A synaptic mechanism for generative inference: Inter-postsynaptic LINKs explain learning-induced redundancy in macaque V4

Kunjumon I. Vadakkan
Neurosearch Center, Toronto, Ontario, Canada
Email: k.vadakkan@gmail.com

Abstract

Task learning in macaque V4 increases information redundancy among orientation-selective neurons while individual-neuron information rises, a pattern that challenges efficient-coding models and supports generative inference. However, the biophysical mechanism implementing this information sharing at the synaptic level remains unspecified. Here, I propose that inter-postsynaptic functional LINKs (IPLs) that are activity-dependent connections between dendritic spines (postsynaptic terminals) of different neurons provide a candidate mechanism. By grounding the computational principles of generative inference in a specific biophysical substrate, the IPL framework offers a path toward integrating systems-level and cellular-level understanding of learning-induced plasticity in sensory cortex.

Introduction

A recent study (Liu et al., 2026) provides compelling evidence that learning reshapes sensory representations in ways that challenge long-held assumptions. Liu et al., report that task learning in macaque V4 increases information redundancy among orientation-selective neurons, while individual-neuron information rises and total population information is largely preserved. This pattern directly contradicts classical efficient-coding models, which predict that learning should reduce redundancy through decorrelation. The authors interpret their findings within a generative-inference framework, where shared prior beliefs introduce correlated variability across neurons through feedback signals from higher cortical areas.

While this framework provides a compelling computational-level account, it leaves open a fundamental question: What is the biophysical mechanism that implements this information sharing at the synaptic and dendritic level? Here, I propose that inter-postsynaptic functional LINKs (IPLs), which are activity-dependent functional connections between dendritic spines (postsynaptic terminals) belonging to different neurons provide a candidate mechanism capable of producing the observed population-

level effects (Vadakkan 2013; 2016). This proposal reframes the generative-inference model's feedback signals as modulatory inputs that bias local spine co-activation, enabling IPL formation and subsequent lateral information sharing within V4.

The mechanistic gap in current interpretations

Liu *et al.* use linear Fisher information to evaluate responses of orientation-selective neurons in macaque V4 during a task learning paradigm. To assess the information carried by each neuron, the authors performed two types of measurements: a) I_{real}, which measures information from the population of neurons, including all correlations and interactions between them; and b) $I_{shuffle}$, which measures the total information contributed by individual neurons alone, obtained by destroying all correlations between neurons by shuffling data across trials.

$$I_{redundancy}=I_{shuffle}-I_{real}$$

If the classical view were correct, and learning simply made the code more efficient by removing correlations, we are expected to see $I_{redundancy}$ decrease and $I_{shuffle}$ stay the same. For this to happen, I_{real}should increase. On the other hand, if $I_{shuffle}$ increases, it means the individual V4 neurons become more informative by a still unknown mechanism that can be hypothesized, raise testable predictions and then verified.

Liu et al., found that $I_{redundancy}$ increases with learning, while I_{real} is preserved or even increases for sampled neuronal subsets implies that correlations become more information-limiting yet behaviorally beneficial. Increases in redundancy of neural responses in macaque visual cortex (area V4) during task learning, reveals an unexpected reorganization of population coding. The authors interpret this primarily in terms of feedback signals from higher areas that distribute task-specific prior beliefs across sensory neurons. Understanding the mechanisms generating redundancy is important because it determines how information is distributed across neural populations and influences the robustness, efficiency, and computational capacity of cortical representations.

The increase in $I_{shuffle}$ with learning indicates that V4 neurons became more informative about the stimulus, whether through sharper tuning, reduced variability, or both. Second, the gap between this individual-neuron information and the information in the actual correlated population $I_{redundancy}=I_{shuffle}-I_{real}$ also increased. As neurons became more informative individually, that information became increasingly shared across the

population. The authors favour a Bayesian/generative interpretation in which feedback broadcasts priors, producing both effects. They support this with simulations showing that such priors can recapitulate the observed patterns and note higher redundancy during active task performance compared to passive viewing. While this feedback-centric view is compelling and aligns with some predictive-coding ideas, several lines of anatomical, dynamical, dimensional, and computational evidence suggest that local cortical circuits likely contribute substantially, and perhaps primarily, to the observed changes.

