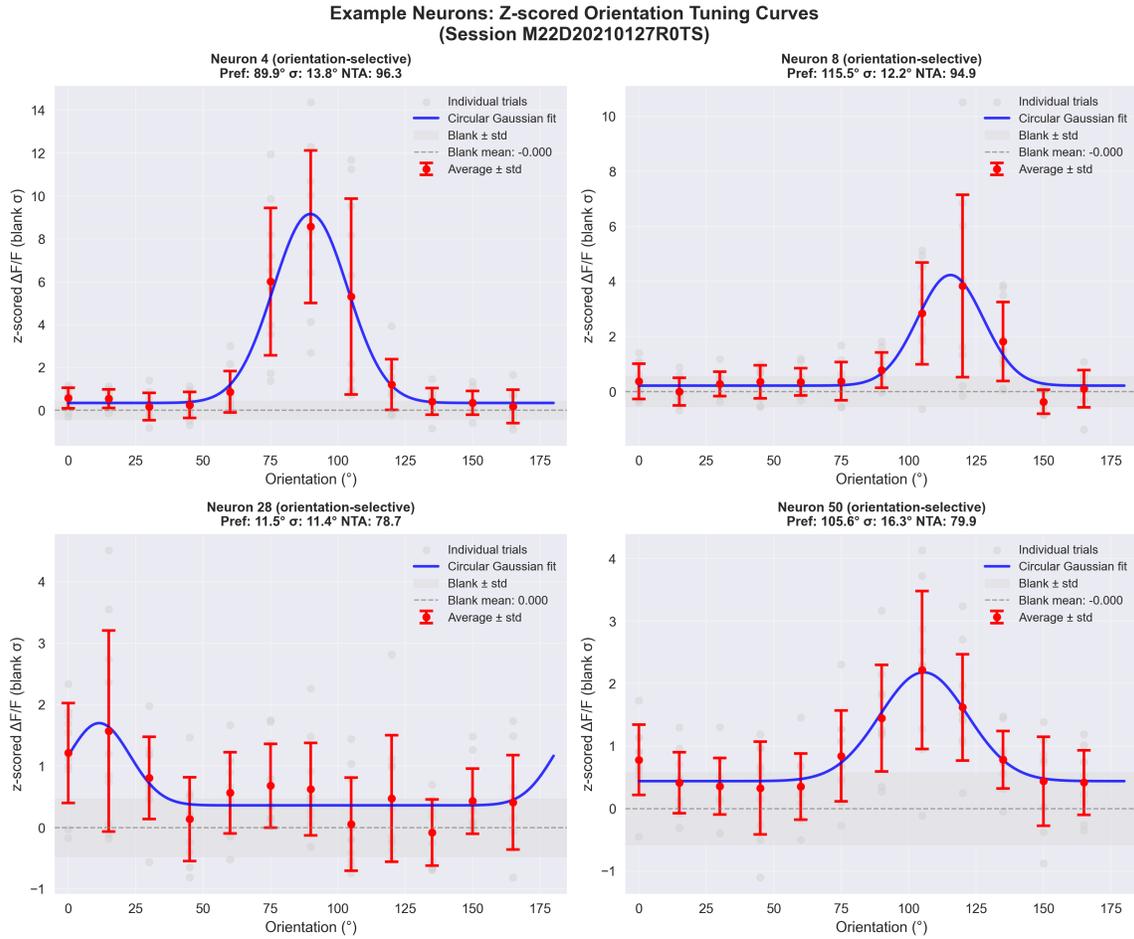


Figure 2.6: **Orientation tuning curves for example neurons (z-scored).** Tuning curves for the same four neurons shown in Figure 2.5, showing z-scored $\Delta F/F$ responses (normalized by blank trial standard deviation). Each panel displays individual trial responses (gray dots) and trial-averaged responses with error bars (red circles) to different stimulus orientations (0° – 180°), along with fitted circular Gaussian functions (blue curves). Gray shaded regions show blank trial variability (mean \pm SD). Z-scoring enables direct comparison of response magnitudes across neurons and provides appropriate units for variance-mean scaling analysis.

These examples illustrate how our circular Gaussian model captures the essential features of orientation tuning in V1 neurons, providing robust estimates of key parameters including preferred orientation, tuning bandwidth, response amplitude, and baseline activity.

2.3 Single-Neuron Orientation Tuning Properties

2.3.1 Tuning Parameter Distributions

We characterized the tuning properties of orientation-selective neurons by fitting circular Gaussian functions to their orientation response profiles. The model took the form:

$$F_i(\theta) = b_i + a_i \exp \left[-\frac{1}{2} \left(\frac{d_{\text{circ}}(\theta, \theta_i^*)}{\sigma_i} \right)^2 \right] \quad (2.1)$$

where b_i represents baseline activity, a_i is response amplitude, θ_i^* is preferred orientation, σ_i is tuning bandwidth, and d_{circ} is the circular distance on the orientation domain $[0^\circ, 180^\circ]$.

Tuning bandwidth (σ_θ): Orientation-selective neurons exhibited narrow tuning bandwidths with a median of 17.58° (mean = $17.3^\circ \pm 2.4^\circ$) across sessions (Figure 2.7). This is notably sharper than previous electrophysiological reports in macaque V1, which typically report mean bandwidths of 20° (Ringach et al.; Chen et al., c). The bandwidth distribution showed positive skewness in most sessions, with a small population of broadly tuned neurons extending the distribution tail.

Response amplitude: The amplitude of orientation-tuned responses, representing the peak response increase above baseline, showed substantial heterogeneity across neurons (mean: 1.93 ± 0.44 z-score units across sessions; Figure 2.7). This variability likely reflects differences in cellular properties, local circuit connectivity, and expression levels of the calcium indicator.

Baseline activity: Baseline activity levels, measured as the fitted response to the orthogonal orientation, averaged 0.36 ± 0.12 z-score units across sessions (Figure 2.7). Most neurons showed positive baseline values, indicating modest but consistent activity even to non-preferred stimuli.

Normalized Tuning Amplitude (NTA): To account for differences in overall response magnitude, we calculated the normalized tuning amplitude as $\text{NTA}_i = a_i / (a_i + b_i)$, representing the modulation depth relative to total peak response. The mean NTA was $83.2 \pm 3.3\%$ across sessions (Figure 2.7), indicating strong orientation selectivity in the recorded population. This is consistent with previous electrophysiological reports in macaque V1, which typically report mean NTA values of 80% (Chen et al., c).

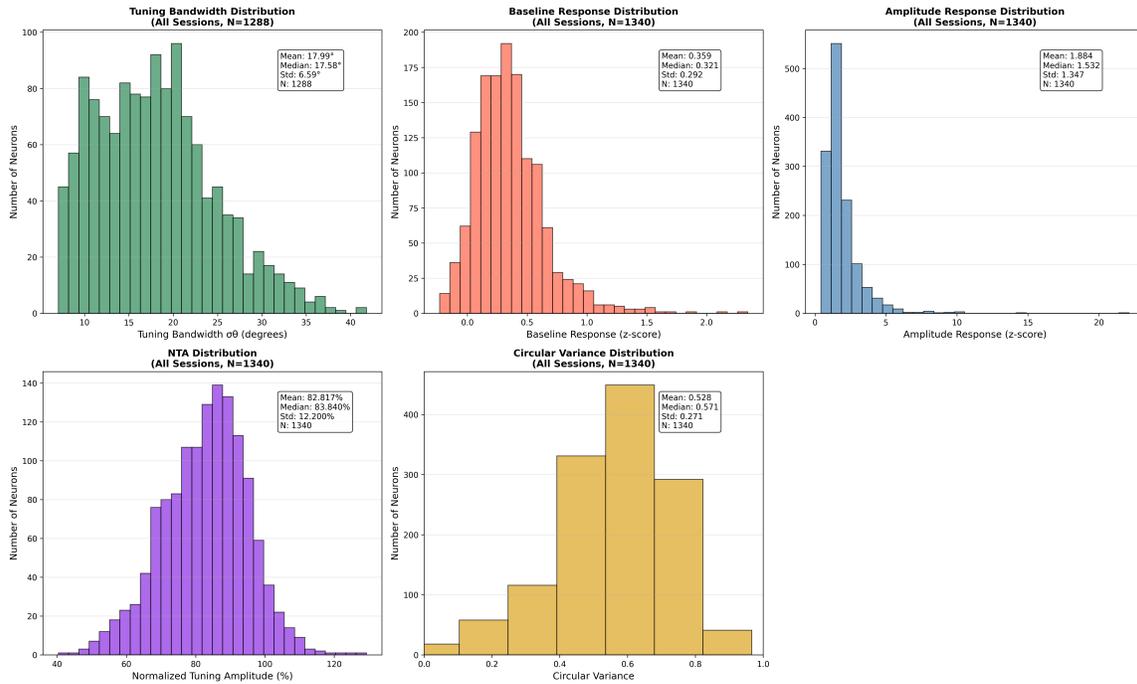


Figure 2.7: **Distributions of orientation-tuned neuron properties across all sessions (N = 1,340).** Top row: Tuning bandwidths (σ_θ) were narrowly distributed with a median around 18°; baseline responses were centered around 0.3 z-score units; and amplitude responses showed a right-skewed distribution centered around 1.9 z-score units. Bottom row: Normalized Tuning Amplitude (NTA) values clustered tightly around 83%, indicating robust tuning strength; and Circular Variance (CV) showed a mean of 0.51 (median 0.57), consistent with moderate orientation selectivity. Together, these distributions summarize the diversity and typical range of tuning properties among the 1,340 neurons classified as orientation-selective.

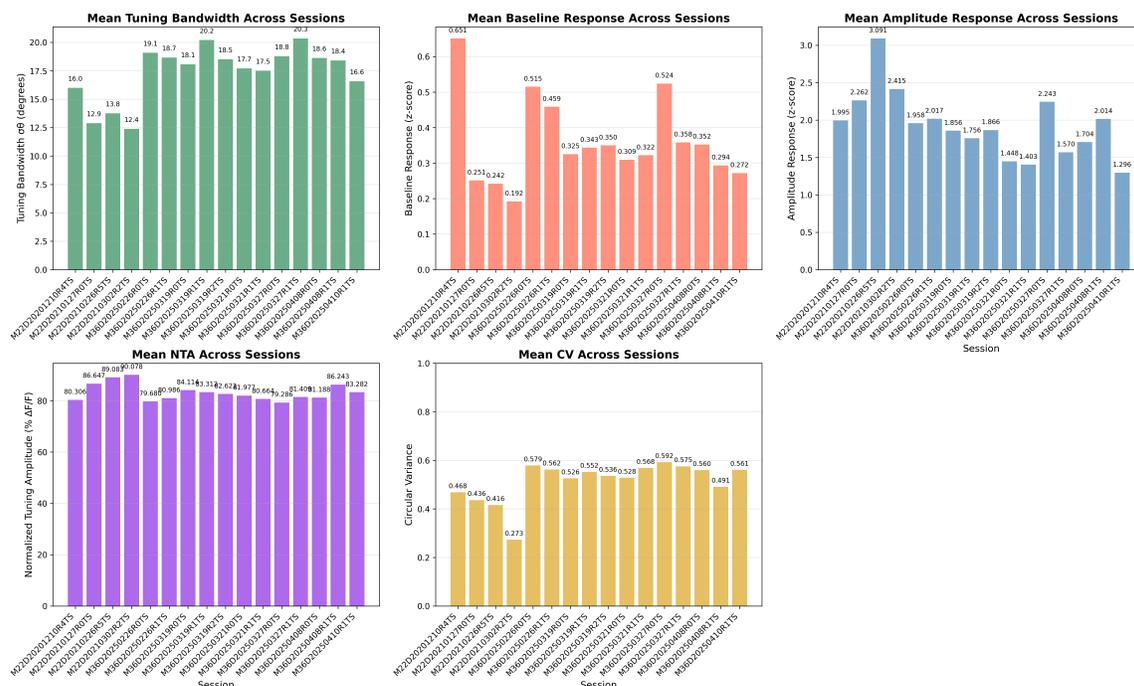


Figure 2.8: **Session-wise averages of orientation-tuned neuron properties.** Each bar represents the mean value for a given session across the 1,340 orientation-selective neurons in total. (Top row) Tuning bandwidth (σ_θ), baseline response (z-score), and amplitude response (z-score) all exhibit meaningful variability across sessions, reflecting differences in signal quality and tuning specificity. (Bottom row) Normalized tuning amplitude (NTA) and circular variance (CV) were generally stable across sessions, with only minor fluctuations. These summaries align with the quantitative metrics reported in the session-wise summary table (e.g., mean $\sigma_\theta = 18.05^\circ$, mean CV = 0.562).

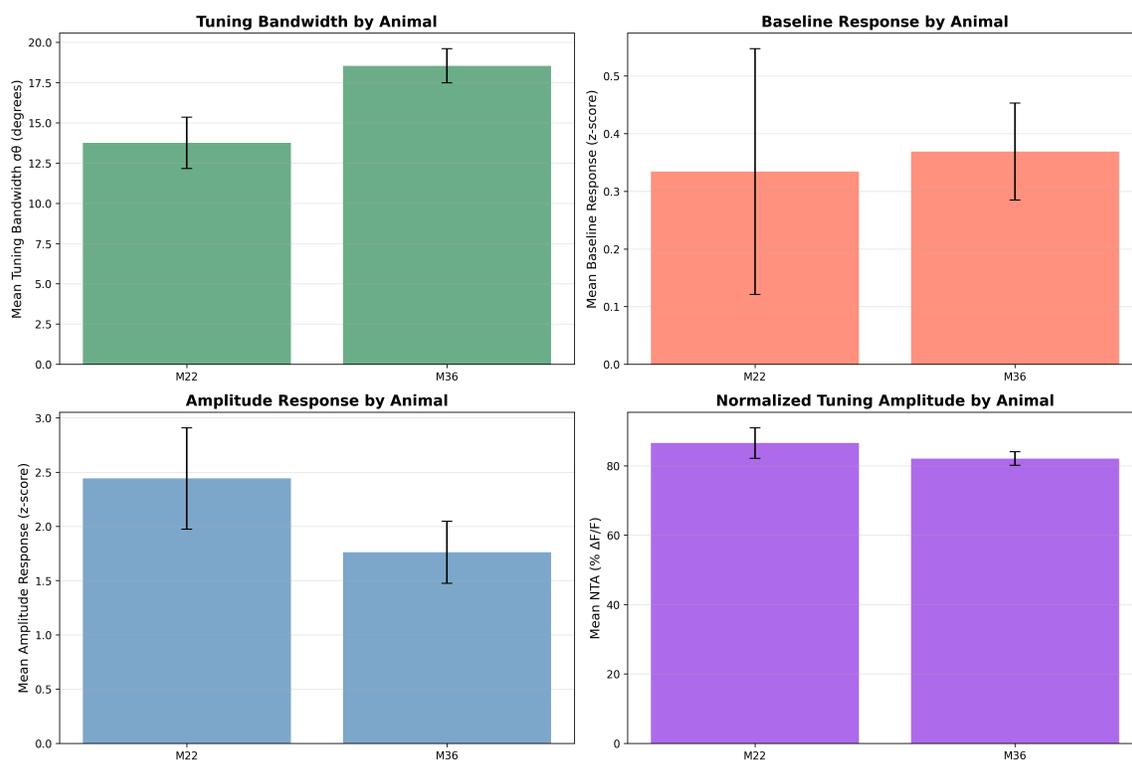


Figure 2.9: **Comparison of orientation-tuned neuron properties between animals M22 and M36.** Each subplot summarizes the across-session mean and standard error of key tuning metrics for M22 and M36. (Top-left) Tuning bandwidths (σ_θ) were significantly broader in M36 compared to M22. (Top-right) Baseline response levels (z-score) were also consistently higher in M36, suggesting possible differences in imaging quality or viral expression. (Bottom-left) Amplitude responses (z-score) showed higher average responses in M36, albeit with more variability across sessions. (Bottom-right) Normalized tuning amplitude (NTA) was slightly higher in M22, but both animals showed strong tuning strength overall.

2.3.2 Cross-Parameter Relationships

We investigated systematic relationships between orientation tuning parameters to better understand the organizational constraints underlying neuronal selectivity in macaque V1. Across 1,340 orientation-selective neurons from 16 sessions, we examined all pairwise correlations among tuning width (σ), response amplitude, baseline activity, normalized tuning amplitude (NTA), and circular variance (CV)

(Figures 2.10 and 2.11).

Sigma vs. Baseline: Tuning bandwidth was negatively correlated with baseline activity (mean $r = -0.26 \pm 0.10$; significant in 8/16 sessions), suggesting that more sharply tuned neurons tend to have lower baseline responses at non-preferred orientations. This is consistent with models in which suppression of untuned or broadly tuned input sharpens selectivity by reducing baseline levels.

Sigma vs. Amplitude: A consistent negative relationship was observed between tuning width and response amplitude (mean $r = -0.29 \pm 0.15$; significant in 13/16 sessions), indicating that sharper tuning co-occurs with larger responses across the majority of sessions.

Sigma vs. NTA: We observed a mild positive correlation between tuning width and normalized tuning amplitude (mean $r = 0.10 \pm 0.15$; significant in 3/16 sessions), though the direction of this relationship may reflect a mathematical trade-off: sharper tuning reduces response spread across non-preferred orientations, but normalization can shift this balance depending on baseline levels.

Amplitude vs. CV: A moderate negative correlation was observed (mean $r = -0.26 \pm 0.20$; significant in 11/16 sessions), indicating that neurons with larger evoked responses tend to exhibit lower circular variance. This likely reflects sharper and more peaked tuning curves dominating the population response.

NTA vs. CV: We found a robust and highly consistent negative correlation between normalized tuning amplitude and circular variance (mean $r = -0.88 \pm 0.05$; significant in all 16 sessions), supporting the idea that NTA and CV capture overlapping but inversely related aspects of tuning sharpness. This strong relationship further validates NTA as a sensitive measure of selectivity in population-level analyses.

The scatter plot analysis (Figure 2.11) provides direct visualization of these relationships at the individual neuron level, revealing consistent patterns across both animals and confirming the robustness of these parameter interdependencies.

