Symmetry constraints on feedback-driven orientation tuning dynamics in translation-invariant models of primate V1

Makoto Fukushima[1¤*]
1 Division of Physical Sciences, The University of Chicago, Chicago, Illinois 60637, USA
¤Current address: Honda Research Institute Japan, 8-1 Honcho, Wako-shi, Saitama 351-0114, Japan
* mfukush@uchicago.edu

Abstract

Orientation selectivity in primate V1 is shaped by feedforward drive, intracortical recurrence, and corticocortical feedback. Yet a basic conceptual question remains: what aspects of orientation tuning can (or cannot) be altered by a recurrent/feedback loop that is rotationally symmetric in cortical space? We develop a minimal translation-invariant linear field model in which stimulus drive and recurrent interactions are analyzed in the spatial Fourier domain (k, ϕ). Our central analytic result is an exact symmetry constraint: for any temporally filtered but isotropic recurrent coupling $W_0(k, \omega)$, the response at fixed (k, ω) is multiplied by an angle-independent complex gain. Consequently, any normalized orientation metric computed at fixed k—including circular variance, orientation selectivity index (OSI), and bandwidth—is invariant under isotropic recurrence. Isotropic recurrence can nevertheless rescale gain and select a spatial band through its k-dependence. We then show that adding a weak anisotropic component, especially when endowed with a slower timescale than a fast isotropic pathway, produces delayed sharpening and modest preferred-angle drifts. A perturbative calculation predicts linear scaling of tuning changes with anisotropy strength, and a simple phasor-sum formula explains $\phi^*(t)$ drift when feedforward and feedback anisotropies compete. We provide stability criteria for the temporally filtered (auxiliary-variable) implementation and verify them numerically. Finally, we identify concrete regimes in which isotropic recurrence can alter *measured* tuning in practice, including finite-k readouts with k-dependent feedforward anisotropy and nonlinear (tuned) divisive-normalization readouts. Together these results sharpen a null hypothesis for perturbation experiments and delineate the minimal feature-specific ingredients required to explain dynamic tuning in macaque V1.

Author summary

Neurons in the primary visual cortex (V1) respond more strongly to some stimulus orientations than others. This "orientation selectivity" is often attributed to feedforward input from the thalamus, but it is also shaped by recurrent and feedback connections within and between cortical areas. A common intuition is that if feedback is not itself tuned for orientation—for example, if it is spatially symmetric and acts like a global gain control—then it should change response strength without changing tuning. Here we turn that intuition into an exact mathematical statement in a simple but widely used approximation: locally translation-invariant, small-signal (linearized) dynamics analyzed in the spatial Fourier domain. We prove that any rotationally symmetric recurrent coupling cannot change normalized tuning metrics at a fixed

spatial frequency, even if it has rich temporal filtering. We then show what is minimally required to change tuning: a weak orientation-biased component, especially if it has a slower timescale than an untuned pathway. This two-timescale mechanism produces delayed sharpening and small drifts in preferred orientation, features reported in macaque V1. We also explain why tuning can still appear to change under untuned feedback in real experiments, for example when measurements combine multiple spatial scales or when normalization pools are orientation selective.

Introduction

Orientation selectivity is a canonical computation of primate V1 [1, 2]. Multiple circuit motifs contribute to this selectivity, including feedforward thalamocortical drive, intracortical recurrence, and long-range corticocortical feedback [3–9]. Classic theoretical accounts emphasize structured recurrence in an orientation feature space, including ring/continuous-attractor models and their variants [10–12]. At the same time, both theory and experiment highlight large untuned components such as global suppression and divisive normalization [13–15]. Orientation tuning is also dynamic: tuning metrics and preferred orientation can evolve over tens of milliseconds after stimulus onset in macaque V1 [15–17].

A useful way to organize these debates is to ask what is forced by symmetry. If recurrent couplings are translationally homogeneous and isotropic in cortical space, their Fourier representation depends only on the magnitude k of the spatial frequency, not on its angle ϕ. Intuitively, such isotropic recurrence should amplify or suppress responses without "choosing" a preferred orientation. However, it is often unclear which tuning measurements this intuition constrains, especially when temporal filtering is present or when experiments quantify tuning after integrating over spatial frequency or applying nonlinear readouts.

Here we formalize these issues in a minimal, analytically tractable framework. We derive an exact invariance theorem for *normalized* orientation metrics at fixed (k, ω) under isotropic recurrent coupling, providing a sharp null hypothesis for perturbation experiments. We then introduce a weak anisotropic component and show how delayed, feature-specific recurrence can generate sharpening and preferred-angle drift. Our analysis complements (i) classical ring-model accounts that assume explicitly feature-specific recurrence [10–12], (ii) linearized network analyses showing that random (feature-unspecific) recurrence amplifies but does not bias tuning [18], and (iii) salt-and-pepper modeling results highlighting the need for feature-specific connectivity to match observed tuning distributions [19, 20]. We further discuss how developmental/plasticity mechanisms can generate anisotropy [21] and how layer-resolved models constrain timescales and delays [22].

Materials and methods

Translation-invariant field model and small-signal interpretation

We consider a linear rate field $u(\mathbf{x}, t)$ on the cortical plane $\mathbf{x} \in \mathbb{R}^2$ driven by a stimulus $I(\mathbf{x}, t)$,

$$\partial_t u(\mathbf{x}, t) = -\alpha u(\mathbf{x}, t) + (w_{\mathrm{ff}} * I)(\mathbf{x}, t) + (w_{\mathrm{rec}} * u)(\mathbf{x}, t), \tag{1}$$

where $\alpha > 0$ is an effective decay rate and $*$ denotes convolution in space (and, when specified, time). Equation (1) can be viewed as a small-signal linearization of nonlinear rate models around a spatially homogeneous operating point. For example, in Wilson–Cowan-type formulations the firing rate is a sigmoidal function of synaptic

input, $r = f(u)$; expanding about a baseline u_0 gives $\delta r \approx f'(u_0)\,\delta u$, so linear dynamics describe small perturbations with effective gain absorbed into the coupling kernels [23]. This interpretation helps delimit the scope of the exact invariance theorem below: nonlinearities that couple orientations (e.g., tuned normalization) can break invariance even when the recurrent kernel is isotropic.