Several challenges arise when considering whether feedback alone can explain the increase in I_{shuffle}. First, long-range feedback projections terminate diffusely in superficial layers and primarily contact apical dendrites (Rockland and Pandya, 1979; Shipp, 2007; Markov et al., 2014). This pattern is better suited for modulatory or contextual roles than for precise, feature-specific routing to targeted V4 subpopulations (Larkum, 2013; Fišek et al., 2023; Marques et al., 2018). If feedback cannot selectively enhance orientation-specific neurons, it is difficult to see how it could be the primary driver of increased I_{shuffle}.

Second, the timescale of the effect poses a challenge. Liu et al. show redundancy increases within a single trial, with buildup occurring over hundreds of milliseconds. Such rapid, moment-to-moment dynamics may outpace reliable long-range feedback loops. Studies of feedback timing indicate that delayed feedback engages distinct (often slower) neural circuits compared to immediate processing (Röhlinger et al., 2025). If long-range signals are too slow or imprecise for within-trial changes, the gradual buildup likely depends on local recurrent processing.

Third, the dimensionality mismatch is problematic. Feedback from prefrontal or parietal areas typically encodes low-dimensional variables (task rules, attention, confidence), whereas sensory populations represent high-dimensional stimulus features (Stringer et al., 2019; Chung and Abbott, 2021). Redistributing orientation-specific information within this high-dimensional space requires precise coordination that low-dimensional signals alone are unlikely to provide (Saxena and Cunningham, 2019). Local circuit interactions offer a more plausible substrate.

Fourth, the credit assignment problem remains unresolved for feedback-driven accounts. Routing the right prior to the right neurons demands that the brain learn and implement precise mappings, a challenge for biological networks (Lillicrap et al., 2020). Local

synaptic plasticity rules within sensory cortex provide a computationally simpler alternative.

The above considerations favor an alternative. Local recurrent circuits within V4 likely drive the rapid, feature-specific changes in I_{shuffle}, with feedback playing a modulatory or contextual role (e.g., gating or broadly facilitating plasticity) rather than directly specifying increases. Local synapses vastly outnumber long-range feedback inputs (Shipp, 2007; Markov et al., 2014; Harris and Shepherd, 2015), and structured correlations arise naturally from recurrent connectivity and shared inputs (Huang et al., 2019; Doiron and Litwin-Kumar, 2019). The scale of redundancy (about half of individual-neuron information shared) aligns with strong local coupling (Zohary et al., 1994; Averbeck et al., 2006; Cohen and Kohn, 2011).

The within-trial gradual buildup resembles attractor convergence in recurrent networks (Hopfield, 1982; Wang, 2002; Litwin-Kumar and Doiron, 2014) more than rapid top-down modulation. Perceptual learning studies show long-lasting plastic changes in sensory cortex via local reorganization (Schoups et al., 2001; Yang and Maunsell, 2004; Law and Gold, 2008). Notably, the task-versus-passive difference (higher redundancy during active performance) can be explained by local IPL-like coupling being gated or amplified by task-related neuromodulation or disinhibition, without requiring precise feature-specific feedback. Lateral interactions increase redundancy by spreading activity among neighbors, creating shared information without necessarily raising total information content (Zohary et al., 1994; Averbeck et al., 2006). Strengthening such coupling could also boost I_{shuffle} by stabilizing responses to preferred stimuli via recurrent amplification or noise reduction.

Ultrastructural evidence supports IPL plausibility. A typical pyramidal neuron in mammalian cortex possesses 10,000 to 30,000 dendritic spines (Heck and Santos, 2023), each representing a potential site where sensory inputs arrive. Firing of pyramidal neurons can be triggered by activation of approximately100 to 200 dendritic spines within a 10-20 ms window, regardless of which specific subset of dendritic spines is activated (Azouz and Gray, 2000; Palmer et al., 2014; Eyal et al., 2018). In dense neuropil, dendritic spines from different neurons can reside within tens of nanometers of each other, forming microdomains for functional interactions (Spacek and Harris, 1997; Shepherd and Harris, 1998; Mishchenko et al., 2010; Kasthuri et al., 2015). Furthermore, connectomic reconstructions reveal intricate interdigitation of synapses, spines, and axons (Helmstaedter et al., 2013; Motta et al., 2019). Learning could

strengthen functional links within these spine networks, expanding coordinated ensembles and producing both more reliable individual responses and higher population redundancy, matching Liu et al.'s observations.

