

**Field Fibers:
Neural States That Share a Mesoscale Electric Field, and
How to Detect Field Stabilization**

PIATRA . INSTITUTE

September 2026

Abstract

The neurons that carry a memory change from day to day while behaviour stays the same, and recent work proposes that mesoscale electric fields act as guard rails that keep the representation stable. We make the proposal geometric. The field fiber of a state is the set of neural states that produce the same field. For current-dipole sources in two laminae beneath a recording plane, the field map is diagonal in spatial frequency; its exact fiber is the redistribution of activity within each patch plus a laminar ambiguity of one direction per spatial mode, 87.5% of all state directions. Sensor noise of 1% on a plane 0.25 mm above the sources hides every spatial mode with wavelength below 0.37 mm and enlarges the fiber to 95.4%. Fine binary patterns become invisible exponentially fast, with field error $e^{-2\pi mh/L}$ for m periods. In a linear stochastic model a stable field over drifting neurons arises without feedback, because the field averages fluctuations away. Field feedback through the lead field cuts field variance to 0.235 of its passive value and leaves neural drift almost unchanged. A structured-noise model with no feedback reproduces the feedback model’s equal-time covariance exactly, so the anticorrelation between variance and field sensitivity cannot detect feedback. Lagged field autocorrelation and matched perturbations can: an equal-norm kick relaxes at rate 51 along a field-sensitive direction and 1 along the fiber. Ephaptic neural fields with Lipschitz rates are well posed, so convex integration can enter only through discontinuous or set-valued dynamics.

Introduction

Neurons that encode a learned variable are replaced by other neurons over days to weeks while the animal’s behaviour stays the same. In mouse hippocampus each day’s representation of a familiar environment used a unique subset of CA1 cells, and the 15–25% of cells shared between days kept their place fields (Ziv et al. 2013). In posterior parietal cortex the cells informative about a navigation task changed across days while the population kept statistically similar properties (Driscoll et al. 2017). Drift has since been reported in visual cortex across areas, layers and cell types (Deitch et al. 2021) and in piriform cortex, where a classifier trained on the first day approached chance after 32 days (Schoonover et al. 2021). Reviews ask how stable behaviour survives a changing code and what drift could be for (Rule et al. 2019; Driscoll et al. 2022). Several answers share a structure: the code is degenerate, so many population states support the same readout, and noise or plasticity moves the state within that degenerate set. Drift in posterior parietal cortex is constrained far above chance to directions that preserve a linear readout, although a component of it degrades a fixed readout (Rule et al. 2020). Hebbian/anti-Hebbian networks whose objective has a degenerate solution space drift by a coordinated random walk through it (Qin et al. 2023). Low-dimensional latent dynamics in motor cortex stay stable for up to two years while recorded neurons turn over (Gallego et al. 2020). Related ideas appear in output-null spaces that let motor cortex prepare without moving (Kaufman et al. 2014), in stable subspaces of heterogeneous prefrontal working-memory activity (Murray et al. 2017; Druckmann and Chklovskii 2012), and, much earlier, in the degeneracy of biological systems and neural circuits (Edelman and Gally 2001; Prinz et al. 2004; Marder and Goaillard 2006).

Earl K. Miller and Dimitris Pinotsis have proposed a physical candidate for the degenerate set. In their analysis of prefrontal recordings, electric fields carried information about working-memory content and could act as guard rails that funnel higher-dimensional variable neural activity along stable lower-dimensional routes (Pinotsis and Miller 2022). Subsequent work reported evidence for in vivo ephaptic coupling in memory representations (Pinotsis and Miller 2023), proposed a cytoelectric coupling hypothesis in which neural fields are causal down to the cytoskeleton (Pinotsis et al. 2023), and found that field-to-neuron interactions were stronger than neuron-to-field interactions and correlated with trial-by-trial changes in oscillatory power (Pinotsis and Miller 2026). The proposal builds on a programme that places working memory in brief oscillatory bursts and layer-specific rhythms (Lundqvist et al. 2016; Miller et al. 2018). That fields act back on neurons is established at the cellular level: extracellular fields changed somatic potentials by less than 0.5 mV yet strongly entrained spiking, particularly for fluctuations slower than 8 Hz (Anastassiou et al. 2011); weak fields within the endogenous range entrained the neocortical slow oscillation, and feedback fields demonstrated a loop between network activity and field (Fröhlich and McCormick 2010); slow hippocampal activity propagated without synaptic transmission or gap junctions by a mechanism consistent with ephaptic coupling (Chiang et al. 2019; see also Jefferys 1995; Anastassiou and Koch 2015).