Let $\hat{u}(\mathbf{k}, \omega)$ denote the spatiotemporal Fourier transform of u. Translation invariance diagonalizes Eq. (1) in the Fourier basis,

$$\hat{u}(\mathbf{k}, \omega) = \frac{\hat{w}_{\mathrm{ff}}(\mathbf{k})\,\hat{I}(\mathbf{k}, \omega)}{\alpha + i\omega - \hat{w}_{\mathrm{rec}}(\mathbf{k}, \omega)}. \tag{2}$$

Write $\mathbf{k} = k(\cos\phi, \sin\phi)$ and define the tuning curve at fixed (k, ω) as $R(\phi) = |\hat{u}(k, \phi, \omega)|$. We define the normalized angular profile $p(\phi) = R(\phi)/\int_0^\pi R(\phi')\,d\phi'$ and quantify selectivity using the second-harmonic orientation selectivity index (OSI)

$$\mathrm{OSI} = \left| \int_0^\pi p(\phi) e^{i2\phi}\,d\phi \right|, \qquad \mathrm{CV} = 1 - \mathrm{OSI}, \tag{3}$$

so $\mathrm{OSI} \in [0, 1]$ and $\mathrm{CV} \in [0, 1]$.

Isotropic and anisotropic decomposition

We decompose the recurrent kernel into an isotropic component plus a weak anisotropic perturbation,

$$\hat{w}_{\mathrm{rec}}(k, \phi, \omega) = \mu W_0(k, \omega) + \mu\varepsilon W_2(k, \omega)\cos 2(\phi - \phi_0), \tag{4}$$

where μ sets overall coupling strength, ε is dimensionless anisotropy, and ϕ_0 sets the anisotropy axis. We also allow a weak feedforward orientation bias at fixed k,

$$\hat{w}_{\mathrm{ff}}(k, \phi) \propto 1 + \eta_{\mathrm{ff}} \cos 2(\phi - \phi_{\mathrm{ff}}), \tag{5}$$

with bias strength η_{ff} and axis ϕ_{ff}.

Spatial kernel and gain parameterization

For the isotropic spatial kernel we use a difference-of-Gaussians (DoG) in Fourier space, $W_0(k) = A\exp(-\frac{1}{2}s_1^2k^2) - B\exp(-\frac{1}{2}s_2^2k^2)$. The DoG parameters determine a characteristic spatial frequency k_c where $W_0(k)$ is maximal. We define a dimensionless closed-loop gain ratio $g_0 \equiv \mu W_0(k_c)/\alpha$. Figure 1 plots the corresponding amplification factor $|\alpha/(\alpha - \mu W_0(k))|$ to visualize k-dependent gain and band selection under isotropic recurrence.

Temporal filtering and auxiliary-variable implementation

To capture onset dynamics we endow isotropic and anisotropic recurrent components with distinct temporal filters. We use an exponential low-pass for the isotropic component, $h_{\mathrm{iso}}(t) = (1/\tau_{\mathrm{iso}})e^{-t/\tau_{\mathrm{iso}}}H(t)$, and a second-order (gamma) filter for the anisotropic component, $h_{\mathrm{aniso}}(t) = (t/\tau_{\mathrm{aniso}}^2)e^{-t/\tau_{\mathrm{aniso}}}H(t)$, with Laplace transforms $\tilde{h}_{\mathrm{iso}}(s) = 1/(1 + s\tau_{\mathrm{iso}})$ and $\tilde{h}_{\mathrm{aniso}}(s) = 1/(1 + s\tau_{\mathrm{aniso}})^2$. These convolutions are implemented exactly via auxiliary variables,

$$\tau_{\mathrm{iso}}\dot{v}_{\mathrm{iso}} = u - v_{\mathrm{iso}}, \tag{6}$$

$$\tau_{\mathrm{aniso}}\dot{w} = u - w, \qquad \tau_{\mathrm{aniso}}\dot{v}_{\mathrm{aniso}} = w - v_{\mathrm{aniso}}. \tag{7}$$

At fixed (k, ϕ) the reduced dynamics take the form

$$\dot{u} = -\alpha u + I + g_0 \alpha \left[v_{\text{iso}} + \varepsilon \cos 2(\phi - \phi_0)\, v_{\text{aniso}} \right], \tag{8}$$

where I is the (Fourier-mode) feedforward drive. We integrate these ODEs with a fixed-step explicit Euler method ($\Delta t = 0.2$ ms in all figures) and verify numerical stability by checking the eigenvalues of the state matrix and the Euler stability region for the chosen step size.

Finite-k readout and nonlinear readouts

The invariance theorem is formulated at fixed (k, ω). To model common measurement pipelines that integrate over spatial frequency, we compute broadband readouts $R_{\text{int}}(\phi) = \int S(k) R(k, \phi)\, dk$ with a specified stimulus spectrum $S(k)$ and k-dependent feedforward anisotropy $\eta_{\text{ff}}(k)$. To probe nonlinear measurement effects, we also compute selectivity metrics after applying (i) rectification and (ii) divisive normalization readouts. In an orientation-channel representation,

$$r(\phi) = \frac{[u(\phi)]_+}{\sigma + \kappa \int_0^\pi w_{\text{norm}}(\phi - \phi')[u(\phi')]_+ \, d\phi'}, \tag{9}$$

with untuned (constant) or tuned (orientation-dependent) normalization pools w_{norm}.