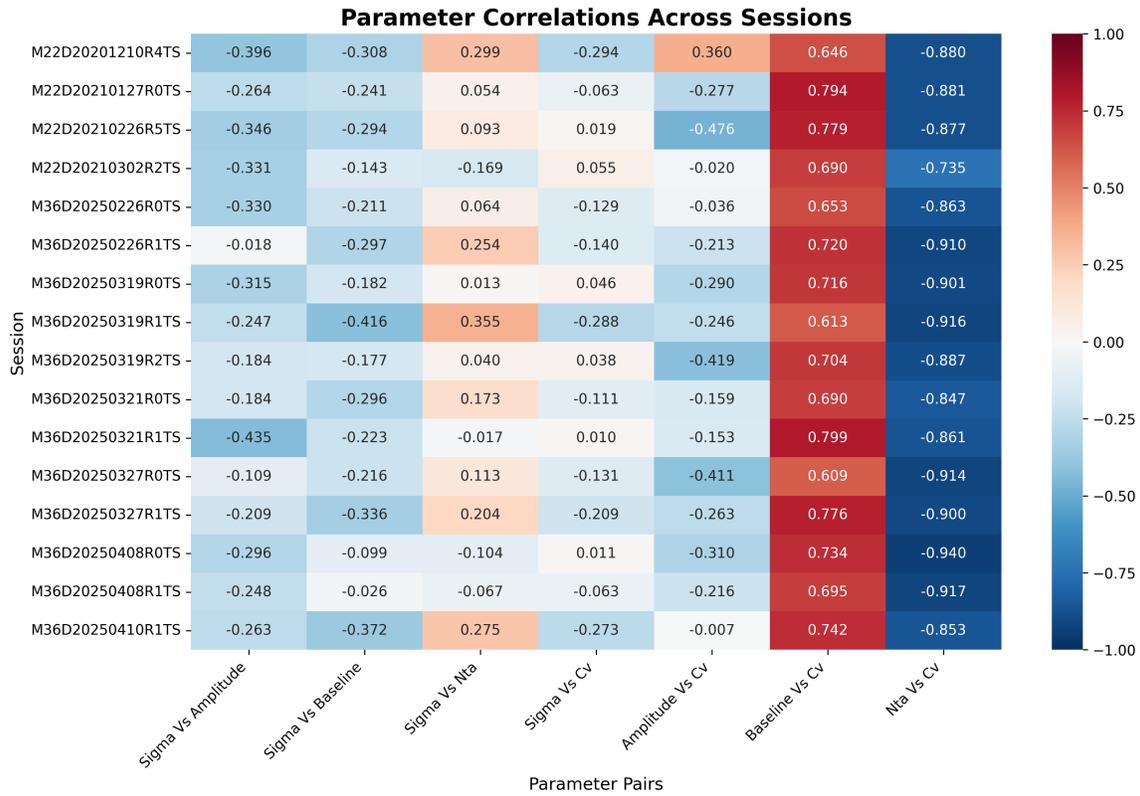


Figure 2.10: **Session-wise correlation matrix of tuning parameters.** Heatmap displaying Pearson correlation coefficients between key orientation tuning parameters across all 16 recording sessions. Each row corresponds to a session, and each column to a unique pair of parameters: tuning bandwidth (σ_θ), baseline response (z-score), response amplitude (z-score), normalized tuning amplitude (NTA), and circular variance (CV). Color intensity represents the strength and sign of the correlation: red for positive, blue for negative, and white for near-zero. Annotated values within each cell show the exact Pearson correlation coefficient. This matrix reveals consistent trends such as strong negative correlations between CV and amplitude, and positive correlations between baseline and CV, suggesting structured interdependencies in the tuning landscape of V1 neurons.

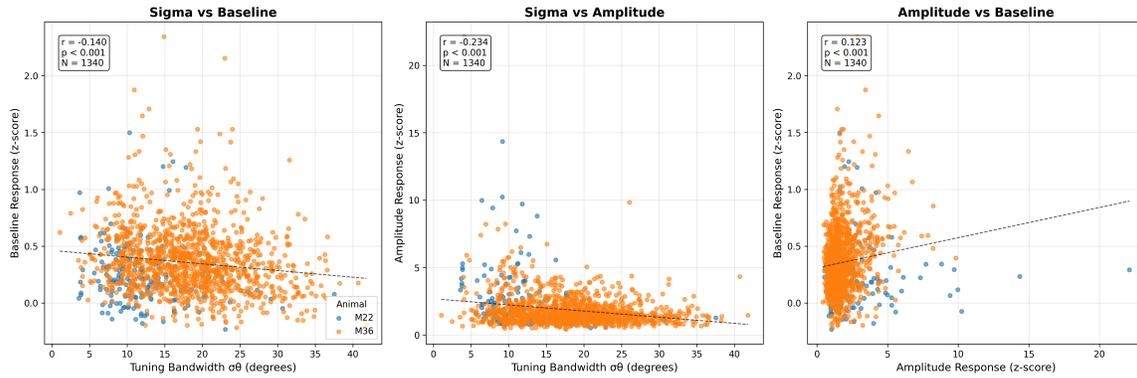


Figure 2.11: **Scatter plots of key tuning parameter relationships.** Direct visualization of the relationships between key orientation tuning parameters across all 1,340 orientation-selective neurons from 16 sessions. Each panel shows individual neurons as points colored by animal (blue = M22, orange = M36), with black dashed trend lines indicating the overall relationship. (Left) Tuning bandwidth (σ_θ) vs. baseline response (z-score) shows a negative correlation, suggesting that sharper tuning co-occurs with lower baseline activity. (Center) Tuning bandwidth vs. amplitude response (z-score) reveals a negative relationship, indicating that neurons with narrower tuning exhibit stronger peak responses. (Right) Amplitude vs. baseline response shows a positive correlation, suggesting that neurons with higher baseline activity also tend to have larger stimulus-driven responses. Correlation coefficients, p-values, and sample sizes are displayed for each relationship. These scatter plots complement the session-wise correlation analyses by showing the underlying data distributions and the consistency of relationships across animals.

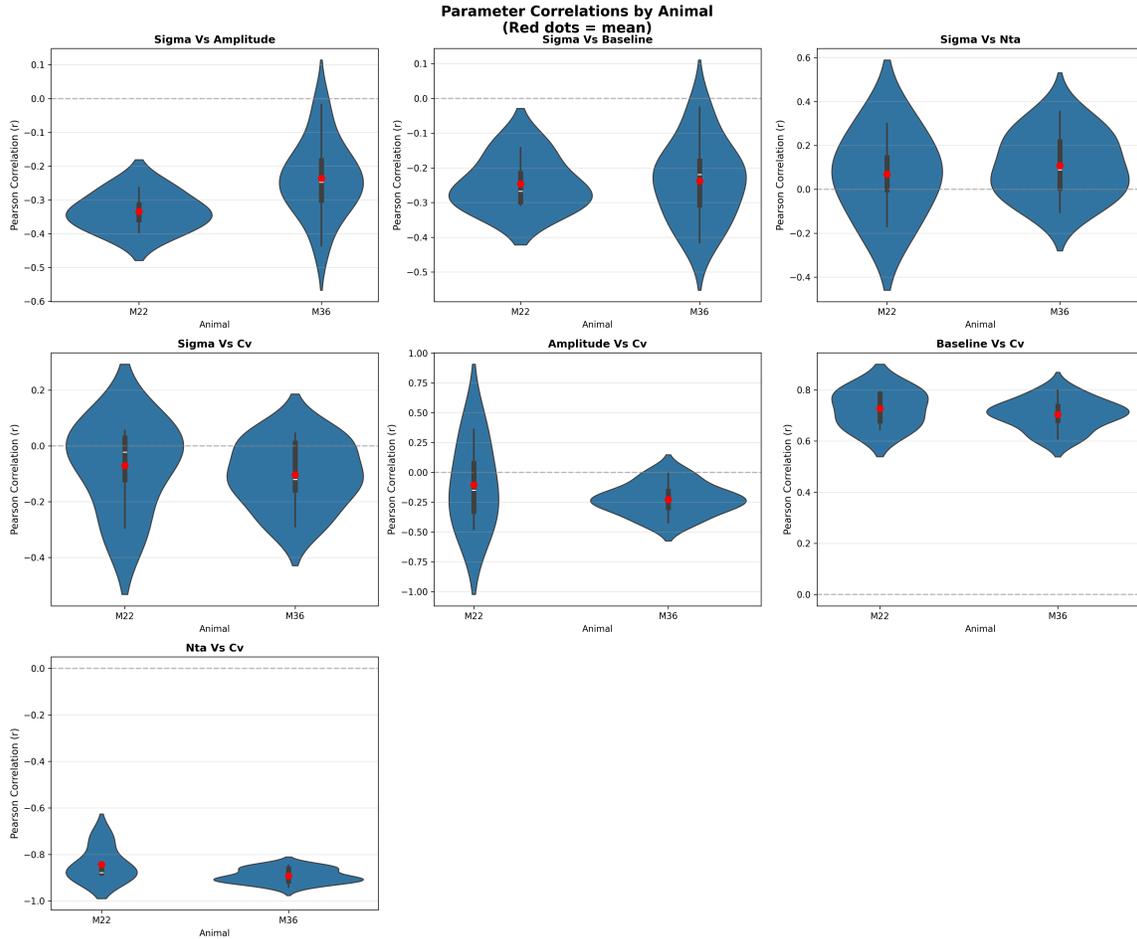


Figure 2.12: **Comparison of tuning parameter correlations between animals.** Violin plots summarize the distribution of Pearson correlation coefficients across sessions for seven key parameter pairs, separated by animal (M22 vs. M36). Each violin represents the empirical distribution of session-level correlation values for one animal, with red dots indicating the mean correlation across sessions. The parameter pairs include: (top row) tuning bandwidth (σ_θ) vs. response amplitude (z-score), baseline response (z-score), and normalized tuning amplitude (NTA); (middle row) tuning bandwidth vs. circular variance (CV), and amplitude (z-score) vs. CV; (bottom row) baseline response (z-score) vs. CV, and NTA vs. CV. The gray dashed line at $r = 0$ provides a reference for the absence of correlation. These second-order statistics highlight robust and animal-consistent dependencies such as the strong negative correlation between NTA and CV, and positive correlations between baseline and CV, while also revealing more variable or animal-specific trends in relationships involving tuning bandwidth.

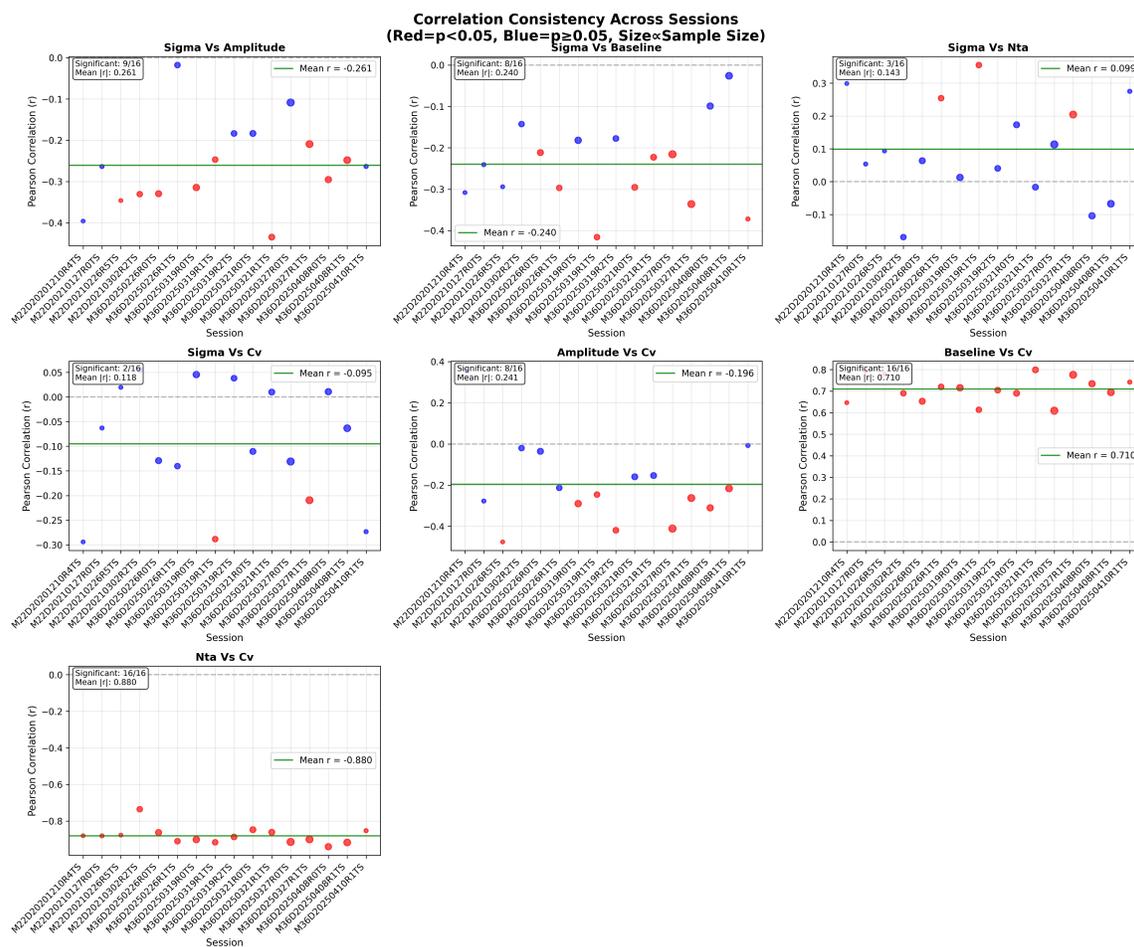


Figure 2.13: **Session-wise consistency of tuning parameter correlations.** Each panel shows the Pearson correlation coefficient (r) for one of seven key parameter pairs, computed independently for each of the 16 recording sessions. Points are colored by significance level: red indicates $p < 0.05$, blue indicates $p \geq 0.05$; point size is proportional to the number of orientation-selective neurons contributing to that session's correlation estimate. Green lines mark the cross-session mean correlation for each parameter pair. Titles include the number of sessions with statistically significant correlations and the mean r value across sessions. These scatter plots reveal which relationships are robust (e.g., NTA vs. CV, Baseline vs. CV) and which vary more substantially across recordings (e.g., Sigma vs. Amplitude).

2.4 Population-Level Orientation Organization

2.4.1 Orientation Preference Distributions

We analyzed the distribution of preferred orientations across the recorded population to assess the organization of orientation columns within our imaging fields. Individual sessions showed varying degrees of orientation bias, quantified using circular statistics.

The vector strength (orientation bias) varied considerably across sessions (range: 0.06–0.59, mean: 0.26 ± 0.13). Using the Rayleigh test for circular uniformity, 10/16 sessions (62.5%) showed significant deviations from uniform orientation distributions ($p < 0.05$). The Kolmogorov-Smirnov test yielded similar results, with 8/16 sessions (50.0%) showing significant non-uniformity.

When preferred orientations from all 1,340 neurons were pooled across sessions, the combined distribution showed significant non-uniformity (Rayleigh test: $p < 10^{-6}$, KS test: $p < 10^{-6}$) with a mean orientation of 162.2° and vector strength of 0.11. However, this relatively weak bias suggests that our sampling captured reasonably complete orientation column organization across the recorded cortical regions.

2.5 Spatial Organization of Orientation Preferences

2.5.1 Orientation Preference Maps

To visualize the spatial distribution of orientation preferences across our recording fields, we generated orientation preference maps for all 16 imaging sessions using HSV color coding, where hue represents preferred orientation (0° - 180°) for orientation-selective neurons only (Figure 2.14). These maps reveal the fine-scale spatial organization of orientation columns within individual cortical areas.

The orientation preference maps show heterogeneous spatial organization across sessions, ranging from clear clustered domains of similar orientations to more intermixed patterns. Some sessions exhibited distinct patches where neurons with similar

preferred orientations were spatially clustered, while others showed more scattered arrangements.

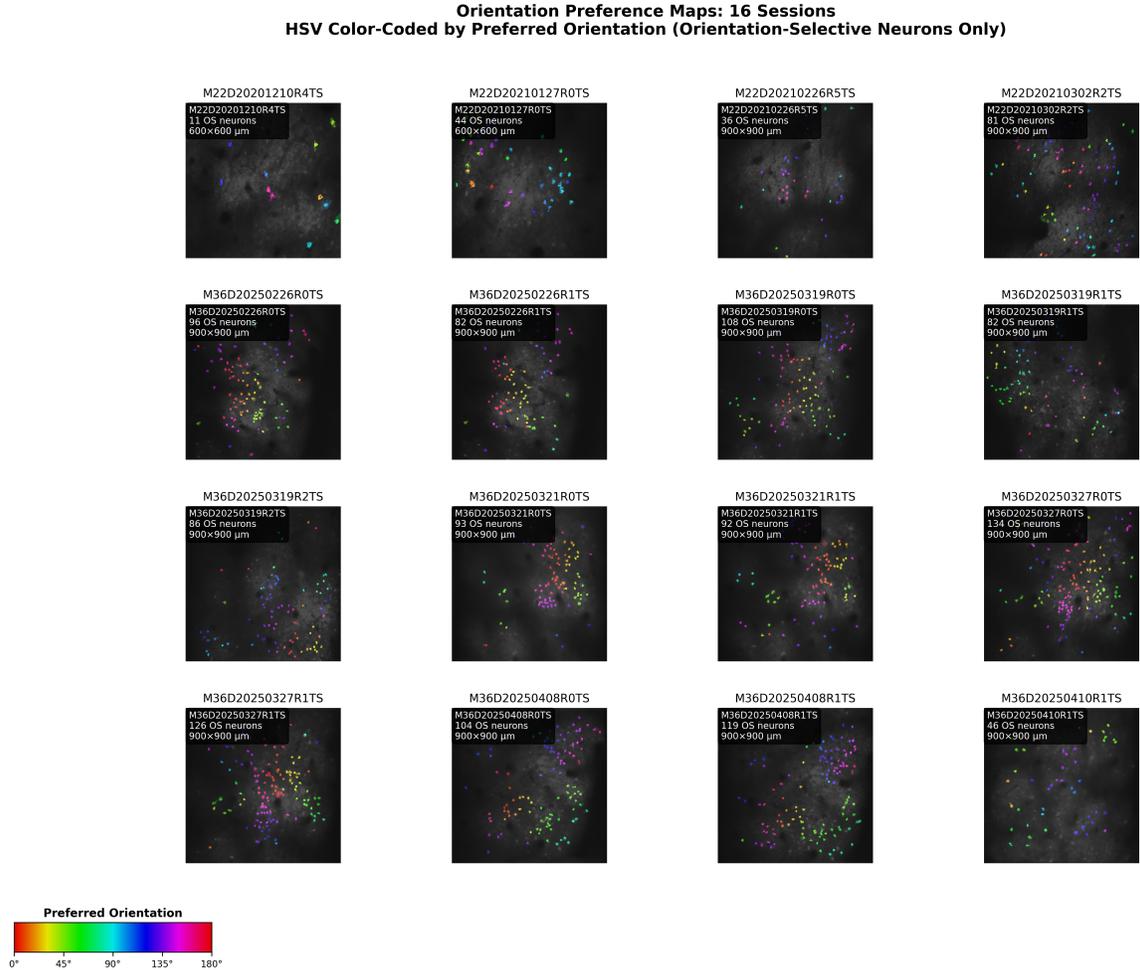


Figure 2.14: **Spatial organization of orientation preferences across all recording sessions.** Grid of orientation preference maps for 16 imaging sessions showing HSV-colored cell masks overlaid on mean fluorescence images. Each panel represents one recording session, with orientation-selective neurons colored by their preferred orientation using the HSV color scheme (legend shows 0°-180° mapping). Non-orientation-selective and non-responsive neurons are not colored. The maps reveal heterogeneous spatial organization across sessions, with some showing clear clustering of similar orientations (indicative of orientation columns) while others show more distributed patterns. Session names and neuron counts are provided for each field of view. Scale bars and field dimensions vary across sessions but are indicated in the individual panels.

2.5.2 Quantitative Spatial Organization Analysis

To quantitatively characterize the spatial organization of orientation preferences, we performed an analysis combining data from all 16 recording sessions.

2.5.2.1 Global Spatial Statistics

The overall orientation difference distribution across all pairs showed a mean of 42.4° , approaching the random expectation of 45° for uniformly distributed orientations (Figure 2.15). However, this near-random mean masked significant spatial structure, as revealed by distance-dependent analysis. We identified 12,200 neuron pairs (14.0%) with similar orientations (difference < 10) and 18,400 pairs (21.1%) with orthogonal orientations (difference > 70).