A multi-stage mechanistic model

I propose a three-stage mechanism that bridges the systems-level observations of Liu *et al.* with spine-level biophysics:

Stage 1: Task-dependent bias (systems level). Top-down signals from decision-making areas (e.g., prefrontal or parietal cortex) bias which orientation-selective V4 neurons are active during task performance. These signals do not themselves carry detailed orientation information, which is consistent with the non-specific nature of many feedback projections and ensures that the specific orientation tuning remains a property of the local V4 microcircuit. However, top-down signals from decision-making areas modulate excitability or attention, increasing the probability that spines tuned to the task-relevant orientation will be co-activated.

Stage 2: IPL formation (dendritic level). When dendritic spines belonging to *different* neurons are activated nearly simultaneously and repeatedly over the course of learning, IPLs form between them (**Fig.1A**). These are reversible, activity-dependent functional connections that allow subthreshold depolarization to spread laterally between neurons without requiring changes in classical synaptic weights. The term “functional” is used to indicate that they form as a function of co-activation of the inter-LINKed spines. The coupling strength $\alpha_{ik} \in [0,1]$ for the connection from spine k to neuron i grows with repeated joint activation.

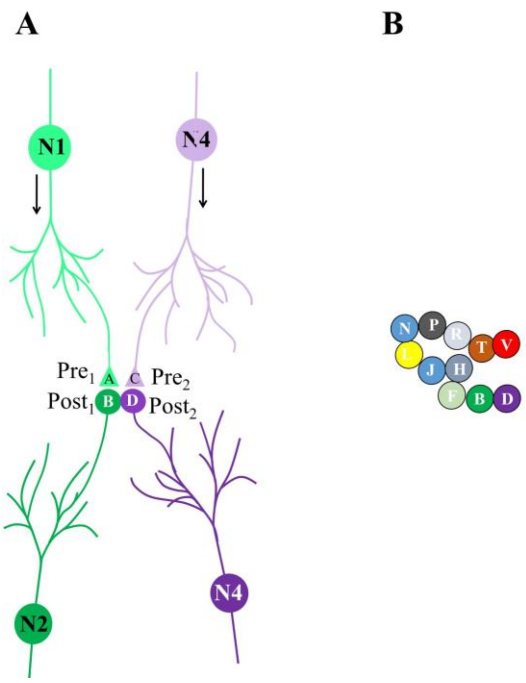

Figure 1. IPL formation between spines of different neurons. ***A)*** *When spines B and D, of two different neurons N2 and N4 respectively, are stimulated together during learning, a reversible, and stabilizable inter-postsynaptic functional LINK (IPL) B-D is formed between them. This allows their postsynaptic neurons to get tuned to the same orientation and is responsible for positive correlation between their correlated firing. The IPL B-D operates bidirectionally.* ***B)*** *Continuation of learning introduces additional IPLs linking to either B or D or both, leading to the formation of islets of inter-LINKed spines (IILSPs). Note that IPL B-D bottom right of the IILSPs (Figure modified from Vadakkan, 2013).*

Stage 3: Lateral information sharing (population level). Once formed, networks of IPLs create islets of inter-LINKed spines (IILSPs) that distribute task-related signals across multiple neurons (**Fig.1B**). The effective postsynaptic depolarization at spine j of neuron i becomes

$$V_{ij}=V_{ij}^{\text{synaptic}}+\sum_{k\in\text{IILSPs}}\alpha_{ik}\cdot s_k$$

V_{ij}^{synaptic} is the classical synaptic potential from afferent input. This coupling term can be thought of as a form of "effective connectivity" that is not captured by standard cross-correlograms of spiking activity alone, as it operates on sub-threshold potentials. The

summation runs over all IILSPs that include spine j of neuron i. α_{ik} = coupling coefficient from IILSPs k to neuron i, s_k represents subthreshold potentials arriving through IPLs from inter-LINKed spine k. Firing rate $r_i \propto V_{ij}$ therefore shows spike-time alignment that increases with learning, exactly the pattern in Liu *et al.* Fig.1A.

This multi-stage model preserves the core insight of generative inference that prior beliefs are shared among neurons while providing a specific biophysical implementation. Crucially, it resolves the scale mismatch between long-range feedback signals (systems level) and local correlations within V4 (population level): top-down signals provide the *bias*, while IPLs provide the *lateral connectivity* that implements information sharing locally (**Fig.2**).