Artificial intelligence tools and technologies

We here report use of artificial intelligence (AI) tools or technologies in preparation of a manuscript and associated files. During the development of this study and article, we used a large language model (LLM) tool (ChatGPT, OpenAI) to assist with: (i) drafting and editing portions of the manuscript text (including rewording for clarity and structure), (ii) generating initial versions of code for numerical simulations, parameter sweeps, and figure generation, and (iii) assisting with identification of relevant background literature. The author evaluated the validity of AI tool outputs by independently evaluating all key analytic steps and numerical results and verifying that cited sources support the corresponding statements. All quantitative results reported in the manuscript are produced by analytic derivations and numerical simulations run using the provided code, and all interpretations and conclusions are the author's responsibility. All python scripts used to generate figures and run parameter sweeps are included with this submission.

Results

Exact symmetry constraint for isotropic recurrence

The key observation is immediate from Eq. (2): if $\hat{w}_{\text{rec}}(k, \phi, \omega) = \mu W_0(k, \omega)$ is isotropic (independent of ϕ), then the denominator of Eq. (2) is ϕ-independent. At fixed (k, ω) the tuning curve can be written as

$$R(\phi) = G(k, \omega)\, R_{\text{ff}}(\phi), \qquad G(k, \omega) = \left| \alpha + i\omega - \mu W_0(k, \omega) \right|^{-1}, \tag{10}$$

so the normalized angular profile $p(\phi)$ is unchanged. Therefore any normalized orientation metric computed from $p(\phi)$ (e.g., OSI, CV, bandwidth) is invariant under isotropic recurrence at fixed (k, ω). This provides a sharp null hypothesis for perturbation experiments: changes in normalized tuning at fixed spatial frequency imply feature-specific (anisotropic) effective coupling.

What isotropic recurrence can do: gain and spatial band selection

Although isotropy constrains normalized metrics at fixed k, isotropic recurrence can strongly affect overall gain and the spatial band emphasized by the circuit because $W_0(k)$ is generally k-dependent. Figure 1 shows the closed-loop amplification factor as a function of k for a DoG kernel, illustrating spatial band selection through isotropic recurrence.

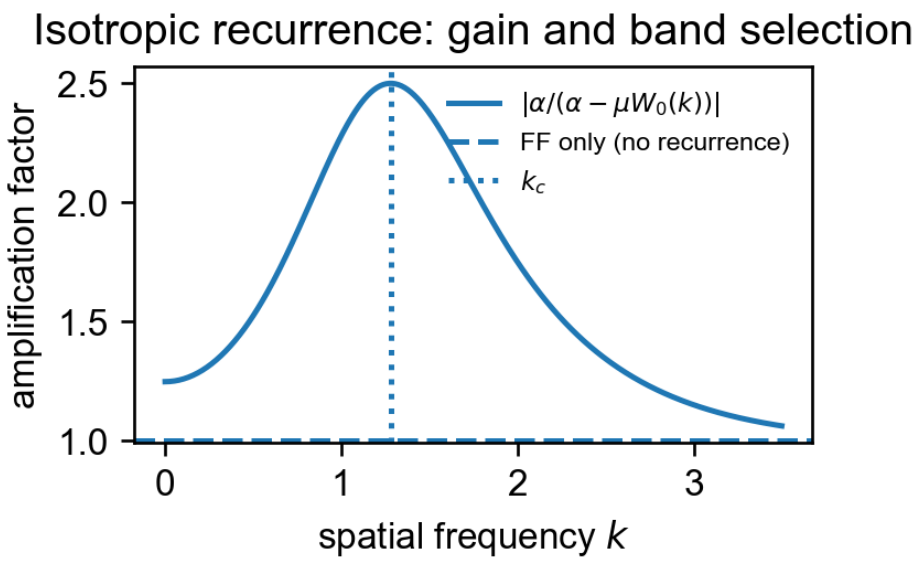

Fig 1. Isotropic recurrence reshapes gain versus spatial frequency. Isotropic recurrence cannot change normalized orientation metrics at fixed (k, ω), but it can strongly reshape the spatial-frequency dependence of response amplitude. Shown is the closed-loop amplification factor $|\alpha/(\alpha - \mu W_0(k))|$ for an isotropic DoG kernel $W_0(k)$, with μ chosen so that $(\mu W_0(k_c))/\alpha = g_0$ at the band-selection point k_c. The horizontal dashed line indicates the feedforward-only baseline.

Weak anisotropy generates tuning and predicts linear-in-ε changes

Let $D_0(k, \omega) \equiv \alpha + i\omega - \mu W_0(k, \omega)$. An exact algebraic rearrangement yields

$$\hat{u}(k, \phi, \omega) = \frac{\hat{u}_{\text{iso}}(k, \phi, \omega)}{1 - \varepsilon \Gamma(k, \omega) \cos 2(\phi - \phi_0)}, \tag{11}$$

where $\hat{u}_{\text{iso}}$ is the solution with $\varepsilon = 0$ and $\Gamma(k, \omega) \equiv \mu W_2(k, \omega)/D_0(k, \omega)$. Equation (11) is exact; the small-ε approximation enters only when expanding $|\hat{u}|$ and normalized metrics for $|\varepsilon\Gamma| \ll 1$. To first order,

$$R(\phi) \approx R_{\text{iso}}(\phi) \left[1 + \varepsilon \operatorname{Re} \Gamma(k, \omega) \cos 2(\phi - \phi_0)\right], \tag{12}$$

implying that changes in smooth normalized metrics scale linearly with ε for sufficiently weak anisotropy.

Two-timescale recurrence produces delayed sharpening and preferred-angle drift

A key conceptual case is $\eta_{\text{ff}} = 0$ (isotropic feedforward drive) with $\varepsilon \neq 0$ (anisotropic recurrence). In this "Case 2" regime, anisotropic recurrence alone generates orientation selectivity, whereas isotropic recurrence cannot. Separating a fast isotropic pathway and a slower anisotropic pathway yields delayed tuning changes and modest preferred-angle drift. Figure 2 shows the time course of circular variance $\text{CV}(t)$ (solid curves) and preferred orientation $\phi^*(t)$ (dotted curve; computed with a weak initial feedforward bias to define the early axis).