The guard-rail picture contains a geometric claim that can be stated exactly. Write the neural state as $z \in \mathbb{R}^N$ and the field it generates as $H(z)$. The neural states that produce a given field form the set $\mathcal{F}_E = \{z : H(z) = E\}$, which we call the field fiber. If the field is the variable that matters and feedback holds it fixed, drift should move the state within the fiber while the field stays put. We ask three questions about this picture. How large is the fiber for a physically specified field map, and how much of it comes from the physics of volume conduction, from measurement noise and from the number of electrodes? Which statistics of neural and field data can distinguish field stabilization from two simpler explanations of a stable field over drifting neurons, passive averaging and structured noise? And where in this geometry could convex integration, a method for constructing many wildly oscillating solutions of a single differential relation (Nash 1954; Gromov 1986; De Lellis and Székelyhidi 2009), have anything to say? The answers come from a quasistatic forward model with explicit singular values, a linear stochastic model with closed-form statistics, and simulations that check every closed form. The models are minimal and fitted to no recording. No result here shows that any neural population drifts along a field fiber, and the published electric-field studies enter only through what their abstracts report.

A forward model and its fibers

Sources, field and sensors

Extracellular potentials in cortex arise mostly from transmembrane currents of aligned pyramidal cells, and at the frequencies of the local field potential gray matter behaves as a homogeneous, frequency-independent resistive conductor (Buzsáki et al. 2012; Einevoll et al. 2013; Logothetis et al. 2007). In the quasistatic approximation the potential is linear in the currents, so each neural state produces a field by superposition (Nunez and Srinivasan 2006; Hämmäläinen et al. 1993). We model

a square patch of cortex of side $L = 4$ mm with periodic boundaries, divided into 32×32 patches of 0.125 mm. Each patch holds two laminae of current-dipole sources separated by $d = 0.6$ mm, standing for superficial and deep current sinks of the same columns, and each lamina in each patch holds $c = 4$ cell groups. The state $z \in \mathbb{R}^N$ with $N = 2 \cdot 32^2 \cdot 4 = 8,192$ lists the dipole amplitudes of all groups. A dipole layer with Fourier amplitude $\hat{p}(k)$ produces at height z_0 above it a potential with Fourier amplitude proportional to $\hat{p}(k)e^{-kz_0}$, where $k = |\mathbf{k}|$ is the spatial wavenumber. The recording plane lies at height h above the superficial lamina, with $h = 0.25$ mm in the reference case.

Proposition 1. Let S_1 and S_2 be the patch sums of the two laminae, scaled by $c^{-1/2}$. The field map is $\hat{\phi}(\mathbf{k}) = (e^{-kh}\hat{S}_1(\mathbf{k}) + e^{-k(h+d)}\hat{S}_2(\mathbf{k}))/\sqrt{2}$, and its nonzero singular values are

$$s(k) = e^{-kh} \sqrt{\frac{1}{2}(1 + e^{-2kd})},$$

one for each spatial mode, normalized so that $s(0) = 1$. The map has rank $32^2 = 1,024$, and its kernel, the exact fiber, has dimension $N - 1,024 = 7,168$, 87.5% of all state directions. The kernel is the direct sum of the contrasts among cell groups within a patch (dimension $2 \cdot 32^2 \cdot 3 = 6,144$) and a laminar ambiguity of one direction per spatial mode, $\hat{S}_1 = -e^{-kd}\hat{S}_2$ (dimension 1,024), which produces zero potential on every plane above the superficial lamina.