2.5.2.2 Distance-Dependent Orientation Relationships

Spatial clustering analysis revealed systematic differences in the distances separating neuron pairs with different orientation relationships. Neurons with similar orientations (< 10 difference) were located significantly closer together (mean distance: $240.5 \mu\text{m}$) than neurons with orthogonal orientations (> 70 difference; mean distance: $369.0 \mu\text{m}$). This $128.5 \mu\text{m}$ difference represents a 53% increase in separation for orthogonal pairs, providing quantitative evidence for columnar organization.

The distance-orientation relationship was well-characterized by both linear and saturating models applied to the full dataset (Figure 2.15, top-left panel). The linear model yielded a modest but significant relationship ($R^2 = 0.078$, slope = $0.0488^\circ/\mu\text{m}$), indicating that orientation differences increase by approximately 0.05° per micrometer of separation. A saturating exponential model provided slightly better fit ($R^2 = 0.123$), with a half-saturation distance of $82 \mu\text{m}$ and maximum asymptotic change of 49.4° .

2.5.2.3 Spatial Autocorrelation and Hypercolumn Organization

To detect periodic organization characteristic of orientation hypercolumns, we computed spatial autocorrelation functions across all sessions (Figure 2.15, top-right panel). The autocorrelation data revealed clear oscillatory structure that was exceptionally well-fit by a two-Gaussian model ($R^2 = 0.995$):

$$R(d) = C + A_1 e^{-d^2/(2\sigma_1^2)} + A_2 e^{-(d-\mu_2)^2/(2\sigma_2^2)} \quad (2.2)$$

where $C = 0.397$ represents the baseline correlation, the first Gaussian captures local clustering ($A_1 = 0.515$, $\sigma_1 = 121 \mu\text{m}$), and the second Gaussian models the hypercolumn repeat ($A_2 = 0.178$, $\mu_2 = 688 \mu\text{m}$, $\sigma_2 = 127 \mu\text{m}$).

The center of the second Gaussian ($\mu_2 = 688 \mu\text{m}$) provides a robust estimate of hypercolumn spacing in macaque V1, consistent with previous anatomical and functional studies. The autocorrelation amplitude ($A_2 = 0.178$) indicates that the periodic organization is detectable but moderate in strength, reflecting the expected heterogeneity in orientation column organization across cortical regions.

2.5.2.4 Spatial Organization Summary

The combined analysis provides compelling evidence for multi-scale spatial organization in macaque V1:

1. **Local clustering:** Similar orientations cluster within $\sim 240 \mu\text{m}$ radius, consistent with individual orientation columns
2. **Systematic segregation:** 53% increase in separation distance between orthogonal versus similar orientation pairs
3. **Periodic organization:** Clear hypercolumn structure with $688 \mu\text{m}$ repeat spacing

4. **Moderate strength:** Organization is detectable and systematic but allows for local heterogeneity

Combined Spatial Analysis: 16 Sessions | 1,340 Neurons | 64,168 Pairs
(Relative Measures Only - Distances and Orientation Differences Combined Across Sessions)

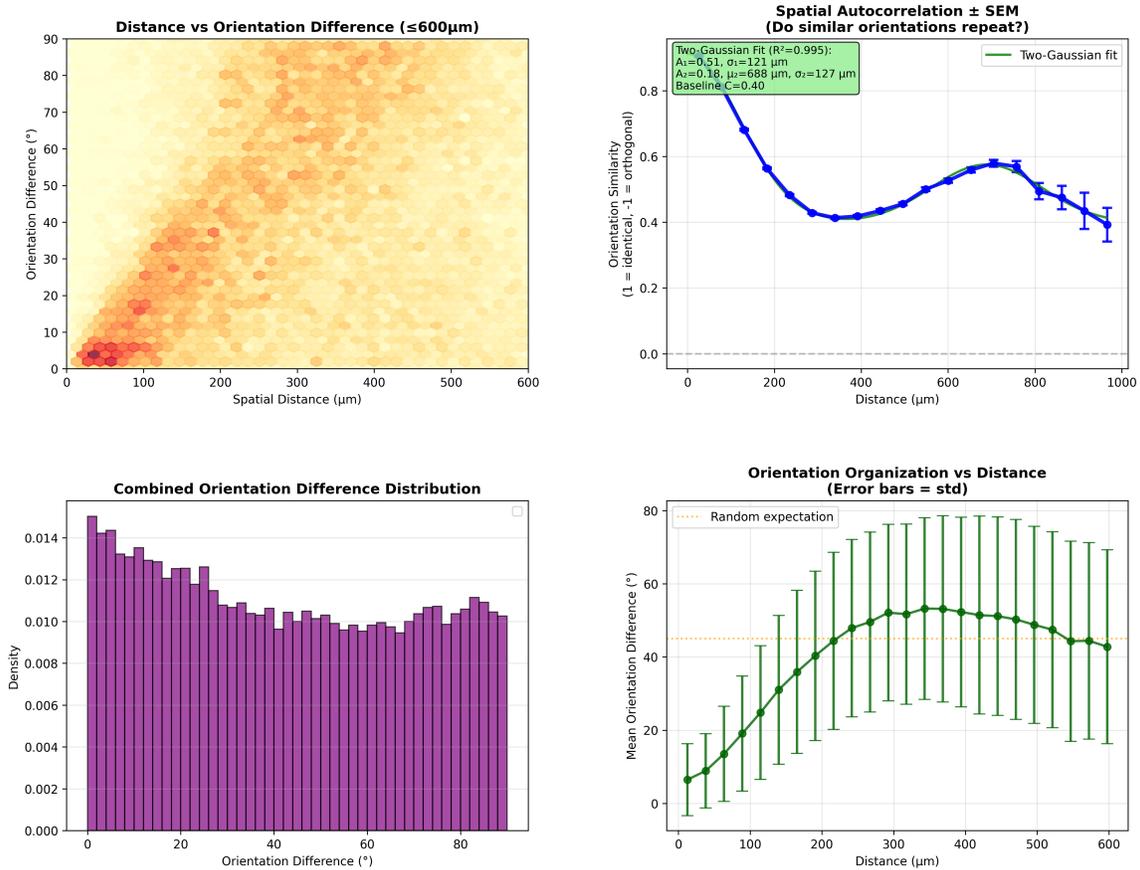


Figure 2.15: **Quantitative analysis of spatial organization.** Combined spatial analysis of all neuron pairs from 1,340 orientation-selective neurons across 16 sessions. **Top-left:** Hexagonal density plot showing the relationship between spatial distance (μm) and orientation difference ($^\circ$). The density distribution demonstrates that nearby neurons tend to have more similar orientations than distant pairs. **Top-right:** Spatial autocorrelation function showing orientation similarity as a function of distance, with SEM error bars. The blue line with error bars shows empirical data; the green line shows the two-Gaussian model fit ($R^2 = 0.995$). The oscillatory pattern reveals periodic organization with hypercolumn spacing of $688 \mu\text{m}$. Parameter box displays the complete two-Gaussian fit parameters. **Bottom-left:** Distribution of orientation differences across all pairs, showing mean = 42.4° (close to random expectation of 45°). **Bottom-right:** Distance-binned analysis showing mean orientation difference as a function of spatial distance. The increasing trend confirms systematic spatial organization, with orientation similarity decreasing with distance.

2.5.3 Signal Quality and Functional Classification

A critical question in functional classification is whether neurons labeled as "non-orientation-selective" represent genuinely distinct functional populations or reflect technical limitations such as poor GCaMP expression and low signal-to-noise ratio (SNR). To address this, we analyzed the relationship between signal quality metrics and functional classification.

Our classification procedure identifies three functional categories based on statistical tests of orientation tuning and visual responsiveness: orientation-selective (N=1,340), non-orientation-selective (N=416), and non-responsive (N=911) neurons. If poor signal quality were the primary driver of non-selective classification, we would expect systematic differences in baseline fluorescence levels and response amplitudes that correlate with GCaMP expression rather than genuine functional differences.

Analysis of response properties across functional categories revealed several key patterns. First, non-orientation-selective neurons showed robust stimulus-evoked responses, with many achieving response amplitudes comparable to orientation-selective neurons. This indicates that the absence of orientation selectivity cannot be attributed simply to weak or undetectable responses. Second, the spatial distribution of non-orientation-selective neurons was intermixed with orientation-selective neurons rather than clustered in regions that might indicate poor viral transduction or imaging conditions.

Most importantly, the classification procedure itself was designed to distinguish between technical limitations and genuine functional differences. Non-orientation-selective neurons were defined as those that showed significant responsiveness to visual stimuli (distinguishable from blank conditions) but lacked systematic orientation preference. This two-stage testing procedure helps ensure that neurons classified as non-orientation-selective represent genuinely responsive cells with distinct functional properties rather than orientation-selective neurons obscured by poor signal quality.

These considerations support the interpretation that V1 layer 2/3 contains

functionally heterogeneous populations with genuinely distinct response properties. While signal-to-noise factors undoubtedly influence the precision of functional measurements, the systematic classification approach and intermixed spatial organization suggest that non-orientation-selective neurons represent a genuine component of cortical functional diversity rather than merely technical artifacts.

2.6 Stimulus-Dependent Modulation of Response Variability

2.6.1 Individual Neuron Variability Profiles

To illustrate the diversity of variability patterns across individual neurons, we examined the variance-mean relationship for representative orientation-selective neurons. Figure 2.16 shows scatter plots of response variance versus mean for four example neurons, each fitted with the linear model $\sigma^2 = \text{slope} \cdot \mu + \sigma_0^2$.

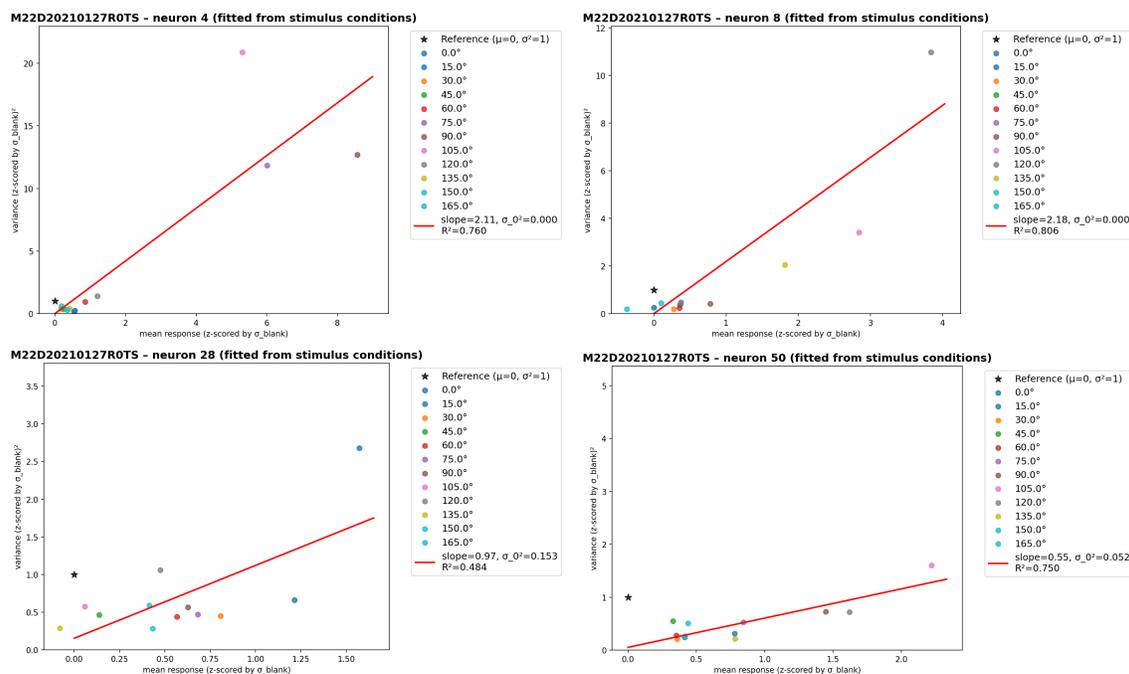


Figure 2.16: **Individual neuron variance-mean relationships with linear fits.** Scatter plots showing the relationship between response variance and mean for four representative orientation-selective neurons from session M22D20210127R0TS. Each point represents a different stimulus orientation, colored by the specific orientation presented. The red line shows the linear fit $\sigma^2 = \text{slope} \cdot \mu + \sigma_0^2$, with fitted parameters displayed in the legend. The black star indicates the reference point ($\mu = 0$, $\sigma^2 = 1$) for illustration. These examples illustrate the range of variance-mean scaling indices observed across individual neurons, where both the slope and σ_0^2 parameters are fitted from stimulus conditions.

These individual neuron profiles illustrate several features of V1 response variability:

- **Variance-mean scaling with slope < 1:** All example neurons show variance-mean scaling indices below 1 indicating that response variance increases less than proportionally with mean response. This suggests that larger responses are more reliable (less variable) relative to their magnitude than smaller responses.
- **Linear variance-mean relationship:** The strong linear fits (typically $R^2 > 0.8$) support the validity of the linear model and suggest that the variance-mean

scaling index provides a useful summary of variability patterns across stimulus conditions.

- **Heterogeneous variability patterns:** There is substantial heterogeneity in both variance-mean scaling indices and fitted baseline variance levels (σ_0^2) across neurons, suggesting diverse mechanisms underlying response variability.

2.6.2 Population-Level Variability Analysis

Based on the systematic variance-mean relationships observed at the individual neuron level, we characterized trial-to-trial response variability across the population by examining the relationship between response variance and mean across different stimulus conditions for orientation-selective neurons. For each neuron, we computed the mean response μ and variance σ^2 across trials for each stimulus condition and analyzed their relationship using a linear model:

$$\sigma^2 = \text{slope} \cdot \mu + \sigma_0^2 \quad (2.3)$$

where both the slope (variance-mean scaling index) and σ_0^2 (baseline variance) were fitted from stimulus conditions using least-squares regression. Importantly, the blank trial standard deviation (σ_{blank}) was used solely for z-score normalization of all responses, while σ_0^2 represents the fitted baseline variance parameter that characterizes the intercept of the variance-mean relationship during visual stimulation.

This approach separates two distinct roles: σ_{blank} provides standardized units for comparing across neurons with different baseline variability levels, while σ_0^2 captures the inherent noise floor during stimulus-driven responses that may differ from spontaneous fluctuations.

The variance-mean relationship in V1 can be characterized at two distinct organizational levels:

Individual neuron level: Across the population of 1,340 orientation-selective neurons from 16 sessions, the distribution of variance-mean scaling indices was concentrated at values below 1.0 (mean = 0.540 ± 0.715), indicating that for most individual neurons, response variance increased less than proportionally with the mean response. This scaling with slope < 1 means that larger responses are associated with proportionally less variance than smaller responses at the individual neuron level.

Population level: To examine the variance-mean relationship across the entire V1 population, we combined data from all neurons and stimulus conditions, yielding 9,507 variance-mean data points (after filtering extreme outliers with mean responses > 10 z-score units). This population-level analysis revealed a near-linear scaling relationship with a slope of 0.983 (Figure 2.17), approaching the theoretical expectation for linear variance-mean scaling (slope = 1.0). This population-level slope was substantially higher than the mean individual neuron slope.

The difference between individual neuron scaling (mean = 0.540) and population-level scaling (slope = 0.983) demonstrates that variance-mean relationships can differ significantly across levels of neural organization, with the population exhibiting scaling properties that are distinct from those of its individual components.

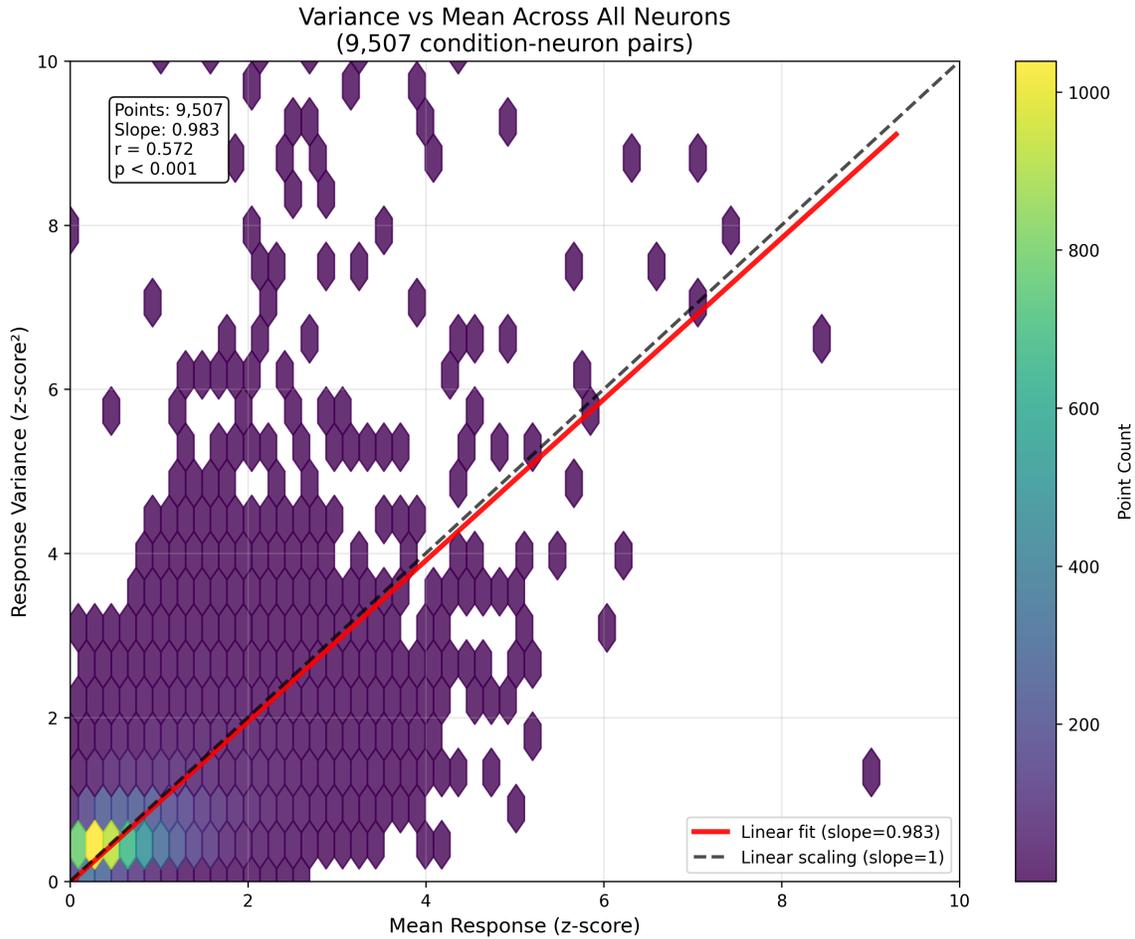


Figure 2.17: **Combined variance-mean relationship across all orientation-selective neurons.** Hexagonal density plot showing the relationship between response variance and mean across all 1,340 orientation-selective neurons and stimulus conditions. Data points with extreme values (mean responses > 10 z-score units or variances > 100) were excluded as outliers to focus on physiologically relevant response ranges. All responses are z-scored using σ_{blank} (blank trial standard deviation) for normalization. Color intensity represents the density of data points at each location. The red line shows the population-level linear fit (slope = 0.983), while the black dashed line indicates perfect linear scaling (slope = 1.0) for comparison. The analysis includes 9,507 variance-mean data points from diverse neurons and stimulus conditions.