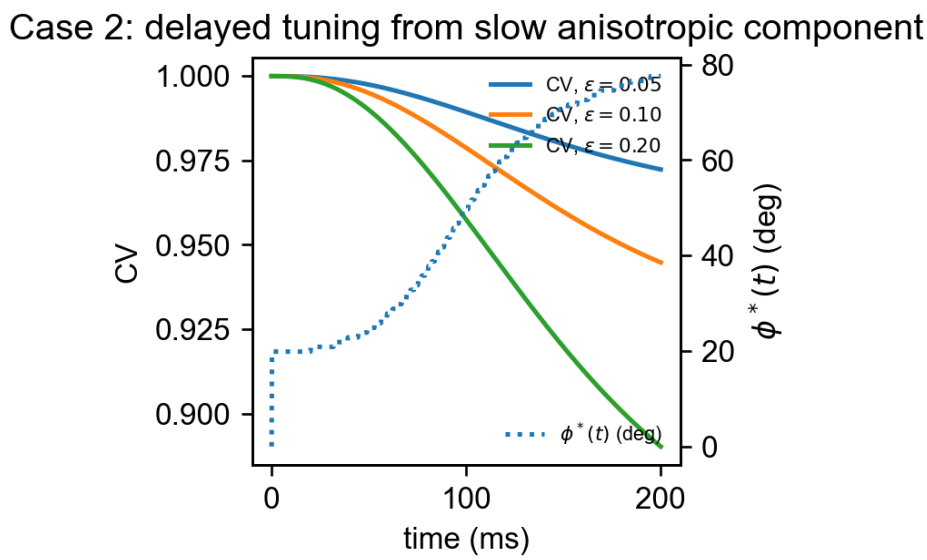

Fig 2. Case 2 dynamics with two timescales.
Solid curves: $\text{CV}(t)$ for several anisotropy strengths ε (feedforward bias set to zero). Dotted curve: preferred orientation $\phi^*(t)$ for a weak feedforward bias ($\eta_{\text{ff}} = 0.05$) and mismatched axes $\phi_{\text{ff}} = 20°$, $\phi_0 = 90°$, demonstrating drift via phasor competition. Parameters: $\alpha = 100\ \text{s}^{-1}$, $g_0 = 0.6$, $\tau_{\text{iso}} = 10$ ms, $\tau_{\text{aniso}} = 40$ ms.

When both feedforward and recurrent components carry (possibly misaligned) anisotropy, their second-harmonic components add as phasors. To leading order the instantaneous second-harmonic coefficient can be written as $A(t) = \eta_{\text{ff}} e^{i2\phi_{\text{ff}}} + \varepsilon\, \gamma(t) e^{i2\phi_0}$, where $\gamma(t)$ is the time-dependent effective gain of the anisotropic pathway. The preferred orientation (modulo π) is then $\phi^*(t) = \frac{1}{2} \arg A(t)$, so $\phi^*(t)$ drifts whenever the relative weight of the two phasors changes over time.

Stability of temporally filtered recurrence

The auxiliary-variable implementation yields a linear ODE system whose stability can be analyzed via eigenvalues. A simple sufficient condition for stability in this two-filter architecture is $g_0(1 + |\varepsilon|) < 1$, corresponding to DC loop gain smaller than unity for the worst-case orientation. Figure 3 verifies numerically that the stability boundary closely tracks this bound over the parameter ranges explored.

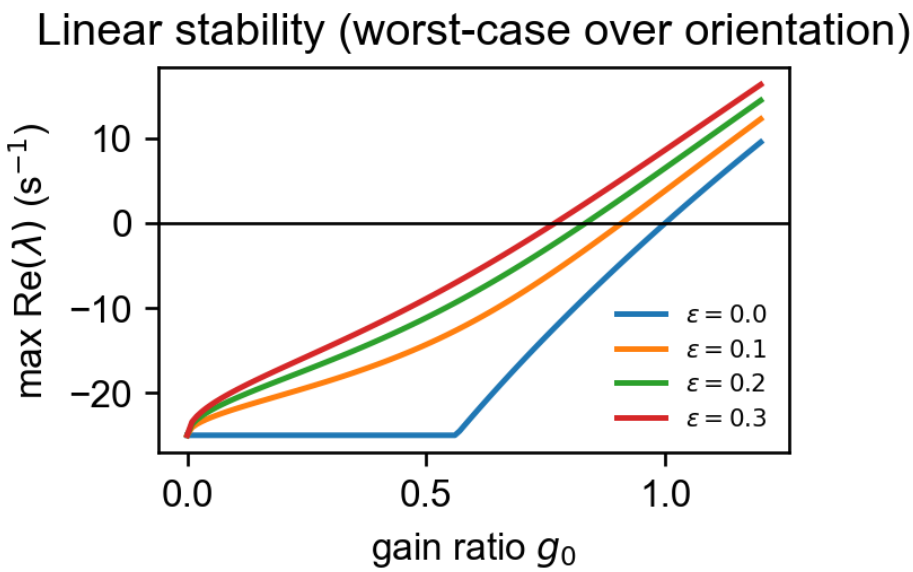

Fig 3. Worst-case linear stability margin versus gain.
Plotted: maximum real part of the eigenvalues of the auxiliary-variable system, maximized over $\cos 2(\phi - \phi_0) \in [-1, 1]$. Stability corresponds to negative values. The observed boundary agrees well with the sufficient condition $g_0(1 + |\varepsilon|) < 1$.