Proof. The patch sums are an orthogonal projection whose kernel is the within-patch contrasts. The Fourier transform diagonalizes the translation-invariant transfer, and for each mode the map from $(\hat{S}_1, \hat{S}_2)$ to $\hat{\phi}$ is the row vector $(e^{-kh}, e^{-k(h+d)})/\sqrt{2}$, of norm $s(k)$, whose kernel is spanned by $(e^{-kd}, -1)$. The potential of that kernel direction at height z_0 is proportional to $e^{-kz_0}(\hat{S}_1 + e^{-kd}\hat{S}_2) = 0$ for every $z_0 > 0$. $\square$

The laminar ambiguity is the planar case of Helmholtz's observation that sources inside a volume conductor cannot be recovered from external potentials (Helmholtz 1853; Hämäläinen et al. 1993). A random deep-lamina source paired with its laminar counterpart produces a potential at most $2.0\text{e-}16$ of the deep source's own potential on every plane tested. The within-patch part is the part that matters for drift: cell groups in the same patch with the same dipole geometry are interchangeable at any distance, so the field cannot tell which of them carries the current. Explicit construction of the 256×256 map for an 8×8 instance confirms the rank, a fiber of dimension 192 and the singular values of Proposition 1 to $1.5\text{e-}15$.

Noise and electrodes

With sensor noise ε relative to the strongest mode, a spatial mode is measurable only if $s(k) \geq \varepsilon$. For $kd \gg 1$ this gives the cutoff $k_\varepsilon \approx \ln(1/\sqrt{2}\varepsilon)/h$, and the number of measurable modes is close to $\pi(k_\varepsilon L/2\pi)^2$ whenever the cutoff lies well below the patch-grid Nyquist wavenumber (within 13% across the cases tested). At $\varepsilon = 0.01$ and $h = 0.25$ mm the cutoff is $k_\varepsilon = 17.03$ rad/mm, a wavelength of 0.37 mm; 373 of the 1,024 field modes remain, and the noise-limited fiber covers

95.4% of state directions (Figure 1a). Raising the plane to $h = 1$ mm leaves 21 modes at the same noise level.

Electrodes add a third layer. A grid of M electrodes resolves at most M directions, so the measurement fiber contains the physical fiber and every field direction the grid misses. A 10×10 grid with 0.4 mm spacing resolves 100 directions and leaves a measurement fiber of 98.8% (Figure 1b). Three nested sets therefore have to be kept apart: the exact fiber fixed by the physics of the field, the larger fiber fixed by noise, and the measurement fiber fixed by electrode count. Only the first is a property of the brain. A demonstration that many neural states fit the same electrode recordings mostly measures the third.

Fine structure the field cannot see

The field map smooths. Its kernel e^{-kh} makes it a compact operator, and compact operators turn weak convergence into strong convergence: if a sequence of states q_m converges weakly to $\bar{q}$, the fields $H(q_m)$ converge to $H(\bar{q})$ in norm. Any microstructure that averages out at fine scales is therefore invisible in the limit, whatever its amplitude.

A concrete case shows how fast this happens. Take a graded coarse profile $\bar{q}(x)$ with values between 0.05 and 0.95 and realize it by a binary microstate $q_m(x) \in \{0, 1\}$ that is active on the first fraction $\bar{q}$ of each of m periods across the domain, the construction of naturally sampled pulse-width modulation (Figure 1c). The binary state has the right coarse average: at $m = 128$ its patch averages differ from those of $\bar{q}$ by at most 0.0081. Its field differs from that of $\bar{q}$ only through the carrier sidebands near wavenumber $2\pi m/L$, so the relative field error decays as $e^{-2\pi m h/L}$. Fitted slopes are -0.15708 per period at $h = 0.1$ mm and -0.39268 at $h = 0.25$ mm, against $-2\pi h/L = -0.15708$ and -0.39270 ; at $h = 0.25$ mm the error is $2.3\text{e-}3$ at $m = 16$ and $4.3\text{e-}6$ at $m = 32$ (Figure 1d). Two binary realizations at $m = 64$ that differ by a half-period shift are far apart as states (relative distance 1.38) and have fields that differ by $3.0\text{e-}11$.

In the language of Young measures (Ball 1989; Müller 1999), a sequence of microstates converges to a family of probability distributions ν_x over local states. Because the quasistatic field is linear in the currents, the limiting field depends on ν_x only through its barycenter: $E = H(\int q d\nu_x)$. If the currents are a nonlinear function S of the local state, the field depends on $\int S(q) d\nu_x$, still a single moment of the distribution. Every pair of microstate distributions with the same mean current produces the same field. In this sense the field fiber of a coarse state is enormous: it contains every fine-scale arrangement of activity, binary or graded, with the right local averages, the fine phase mixtures familiar from the theory of microstructure (Ball and James 1987).