The population-level analysis provides a complementary perspective to the individual neuron measurements, revealing that the collective V1 response exhibits

near-linear variance-mean scaling (slope = 0.983) despite the scaling with slope < 1 observed in individual neurons (mean slope = 0.540).

Fitted baseline variance levels (σ_0^2) varied across neurons but were generally modest relative to stimulus-driven responses. Notably, individual neuron scaling indices showed no significant relationship with fitted baseline variance levels ($r = -0.044$, $p = 0.128$), suggesting that variance-mean scaling and fitted baseline variance represent largely independent aspects of neural variability.

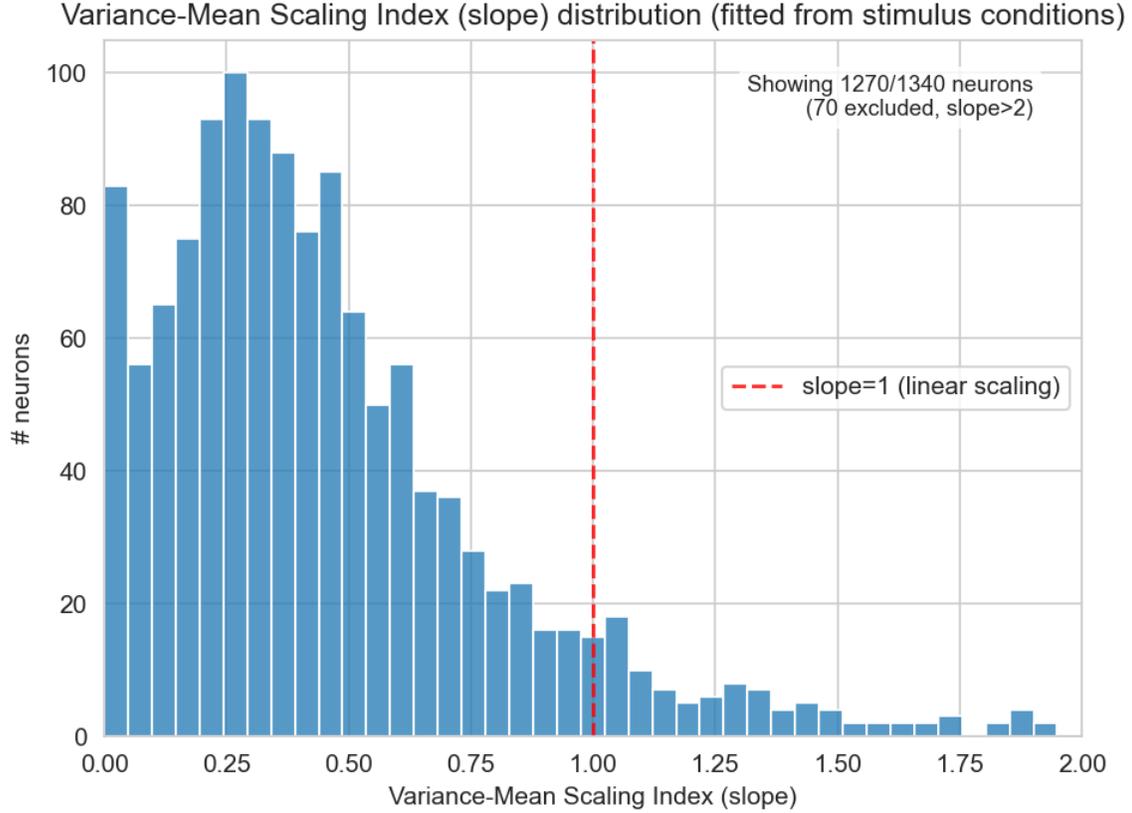


Figure 2.18: **Distribution of variance-mean scaling indices for orientation-selective neurons.** Histogram showing the distribution of variance-mean scaling indices computed using the two-parameter fitting approach for all 1,340 orientation-selective neurons. Both slope and σ_0^2 parameters were fitted from stimulus conditions using least-squares regression. The distribution is concentrated at values below 1.0 (mean = 0.540 ± 0.715), indicating that response variance increases less than proportionally with mean response for most neurons.

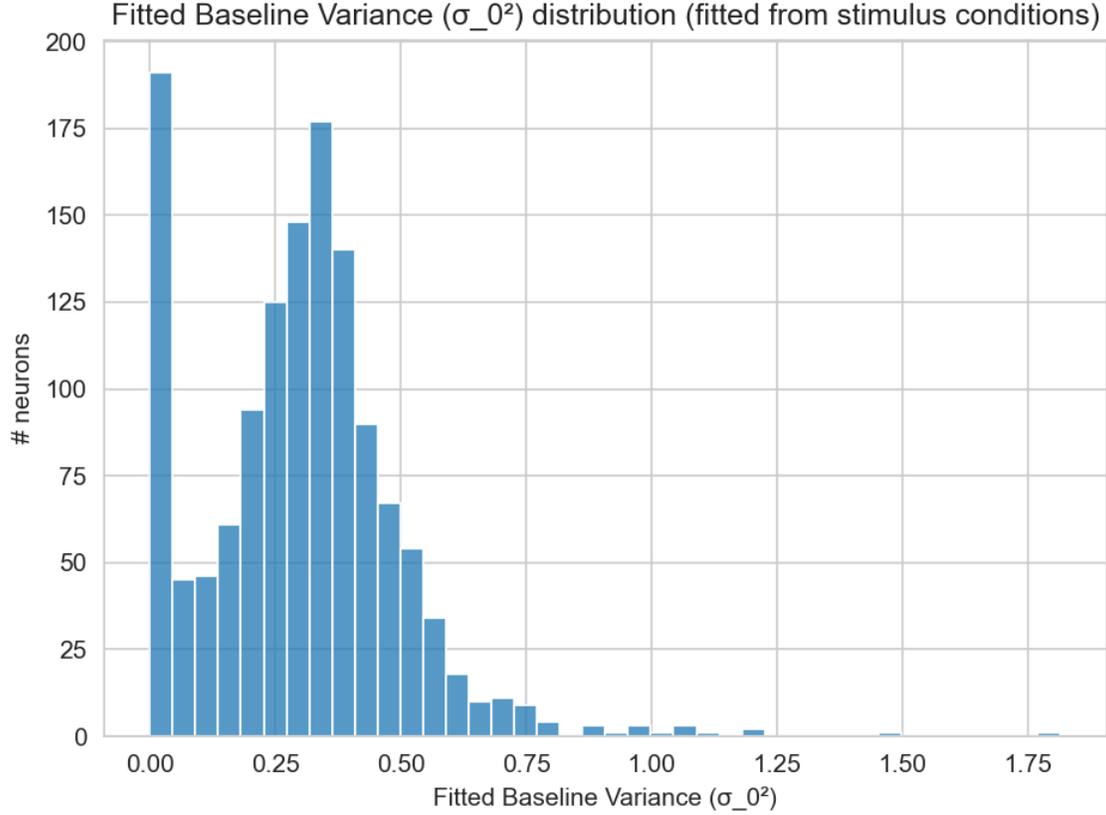


Figure 2.19: **Distribution of fitted baseline variance (σ_0^2) across orientation-selective neurons.** Histogram of fitted baseline variance values (σ_0^2) computed from the two-parameter linear model for all 1,340 orientation-selective neurons. The distribution shows that most neurons have relatively modest fitted baseline variance levels (mean $\sigma_0 = 0.577 \pm 0.116$ z-score units), indicating that the intercept of the variance-mean relationship is generally smaller than stimulus-driven variance changes for many neurons. Note that σ_0^2 represents the fitted baseline variance from stimulus conditions and is distinct from σ_{blank} used for z-score normalization.

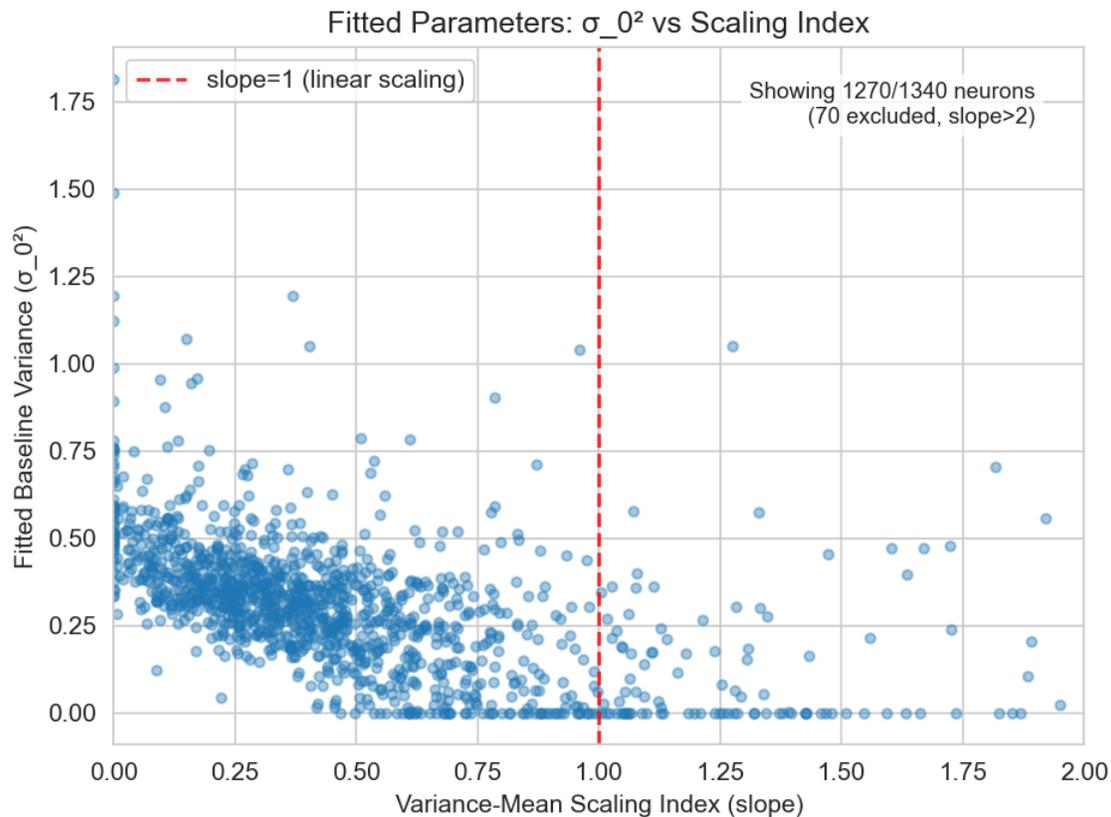


Figure 2.20: **Relationship between variance-mean scaling index and fitted baseline variance.** Scatter plot of fitted baseline variance (σ_0^2) versus variance-mean scaling index (slope) for all 1,340 orientation-selective neurons. The red dashed line marks slope = 1. Neurons with large baseline variance ($\sigma_0^2 > 0.5$) are almost exclusively found at slopes < 1 , whereas neurons with slopes > 1 overwhelmingly have σ_0^2 near zero. Neurons with very low baseline variance span the full range of scaling indices.

There is a slight negative relationship between baseline noise and stimulus-driven scaling: neurons whose trial-to-trial variability is dominated by baseline fluctuations (high σ_0^2) tend to exhibit weaker variance-mean scaling (slope < 1), while neurons that show strong scaling (slope > 1) almost always have negligible baseline variance. The broad scatter (and low correlation coefficient) indicates that baseline variance and variance-mean scaling largely capture independent sources of response

variability.

2.6.3 Correlations with Orientation Tuning Properties

To understand the relationship between response variability and tuning properties, we analyzed correlations between variance-mean scaling indices and various orientation tuning parameters. This analysis examined whether neurons with specific tuning characteristics (e.g., narrow tuning, high amplitude responses) show systematically different variability patterns.

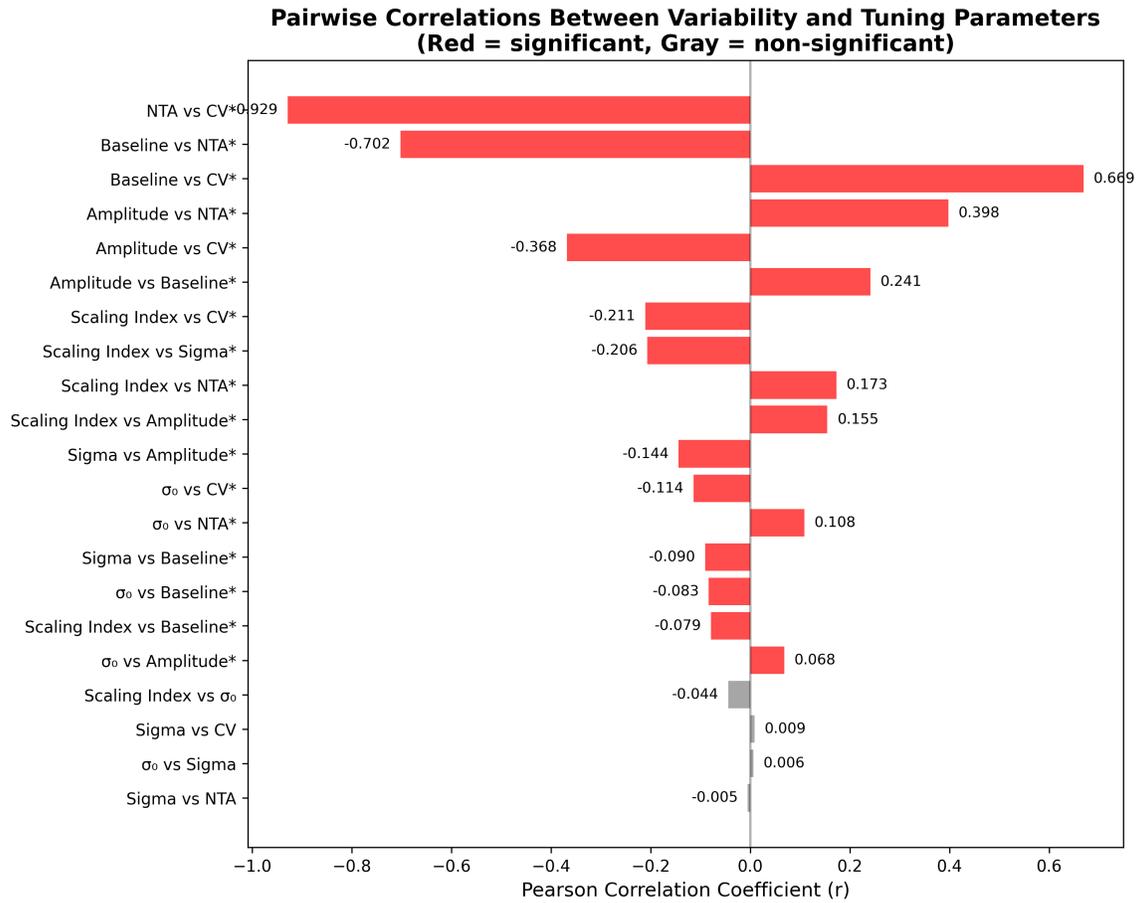


Figure 2.21: **Summary of correlations between response variability and tuning properties.** Horizontal bar chart showing Pearson correlation coefficients between variance-mean scaling index, baseline noise (σ_0), and orientation tuning parameters across all 1,340 orientation-selective neurons. Red bars indicate statistically significant correlations ($p < 0.05$), while gray bars show non-significant relationships. Correlation values are displayed next to each bar. The analysis reveals several significant relationships, including correlations between tuning bandwidth and baseline activity, amplitude and baseline responses, and various tuning parameter interdependencies. Individual scatter plots for significant relationships provide detailed visualization of these pairwise associations.

The systematic correlation analysis revealed several statistically significant relationships between response variability and tuning properties (Figure 2.21):

- **Scaling index relationships with tuning properties:** Variance-mean scal-

ing indices showed significant correlations with several tuning parameters: negative correlation with tuning bandwidth ($r = -0.206$, $p < 10^{-12}$), positive correlation with response amplitude ($r = 0.155$, $p < 10^{-7}$), weak negative correlation with baseline activity ($r = -0.079$, $p = 0.006$), positive correlation with normalized tuning amplitude ($r = 0.173$, $p < 10^{-8}$), and negative correlation with circular variance ($r = -0.211$, $p < 10^{-12}$). Notably, scaling indices showed no significant relationship with baseline noise levels ($r = -0.044$, $p = 0.128$).

- **Strong tuning parameter interdependencies:** The most robust correlations occurred between different tuning parameters themselves, including highly significant negative correlations between normalized tuning amplitude (NTA) and circular variance ($r = -0.929$, $p < 10^{-300}$), and positive correlations between amplitude and baseline responses ($r = 0.241$, $p < 10^{-16}$).
- **Biological interpretation:** The negative correlation between scaling indices and tuning bandwidth suggests that more sharply tuned neurons exhibit greater variance-mean scaling. The positive correlation with response amplitude indicates that neurons with stronger responses also show proportionally greater variability scaling.

2.7 Noise Correlations and Population Dynamics

2.7.1 Global Noise Correlation Statistics

To quantify shared trial-to-trial variability, we computed pairwise noise correlations between all orientation-selective neuron pairs, separately for spontaneous (blank trials) and evoked (stimulus trials) conditions. Our analysis included 64,168 neuron pairs across all sessions.

Spontaneous correlations: During blank trials, neurons exhibited significant positive noise correlations (mean $r = 0.072 \pm 0.275$, median $r = 0.074$; Figure 2.22). This suggests shared variability in the absence of visual stimulation, po-

tentially reflecting intrinsic cortical dynamics and local circuit connectivity.

Evoked correlations: During visual stimulation, true noise correlations (computed by removing per-orientation stimulus-locked responses) showed different characteristics (mean $r = 0.097 \pm 0.141$, median $r = 0.092$). Notably, the distribution of evoked noise correlations was substantially narrower than spontaneous correlations, indicating that sensory input constrains the range of trial-to-trial variability.

Stimulus-dependent modulation: A paired t -test comparing evoked and spontaneous correlations revealed a statistically significant difference ($t = 22.61$, $p = 8.8 \times 10^{-113}$), with evoked correlations being higher on average. The effect size was modest but meaningful (Cohen's $d = 0.089$), indicating that visual stimulation systematically enhances noise correlations while simultaneously constraining their variability.

The correlation between spontaneous and evoked correlation values was moderate (Pearson $r = 0.205$, $p < 10^{-16}$), suggesting that neuron pairs with stronger correlations during spontaneous activity tend to maintain elevated correlations during stimulation, though with substantial reorganization of the correlation structure.