Parameter sweeps and latency mapping to macaque dynamics

Parameter sweeps over the slow anisotropic timescale τ_{aniso}, the fast isotropic timescale τ_{iso}, the closed-loop gain ratio g_0, and anisotropy strength ε quantify how tuning latency and amplitude depend on physiologically interpretable quantities. Figure 4 summarizes representative sweeps at $k = k_c$: (A) the tuning half-time t_{50} (time to reach 50% of the total ΔCV) and (B) the total tuning change ΔCV as functions of $(\tau_{\text{aniso}}, g_0)$ (tabulated values are provided in **S1 Data**). To connect to macaque tuning dynamics, we focus on an onset metric t_{10} (time to reach 10% of the total ΔCV).

Figure 5 reports the onset metric t_{10} (time to reach 10% of the total ΔCV) as a function of the slow anisotropic timescale τ_{aniso} and closed-loop gain ratio g_0 (computed at $k = k_c$). The horizontal band denotes a representative macaque V1 onset window (30–45 ms) reported in Refs. [15, 16], illustrating how τ_{aniso} in the tens-of-milliseconds range can account for observed tuning latencies depending on g_0.

When isotropic recurrence can affect measured tuning: finite-k readout

The invariance theorem applies at fixed (k, ω). In practice, measurements (and nonlinear readouts) often integrate responses over a finite spatial-frequency bandwidth. If feedforward anisotropy varies with k, isotropic recurrence can change apparent tuning by reweighting spatial frequencies. Figure 6 demonstrates this effect for a broadband readout $R_{\text{int}}(\phi) = \int S(k) R(k, \phi)\, dk$ under two scenarios: (i) factorable anisotropy $\eta(k) = \eta_0$ (invariance preserved) and (ii) k-dependent anisotropy peaked away from k_c (invariance broken by k-reweighting).

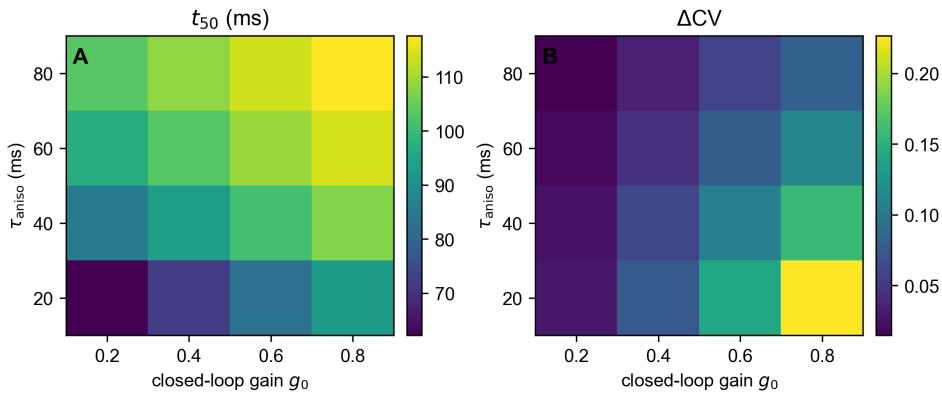

Fig 4. Parameter sweeps for tuning latency and amplitude (computed at $k = k_c$).
(A) Heatmap of tuning latency t_{50} (time to reach 50% of the total ΔCV) versus the slow anisotropic timescale τ_{aniso} and closed-loop gain ratio g_0. (B) Heatmap of total tuning change ΔCV versus the same parameters. Other parameters are fixed as in Figure 2, with $\varepsilon = 0.25$ for visualization.

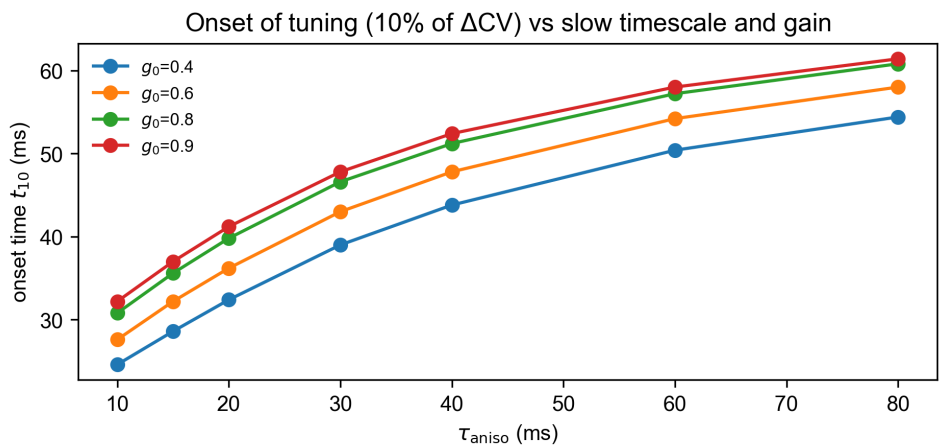

Fig 5. Onset time of tuning versus slow anisotropic timescale and gain.
Onset time t_{10} (time to reach 10% of the total ΔCV) versus the anisotropic timescale τ_{aniso} for several closed-loop gains g_0 (computed at $k = k_c$). The horizontal shaded band indicates a representative macaque V1 onset range (30–45 ms) from Refs. [15,16].

Nonlinear readouts: rectification and divisive normalization

The exact invariance theorem applies to the linear field variable $u(k, \phi, \omega)$. In experiments, tuning curves are estimated from firing rates and/or pooled population signals, which introduce nonlinear transformations. For a pointwise static nonlinearity $r = f(u)$ (e.g., rectification or a saturating sigmoid), isotropic recurrence can in principle change normalized metrics at fixed (k, ω) because a global gain change moves different orientations through different parts of the nonlinear input–output curve. In the small-signal regime $u \approx u_0$ where f is well approximated by its derivative, $r \approx f'(u_0)u$, and the linear invariance theorem is recovered.