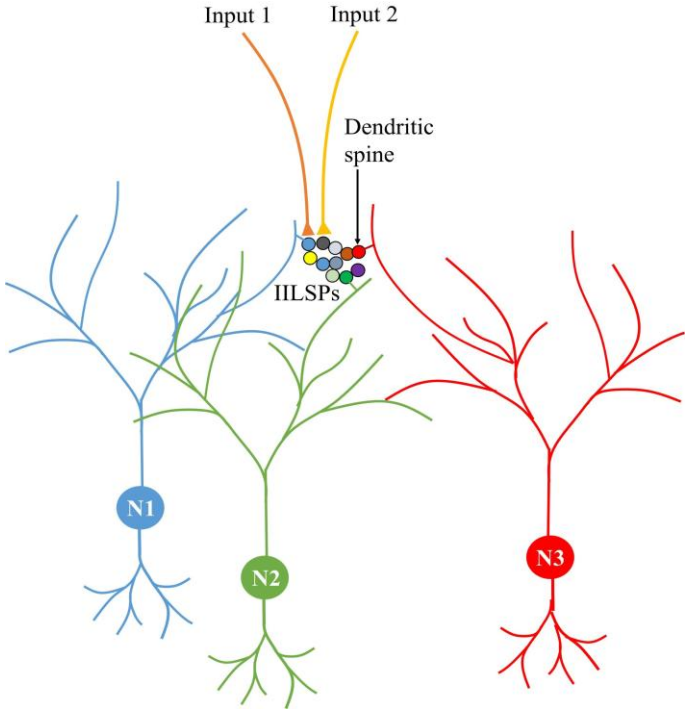

Figure 2. Population model accompanying lateral information sharing through IPL coupling*. A population of three neurons N1, N2 and N3 are shown with one each of their spines inter-LINKed with spines of seven other neurons within an islet of inter-LINKed spines (IILSPs). Even though input arrives only to the spine of neuron N1 and to another spine within the IILSPs, postsynaptic potentials propagate within the IILSPs and may allow either one or two or all the neurons N1, N2 and N3 to cross their thresholds for firing. The increase in I_{shuffle} with learning in Liu et al.,'s work indicates that each individual neuron becomes more informative about the stimulus, even when all the correlation were destroyed by shuffling data across trials. As neurons become more*

informative individually, that information become increasingly shared across the population. This necessitates the presence of a yet undiscovered connections between the neurons capable of sharing information. Connection between the neurons across the IILSPs offer a testable mechanism.

Mapping empirical observations to the IPL mechanism

The observations in Liu *et al.* (2026) follow naturally from this multi-stage model:

Figure 1 (Liu *et al.*) motivates the experimental test by contrasting two theoretical predictions: classical feedforward models predict that learning reduces redundancy among sensory neurons, whereas generative inference models predict that learning increases it. The authors therefore propose to test this by measuring redundancy, defined as the difference between Fisher information in the real population activity and that obtained after removing correlations. Positive correlations between neurons tuned to the same orientation arise because IPLs connect spines representing shared features. The coupling term $\sum \alpha_{ik} s_k$ introduces correlated variability without requiring changes in feedforward synaptic weights.

Figure 2 (Liu *et al.*) describes the behavioral experiment in which monkeys learned to discriminate visual orientations embedded in noise and reported their choice with a saccade. Behavioral analyses show that the animals gradually learned the task over weeks and performed close to the ideal observer in late training. The 1.6-s oriented-noise stimuli repeatedly co-activate orientation-selective spines, satisfying the requirement for IPL formation (α_{ik} grows with repeated joint depolarization). Psychometric curves approach ideal-observer performance because IILSPs distribute the same orientation "belief" across many neurons, enabling consistent readout.

Figure 3 (Liu *et al.*) shows that as monkeys got better at a visual task, the neurons in their visual cortex started to share more information with each other showing increased redundancy of their responses. The monotonic rise in $I_{\text{redundancy}}$ across training sessions reflects cumulative IPL formation

$$\frac{dI_{\text{redundancy}}}{dt} \propto \frac{d}{dt}\left(\sum \alpha_{ik}\right)$$

Critically, the observation that redundancy is far higher during *active task performance* than during *passive viewing* on the same day (Fig. 3C) aligns with the reversible, state-dependent nature of IPLs. This provides more explanatory power than the generative inference model, which might imply a more static "prior." Passive viewing lacks the top-

down biasing signals that maintain task-appropriate spine co-activation; consequently, coupling strength α drops rapidly. This provides a mechanistic explanation for why the same neuronal population shows different redundancy levels depending on behavioral context, a finding that pure feedforward models struggle to explain.