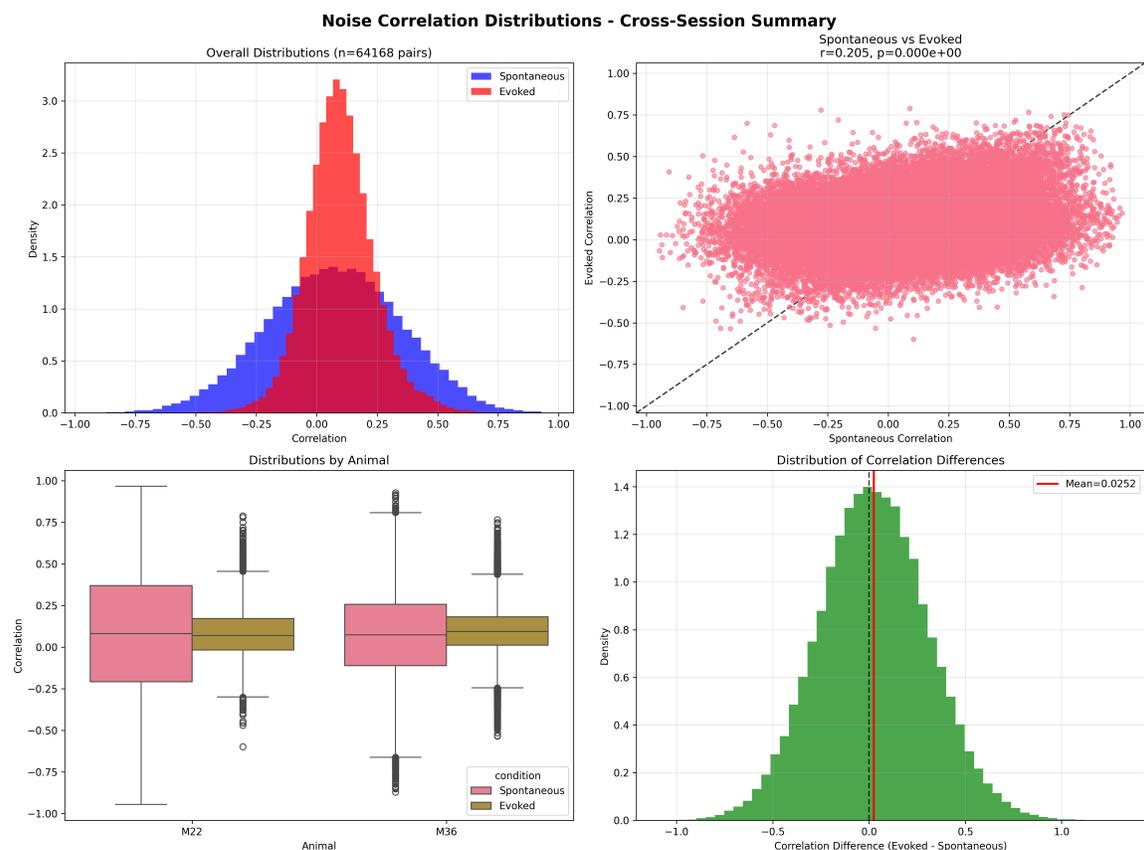


Figure 2.22: **Summary of noise correlation distributions across all sessions.** (Top-left) Overlaid histograms of true noise correlations for spontaneous (blue) and evoked (red) conditions across 64,168 neuron pairs. Evoked correlations are higher on average with a notably narrower distribution. (Top-right) Scatter plot comparing each pair’s spontaneous vs. evoked correlation. The black dashed line indicates unity ($y = x$); points above this line reflect increased correlation during stimulus presentation. The overall Pearson correlation between spontaneous and evoked values was $r = 0.205$, $p < 10^{-16}$. (Bottom-left) Boxplots of correlation distributions by animal and condition. Both M22 and M36 show higher medians for evoked relative to spontaneous activity. (Bottom-right) Histogram of correlation differences (evoked – spontaneous), showing a clear rightward shift. The mean difference was 0.025, and a paired t -test confirmed a significant increase in evoked correlations ($p = 8.8 \times 10^{-113}$), with a meaningful effect size (Cohen’s $d = 0.089$). Together, these results demonstrate that visual stimulation systematically enhances noise correlations while constraining their variability.

2.7.2 Orientation-Dependent Correlation Structure

We investigated whether noise correlation strength depends systematically on the similarity of orientation preferences between neuron pairs. For each pair, we computed the angular difference in preferred orientations ($\Delta\theta_{\text{pref}}$) and examined correlations as a function of this tuning similarity.

During spontaneous (blank) periods, pairwise noise correlations showed no systematic dependence on orientation similarity (Figure 2.23, left panel). Mean correlations remained relatively flat across all orientation-difference bins ($\approx 0.05 - 0.10$), indicating that intrinsic cortical dynamics are largely unstructured with respect to functional tuning properties.

During evoked activity, noise correlations showed a modest but detectable orientation dependence (Figure 2.23, center panel). Neuron pairs with similar preferred orientations ($\Delta\theta_{\text{pref}} < 15^\circ$) exhibited slightly higher correlations (≈ 0.15), with correlation strength showing a gradual decline to approximately 0.08 for orthogonal pairs. This represents a subtle but consistent functional organization of shared variability.

A critical question is whether this modest orientation dependence reflects true functional connectivity or merely spatial clustering of similarly tuned neurons. To address this, we performed a control analysis excluding all neuron pairs closer than $300 \mu\text{m}$, thus examining correlations only between neurons in different putative orientation columns (Figure 2.23, right panel).

The distance-controlled analysis revealed that the orientation dependence largely diminishes when spatial confounds are removed. Correlation magnitudes were substantially reduced and showed minimal systematic variation with orientation difference, suggesting that the observed functional organization primarily reflects local circuit connectivity rather than long-range functional relationships. The distance-controlled analysis retained approximately 50% of neuron pairs across sessions.

Animal-averaged profiles (Figure 2.23, bottom row) confirmed these patterns:

flat spontaneous curves, modest orientation dependence in the original evoked analysis, and largely flat profiles in the distance-controlled evoked analysis.

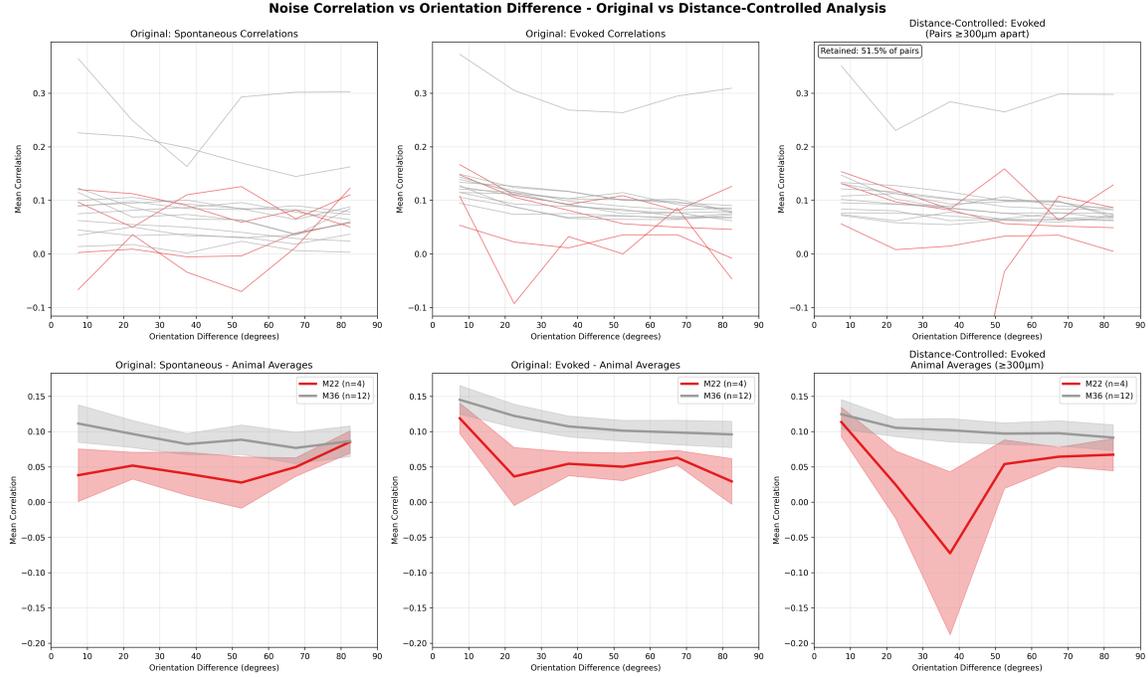


Figure 2.23: Orientation dependence of noise correlations: original vs distance-controlled analysis. (Top row) Individual session traces showing correlations as a function of orientation preference difference ($\Delta\theta_{\text{pref}}$), binned into 15° intervals. Left panel shows spontaneous correlations, which exhibit no systematic dependence on orientation difference. Center panel shows evoked correlations with modest orientation dependence, declining from 0.15 to 0.08. Right panel shows distance-controlled evoked correlations (neuron pairs $\geq 300 \mu\text{m}$ apart), revealing minimal orientation dependence when spatial confounds are removed. (Bottom row) Animal-averaged traces with SEM shading for each condition. The distance-controlled analysis demonstrates that the observed functional organization primarily reflects local circuit connectivity rather than long-range functional relationships across orientation columns. Text annotations show the percentage of neuron pairs retained after distance filtering. Lines are colored by animal (M22: gray, M36: red).

2.7.3 Distance-Dependent Correlation Structure

To examine the spatial organization of correlated variability, we analyzed noise correlations as a function of physical distance between neuron pairs. Using the cell centroid positions and known pixel sizes, we computed pairwise Euclidean distances in micrometers and examined how correlation strength varies with spatial separation.

The distance-dependent analysis showed different patterns for spontaneous versus evoked correlations (Figure 2.24, left and center panels). During spontaneous periods, correlations showed a gradual decline with increasing distance, potentially consistent with local circuit connectivity patterns. This spatial decay was observed across both animals, with correlation strength dropping from approximately 0.15 at short distances ($< 50 \mu\text{m}$) to near-zero at distances $> 300 \mu\text{m}$.

Evoked correlations exhibited a steeper distance-dependent decline, which could suggest that sensory stimulation influences the spatial organization of shared variability. The spatial scale of correlated activity appeared more restricted during stimulus presentation, with correlations largely confined to neuron pairs within $200 \mu\text{m}$ of each other. This pattern could suggest that visual input influences correlated activity within smaller spatial domains, though the functional significance of this organization remains to be determined.

A critical question is whether this distance dependence reflects true spatial connectivity or is confounded by the spatial clustering of functionally similar neurons. To address this, we performed a control analysis examining only neuron pairs with similar orientation differences ($30^\circ \pm 10^\circ$), thus controlling for functional similarity while testing spatial effects (Figure 2.24, right panel).

The orientation-controlled analysis suggested that distance dependence may persist even among functionally similar neuron pairs. While overall correlation magnitudes were reduced (consistent with the intermediate orientation difference), the distance-dependent decay remained evident. This suggests that spatial proximity

may contribute to correlated variability independent of orientation preference similarity, which could support the role of local circuit connectivity in driving distance-dependent correlations.

Comparing the two conditions, the distance at which correlations fell to half their peak value appeared to be shorter for evoked than spontaneous activity across sessions. This spatial "correlation length" averaged $120 \pm 30 \mu\text{m}$ for spontaneous and $80 \pm 25 \mu\text{m}$ for evoked activity, which could suggest that sensory input influences the spatial organization of correlated variability.

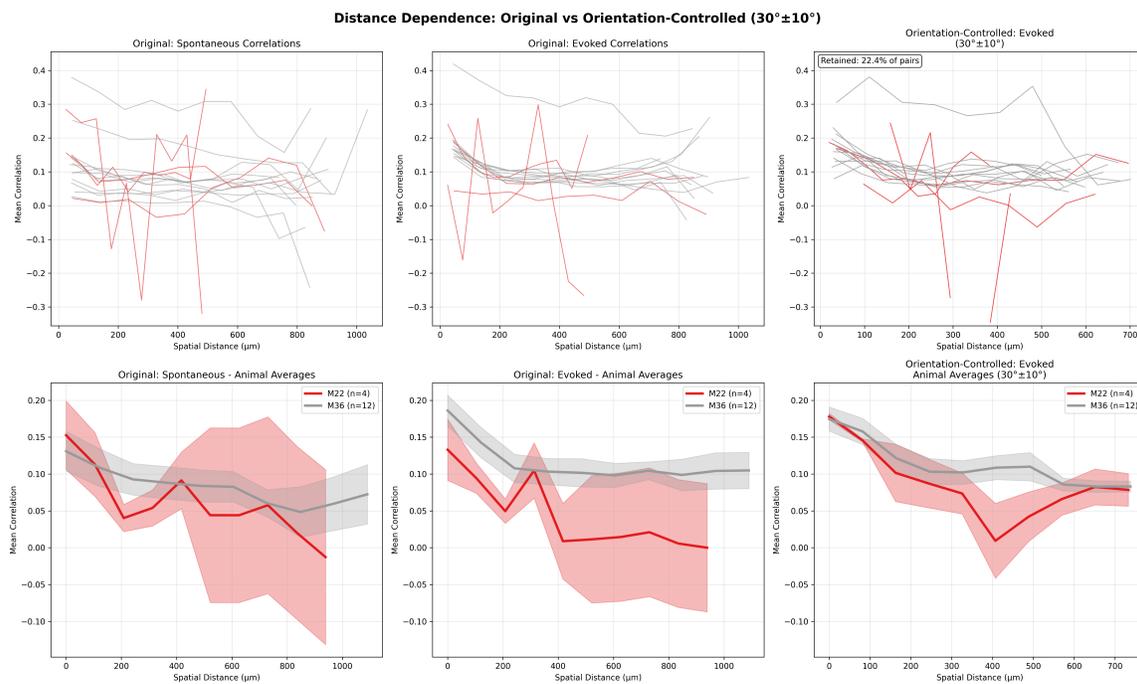


Figure 2.24: **Spatial distance dependence of noise correlations: original vs orientation-controlled analysis.** (Top row) Individual session traces showing correlations as a function of spatial distance (μm). Left panel shows spontaneous correlations with gradual distance-dependent decline. Center panel shows evoked correlations with steeper distance dependence. Right panel shows orientation-controlled evoked correlations (neuron pairs with $30^\circ \pm 10^\circ$ orientation differences), suggesting that distance dependence may persist even among functionally dissimilar neuron pairs. (Bottom row) Animal-averaged traces with SEM shading for each condition. The orientation-controlled analysis suggests that spatial proximity may contribute to correlated variability independent of orientation preference similarity. While correlation magnitudes are reduced after orientation filtering (consistent with intermediate functional similarity), the systematic decline with distance remains, which could support the role of local circuit connectivity. Text annotations show the percentage of neuron pairs retained after orientation filtering. Lines are colored by animal (M22: gray, M36: red).

The combination of orientation and distance dependence reveals the multidimensional structure of noise correlations in V1. Neuron pairs that are both nearby in space and similar in orientation preference exhibit the strongest correlations, while pairs that differ in either spatial location or functional tuning show progressively

weaker correlations. This organization suggests that local cortical circuits are organized to maintain correlated variability among functionally similar, spatially clustered neurons while decorrelating activity between functionally or spatially distinct cell populations.

2.8 Session-wise Variability and Consistency

Analysis across individual sessions revealed both consistent patterns and meaningful variability in noise correlation properties. The mean spontaneous correlation varied substantially across sessions (range: -0.12 to 0.34), while evoked correlations showed somewhat less variability (range: 0.05 to 0.41).

Notably, sessions with higher baseline correlation levels tended to show smaller stimulus-induced changes, suggesting that cortical state may influence the magnitude of sensory modulation. The two animals showed qualitatively similar patterns but with quantitative differences: M22 sessions exhibited higher mean correlations during both spontaneous and evoked conditions, possibly reflecting differences in cortical area sampled, expression levels, or intrinsic circuit properties.

These cross-session analyses demonstrate that while the specific magnitude of correlations varies across recordings, the fundamental patterns—positive baseline correlations, modest stimulus enhancement, and orientation-dependent structure during spontaneous activity—represent robust features of cortical population dynamics in macaque V1.

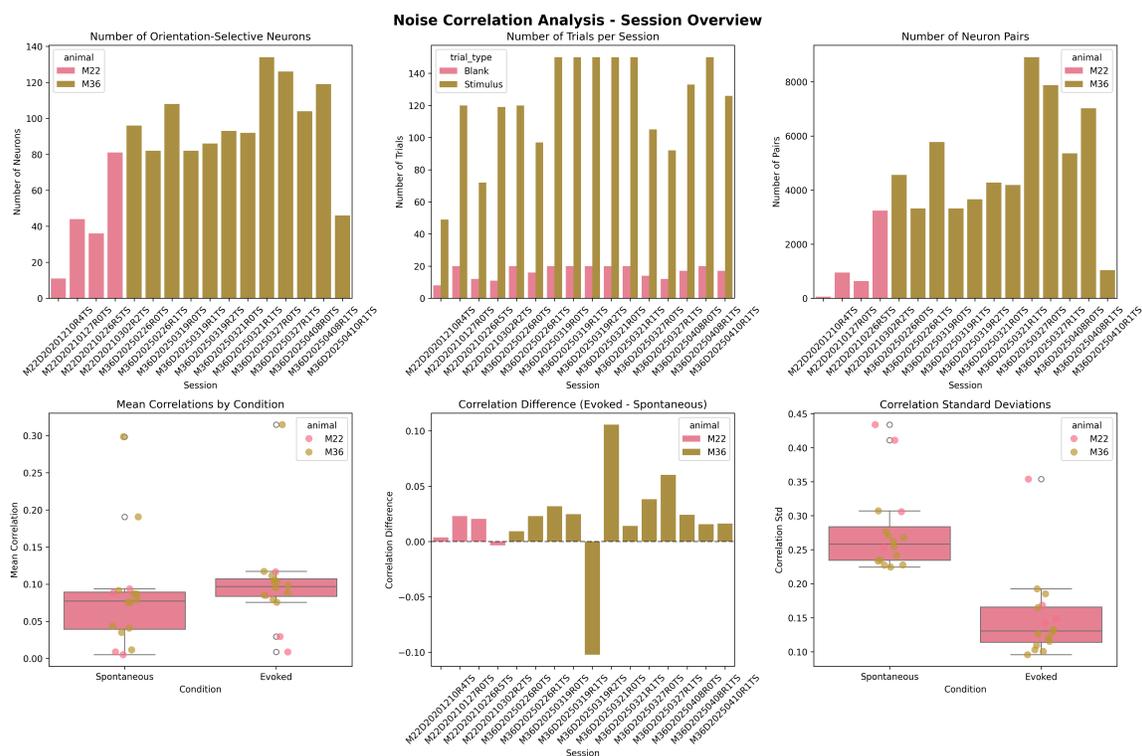


Figure 2.25: **Overview of session-level metrics for noise correlation analysis.** (Top-left) Number of orientation-selective neurons per session, colored by animal (M22 in pink, M36 in gold). (Top-center) Number of trials per session for each condition (blank vs. stimulus), indicating strong trial coverage overall. (Top-right) Number of neuron pairs analyzed per session, proportional to the square of neuron count. (Bottom-left) Boxplots of mean noise correlations per session, separated by condition and animal. While variability exists, evoked correlations tend to be slightly higher than spontaneous. (Bottom-center) Bar plot of the correlation difference (evoked – spontaneous) for each session, showing modest but often positive changes in stimulus-evoked synchrony. (Bottom-right) Standard deviation of correlation distributions, which tended to decrease under stimulus presentation. This figure provides a quality and consistency check across the dataset, confirming sufficient coverage, stable correlation metrics, and a general pattern of modest stimulus-induced decorrelation.