To make the scope of this approximation explicit, write the firing rate as a smooth nonlinearity $r = f(u)$ and expand around an operating point u_0:
$f(u) \approx f(u_0) + f'(u_0)(u - u_0)$. A simple sufficient *small-signal* criterion is that the

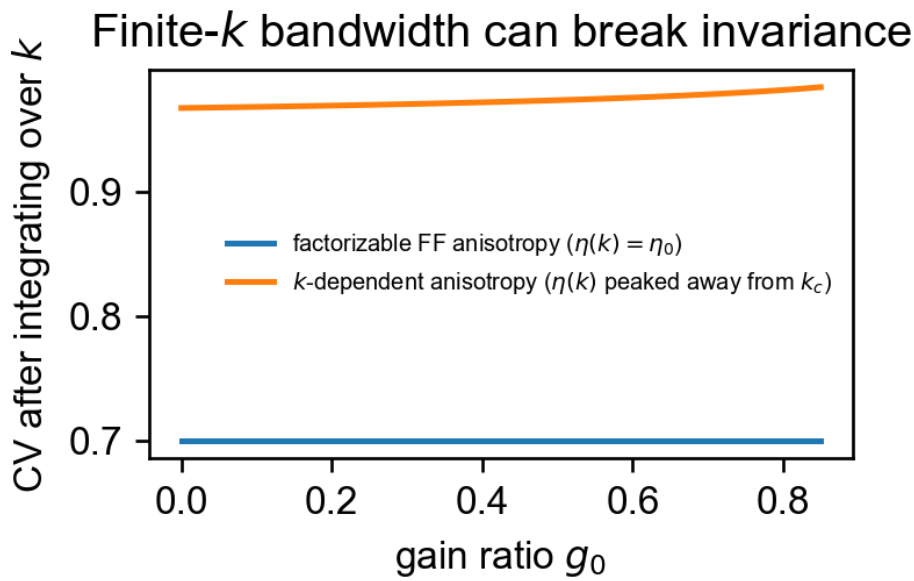

Fig 6. Finite-k readout can break invariance through k-reweighting. We compute CV from a broadband readout that integrates responses over a stimulus spectrum $S(k)$. If feedforward anisotropy factorizes ($\eta(k) = \eta_0$), CV remains independent of isotropic gain g_0. If $\eta(k)$ varies with k, isotropic recurrence changes the effective k-weighting and modifies CV.

curvature term is negligible over the relevant range of u,

$$|f''(u_0)| \max_t |u(t) - u_0| \ll |f'(u_0)|. \tag{13}$$

For sigmoidal nonlinearities, f'' is smallest near the inflection point (often near threshold), so Eq. (13) is most likely to hold when responses are weak (e.g., low contrast and/or early post-onset transients) and least likely to hold when responses explore the saturating regime. Operationally, the small-signal regime can be assessed by checking whether tuning curves at nearby contrasts (or nearby operating points) are approximately related by a scalar gain, in which case normalized metrics are expected to remain nearly invariant under isotropic gain modulation.

Divisive normalization is ubiquitous in V1 [13, 14] and can be implemented dynamically by recurrent circuits [24, 25]. If the normalization pool is untuned (constant in ϕ), the denominator in Eq. (9) is ϕ-independent and isotropic recurrence still cannot change normalized tuning metrics: any isotropic gain change rescales $r(\phi)$ by a ϕ-independent factor that cancels in normalized metrics. In contrast, a tuned normalization pool can produce *gain-dependent* changes in the angular profile when the normalization includes a non-homogeneous term such as the semi-saturation constant σ (or a baseline). In the idealized $\sigma \to 0$ limit (and for nonnegative responses), divisive normalization is homogeneous in gain and even a tuned pool becomes invariant to isotropic gain. Figure 7 illustrates these regimes.

Higher-harmonic anisotropy and generalized metrics

Standard OSI and CV are tied to the second angular Fourier component. If anisotropy is dominated by higher harmonics (e.g., $\cos 4\phi$), second-harmonic OSI can remain small

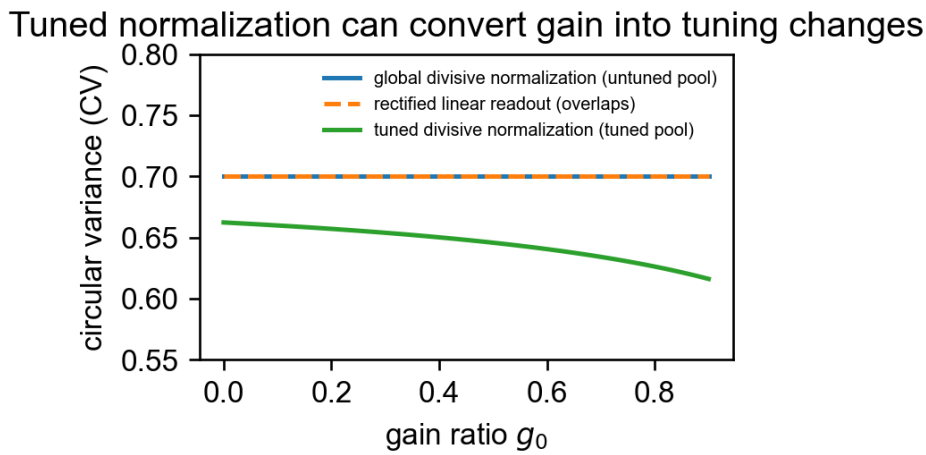

Fig 7. Tuned divisive normalization converts isotropic gain into tuning changes.
CV computed after applying a rectified readout (dashed) and global divisive normalization with an untuned pool (solid), which overlap exactly, and after divisive normalization with a tuned pool (green), as a function of isotropic gain ratio g_0. Gain-dependence in the tuned case arises from the semi-saturation constant $\sigma \neq 0$; parameters: $\eta = 0.6$, $\sigma = 5.0$, $\kappa = 1.0$, and $w_{\text{norm}}(\Delta\phi) = 1 + q\cos 2\Delta\phi$ with $q = -0.8$.

even when tuning is multi-peaked. Figure 8 illustrates this effect with a toy response and motivates harmonic-specific metrics when needed.