Dynamics with and without field feedback

The model

Let the state fluctuate around a memory pattern z^* with leak rate a and white noise of intensity q , and let the field feed back:




Figure 1: Geometry of the field fiber. (a) Singular value $s(k)$ of the field map for recording planes at four heights above the superficial lamina, with noise levels ε of 0.1, 0.01 and 0.001 and the Nyquist wavenumber of the patch grid. (b) Number of state directions resolved by grids of 16 to 1,024 electrodes at two noise levels, with the exact field rank and the noise-limited count for $\varepsilon = 0.01$. (c) A binary microstate with 16 periods across the domain whose local duty cycle follows a graded profile $\bar{q}$. (d) Relative field error of binary microstates against the number of periods m , at two plane heights, with the predicted slope $-2\pi h/L$.

$$dz = -[a(z - z^*) + gH^T H(z - z^*)] dt + \sqrt{q} dW.$$

The feedback term is the gradient of the field error $\frac{1}{2}g|H(z - z^*)|^2$: it corrects deviations that change the field and leaves the fiber free, the structure that optimal feedback control predicts for redundant motor tasks (Todorov and Jordan 2002) and that the uncontrolled-manifold analysis looks for in motor variability (Scholz and Schöner 1999). Using H^T for the coupling from field to sources is natural by reciprocity, since the adjoint of a lead field describes how a field pattern at the sensors maps back onto the sources (Plonsey 1963), but none of the results below depends on its biophysical exactness. We take $a = 1$, $g = 50$ and $q = 2$, so that unstabilized directions have unit stationary variance, and a smooth memory pattern (a Gaussian bump of width 0.5 mm in both laminae) with $|z^*|^2$ equal to 0.3 of the total stationary variance.

Proposition 2. The fluctuation $z - z^*$ is an Ornstein–Uhlenbeck process (Gardiner 2009). Along a right singular direction of H with singular value s it has relaxation rate $a + gs^2$ and stationary variance $q/[2(a + gs^2)]$; along the fiber the rate is a and the variance $q/2a$. The stationary covariance is $\Sigma = \frac{q}{2}(aI + gH^T H)^{-1}$. The neural similarity between states separated by a lag Δ ,

$$S_z(\Delta) = \frac{|z^*|^2 + \sum_j \frac{q}{2\lambda_j} e^{-\lambda_j \Delta}}{|z^*|^2 + \sum_j \frac{q}{2\lambda_j}},$$

with λ_j the rates of all N directions, and the analogous field similarity $S_E(\Delta)$ with each term weighted by s_j^2 , follow in closed form.

A stable field over drifting neurons needs no feedback

Without feedback ($g = 0$), the neural similarity falls from 0.977 at a lag of 0.03 to 0.231 at long lags, while the field similarity levels off at 0.993 (Figure 2a). With feedback the neural similarity at long lags is 0.232 and the field similarity 0.998. Feedback cuts the variance of field fluctuations to 0.235 of its passive value and changes neural drift by less than one part in a hundred. The stable field of the passive model comes from averaging: the field map collapses 8,192 independent fluctuations onto a handful of strongly weighted spatial modes, the field variance is the sum of squared singular values, 5.89, against a squared field norm of 807 for the memory pattern. Sampled trajectories reproduce both similarity curves to 7.0e-4. “Stable field, drifting neurons” is therefore what volume conduction produces from any population whose fluctuations are spread over many neurons, and observing it says little about feedback.

Statistics that separate field stabilization from its rivals

Equal-time covariance cannot

Field stabilization predicts that directions with high field sensitivity carry little variance. In the feedback model the stationary variance along a direction falls monotonically with its sensitivity s ,

a rank correlation of -0.99999 across field modes, the shape the uncontrolled-manifold analysis tests in motor data (Scholz and Schöner 1999).

Proposition 3. The model with no feedback and structured noise, $dz = -a(z - z^*) dt + \sqrt{2a\Sigma} dW$ with Σ the feedback model’s stationary covariance, has the same stationary covariance as the feedback model.