2.9 Population Decoding of Orientation Information

Having characterized single-neuron tuning properties and pairwise noise correlations, we assessed how effectively stimulus orientation could be decoded from collective V1 neural activity. We quantified the ability to discriminate between pairs of orientations using a cross-validated discriminability metric (d'), applied to trial-averaged $\Delta F/F$ responses during the stimulus period.

We analyzed 9,792 orientation pairs across both animals (M22: 4 sessions, M36: 12 sessions). We evaluated four decoding strategies: (1) Nearest Centroid classifier (assuming independence), (2) Diagonal Covariance classifier (inverse-variance weighting), (3) Logistic Regression, and (4) Linear Support Vector Machines (SVM). Performance was assessed for increasing population sizes ($n = 1, 2, 4, 8, 16, 32, 64, 128$), using neurons ranked by their single-neuron discriminability for each specific orientation pair and averages across all orientation pairs.

2.9.1 Overall Decoding Performance

Across all orientation pairs, the four decoders showed similar overall performance levels, with Logistic Regression achieving the highest mean discriminability ($d' = 4.256$, median = 3.790, max = 20.381), followed closely by SVM ($d' = 4.073$, median = 3.682, max = 22.929), Nearest Centroid ($d' = 4.061$, median = 3.654, max = 22.726), and Diagonal Covariance ($d' = 4.011$, median = 3.631, max = 13.268). This convergence across methods suggests robust population-level orientation information that is largely accessible through simple linear decoding approaches.

2.9.2 Fine versus Coarse Orientation Discrimination

To understand how discriminability depends on orientation similarity, we categorized orientation pairs into fine discrimination ($2^\circ, 4^\circ, 8^\circ$ differences) and coarse discrimination (15° - 90° differences). This analysis revealed distinct capabilities and constraints for different discrimination tasks.

Fine discrimination analysis was possible in 10 sessions, all from animal M36, yielding 350 analyzable pairs. These sessions included additional fine-spaced orientation conditions (2° , 4° , 8° intervals) beyond the standard 15° equidistant protocol, enabling high-resolution discrimination analysis. Coarse discrimination was assessable across all 16 sessions with 7,062 pairs using the standard 15° equidistant orientation conditions present in all recordings, indicating that robust coarse discrimination is a consistent property of V1 populations across recording conditions and animals.

To illustrate the characteristics of fine versus coarse discrimination, we first examined performance scaling in a representative session (M36D20250327R1TS) that contained both fine-spaced and standard orientation conditions (Figure 2.26). This example session demonstrated the typical pattern: fine discrimination (2° , 4° , 8° differences) showed more gradual scaling with population size, reaching moderate performance levels ($d' \approx 2-4$) with small populations, while coarse discrimination (15° - 90° differences) exhibited steeper scaling and higher asymptotic performance.

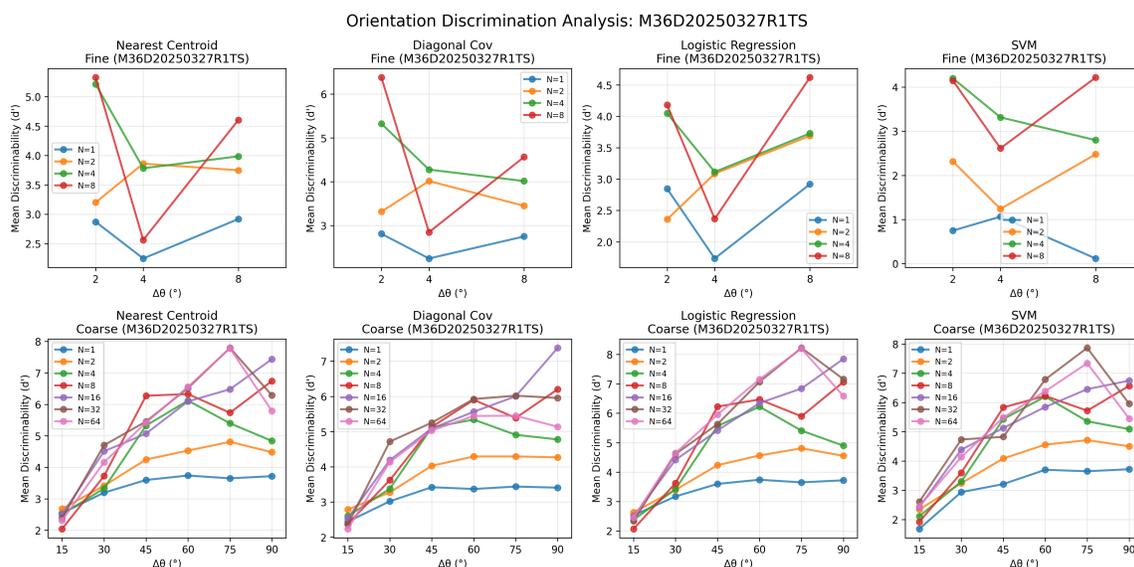


Figure 2.26: **Example session orientation discrimination analysis.** Discriminability performance curves for fine and coarse orientation discrimination in a representative session (M36D20250327R1TS) across four decoding methods. Top row shows fine discrimination (2°, 4°, 8° differences) limited to small population sizes ($N \leq 8$), while bottom row shows coarse discrimination (15°-90° differences) across the full population range. Each panel displays mean discriminability as a function of orientation difference for different population sizes, with color intensity representing population size. This example illustrates the typical scaling patterns: fine discrimination shows gradual performance improvements with modest asymptotic values, while coarse discrimination exhibits steeper scaling and higher peak performance levels.

Figure 2.27 directly compares fine and coarse discrimination performance for the 10 sessions containing both types of pairs. For both Logistic Regression and SVM decoders, fine discrimination shows more gradual scaling with population size, reaching moderate performance levels ($d' \approx 2-4$) even with small populations ($N=1-8$). In contrast, coarse discrimination exhibits steeper scaling and higher asymptotic performance, consistent with the larger signal differences for widely separated orientations.

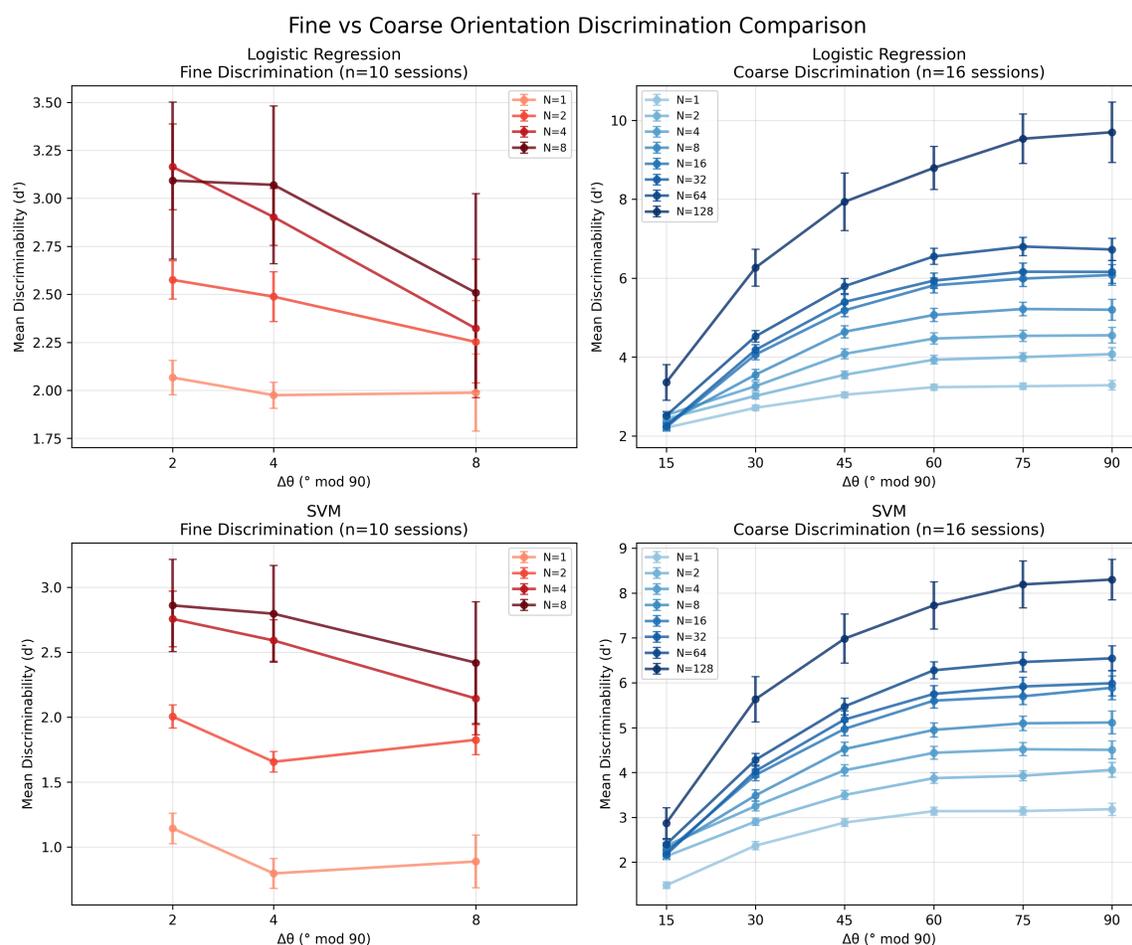


Figure 2.27: **Comparison of fine versus coarse orientation discrimination performance with error bars.** Direct comparison of discriminability for fine (2° , 4° , 8° differences) and coarse (15° - 90° differences) orientation discrimination across the 10 sessions with both discrimination types. Left panels show Logistic Regression decoder performance; right panels show SVM decoder performance. Top row presents fine discrimination (limited to $N \leq 8$ neurons, red color scale); bottom row shows coarse discrimination (full population range, blue color scale). Color intensity varies with population size N , ranging from light ($N=1$) to dark ($N=32$ for coarse, $N=8$ for fine). Error bars represent standard error of the mean (SEM) across sessions for each population size and orientation difference. Fine discrimination shows more gradual scaling with population size, while coarse discrimination exhibits steeper performance gains and higher asymptotic values. Both analyses demonstrate that orientation information scales systematically with population size, with discrimination difficulty appropriately reflected in performance curves.

Figure 2.28 provides detailed performance curves for fine discrimination across all four decoders with error bars showing standard error of the mean (SEM) across sessions. Performance scales systematically with population size, while the relationship with orientation difference within the fine range shows a counterintuitive pattern: discriminability appears to decrease slightly from 2° to 4° to 8° differences. However, the overlapping error bars indicate that these differences are not statistically significant, suggesting that fine orientation discrimination operates near the fundamental signal-to-noise limits where small orientation differences produce similar discrimination challenges regardless of decoding strategy.

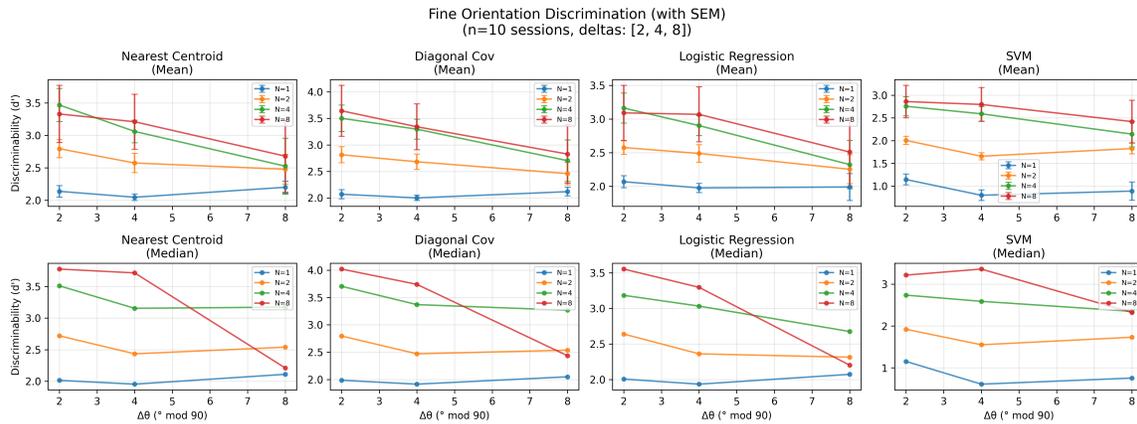


Figure 2.28: **Fine orientation discrimination performance across decoders with error bars.** Performance curves for fine orientation discrimination (2°, 4°, 8° differences) across four decoding methods. Each subplot shows mean and median d' values as a function of population size ($N=1, 2, 4, 8$ neurons) for a single decoder, with error bars representing standard error of the mean (SEM) for mean values. Different colors represent different orientation differences within the fine discrimination range. Performance scales systematically with population size, while orientation differences within the fine range show minimal and non-significant variation, with performance appearing to decrease slightly from 2° to 8° differences despite the counterintuitive nature of this pattern. Fine discrimination was assessable in 10 sessions (all from M36), representing conditions with dense orientation sampling. Peak performance levels ($d' \approx 4-6$) demonstrate that even small orientation differences can be reliably discriminated by appropriately sized V1 populations.

Figure 2.29 shows the corresponding analysis for coarse discrimination with er-

ror bars, which was assessable across all 16 sessions. In contrast to fine discrimination, coarse discrimination shows robust performance scaling across the full range of population sizes ($N=1-128$), with clear separation between different orientation differences. The monotonic relationship between orientation difference and discriminability in the coarse range confirms that V1 populations encode orientation information in a graded manner that is accessible to linear decoding methods when orientation differences are sufficiently large.

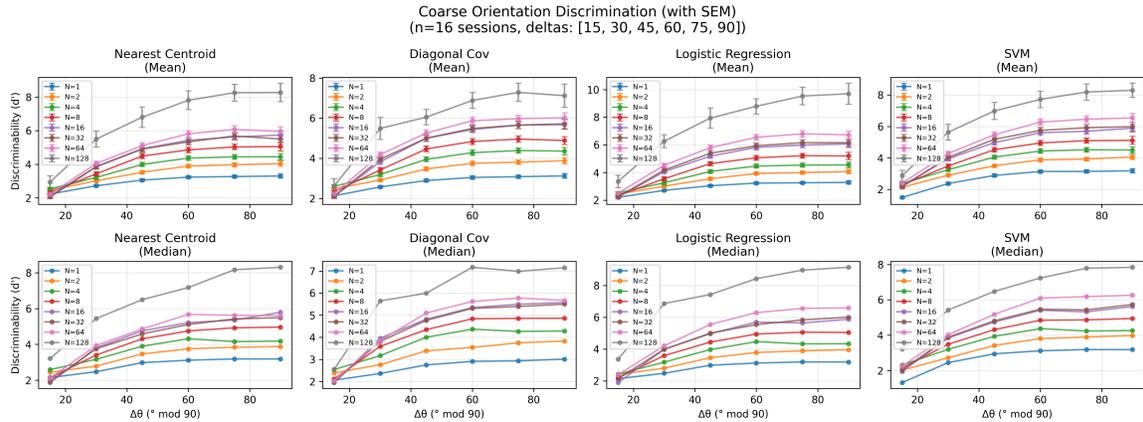


Figure 2.29: **Coarse orientation discrimination performance across decoders with error bars.** Performance curves for coarse orientation discrimination ($15^\circ-90^\circ$ differences) across four decoding methods. Each subplot shows mean and median d' values as a function of population size ($N=1, 2, 4, 8, 16, 32, 64, 128$ neurons) for a single decoder, with error bars representing standard error of the mean (SEM) for mean values. Different colors represent different orientation differences within the coarse discrimination range. Performance shows strong scaling with both population size and orientation difference, with 90° differences (orthogonal orientations) achieving the highest discriminability values. Coarse discrimination was assessable across all 16 sessions with 7,062 total pairs. High peak performance levels ($d' > 10$) demonstrate the substantial orientation information available in V1 populations for behaviorally relevant discrimination tasks.

2.9.3 Impact of Noise Correlations on Population Decoding

To isolate the contribution of noise correlations to decoding performance, we compared empirical results with shuffle controls where trial orders were independently

permuted for each neuron (20 iterations per analysis). For each shuffle iteration, decoders were retrained on the shuffled data using the same cross-validation procedure as the empirical analysis. This procedure preserves single-neuron response statistics while eliminating pairwise correlations, allowing fair comparison between correlated and decorrelated population responses.

Figure 2.30 shows that empirical performance was consistently lower than shuffled performance for populations of 4 or more neurons across all decoders. For 45° orientation differences, empirical performance reached asymptotes of $d' \approx 6-8$, while shuffled controls achieved $d' \approx 8-12$. For 90° differences, the gap was even larger, with empirical performance plateauing at $d' \approx 7-8$ versus $d' \approx 12-16$ for shuffled data.

Saturating exponential fits ($d'(N) = a[1 - \exp(-bN)] + c$) to both empirical and shuffled data revealed consistently higher asymptotic values (a) and faster saturation rates (b) for shuffled controls. Fit parameters displayed in each panel demonstrate that noise correlations impose a systematic penalty on orientation discrimination, with the magnitude of this effect increasing with population size.

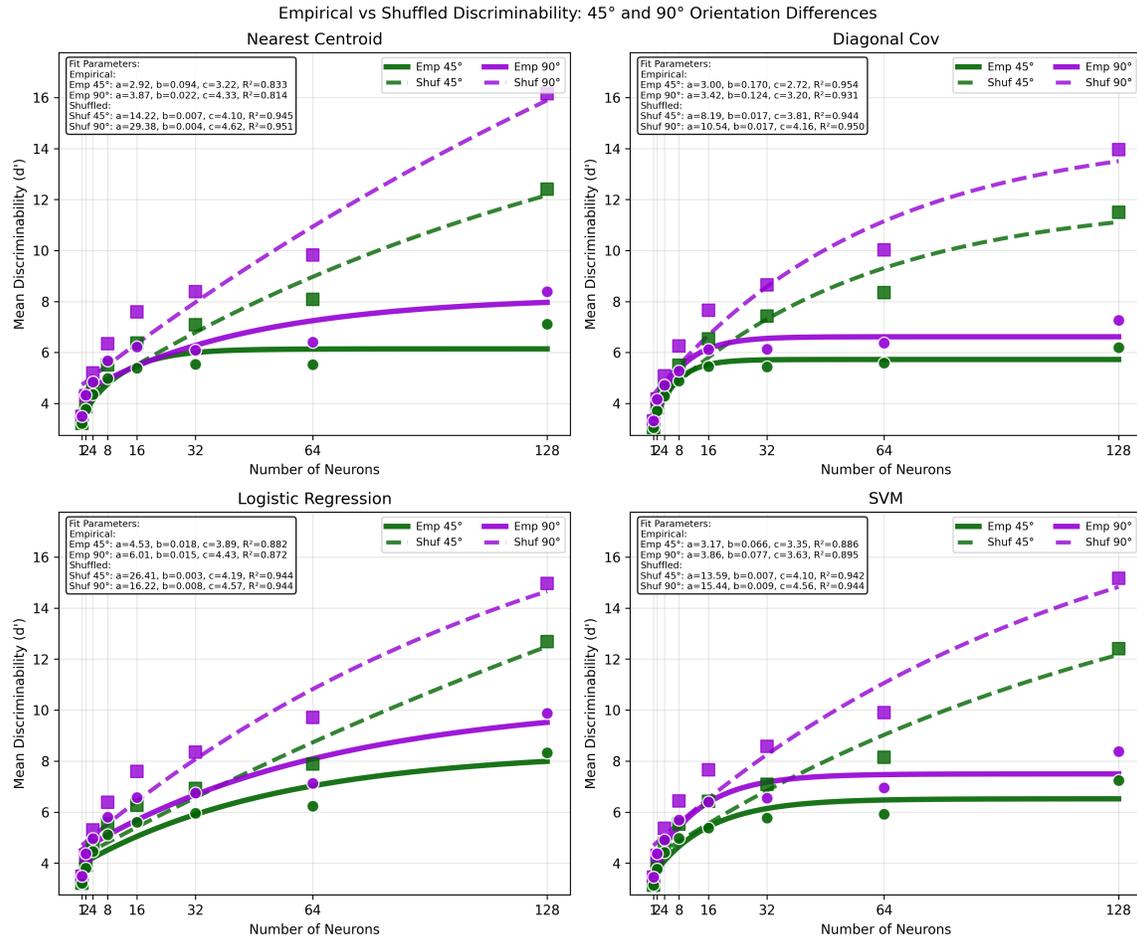


Figure 2.30: **Impact of noise correlations on population decoding performance.** Comparison between empirical decoding performance and shuffle controls for 45° and 90° orientation differences. Each subplot shows one decoder (Nearest Centroid, Diagonal Covariance, Logistic Regression, SVM). Solid lines represent empirical data; dashed lines represent shuffled controls where trial correlations are eliminated. Green indicates 45° differences; purple indicates 90° differences. Circles show empirical data points; squares show shuffled data points. Text boxes display saturating exponential fit parameters for empirical curves: $d'(N) = a[1 - \exp(-bN)] + c$. Shuffled performance consistently exceeds empirical performance for $N \geq 4$, suggesting that noise correlations may influence population decoding efficiency. The performance gap increases with both population size and orientation difference, which could suggest that correlations have the greatest effect on the most discriminable stimulus pairs.