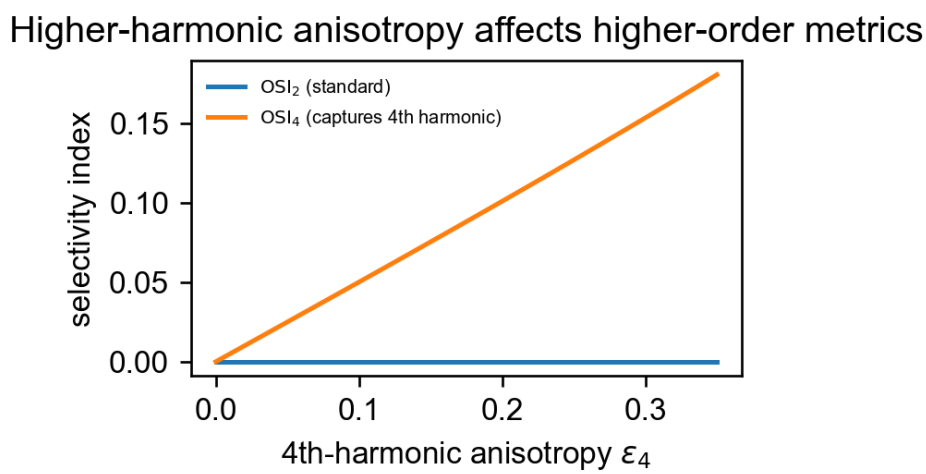

Fig 8. Higher-harmonic anisotropy can evade second-harmonic OSI.
For a toy response $R(\phi) \propto [1 - \varepsilon_4 \cos 4\phi]^{-1}$, the standard second-harmonic OSI remains near zero while a fourth-harmonic selectivity measure increases with ε_4.

Discussion

Our results provide a symmetry-based baseline for interpreting feedback/recurrent effects on orientation tuning dynamics in primate V1.

Exact null hypothesis for isotropic recurrence. Within a translation-invariant small-signal regime, isotropic recurrent coupling cannot change any normalized orientation metric at fixed (k, ω). This formalizes a common intuition and aligns with linearized spiking-network results showing that feature-unspecific recurrence amplifies but does not bias tuning [18]. It also clarifies why manipulations that primarily modulate gain or spatial summation may leave bandwidth unchanged at fixed k [6, 7].

Feature-specific coupling is necessary for tuning changes at fixed k. Any sharpening of normalized tuning at fixed k must arise from anisotropy in the effective recurrent kernel—whether due to corticocortical feedback, tuned suppression, or structured intracortical connectivity. This perspective is consistent with classical ring-model accounts in feature space [10–12] and with salt-and-pepper modeling studies arguing that feature-specific connectivity is required to match observed tuning distributions [19, 20].

Connection to classical ring models. In ring models, tuned recurrence in an orientation feature space is often written as a low-order harmonic expansion, $J(\Delta\theta) = J_0 + J_2 \cos 2\Delta\theta + \cdots$ [10–12]. In our spatial formulation, an analogous harmonic structure appears in the *angular dependence* of the Fourier-space kernel $\hat{w}_{\mathrm{rec}}(k, \phi, \omega)$: the isotropic part is independent of ϕ (the analogue of an untuned component), while the leading anisotropic contribution is a $\cos 2(\phi - \phi_0)$ term. The invariance theorem shows that a purely ϕ-independent loop cannot reshape normalized tuning at fixed k in the translation-invariant linear regime; any sharpening at fixed k therefore requires ϕ-dependent (feature-specific) coupling, consistent with the key assumption built into ring-model mechanisms.

Two-timescale mechanism for delayed tuning and drift. By separating a fast isotropic pathway and a slower anisotropic pathway, the model naturally produces delayed sharpening and modest preferred-angle drift. The phasor-sum formula makes drift predictions transparent: drift is expected when early tuning is set by feedforward bias but later tuning is dominated by slow anisotropic recurrence. This offers an analytically grounded mechanism for aspects of the tuning dynamics reported in macaque V1 [15–17].

When invariance may fail in practice. The invariance theorem is exact under its assumptions, but both biology and experimental readout can violate those assumptions in concrete ways. First, common tuning measurements effectively integrate responses over a finite spatial-frequency bandwidth; if feedforward anisotropy varies with k, isotropic recurrence can change apparent tuning by k-reweighting (Fig 6). Second, nonlinearities can couple orientations and convert a gain change into a shape change. Pointwise nonlinearities can break invariance when responses traverse a nonlinear operating regime; near a baseline operating point they reduce to the linearized regime of Eq. (1). Divisive normalization is ubiquitous [13, 14] and can be implemented by recurrent "normalization circuits" [24, 25]; as shown in Fig 7, untuned pools preserve invariance while tuned pools [26] can break it even at a fixed Fourier mode. Third, strong nonlinear recurrent regimes (e.g., stabilized supralinear networks) combine tuned and untuned components and can lie beyond the small-signal regime emphasized here [27]. Finally, spatial inhomogeneities (orientation maps, pinwheels) couple Fourier modes and can invalidate strict translation invariance.

Relation to modern data-driven V1 models. Recent system-identification models that incorporate explicit contextual modulation mechanisms (e.g., self-attention modules) can improve prediction of V1 responses compared to purely local feedforward baselines [28]. While these models are not designed to isolate Euclidean symmetry, their success is consistent with the central message here: capturing dynamic changes in tuning and contextual effects generally requires interactions that are feature-specific and/or nonlocal, rather than purely isotropic gain control.

Conclusion

We derived an exact symmetry constraint on what isotropic recurrence can do to orientation tuning at fixed (k, ω) in translation-invariant small-signal models: it rescales responses by an angle-independent gain and therefore cannot change any normalized tuning metric at fixed k. We showed that weak anisotropy, especially with a slower timescale than an untuned pathway, is sufficient to generate delayed tuning changes and preferred-angle drift. We also identified practical routes by which isotropic recurrence can affect measured tuning—finite-k readouts with k-dependent feedforward anisotropy and tuned divisive normalization—thereby connecting the theorem to common experimental analyses.