Figure 4 (Liu *et al.*). To assess the information carried by each neuron, the authors performed two types of measurements: a) I_{real}, which measures information from the population of neurons, including all correlations and interactions between them; and b) I_{shuffle}, which measures the total information contributed by individual neurons alone, obtained by destroying all correlations between neurons by shuffling data across trials. If I_{shuffle} increases, it means the individual V4 neurons become more informative by a still unknown mechanism that can be hypothesized, raise testable predictions and then verified. The model simulations in Fig.4 of the study show I_{shuffle} rising faster than I_{real}, with redundancy increasing without net information loss.

The IPL framework can reproduce this pattern: each neuron gains extra depolarization from IPL-linked partners, increasing marginal information per neuron (reflected in I_{shuffle}) while simultaneously introducing correlated variability that tempers the growth of population information I_{real}. Lateral links between spines can make a neuron's output more informative by allowing depolarization from inter-LINKed neighbours to boost local signal-to-noise, enabling the neuron to carry information beyond that provided by its own synapses. The rise in I_{shuffle}) reflects the contribution of the diagonal elements of the coupling matrix A and the increase in $f'(s)$ as each neuron's effective signal is boosted by its IILSPs.

Figure 5 (Liu *et al.*) shows that I_{shuffle} increases much faster than I_{real} moment-by-moment within a trial, meaning that the gap between them ($I_{\text{redundancy}}$) grows, implying that information is continuously shared across neurons by a still unknown mechanism. The actual neurobiological or biophysical mechanism should be able to explain this information transfer route. The within-trial rise in redundancy (steeper late in learning) reflects the rapid kinetics of IPL formation or efficacy modulation over hundreds of milliseconds. The compensation for adaptation-related firing rate drops can be explained by propagation of potentials across IPLs sustaining the signal.

$$r_i(t) = r_i^{\text{synaptic}}(t) + \sum \alpha_{ik}(t) \cdot s_k(t\text{-}\Delta t)$$

The single exception (monkey R cardinal task, where coarse 200-ms bins showed no increase) actually supports the model: when analyzed with 50-ms bins, the predicted rise appears (Fig. S19), consistent with sub-second IPL dynamics that can be missed with coarse temporal sampling.

A population model linking IPL coupling to information metrics

In population recording studies the activity of many neurons is often summarized by their covariant matrix, which captures how much neurons' activities co-vary, reflecting shared inputs, common noise, or functional connectivity) yields exactly the simulated trajectory (the time-evolving path of neural population activity in a low-dimensional space). To formalize the link between IPLs and the observed information metrics, consider a population of (N) neurons responding to stimulus (s). In standard models, neuronal responses are $r_i = f_i(s) + \epsilon_i$ with independent noise ϵ_i. When IPLs form, activity can spread between neurons, giving,

$$r = f(s) + \epsilon + Au$$

Here, $f(s)$ represents the feedforward tuning, while the Au term captures the additional IPL-mediated contribution from other spines. Assuming independent synaptic noise $\epsilon \sim (0, \sigma^2 I)$ and spine activity $u \sim (0, C_u)$, the covariance matrix of the population response follows directly.

$$C = \mathrm{Cov}(r) = \mathrm{Cov}(\epsilon + Au) = \sigma^2 I + AC_u A^T$$

where the second term represents correlated variability introduced by IPL coupling. Linear Fisher information for the population is $I_{\mathrm{real}} = f'(s)^T C^{-1} f'(s)$. after shuffling trials to destroy correlations the covariance becomes $\sigma^2 I$, yielding

$$I_{\mathrm{shuffle}} = \left|\left|f'(s)\right|\right|^2 / \sigma^2$$

The IPL coupling matrix A affects these quantities in specific ways. 1) Diagonal entries (a neuron's self-coupling through its own spines) contribute to $f'(s)$, increasing the marginal information each neuron carries. This drives the increase in I_{shuffle}. 2) Off-diagonal entries (cross-neuron coupling) contribute to the correlated term $AC_u A^T$ which increases redundancy and tempers the growth of I_{real}.

The observed pattern I_{shuffle} rising faster than I_{real}, with $I_{\mathrm{redundancy}} > 0$, emerges naturally when IPLs form during learning. Because the number of possible pairwise interactions

grow as $O(N^2)$, even modest increases in IPL formation probability can generate substantial increases in population redundancy once learning stabilizes.