Proof. Σ solves the Lyapunov equation $-a\Sigma - a\Sigma + 2a\Sigma = 0$ of the structured model. $\square$

Any statistic of single time points, including the anticorrelation of variance with field sensitivity, the principal components of neural variability and their alignment with the field map, is identical in the two models (Figure 2c). The structured model needs no field at all: its noise is merely smaller in field-sensitive directions, as it would be if upstream inputs were themselves organized around the field’s spatial modes. Simulated with an exact discretization on the explicit 8×8 instance, the two models give variances within 1.7% of the closed forms along all 64 field-sensitive directions and within 1.3% along the fiber.

Lagged field statistics and matched perturbations can

The models differ in how fast field-sensitive fluctuations relax. In the feedback model the uniform field mode relaxes at $a + g = 51$, and in the structured and passive models every direction relaxes at $a = 1$. The simulations recover a rate of 51.1 for the uniform mode with feedback and a mean of 1.04 over the field-sensitive directions of the structured model. At the level of the measured field, the autocorrelation of field fluctuations at a lag of 0.1 is 0.905 in the passive and structured models and 0.721 with feedback (Figure 2b). The contrast is largest for the best-measured, long-wavelength modes, since feedback acts in proportion to s^2 and fine modes relax almost as slowly as the fiber.

A perturbation experiment sharpens the test. Two kicks of equal norm, one along the uniform field mode and one along a within-patch contrast, decay at rates 51.0 and 1.0 with feedback and at the same rate without it. A perturbation that changes the field and one that does not, matched in neural magnitude, are the most direct probe of a guard rail, and the prediction is a ratio of relaxation rates set by the feedback gain.

Nullspace alignment needs its baseline

A natural test asks what fraction of drift lies in directions the recordings cannot see. That fraction is large under any hypothesis. For increments over a unit lag and a grid of 100 electrodes, the passive model puts 0.988 of drift energy in the measurement nullspace, exactly the fraction of state directions the grid cannot resolve, and the feedback model puts 0.992 there (Figure 2d). The informative quantity is the excess over the isotropic baseline, $(A - A_0)/(1 - A_0)$, which is 0 without feedback and 0.369 with it at 100 electrodes. Against the exact physical fiber the excess with feedback is only 0.040, because most field modes have small singular values and feedback barely stabilizes them. Drift in posterior parietal cortex was found to be constrained toward readout-

preserving directions while still degrading a fixed readout (Rule et al. 2020), a pattern compatible with partial stabilization of this kind.

Sufficiency of the field depends on the readout

The strong field hypothesis predicts that once the field is known the remaining neural detail adds nothing about behaviour, $I(B; z | E) \approx 0$. For a Gaussian readout $B = w^T z + \text{noise}$ with noise standard deviation σ_B , $I(B; z | E) = \frac{1}{2} \ln[(w^T \Sigma_z w + \sigma_B^2) / \sigma_B^2]$. A readout along the uniform field mode gives zero; a readout with half its weight on a within-patch contrast gives 0.549 nats at $\sigma_B = 0.5$, with or without feedback, and a regression estimate from 40,000 samples gives 0.546. Conditional information measures where the readout lies relative to the field's row space. It is silent on whether the field is stabilized, and it should be paired with the dynamical tests above.




Figure 2: Field stabilization and its rivals in the linear model. (a) Similarity of neural states and of fields between sessions separated by a lag, with and without field feedback; the neural curves of the two models nearly coincide. (b) Autocorrelation of field fluctuations against lag for the passive model, the structured-noise model with the feedback model's covariance, and the feedback model. (c) Stationary variance along a direction against its field sensitivity s : the feedback and structured-noise models coincide. (d) Fraction of drift energy over a unit lag lying in the measurement nullspace of electrode grids, with the isotropic baseline.