Supporting information

S1 Data. This includes all python scripts used to generate figures and run parameter sweeps, CSV files containing the simulated time courses, stability curves, and parameter-sweep summaries used to generate all main-text figures.

Acknowledgments

A portion of this work was originally submitted as the master's thesis of the first author in the Division of Physical Sciences at the University of Chicago. We gratefully acknowledge Prof. Jack Cowan for supervising this project.

References

1. Hubel DH, Wiesel TN. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. Journal of Physiology. 1962;160:106-54.
2. Ferster D, Miller KD. Neural mechanisms of orientation selectivity in the visual cortex. Annual Review of Neuroscience. 2000;23:441-71.
3. Chance FS, Nelson SB, Abbott LF. Complex cells as cortically amplified simple cells. Nature Neuroscience. 1999;2:277-82.
4. Angelucci A, Bullier J. Reaching beyond the classical receptive field of V1 neurons: horizontal or feedback axons? Journal of Physiology Paris. 2003;97:141-54.
5. Angelucci A, et al. Circuits and mechanisms for surround modulation in visual cortex. Annual Review of Neuroscience. 2017;40:425-51.
6. Nurminen L, Merlin S, Bijanzadeh M, Federer F, Angelucci A. Top-down feedback controls spatial summation and response amplitude in primate visual cortex. Nature Communications. 2018;9:2281. doi:10.1038/s41467-018-04500-5.
7. Federer F, Ta' afua SF, Merlin S, Hassanpour MS, Angelucci A. Stream-specific feedback inputs to the primate primary visual cortex. Nature Communications. 2021;12:228. doi:10.1038/s41467-020-20505-5.
8. Korkian Y, Nakhla N, Pack CC. Feature selectivity of corticocortical feedback along the primate dorsal visual pathway. Journal of Neurophysiology. 2025;133:799-814.

9. Xin Y, Yan Y, Li W. A central and unified role of corticocortical feedback in parsing visual scenes. Nature Communications. 2025;16:6930.

10. Ben-Yishai R, Bar-Or RL, Sompolinsky H. Theory of orientation tuning in visual cortex. Proceedings of the National Academy of Sciences of the USA. 1995;92:3844-8.

11. Somers DC, Nelson SB, Sur M. An emergent model of orientation selectivity in cat visual cortical simple cells. Journal of Neuroscience. 1995;15:5448-65.

12. Hansel D, Sompolinsky H. Modeling feature selectivity in local cortical circuits. In: Koch C, Segev I, editors. Methods in Neuronal Modeling: From Ions to Networks. 2nd ed. Cambridge: MIT Press; 1998. p. 499-566.

13. Heeger DJ. Normalization of cell responses in cat striate cortex. Visual Neuroscience. 1992;9:181-97.

14. Carandini M, Heeger DJ. Normalization as a canonical neural computation. Nature Reviews Neuroscience. 2012;13:51-62.

15. Ringach DL, Hawken MJ, Shapley RM. Dynamics of orientation tuning in macaque V1: the role of global and tuned suppression. Journal of Neurophysiology. 2003;90:342-52. doi:10.1152/jn.01018.2002.

16. Ringach DL, Hawken MJ, Shapley RM. Dynamics of orientation tuning in macaque primary visual cortex. Nature. 1997;387:281-4. doi:10.1038/387281a0.

17. Wang T, et al. Laminar subnetworks of response suppression in macaque primary visual cortex. Journal of Neuroscience. 2020;40:7436-50. doi:10.1523/JNEUROSCI.1129-20.2020.

18. Sadeh S, Rotter S. Distribution of orientation selectivity in recurrent networks of spiking neurons with different random topologies. PLOS ONE. 2014;9:e114237. doi:10.1371/journal.pone.0114237.

19. Roy D, et al.. Afferent specificity, feature specific connectivity influence orientation selectivity: A computational study in mouse primary visual cortex; 2013. arXiv:1301.0996.

20. Fu J, et al. Heterogeneous orientation tuning in the primary visual cortex of mice diverges from Gabor-like receptive fields in primates. Cell Reports. 2024;43:114639.

21. Cho MW, Abbott RJ, Sederberg AM. A symmetries approach to understanding the emergence of orientation selectivity in visual cortex: a model with spike timing dependent plasticity; 2015. arXiv:1503.04620.

22. Schmidt M, et al. A multi-scale layer-resolved spiking network model of macaque visual cortex. PLOS Computational Biology. 2018;14:e1006359.

23. Wilson HR, Cowan JD. Excitatory and inhibitory interactions in localized populations of model neurons. Biophysical Journal. 1972;12:1-24.

24. Heeger DJ, Zemlianova KO. A recurrent circuit implements normalization, simulating the dynamics of V1 activity. Proceedings of the National Academy of Sciences of the USA. 2020;117:22494-505.

25. Heeger DJ, Mackey WE. Oscillatory recurrent gated neural integrator circuits (ORGaNICs), a unifying theoretical framework for neural dynamics. Proceedings of the National Academy of Sciences of the USA. 2019;116:22783-94.

26. Fang Z, Bloem IM, Olsson C, Ma WJ, Winawer J. Normalization by orientation-tuned surround in human V1–V3. PLOS Computational Biology. 2023;19:e1011704. doi:10.1371/journal.pcbi.1011704.

27. Rubin DB, Van Hooser SD, Miller KD. The stabilized supralinear network: a unifying circuit motif underlying multi-input integration in sensory cortex. Neuron. 2015;85:402-17.

28. Lin I, Wang T, Gao S, Tang S, Lee TS. Self-attention-based contextual modulation improves neural system identification; 2025. arXiv:2406.07843.