Specific predictions of the IPL model

A mechanistic model gains credibility when it generates predictions that distinguish it from alternative explanations. The IPL framework offers several testable predictions that go beyond those of the generative-inference model alone:

1. Spatial scale of redundancy increases. IPLs form locally between nearby spines (within dendritic arbors and between neighboring neurons). Liu *et al.* used Utah arrays with 400 μm inter-electrode spacing; re-analyzing their data to compute $I_{\text{redundancy}}$ as a function of inter-neuron distance would provide a direct test. The generative-inference model, which invokes broadcast feedback signals, does not necessarily predict such spatial specificity. The model predicts that the within-trial slope of $I_{\text{redundancy}}$ should be proportional to the *consistency* of the orientation signal in the preceding ~100-200ms, reflecting the integration time constant of the spine-spine interaction.

2. Rapid state-dependent reversibility. The drop in redundancy during passive viewing (Fig. 3C) reflects rapid reduction in IPL coupling strength α, not dismantling of physical connections. This predicts that switching between task and passive viewing should produce changes in $I_{\text{redundancy}}$ on the order of seconds to minutes, tracking the presence of top-down biasing signals. This could be tested by interleaving task and passive blocks within single sessions.

3. Temporal dynamics of redundancy buildup. If IPLs require repeated co-activation to increase α , the within-trial rise in redundancy (Fig. 5) should show a **s**pecific dependence on stimulus history. For example, redundancy should build more rapidly when successive stimulus frames contain consistent orientation information (as in Liu *et al.*'s dynamic noise stimuli) compared to when orientation varies unpredictably. The generative-inference model might also predict this, but the IPL framework provides a mechanistic rationale based on spine-level integration over tens to hundreds of milliseconds. The model predicts that learning-induced increases in $I_{\text{redundancy}}$ should be strongest for neuron pairs recorded on nearby electrodes (within ~200-400 μm) and weaker for pairs separated by larger distances.

4. Dissociability through pharmacological or genetic manipulation. Most distinctively, the IPL model predicts that interfering with mechanisms required for spine-

spine interaction should reduce or prevent learning-induced redundancy increases, even when top-down feedback signals remain intact. Potential manipulations include 1) Disrupting membrane dynamics that enable IPL formation, 2) Interfering with spine enlargement mechanisms, and 3) Selectively blocking lateral spread of subthreshold depolarization. For example, agents that interfere with membrane fluidity (e.g., cholesterol sequestering agents) or spine actin dynamics (e.g., latrunculin) applied locally to V4 should abolish the learning-induced increase in $I_{\text{redundancy}}$ without impairing basic visual responses or top-down attentional signals.

Such manipulations would be expected to reduce the coupling matrix A, thereby decreasing the correlated component AC_uA^T and reducing $I_{\text{redundancy}}$. The generative-inference model, which does not specify a biophysical implementation, would not predict this specific dependency.

Reconciling with classical findings

The rise in information-limiting correlations during learning seems to contradict studies showing that attention and learning reduce noise correlations. The IPL framework resolves this by distinguishing average noise correlations from information-limiting ones, as emphasized by Moreno-Bote et al. (2014). This fits Liu et al.'s data, where average correlations show no systematic change (Fig. S7), yet redundancy increases. The added redundancy aligns with non-information-limiting correlations in the bulk population (Moreno-Bote et al., 2014), since total population information does not saturate despite growing shared variability. This pattern supports local recurrent origins over rigid, differential structure from feedback. While IPLs generate information-limiting correlations, they function as certainty-enhancing mechanisms in a task context.

IPLs introduce correlations specifically along the task-relevant tuning axis (both neurons preferring the same orientation), creating information-limiting structure. They need not increase average pairwise correlations across all neuron pairs. Indeed, if IPLs form selectively between similarly tuned neurons while other correlations decrease (e.g., through sharper tuning or reduced shared input from irrelevant sources), average correlations could remain stable or even decrease while information-limiting correlations grow. This aligns precisely with Liu *et al.*'s findings and resolves the apparent contradiction with prior work.

Conclusion

Liu et al. provide compelling evidence that learning increases information redundancy in macaque V4, challenging classical efficient-coding accounts and supporting generative-inference frameworks. However, the generative-inference model, while computationally powerful, leaves unspecified the biophysical mechanisms that implement information sharing at the microcircuit level.

The IPL framework proposed here offers a candidate mechanism that provides the "how" for the generative inference "what": activity-dependent functional connections between dendritic spines that form during learning, enabling lateral spread of subthreshold depolarization and creating correlated variability along task-relevant dimensions. This model 1) provides a spine-level substrate for the observed population-level effects, 2) resolves the scale mismatch between long-range feedback and local correlations, 3) explains the rapid task/passive viewing differences through reversible coupling, 4) generates distinguishing predictions, and 5) reconciles the apparent conflict with studies showing reduced average noise correlations.