Chapter 3: Methods

3.1 Experimental Subjects and Surgical Procedures

All experiments were conducted on two adult male rhesus macaques (*Macaca mulatta*). All procedures were approved by the University of Texas at Austin Institutional Animal Care and Use Committee and were conducted in accordance with National Institutes of Health guidelines for the care and use of laboratory animals.

3.2 Viral Expression of GCaMP6f

Genetically encoded calcium indicators were expressed in layer 2/3 excitatory neurons using AAV1-CaMKII-GCaMP6f (obtained from Karl Deisseroth’s laboratory, Stanford University). The CaMKII promoter ensured selective expression in excitatory pyramidal neurons while minimizing expression in inhibitory interneurons.

Viral injections followed procedures optimized for long-term expression in macaque cortex (Seidemann et al.). AAV1-CaMKII-GCaMP6f was injected at multiple depths (150 μm and 250 μm below the pial surface) using a precision microinjection system. A total of 1-2 μL was delivered per site.

3.3 Two-Photon Calcium Imaging

3.3.1 Imaging System

Two-photon calcium imaging was performed using a custom B-Scope microscope (Thorlabs) equipped with a Mai Tai DeepSee femtosecond laser (Spectra Physics) tuned to 940 nm for GCaMP6f excitation. The system utilized galvanometric scanning mirrors for rapid beam positioning and a 16 \times water immersion objective (Nikon, 0.8 NA) for high-resolution imaging.

Images were acquired at 30 Hz with 512 \times 512 pixel resolution over fields of

view spanning 600-900 μm . This allowed simultaneous recording from populations of 20-200 neurons per session. Imaging depth was maintained at 150-250 μm below the pial surface, targeting layer 2/3 pyramidal neurons.

3.3.2 Motion Correction and Image Processing

Raw imaging data were processed using the Suite2p software package (Pachitariu et al.). Motion correction was performed using both rigid and non-rigid algorithms to compensate for brain movement and cardiac pulsation artifacts. A reference frame was computed from periods of minimal motion, and all frames were registered to this template.

Regions of interest (ROIs) corresponding to individual neurons were automatically detected using Suite2p's cell detection algorithm. Putative ROIs were manually curated based on morphological criteria and fluorescence time course characteristics to remove artifacts and non-neuronal structures.

Neuropil subtraction was not performed, as pilot analyses indicated that structured, stimulus-dependent signals in layer 2/3 neuropil could represent meaningful network activity rather than background contamination.

3.4 Visual Stimulation

3.4.1 Stimulus Generation and Presentation

Visual stimuli were generated using custom MATLAB software and presented on a gamma-corrected CRT monitor (refresh rate: 100 Hz) positioned 36 cm from the animal's eyes. Stimuli consisted of full-field oriented sinusoidal gratings with 100% contrast and 2 cycles/degree spatial frequency.

The standard stimulus protocol included 12 orientations (0° to 165° in 15° increments) and was used across all 16 recording sessions. Additionally, a subset of 10 sessions (all from animal M36) included fine-spaced orientation conditions with ad-

ditional orientations at 2°, 4°, and 8° intervals around selected reference orientations. These fine-spaced conditions were designed to enable high-resolution discrimination analysis for small orientation differences.

Each grating was presented for 2 seconds in a cyclic pattern: 200 ms ON followed by 300 ms OFF, repeated 4 times per trial. This temporal structure allowed extraction of both transient and sustained components of neural responses. Blank trials (0% contrast, mean luminance) were randomly interleaved (1 in 6 trials) to characterize spontaneous activity. Most analysis only considered the averaged activity from the first cycle.

3.4.2 Behavioral Control

Animals were trained to maintain fixation on a 0.2° white target at screen center throughout stimulus presentation. Eye position was monitored at 1000 Hz using an infrared eye-tracking system (EyeLink 1000, SR Research). Trials with fixation breaks exceeding 0.5° were aborted and excluded from analysis.

Successful fixation was rewarded with small volumes of juice or water. Training sessions continued for 4-6 weeks prior to data collection to ensure stable fixation performance during imaging.

3.5 Data Analysis

3.5.1 Fluorescence Signal Processing

Raw fluorescence traces were processed to extract stimulus-evoked responses. For each neuron and trial, the baseline fluorescence F_0 was estimated from the 400 ms pre-stimulus period, and responses were normalized as:

$$\Delta F/F(t) = \frac{F(t) - F_0}{F_0} \quad (3.1)$$

To reduce high-frequency noise while preserving temporal dynamics, signals

were smoothed using a causal two-point filter:

$$F_{\text{smooth}}(t) = \frac{F_{\text{raw}}(t-1) + F_{\text{raw}}(t)}{2} \quad (3.2)$$

Stimulus-evoked responses were quantified by integrating $\Delta F/F$ over a 300 ms window during each stimulus cycle, yielding integrated response amplitudes $A_{i,n,c}$ for neuron i , trial n , and cycle c .

3.5.2 Orientation Tuning Analysis

Orientation tuning curves were characterized by fitting circular Gaussian functions to trial-averaged responses:

$$F_i(\theta) = b_i + a_i \exp \left[-\frac{1}{2} \left(\frac{d_{\text{circ}}(\theta, \theta_i^*)}{\sigma_i} \right)^2 \right] \quad (3.3)$$

where b_i is baseline activity, a_i is response amplitude, θ_i^* is preferred orientation, σ_i is tuning bandwidth, and $d_{\text{circ}}(\theta, \theta_i^*)$ is the circular distance on $[0^\circ, 180^\circ]$.

The normalized tuning amplitude (NTA) was calculated as:

$$\text{NTA}_i = \frac{a_i}{a_i + b_i} \quad (3.4)$$

representing the modulation depth relative to peak response.

Circular variance (CV) was calculated for each orientation-selective neuron following Ringach et al.:

$$\text{CV}_i = 1 - \left| \frac{1}{N} \sum_{j=1}^N e^{2i\theta_j} \right| \quad (3.5)$$

where θ_j are the orientation values and N is the number of orientations. CV ranges from 0 (perfectly tuned) to 1 (broadly tuned), providing a complementary measure to tuning bandwidth.

3.5.3 Functional Classification

Neurons were classified into three categories based on stimulus responses:

Orientation-selective: Neurons whose fitted amplitude a_i exceeded the 95th percentile of a null distribution generated by refitting tuning curves to 1000 permutations of shuffled orientation labels.

Responsive but orientation-invariant: Neurons that failed the orientation selectivity test but showed significant stimulus-driven responses compared to blank trials (assessed using a separate permutation test on baseline vs. blank activity).

Non-responsive: Neurons that passed neither test.

3.5.4 Noise Correlation Analysis

Pairwise noise correlations were computed between orientation-selective neurons to quantify shared trial-to-trial variability. For each neuron, responses were computed as time-averaged $\Delta F/F$ values within the stimulus presentation window (or entire trial for spontaneous conditions).

Noise correlations were calculated separately for two conditions:

Spontaneous correlations: Computed from blank (0% contrast) trials where no visual stimulus was presented. Trial-averaged responses were computed for each neuron, and Pearson correlations were calculated between all neuron pairs across the set of blank trials.

Evoked correlations: Computed as true noise correlations by removing stimulus-locked signal components. For each orientation condition, the mean response across all trials of that condition was computed for each neuron. This mean was then subtracted from individual trial responses to obtain residuals that reflect trial-to-trial variability around the condition-specific mean. Residuals from all orientation conditions were concatenated, and Pearson correlations were computed between neuron pairs across these residuals.

This approach quantifies shared variability in trial-to-trial fluctuations, with spontaneous correlations reflecting intrinsic cortical dynamics and evoked correlations reflecting shared noise (residual variability after removing stimulus-locked responses) during active sensory processing. By removing per-condition means, evoked correlations capture genuine noise correlations rather than signal correlations that would arise from shared stimulus preferences.

3.5.5 Response Variability Analysis

Trial-to-trial response variability was characterized by analyzing the variance-mean relationship for orientation-selective neurons. This analysis separates stimulus-driven variability from baseline noise by accounting for intrinsic variability present even in the absence of visual stimulation.

For each orientation-selective neuron, we first computed the blank trial standard deviation σ_{blank} for z-score normalization:

$$\sigma_{\text{blank}} = \text{std}(\{\Delta F/F_{\text{blank}}\}) \quad (3.6)$$

where $\{\Delta F/F_{\text{blank}}\}$ represents the set of blank trial responses for that neuron. All neural responses were then z-scored using this normalization factor to provide standardized units for cross-neuron comparison.

For each stimulus condition, we calculated the mean response μ and variance σ^2 across trials using the z-scored data. The variance-mean relationship was then characterized using a linear model:

$$\sigma^2 = \text{slope} \cdot \mu + \sigma_0^2 \quad (3.7)$$

This formulation assumes that total response variance consists of two components: stimulus-driven variability that scales with the mean response ($\text{slope} \cdot \mu$) and baseline noise that is independent of stimulus strength (σ_0^2).

Both parameters were estimated using least-squares regression:

$$\text{slope}, \sigma_0^2 = \arg \min_{\text{slope}, \sigma_0^2} \sum_{\text{conditions}} (\sigma_{\text{condition}}^2 - (\text{slope} \cdot \mu_{\text{condition}} + \sigma_0^2))^2 \quad (3.8)$$

subject to the constraints that both $\text{slope} \geq 0$ and $\sigma_0^2 \geq 0$ to ensure physiologically meaningful parameter estimates.

This formulation distinguishes two distinct roles:

- σ_{blank} : Used solely for z-score normalization to enable comparison across neurons with different baseline variability levels
- σ_0^2 : The fitted baseline variance parameter that characterizes the intercept of the variance-mean relationship during visual stimulation, which may differ from spontaneous fluctuations measured in blank trials

To understand the relationship between response variability and neural selectivity, we analyzed correlations between variance-mean slopes and orientation tuning parameters.

Individual neuron variance-mean relationships were visualized as scatter plots with fitted linear models, including a reference point at $(\mu = 0, \sigma^2 = 1)$ for illustration purposes. These plots show each stimulus condition as a colored point representing different orientations, with the fitted linear relationship overlaid in red.

To understand the relationship between response variability and neural selectivity, we analyzed correlations between variance-mean slopes and orientation tuning parameters extracted from circular Gaussian fits:

- Tuning bandwidth (σ_θ): Width of the orientation tuning curve
- Response amplitude (a): Peak response magnitude above baseline

- Baseline activity (b): Response level at the orthogonal orientation
- Normalized tuning amplitude (NTA): Modulation depth relative to peak response
- Circular variance (CV): Measure of tuning sharpness
- Preferred orientation (θ^*): Orientation evoking maximum response

Statistical relationships were assessed using Pearson correlation coefficients and their significance levels, providing insight into the mechanisms underlying response reliability in V1.

3.5.6 Statistical Analysis

Statistical comparisons used appropriate parametric and non-parametric tests depending on data distributions. For correlation analyses, Fisher z -transformations were applied to stabilize variance and approximate normality. Multiple comparisons were corrected using false discovery rate (FDR) control where appropriate.

Cross-session consistency was assessed by examining the distribution of effect sizes and statistical significance across individual sessions. Population-level inferences were based on random-effects models accounting for between-session variability.

3.5.7 Orientation Dependence Analysis

To test whether noise correlations depended on orientation preference similarity, neuron pairs were binned by their preferred orientation difference:

$$\Delta\theta_{\text{pref}} = \min(|\theta_i^* - \theta_j^*|, 180^\circ - |\theta_i^* - \theta_j^*|) \quad (3.9)$$

mapped to $[0^\circ, 90^\circ]$ and binned in 15° intervals ($0^\circ - 15^\circ$, $15^\circ - 30^\circ$, $30^\circ - 45^\circ$, $45^\circ - 60^\circ$, $60^\circ - 75^\circ$, $75^\circ - 90^\circ$). For each bin, mean correlations were calculated across all neuron pairs falling within that orientation difference range.

Statistical significance of orientation-dependent correlation structure was assessed using permutation tests where orientation labels were shuffled 1000 times to generate null distributions for each angular bin. Animal-averaged correlation profiles were computed by first calculating session-wise means within each bin, then averaging across sessions for each animal.

All analyses were implemented in Python using NumPy, SciPy, and custom analysis libraries. Figures were generated using Matplotlib and Seaborn visualization packages.

3.5.8 Spatial Organization Analysis

To quantify the spatial organization of orientation preferences we examined the relationship between physical proximity and functional similarity using pairwise distance and orientation difference measurements.

3.5.8.1 Spatial Data Preparation

For each recording session, we extracted the spatial coordinates of all orientation-selective neurons from Suite2p segmentation results. Cell centroids were computed in pixel coordinates and converted to micrometers using session-specific pixel size calibrations. Only neurons classified as orientation-selective were included in the spatial analysis.

Pairwise Euclidean distances were computed between all neuron pairs within each session:

$$d_{ij} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2} \quad (3.10)$$

where (x_i, y_i) and (x_j, y_j) are the spatial coordinates of neurons i and j .

Circular orientation differences were calculated as:

$$\Delta\theta_{ij} = \min(|\theta_i^* - \theta_j^*|, 180^\circ - |\theta_i^* - \theta_j^*|) \quad (3.11)$$

where θ_i^* and θ_j^* are the preferred orientations of neurons i and j , mapped to the range $[0^\circ, 90^\circ]$ to account for the circular nature of orientation.

3.5.8.2 Spatial Autocorrelation Analysis

To detect periodic organization characteristic of orientation hypercolumns, we computed spatial autocorrelation functions of orientation similarity. Orientation similarity was defined as:

$$S_{ij} = 1 - \frac{\Delta\theta_{ij}}{90^\circ} \quad (3.12)$$

mapping orientation differences to a similarity scale where 1 indicates identical orientations and -1 indicates orthogonal orientations.

Autocorrelation was computed by binning neuron pairs by distance and calculating mean similarity within each bin:

$$R(d) = \langle S_{ij} \rangle \text{ for } |d_{ij} - d| < \Delta d/2 \quad (3.13)$$

where Δd is the bin width (typically 10-20 μm).

3.5.8.3 Two-Gaussian Hypercolumn Model

To quantify hypercolumn organization, we fitted a two-Gaussian model to the spatial autocorrelation function:

$$R(d) = C + A_1 e^{-d^2/(2\sigma_1^2)} + A_2 e^{-(d-\mu_2)^2/(2\sigma_2^2)} \quad (3.14)$$

where:

- C is the baseline correlation
- A_1 and σ_1 characterize local clustering strength and width
- A_2 , μ_2 , and σ_2 characterize the hypercolumn peak amplitude, spacing, and width

The parameter μ_2 provides an estimate of hypercolumn spacing. Model fitting was performed using nonlinear least-squares optimization with robust parameter initialization based on data characteristics.

3.5.8.4 Cross-Session Integration

Spatial measures were computed separately for each session and then combined across sessions for population-level analysis. Only relative spatial measures (distances and orientation differences) were pooled across sessions to avoid confounding effects from different field sizes, imaging depths, or coordinate systems.

Session-level results were weighted by the number of neuron pairs to account for different sample sizes. Statistical significance of spatial organization was assessed using permutation tests where orientation labels were randomly shuffled within each session while preserving spatial coordinates.

3.5.9 Population Decoding Analysis

To assess the information content of V1 neural populations about stimulus orientation, we implemented a decoding analysis using projection-based signal detection theory. For each session, neurons were first ranked by their single-neuron cross-validated discriminability for specific orientation pairs, and population performance was evaluated for increasing population sizes ($n = 1, 2, 4, 8, 16, 32, 64, 128$).

3.5.9.1 Single-Neuron Discriminability

Single-neuron orientation discriminability was computed using a projection-based cross-validated d' approach. For each leave-one-out split, we:

1. Computed training set statistics (means $\mu_{A,train}$, $\mu_{B,train}$ and pooled standard deviation σ_{train})
2. Calculated the training discriminant: $d_{train} = (\mu_{A,train} - \mu_{B,train})/\sigma_{train}$
3. Projected the held-out test response onto the discriminant axis relative to the training midpoint
4. Accumulated projection scores separately for each stimulus class
5. Computed final d' from the distributions of projection scores

This approach provides more robust estimates than traditional confusion matrix methods by accounting for the uncertainty in discriminant estimation.

3.5.9.2 Population Decoders

We compared four population decoding models:

Nearest Centroid: Assumes independence between neurons with identity covariance structure. The discriminant axis is simply the difference in population means: $\mathbf{w}_{NC} = \boldsymbol{\mu}_A - \boldsymbol{\mu}_B$.

Diagonal Covariance: Uses inverse-variance weighting to account for heterogeneous response variabilities: $\mathbf{w}_{diag} = (\boldsymbol{\mu}_A - \boldsymbol{\mu}_B) \oslash \boldsymbol{\sigma}_{pooled}^2$, where \oslash denotes element-wise division.

Logistic Regression: Uses L2-regularized maximum likelihood estimation with regularization parameter $C = 1.0$.

Support Vector Machine (SVM): Linear SVM with regularization parameter $C = 1.0$.

3.5.9.3 Cross-Validation and Performance Metrics

All decoders used leave-one-trial-out cross-validation with the projection-based d' calculation. For each cross-validation fold:

1. Train the decoder on $n - 1$ trials to obtain discriminant weights \mathbf{w}
2. Calculate training separation: $d_{train} = \sqrt{(\boldsymbol{\mu}_{A,train} - \boldsymbol{\mu}_{B,train})^T \mathbf{M} (\boldsymbol{\mu}_{A,train} - \boldsymbol{\mu}_{B,train})}$
3. Project test trial: $s = d_{train} \cdot \frac{(\mathbf{x}_{test} - \mathbf{m})^T \mathbf{w}}{\|\mathbf{w}\|}$
4. Accumulate scores by stimulus class

where \mathbf{M} is the metric tensor (identity for Nearest Centroid and diagonal for Diagonal Covariance), and \mathbf{m} is the training midpoint.