Where convex integration could enter

Three kinds of nonuniqueness have appeared. Measurement nonuniqueness belongs to the electrodes. Coarse-graining degeneracy, the exact and noise-limited fibers, belongs to volume conduction. The third kind, dynamical nonuniqueness, would mean that the same initial state and the same inputs admit many solutions of the same equations. Convex integration constructs solutions of this third kind. It began with Nash's C^1 isometric embeddings (Nash 1954) and Gromov's h-principle (Gromov 1986); applied to the Euler equations written as a differential inclusion, it produces weak solutions with compact support in time (De Lellis and Székelyhidi 2009), continuous dissipative flows (De Lellis and Székelyhidi 2013) and, in the end, dissipative solutions below the Onsager exponent $1/3$ (Isett 2018), together with nonunique weak solutions of the Navier–Stokes equations (Buckmaster and Vicol 2019; see the survey Buckmaster and Vicol 2020). Related constructions give Lipschitz solutions of elliptic systems with no regularity (Müller and Šverák 2003). In every case the equations are underdetermined pointwise: a linear differential constraint coupled to a nonconvex set of admissible local states, with room for oscillations along the directions of the wave cone.

Proposition 4. Consider the ephaptic neural field $\tau \partial_t V = -V + w * f(V + \alpha \phi) + I$ with $\phi = \mathcal{G}[V]$, where w is an integrable kernel, $\mathcal{G}$ is a bounded linear field operator and f is Lipschitz. For every bounded initial condition the equation has a unique global solution in L^∞ .

Proof. The right-hand side is globally Lipschitz on L^∞ , with constant $(1 + \|w\|_1 \text{Lip}(f)(1 + |\alpha| \|\mathcal{G}\|))/\tau$, and the Picard–Lindelöf theorem in a Banach space gives a unique global solution. $\square$

The standard Wilson–Cowan and Amari neural fields with smooth rate functions (Wilson and Cowan 1972; Amari 1977; Coombes 2005; Bressloff 2012) are therefore well posed with or without ephaptic coupling, and the dynamical nonuniqueness that convex integration produces cannot occur in them. For a Heaviside rate function the fixed-point argument fails, and existence needs compactness arguments in a weaker space (Potthast and Beim Graben 2010); this is where set-valued dynamics, and with them convex-integration constructions, become conceivable. Kinetic descriptions of population states and hybrid spiking models whose local states are restricted to rest, spike and refractory sets have the same structure of a linear constraint and a nonconvex admissible set.

The field constraint adds nothing to this question. A natural way to bring convex integration to the field hypothesis is to restrict the oscillations of the construction to those that leave the field unchanged. The previous section shows that for a smoothing field map this restriction is empty in the limit: oscillations at wavenumber k change the field by a factor e^{-kh} , so all sufficiently fine oscillations are field-neutral. Whether a neural h-principle holds, in the sense that every admissible coarse field state is realized by many microscopically different exact solutions, depends entirely on the dynamics. It is a question about the wave cone of the neural equations, and the field's only role is to make every answer to it invisible at the electrodes.

What the electric-field hypothesis predicts

The results sort the available explanations of a stable field over drifting neurons by what each predicts.

- **Passive averaging.** The field is a smoothing map of a high-dimensional state, and any population with spread-out fluctuations has a stable field. Prediction: field fluctuations relax no faster than neural fluctuations in the fiber, and matched perturbations relax at equal rates.
- **Measurement limitation.** Apparent degeneracy is mostly the measurement fiber. Prediction: the fraction of drift in the unobserved directions equals the fraction of unobserved directions, and it shrinks as electrode density grows toward the noise-limited count.
- **Structured noise.** Variability is small in field-sensitive directions because inputs are organized that way. Prediction: the equal-time covariance shows the guard-rail signature, while lagged field statistics and perturbation responses are those of the passive model.
- **Field stabilization.** Feedback through the field restores field-changing deviations. Prediction: long-wavelength field fluctuations relax faster than fiber fluctuations, a perturbation along a field-sensitive direction relaxes faster than an equal-norm perturbation along the fiber, and drift exceeds the isotropic nullspace baseline.

The published evidence fits more than one of these. Information about memory content in the field, a stable latent space and a strong field-to-neuron coupling estimate (Pinotsis and Miller 2022, 2026; Pinotsis and Miller 2023) are what field stabilization predicts, and the first two are also what passive averaging and structured noise predict. The discriminating measurements are dynamical. Estimating the lead field from current-source density or a biophysical forward model (Einevoll et al. 2013; Lindén et al. 2011) gives the field-sensitive directions; lagged autocorrelations of field and spiking activity projected on those directions, and closed-loop field perturbations in the style of the feedback fields of slice experiments (Fröhlich and McCormick 2010), then test the relaxation asymmetry directly. The cellular evidence that fields entrain spiking (Anastassiou et al. 2011; Chiang et al. 2019) makes the asymmetry plausible. Whether it has the size needed to make the field the variable that a memory occupies is the open quantity, and it is measurable.