By grounding both the computational principles of generative inference and the generation of first-person inner sensations in the same spine-level IPL mechanism, the framework offers a testable path toward integrating systems-level observations with cellular-level understanding of learning-induced plasticity and subjective experience in sensory cortex. IPLs are a testable, mechanistic hypothesis that bridges levels of analysis: if learning induces IPLs, coordinated activity should rise among spatially proximal neurons; high-resolution imaging and connectomics could test whether high-redundancy neurons show enhanced nanoscale spine proximity; manipulating extracellular matrix or inter-spine interactions should alter redundancy levels. These predictions offer a roadmap for determining whether local microcircuit interactions drive information redistribution during learning. Together, the evidence suggests that learning-related redundancy arises from an interplay of local circuits and modulatory feedback, with IPL-like local interactions likely playing a major role. The ability of a millisecond time-scale IPL mechanism to account for the findings of Liu et al. compels further exploration of this still falsifiable hypothesis.

References

Aitchison L, Lengyel M (2017) With or without you: Predictive coding and Bayesian inference in the brain. Curr Opin Neurobiol. 46:219-227

Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. Nat Rev Neurosci. 7(5):358-366

Azouz R, Gray CM (2000) Dynamic spike threshold reveals a mechanism for synaptic coincidence detection in cortical neurons in vivo. Proc Natl Acad Sci U S A. 97(14): 8110-8115

Chaudhuri R, Knoblauch K, Gariel MA, Kennedy H, Wang XJ (2015) A large-scale circuit mechanism for hierarchical dynamical processing in the primate cortex. Neuron. 88(2):419-431

Chung S, Abbott LF (2021) Neural population geometry: An approach for understanding biological and artificial neural networks. Curr Opin Neurobiol. 70:137-144

Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. Nat Neurosci. 14(7):811-819

Cohen MR, Maunsell JH (2009) Attention improves performance primarily by reducing interneuronal correlations. Nat Neurosci. 12(12):1594-1600

Doiron B, Litwin-Kumar A (2019) Balanced neural architecture and the idling brain. Front Comput Neurosci. 13:56

Eyal G, Verhoog MB, Testa-Silva G, Deitcher Y, Benavides-Piccione R, DeFelipe J et al. (2018) Human cortical pyramidal neurons: From spines to spikes via models. Front Cell Neurosci. 12:181

Fišek M, Herrmann D, Egea-Weiss A, Cloves M, Bauer L, Lee TY, Russell LE, Häusser M (2023) Cortico-cortical feedback engages active dendrites in visual cortex. Nature. 617(7962):769-776

Friston K (2005) A theory of cortical responses. Philos Trans R Soc Lond B Biol Sci. 360(1456):815-836

Harris KD, Shepherd GM (2015) The neocortical circuit: themes and variations. Nat Neurosci. 18(2):170-181

Heck N, Santos MD (2023) Dendritic spines in learning and memory: From first discoveries to current insights. Adv Neurobiol. 34: 311-348

Helmstaedter M, Briggman KL, Turaga SC, Jain V, Seung HS, Denk W (2013) Connectomic reconstruction of the inner plexiform layer in the mouse retina. Nature. 500:168-174

Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. Proc Natl Acad Sci U S A. 79(8):2554-2558

Huang C, Ruff DA, Pyle R, Rosenbaum R, Cohen MR, Doiron B (2019) Circuit models of low-dimensional shared variability in cortical networks. Neuron. 101(2):337-348.e4

Kasthuri N, Hayworth KJ, Berger DR, Schalek RL, Conchello JA, Knowles-Barley S, Lichtman JW (2015) Saturated reconstruction of a volume of neocortex. Cell. 162(3):648-661

Keller GB, Mrsic-Flogel TD (2018) Predictive processing: A canonical cortical computation. Neuron. 100(2):424-435

Larkum M (2013) A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. Trends Neurosci. 36(3):141-151

Law CT, Gold JI (2008) Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. Nat Neurosci. 11(4):505-513

Lillicrap TP, Santoro A, Marris L, Akerman CJ, Hinton G (2020) Backpropagation and the brain. Nat Rev Neurosci. 21(6):335-346

Litwin-Kumar A, Doiron B (2014) Slow dynamics and high variability in balanced cortical networks with clustered connections. Nat Neurosci. 17(11):1506-1515