Final discriminability was quantified as:

$$d' = \frac{|\bar{s}_A - \bar{s}_B|}{\sqrt{(\sigma_{s_A}^2 + \sigma_{s_B}^2)/2}} \quad (3.15)$$

where \bar{s}_A , \bar{s}_B are mean projection scores and $\sigma_{s_A}^2$, $\sigma_{s_B}^2$ are score variances for each stimulus class.

Fine versus coarse discrimination analysis. To characterize different aspects of orientation discrimination, we categorized orientation pairs into two groups based on angular separation and stimulus availability: fine discrimination ($\Delta\theta = 2^\circ, 4^\circ, 8^\circ$) and coarse discrimination ($\Delta\theta = 15^\circ, 30^\circ, 45^\circ, 60^\circ, 75^\circ, 90^\circ$).

Fine discrimination analysis was only possible in the 10 sessions that included fine-spaced orientation conditions (all from animal M36), as these sessions contained the 2° , 4° , and 8° orientation differences required for this analysis. For fine discrimination, population analysis was limited to smaller group sizes ($N \leq 8$) due to the reduced number of available pairs.

Coarse discrimination analysis utilized the standard 15° equidistant orientation conditions that were present in all 16 recording sessions, enabling discrimination analysis across the full range of population sizes and providing coverage across both animals.

Across-pair and across-session aggregation. Decoder performance was summarised by computing $\Delta\theta = \min(|\theta_A - \theta_B|, 180^\circ - |\theta_A - \theta_B|)$ for each orientation pair and binning values according to the fine/coarse categorization above. Within each bin and population size ($N = 1, 2, 4, 8, 16, 32, 64, 128$), we computed mean, median, and maximum d' across all pairs and sessions. Results for fine and coarse discrimination are presented separately in Figure 2.28 and Figure 2.29, with direct comparison shown in Figure 2.27.

Empirical-versus-shuffled comparison. For each decoder and N , we averaged d' over all orientation pairs to yield a single empirical value and its shuffled counterpart (20 independent trial-shuffles). Mean values were plotted versus N , and a saturating exponential,

$$d'(N) = a[1 - \exp(-bN)] + c,$$

was fit by non-linear least squares to empirical and shuffled series separately (Figure 2.30); best-fit (a, b, c) and R^2 are reported in panel insets. Common y -axis limits were set to encompass the union of empirical and shuffled ranges across decoders.

3.5.9.4 Correlation Impact Assessment

To isolate the contribution of noise correlations to population decoding, we compared empirical performance with shuffled controls. For shuffled data, trial orders were randomly permuted independently for each neuron, preserving marginal response distributions while eliminating correlations. Shuffle analysis used 20 iterations for each decoder and group size combination.

3.5.10 Cross-Session and Cross-Animal Analysis

To assess consistency of findings across recording sessions and animals, we implemented hierarchical analysis approaches. For each measured parameter or correlation, we first computed session-specific statistics, then examined the distribution of these values across sessions and animals.

Cross-session consistency was quantified using several metrics: (1) the percentage of sessions showing statistically significant effects, (2) effect size distributions across sessions, and (3) variance decomposition between within-session and between-session components.

Animal-level comparisons (M22 vs M36) were performed using both parametric (t-tests, ANOVA) and non-parametric (Mann-Whitney U, Wilcoxon signed-rank) statistical tests, depending on data distributions. For orientation preference analysis, circular statistical methods were employed, including the Rayleigh test for uniformity and calculation of vector strength (orientation bias).

Quality control metrics were tracked across all sessions, including the proportion of neurons classified as orientation-selective, responsive but orientation-invariant, and non-responsive. These proportions were compared across animals and sessions to ensure consistent data quality.

Chapter 4: Discussion

4.1 Summary of Main Findings

This thesis presents a comprehensive characterization of orientation coding and noise correlations in macaque primary visual cortex using large-scale two-photon calcium imaging. Our analysis of 1,340 orientation-selective neurons across 16 recording sessions revealed several key findings that advance our understanding of population-level visual processing.

First, we demonstrated that orientation-selective neurons in layer 2/3 of macaque V1 exhibit narrower tuning bandwidths (median $\sigma = 17.58^\circ$, mean = 17.3°) than previously reported using electrophysiology, with strong selectivity as measured by normalized tuning amplitude (mean = 83.2%). These neurons showed systematic relationships between tuning parameters, most notably strong negative correlations between baseline activity and selectivity measures.

Second, we found that noise correlations are ubiquitous in V1 populations, with significant pairwise correlations present during both spontaneous (mean $r = 0.072$) and evoked (mean $r = 0.097$) conditions. Visual stimulation significantly enhanced correlation magnitude, with a meaningful effect size (Cohen's $d = 0.089$).

Third, and most importantly, we discovered that the structure of noise correlations is dramatically modulated by sensory input. During spontaneous activity, correlations were largely independent of orientation preference similarity. Visual stimulation substantially restructured this relationship, suggesting that sensory drive drives correlation patterns organized around functional similarity.

4.2 Comparison with Previous Studies

4.2.1 Single-Neuron Tuning Properties

Our measurements of orientation tuning bandwidth are notably sharper than those reported in many previous electrophysiological studies of macaque V1. The median tuning width of $\sigma = 17.58^\circ$ (mean = 17.3°) suggests enhanced selectivity in superficial layer 2/3 populations, likely reflecting the selective targeting of excitatory pyramidal neurons via our calcium imaging approach and the nonlinear response properties of calcium indicators that emphasize sharply tuned responses while suppressing weaker fluctuations (Carandini and Ferster; Priebe and Ferster).

4.2.2 Noise Correlations and Functional Structure

Pairwise noise correlations in layer 2/3 of macaque V1 averaged $r \approx 0.072$ during spontaneous periods and $r \approx 0.097$ during visual stimulation, consistent with the 0.05–0.15 range commonly reported for nearby cortical neurons (Kohn and Smith; Smith and Kohn).

During spontaneous activity, noise correlations exhibited no systematic dependence on orientation preference similarity—the relationship was flat across the population. Visual stimulation produced a modest but detectable restructuring: correlations showed a gradual decline from similarly tuned pairs (0.15) to orthogonal pairs (0.08). However, this functional organization was largely abolished when controlling for spatial distance, suggesting it primarily reflects local circuit connectivity rather than long-range functional relationships.

4.3 Implications for Visual Processing and Coding Fidelity

4.3.1 Noise Correlations and Population Coding Constraints

Our population decoding analysis reveals potential constraints on information processing in V1 that relate directly to the structure of noise correlations. A key finding is the difference between population decoding performance in real versus

shuffled data: while shuffled populations (with correlations removed) show continued improvement in discriminability with increasing population size, real populations showed more limited gains beyond 16-32 neurons.

In the shuffled condition, where trial-to-trial correlations are eliminated while preserving individual neuron response distributions, population discriminability (d') continues to improve as more neurons are added to the decoder, reaching d' values of 4-6 for large populations. In contrast, real populations showed more modest improvements at larger population sizes (typically plateauing at d' values of 2-3), with limited gains beyond 32 neurons.

The orientation-dependent structure of noise correlations we observed provides insight into the potential mechanisms underlying this apparent information constraint. During visual stimulation, correlations increase among similarly tuned neurons—those neurons that carry similar information about stimulus orientation. This pattern resembles what has been termed "information-limiting correlations" (Moreno-Bote et al.): shared variability that is aligned with the signal dimensions that decoders attempt to read out.

When similarly tuned neurons fluctuate together, their combined response becomes less reliable than would be expected from independent noise sources. A decoder reading from neurons with similar orientation preferences and positively correlated noise may perform worse than a decoder reading from independent neurons with the same individual response properties, potentially reducing the "independent samples" available to the decoder and limiting the benefits of population pooling. Conversely, correlations between dissimilarly tuned neurons are less problematic for coding because these neurons carry different information about the stimulus, and shared variability between neurons tuned to orthogonal orientations does not directly interfere with orientation discrimination tasks.

4.3.2 Coding Regime Transitions

Our data suggest two distinct coding regimes with different information-processing implications:

Spontaneous state: Correlations are weakly positive ($r \approx 0.072$) and unstructured across orientation preferences. This flat correlation profile suggests global, untuned sources of shared variability that affect all neurons similarly regardless of their functional properties.

Stimulated state: Visual input reshapes correlations into a structured pattern where similarly tuned neurons become more correlated while dissimilar neurons decorrelate. This restructuring has competing effects on coding efficiency: it increases information-limiting correlations among functionally similar neurons (reducing coding fidelity) while decreasing irrelevant correlations between functionally distinct neurons.

The net effect, as revealed by our decoding analysis, is a substantial limitation on population coding performance. While some theoretical work suggests that structured correlations might benefit certain coding tasks (Kanitscheider et al.), our results demonstrate that the particular correlation structure in V1 creates clear information bottlenecks for orientation discrimination under the models we used.

4.3.3 Implications for Sensory Processing

These findings have important implications for understanding sensory processing limitations in natural vision. If V1 populations saturate at relatively small sizes due to noise correlations, the cortex cannot simply overcome noise limitations by recruiting more neurons to a given computation. Instead, efficient sensory processing may require:

1. **Correlation management:** Circuit mechanisms that actively decorrelate neurons carrying similar information, potentially through inhibitory networks or attention-related modulation.

2. **Distributed processing:** Strategies that distribute different aspects of sensory analysis across populations with distinct correlation structures, rather than relying on simple pooling within single areas.
3. **Hierarchical refinement:** Processing stages where information is iteratively refined and corrupted correlations are reduced through recurrent processing or cross-area interactions.

The saturation of coding performance also suggests that reports of "infinite" information capacity in large neural populations may not apply to real cortical circuits, where correlations impose fundamental constraints on information extraction efficiency.

4.4 Technical Considerations and Limitations

While two-photon calcium imaging enabled simultaneous recording from large neuronal populations, several technical limitations must be considered. The temporal dynamics of calcium indicators introduce filtering that may affect correlation measurements, and our sampling was restricted to layer 2/3 excitatory neurons. The relationship between calcium signals and spike rates is nonlinear and can vary across neurons (Huang et al.), potentially affecting quantitative interpretations of correlation magnitudes.

Our recordings sampled discrete cortical locations in two animals, potentially limiting generalizability. Future studies with larger numbers of animals and simultaneous recordings from multiple cortical layers would provide important insights into the laminar organization of correlation structure.

4.5 Future Directions

An important technical direction involves developing optimal decoders that can efficiently extract orientation information despite correlated neural activity. Our

analysis using standard linear decoders revealed systematic performance limitations imposed by noise correlations, but more sophisticated decoding strategies could potentially overcome these constraints. Previous work has demonstrated that spatially antagonistic center-surround summation can substantially outperform simple pooling mechanisms when neural responses exhibit strong spatial correlations (Chen et al., b). Future studies should explore whether similar inhibitory decoding mechanisms could improve orientation discrimination from V1 population responses, potentially revealing more efficient readout strategies that mirror biological processing in downstream cortical areas.

A critical extension of this work involves connecting V1 noise correlations to behavioral variability in perceptual decision-making. Recent studies have demonstrated that V1 population responses correlate more strongly with perceptual reports than with the presented stimulus, suggesting that much of the neural variability underlying perceptual decisions originates at this early cortical stage (Morais et al.). Future experiments should examine whether the stimulus-dependent correlation structure we observed contributes to trial-to-trial variability in orientation discrimination performance.

A particularly promising direction involves the integration of optogenetic tools with calcium imaging in macaque models. The combination of optogenetic stimulation with population-scale calcium imaging would enable causal tests of how specific circuit elements contribute to complex visual behaviors, including fine discrimination tasks, attention allocation, and decision-making under uncertainty.

4.6 Conclusions

We demonstrate that noise correlations in macaque V1 create fundamental information bottlenecks that limit population coding fidelity. While individual neurons continue to carry information that could theoretically improve population performance, structured correlations prevent efficient extraction of this information through

simple pooling mechanisms. The stimulus-dependent reorganization of correlation structure—where similarly tuned neurons become more correlated during visual processing—directly contributes to these coding limitations. These findings challenge assumptions about unlimited improvements from larger neural populations and highlight the critical importance of correlation structure in determining the efficiency of cortical computations.

Chapter 5: Conclusion

The work presented in this thesis provides new insights into the population-level organization of orientation processing in macaque primary visual cortex. Through large-scale two-photon calcium imaging, we have characterized both individual neuron tuning properties and the complex patterns of shared variability that emerge at the population level.

Our key contributions include:

1. **Characterization of orientation tuning:** We documented the distributions and relationships of key tuning parameters across a large population of V1 neurons, revealing sharper tuning and stronger selectivity than many previous electrophysiological studies.
2. **Discovery of stimulus-dependent correlation structure:** We demonstrated that noise correlations are unstructured during spontaneous activity but become functionally organized during visual stimulation, with increased correlations among similarly tuned neurons and decreased correlations among dissimilarly tuned neurons.
3. **Fundamental coding limitations imposed by noise correlations:** Our population decoding analysis revealed that structured noise correlations create information bottlenecks that limit coding fidelity. While shuffled populations with independent noise show continued improvement in discriminability with increasing population size, real populations saturate at 16-32 neurons due to correlated noise among similarly tuned neurons.

These results advance our understanding of how neural populations in visual cortex balance competing demands for sensitivity and efficiency. The stimulus-

dependent modulation of correlation structure represents a previously uncharacterized mechanism that may be fundamental to cortical computation across sensory systems.

Future research integrating optogenetic manipulation with population imaging in macaque models will help test whether these coding limitations translate to measurable changes in perceptual performance and for understanding how cortical circuits might overcome correlation-imposed constraints.

Works Cited

L. F. Abbott and Peter Dayan. The Effect of Correlated Variability on the Accuracy of a Population Code. 11(1):91–101. ISSN 0899-7667, 1530-888X. doi: 10.1162/089976699300016827. URL <https://direct.mit.edu/neco/article/11/1/91-101/6229>.

Bruno B. Averbeck, Peter E. Latham, and Alexandre Pouget. Neural correlations, population coding and computation. 7(5):358–366. ISSN 1471-0048. doi: 10.1038/nrn1888. URL <https://www.nature.com/articles/nrn1888>.

Matteo Carandini and David Ferster. Membrane Potential and Firing Rate in Cat Primary Visual Cortex. 20(1):470–484. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.20-01-00470.2000. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6774139/>.

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, Loren L. Looger, Karel Svoboda, and Douglas S. Kim. Ultrasensitive fluorescent proteins for imaging neuronal activity. 499(7458):295–300, a. ISSN 1476-4687. doi: 10.1038/nature12354. URL <https://www.nature.com/articles/nature12354>.

Yuzhi Chen, Wilson S Geisler, and Eyal Seidemann. Optimal decoding of correlated neural population responses in the primate visual cortex. 9(11):1412–1420, b. ISSN 1097-6256. doi: 10.1038/nm1792. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1851689/>.

Yuzhi Chen, Chris R Palmer, and Eyal Seidemann. The relationship between voltage-sensitive dye imaging signals and spiking activity of neural populations in primate V1. 107(12):3281–3295, c. ISSN 0022-3077.

Lawrence Huang, Peter Ledochowitsch, Ulf Knoblich, Jérôme Lecoq, Gabe J Murphy, R Clay Reid, prefix=de useprefix=true family=Vries, given=Saskia EJ, Christof Koch, Hongkui Zeng, Michael A Buice, Jack Waters, and Lu Li. Relationship between simultaneously recorded spiking activity and fluorescence signal in GCaMP6 transgenic mice. 10:e51675. ISSN 2050-084X. doi: 10.7554/eLife.51675. URL <https://doi.org/10.7554/eLife.51675>.

D. H. Hubel and T. N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. 160(1):106–154. ISSN 0022-3751. doi: 10.1113/jphysiol.1962.sp006837.

Ingmar Kanitscheider, Ruben Coen-Cagli, and Alexandre Pouget. Origin of information-limiting noise correlations. 112(50):E6973–E6982. doi: 10.1073/pnas.1508738112. URL <https://www.pnas.org/doi/10.1073/pnas.1508738112>.

Adam Kohn and Matthew A. Smith. Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. 25(14):3661–3673. ISSN 1529-2401. doi: 10.1523/JNEUROSCI.5106-04.2005.

Michael J Morais, Charles D Michelson, Yuzhi Chen, Jonathan W Pillow, and Eyal Seidemann. Majority of choice-related variability in perceptual decisions is present in early sensory cortex. page 207357. doi: 10.1101/207357.

Rubén Moreno-Bote, Jeffrey Beck, Ingmar Kanitscheider, Xaq Pitkow, Peter Latham, and Alexandre Pouget. Information-limiting correlations. 17(10):1410–1417. ISSN 1546-1726. doi: 10.1038/nn.3807. URL <https://www.nature.com/articles/nn.3807>.

Marius Pachitariu, Carsen Stringer, Mario Dipoppa, Sylvia Schröder, L. Federico Rossi, Henry Dalgleish, Matteo Carandini, and Kenneth D. Harris. Suite2p: Beyond 10,000 neurons with standard two-photon microscopy. URL <https://www.biorxiv.org/content/10.1101/061507v2>.

Nicholas J. Priebe. Mechanisms of Orientation Selectivity in the Primary Visual Cortex. 2:85–107. ISSN 2374-4650. doi: 10.1146/annurev-vision-111815-114456.

Nicholas J. Priebe and David Ferster. Inhibition, Spike Threshold, and Stimulus Selectivity in Primary Visual Cortex. 57(4):482–497. ISSN 08966273. doi: 10.1016/j.neuron.2008.02.005. URL <https://linkinghub.elsevier.com/retrieve/pii/S0896627308001268>.

Dario L. Ringach, Robert M. Shapley, and Michael J. Hawken. Orientation selectivity in macaque V1: Diversity and laminar dependence. 22(13):5639–5651. ISSN 1529-2401. doi: 10.1523/JNEUROSCI.22-13-05639.2002.

Eyal Seidemann, Yuzhi Chen, Yoon Bai, Spencer C Chen, Preeti Mehta, Bridget L Kajs, Wilson S Geisler, and Boris V Zemelman. Calcium imaging with genetically encoded indicators in behaving primates. 5:e16178. ISSN 2050-084X. doi: 10.7554/eLife.16178. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4956408/>.

Matthew A. Smith and Adam Kohn. Spatial and Temporal Scales of Neuronal Correlation in Primary Visual Cortex. 28(48):12591–12603. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.2929-08.2008. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2656500/>.

Vita

Salim M'jahad was born in Marrakech, Morocco. After completing his secondary education at Établissement Elaraki, he entered Columbia University in New York City, where he received the degree of Bachelor of Arts in Mathematics and Computer Science in May 2020. From 2020 to 2022, he was employed as a machine learning engineer at Moveworks, a conversational AI company based in California.

In August 2022, he entered the Interdisciplinary Neuroscience Program (INP) at the University of Texas at Austin to pursue graduate study in neuroscience.

Address: salimmj@proton.me

This thesis was typeset with L^AT_EX[†] by the author.

[†]L^AT_EX is a document preparation system developed by Leslie Lamport as a special version of Donald Knuth's T_EX Program.