Limitations

The forward model is planar, homogeneous and periodic, with two laminae and identical dipole geometry within a patch; real cortex has folding, heterogeneous conductivity and morphologically diverse sources, all of which shrink the exact fiber by making some within-patch contrasts visible. The dynamics are linear, the feedback acts through the adjoint of the field map, and the noise is white; nonlinear attractor dynamics, plasticity and slow drift of synaptic weights, the mechanisms emphasized in the drift literature (Rule et al. 2019; Qin et al. 2023), are absent. The memory pattern and all parameters are illustrative. The convex-integration discussion establishes where such constructions are excluded and where they could begin; it constructs nothing. The synergetic reading of field-neuron interaction used in some of the electric-field work (Haken 1983; Pinotsis and Miller 2023) is compatible with the feedback model but was not modelled separately.

Reproducibility

The simulation in `simulation/` regenerates every number and figure: `uv run run_all.py` writes `output/results.json` and both figures and evaluates 28 internal checks (closed forms against explicit matrices, sampling and exact stochastic simulation). The claims ledger `claims.yaml` binds each reported number to its location in `results.json` through a recorded execution.

References

- Amari, Shun-ichi. 1977. "Dynamics of Pattern Formation in Lateral-Inhibition Type Neural Fields." *Biological Cybernetics* 27 (2): 77–87. <https://doi.org/10.1007/bf00337259>.
- Anastassiou, Costas A, and Christof Koch. 2015. "Ephaptic Coupling to Endogenous Electric Field Activity: Why Bother?" *Current Opinion in Neurobiology* 31: 95–103. <https://doi.org/10.1016/j.conb.2014.09.002>.
- Anastassiou, Costas A, Rodrigo Perin, Henry Markram, and Christof Koch. 2011. "Ephaptic Coupling of Cortical Neurons." *Nature Neuroscience* 14 (2): 217–23. <https://doi.org/10.1038/nn.2727>.
- Ball, J. M. 1989. "A Version of the Fundamental Theorem for Young Measures." In *PDEs and Continuum Models of Phase Transitions*, edited by M. Rascle, D. Serre, and M. Slemrod, vol. 344. Lecture Notes in Physics. Springer. <https://doi.org/10.1007/bfb0024945>.
- Ball, J. M., and R. D. James. 1987. "Fine Phase Mixtures as Minimizers of Energy." *Archive for Rational Mechanics and Analysis* 100 (1): 13–52. <https://doi.org/10.1007/bf00281246>.
- Bressloff, Paul C. 2012. "Spatiotemporal Dynamics of Continuum Neural Fields." *Journal of Physics A: Mathematical and Theoretical* 45 (3): 033001. <https://doi.org/10.1088/1751-8113/45/3/033001>.
- Buckmaster, Tristan, and Vlad Vicol. 2019. "Nonuniqueness of Weak Solutions to the Navier-Stokes Equation." *Annals of Mathematics* 189 (1). <https://doi.org/10.4007/annals.2019.189.1.3>.
- Buckmaster, Tristan, and Vlad Vicol. 2020. "Convex Integration and Phenomenologies in Turbulence." *EMS Surveys in Mathematical Sciences* 6 (1): 173–263. <https://doi.org/10.4171/emss/34>.
- Buzsáki, György, Costas A. Anastassiou, and Christof Koch. 2012. "The Origin of Extracellular Fields and Currents — EEG, ECoG, LFP and Spikes." *Nature Reviews Neuroscience* 13 (6): 407–20. <https://doi.org/10.1038/nrn3241>.
- Chiang, Chia-Chu, Rajat S. Shivacharan, Xile Wei, Luis E. Gonzalez-Reyes, and Dominique M. Durand. 2019. "Slow Periodic Activity in the Longitudinal Hippocampal Slice Can Self-propagate Non-synaptically by a Mechanism Consistent with Ephaptic Coupling." *The Journal of Physiology* 597 (1): 249–69. <https://doi.org/10.1113/jp276904>.