Liu S, Pletenev A, Haefner RM, Snyder AC (2026) Task learning increases information redundancy of neural responses in macaque visual cortex. Science. 391(6789):1029-1035

Markov NT, Ercsey-Ravasz MM, Ribeiro Gomes AR, Lamy C, Magrou L, Vezoli J, Misery P, Falchier A, Quilodran R, Gariel MA, Sallet J, Gamanut R, Huissoud C, Clavagnier S, Giroud P, Sappey-Marinier D, Barone P, Dehay C, Toroczkai Z, Knoblauch K, Van Essen DC, Kennedy H (2014) A weighted and directed interareal connectivity matrix for macaque cerebral cortex. Cereb Cortex. 24(1):17-36

Marques T, Nguyen J, Fioreze G, Petreanu L (2018) The functional organization of cortical feedback inputs to primary visual cortex. Nat Neurosci. 21(5):757-764

Mishchenko Y, Hu T, Spacek J, Mendenhall J, Chklovskii DB (2010) Ultrastructural analysis of hippocampal neuropil from the connectomics perspective. Neuron. 67(6):1009-1020

Moreno-Bote R, Beck J, Kanitscheider I, Pitkow X, Latham P, Pouget A (2014) Information-limiting correlations. Nat Neurosci. 17(10):1410-1417.

Motta A, Berning M, Boergens KM, Staffler B, Beining M, Loomba S, Helmstaedter M (2019) Dense connectomic reconstruction in layer 4 of the somatosensory cortex. Science. 366(6469):eaay3134

Okun M, Steinmetz N, Cossell L, Iacaruso MF, Ko H, Barthó P, Moore T, Hofer SB, Mrsic-Flogel TD, Carandini M, Harris KD (2015) Diverse coupling of neurons to populations in sensory cortex. Nature. 521(7553):511-515

Palmer LM, Shai AS, Reeve JE, Anderson HL, Paulsen O, Larkum ME (2014) NMDA spikes enhance action potential generation during sensory input. Nat Neurosci. 17(3): 383-390

Rao RP, Ballard DH (1999) Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nat Neurosci. 2(1):79-87

Renart A, de la Rocha J, Bartho P, Hollender L, Parga N, Reyes A, Harris KD (2010) The asynchronous state in cortical circuits. Science. 327(5965):587-590

Rockland KS, Pandya DN (1979) Laminar origins and terminations of cortical connections of the occipital lobe in the rhesus monkey. Brain Res. 179(1):3-20

Röhlinger M, Albrecht T, Ghio M, Bellebaum C (2025) Neural processing of immediate versus delayed feedback in action-feedback and stimulus-feedback associations. J Cogn Neurosci 37(11):2225-2259.

Ruff DA, Cohen MR (2014) Global cognitive factors modulate correlated response variability between V4 neurons. J Neurosci. 34(49):16408-16416

Saxena S, Cunningham JP (2019) Towards the neural population doctrine. Curr Opin Neurobiol. 55:103-111

Schoups A, Vogels R, Qian N, Orban G (2001) Practising orientation identification improves orientation coding in V1 neurons. Nature. 412(6846):549-553

Shepherd GM, Harris KM (1998) Three-dimensional structure and composition of CA3→CA1 axons in rat hippocampus: Implications for presynaptic connectivity and compartmental organization. J Neurosci. 18(20):8300-8310

Shipp S (2007) Structure and function of the cerebral cortex. Curr Biol. 17(12):R443-R449

Spacek J, Harris KM (1997) Three-dimensional organization of smooth endoplasmic reticulum in hippocampal dendrites and dendritic spines of the rat. J Neurosci. 17(1):190-203

Spratling MW (2017) A review of predictive coding algorithms. Brain Cogn. 112:92-97

Stringer C, Pachitariu M, Steinmetz N, Carandini M, Harris KD (2019) High-dimensional geometry of population responses in visual cortex. Nature. 571(7765):361-365

Vadakkan KI (2013) A supplementary circuit rule-set for the neuronal wiring. Front Hum Neurosci. 7:170.

Vadakkan KI (2016) The functional role of all postsynaptic potentials examined from a first-person frame of reference. Rev Neurosci. 27(2):159-184.

Wang XJ (2002) Probabilistic decision making by slow reverberation in cortical circuits. Neuron. 36(5):955-968

Yang T, Maunsell JH (2004) The effect of perceptual learning on neuronal responses in monkey visual area V4. J Neurosci. 24(7):1617-1626

Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. Nature. 370(6485):140-143