- Coombes, S. 2005. "Waves, Bumps, and Patterns in Neural Field Theories." *Biological Cybernetics* 93 (2): 91–108. <https://doi.org/10.1007/s00422-005-0574-y>.
- De Lellis, Camillo, and László Székelyhidi. 2009. "The Euler Equations as a Differential Inclusion." *Annals of Mathematics* 170 (3): 1417–36. <https://doi.org/10.4007/annals.2009.170.1417>.
- De Lellis, Camillo, and László Székelyhidi. 2013. "Dissipative Continuous Euler Flows." *Inventiones Mathematicae* 193 (2): 377–407. <https://doi.org/10.1007/s00222-012-0429-9>.
- Deitch, Daniel, Alon Rubin, and Yaniv Ziv. 2021. "Representational Drift in the Mouse Visual Cortex." *Current Biology* 31 (19): 4327–4339.e6. <https://doi.org/10.1016/j.cub.2021.07.062>.
- Driscoll, Laura N., Lea Duncker, and Christopher D. Harvey. 2022. "Representational Drift: Emerging Theories for Continual Learning and Experimental Future Directions." *Current Opinion in Neurobiology* 76: 102609. <https://doi.org/10.1016/j.conb.2022.102609>.
- Driscoll, Laura N., Noah L. Pettit, Matthias Minderer, Selmaan N. Chettih, and Christopher D. Harvey. 2017. "Dynamic Reorganization of Neuronal Activity Patterns in Parietal Cortex." *Cell* 170 (5): 986–999.e16. <https://doi.org/10.1016/j.cell.2017.07.021>.
- Druckmann, Shaul, and Dmitri B. Chklovskii. 2012. "Neuronal Circuits Underlying Persistent Representations Despite Time Varying Activity." *Current Biology* 22 (22): 2095–103. <https://doi.org/10.1016/j.cub.2012.08.058>.
- Edelman, Gerald M., and Joseph A. Gally. 2001. "Degeneracy and Complexity in Biological Systems." *Proceedings of the National Academy of Sciences* 98 (24): 13763–68. <https://doi.org/10.1073/pnas.231499798>.
- Einevoll, Gaute T., Christoph Kayser, Nikos K. Logothetis, and Stefano Panzeri. 2013. "Modelling and Analysis of Local Field Potentials for Studying the Function of Cortical Circuits." *Nature Reviews Neuroscience* 14 (11): 770–85. <https://doi.org/10.1038/nrn3599>.
- Fröhlich, Flavio, and David A. McCormick. 2010. "Endogenous Electric Fields May Guide Neocortical Network Activity." *Neuron* 67 (1): 129–43. <https://doi.org/10.1016/j.neuron.2010.06.005>.
- Gallego, Juan A., Matthew G. Perich, Raeed H. Chowdhury, Sara A. Solla, and Lee E. Miller. 2020. "Long-Term Stability of Cortical Population Dynamics Underlying Consistent Behavior." *Nature Neuroscience* 23 (2): 260–70. <https://doi.org/10.1038/s41593-019-0555-4>.
- Gardiner, Crispin. 2009. *Stochastic Methods: A Handbook for the Natural and Social Sciences*. 4th ed. Springer Berlin Heidelberg. <https://doi.org/10.1007/978-3-540-70713-4>.
- Gromov, Mikhael. 1986. *Partial Differential Relations*. Springer Berlin Heidelberg. <https://doi.org/10.1007/978-3-662-02267-2>.
- Haken, Hermann. 1983. *Synergetics: An Introduction*. 3rd ed. Springer Series in Synergetics. Springer Berlin Heidelberg. <https://doi.org/10.1007/978-3-642-88338-5>.

- Hämäläinen, Matti, Riitta Hari, Risto J. Ilmoniemi, Jukka Knuutila, and Olli V. Lounasmaa. 1993. "Magnetoencephalography—Theory, Instrumentation, and Applications to Noninvasive Studies of the Working Human Brain." *Reviews of Modern Physics* 65 (2): 413–97. <https://doi.org/10.1103/revmodphys.65.413>.
- Helmholtz, H. 1853. "Ueber einige Gesetze der Vertheilung elektrischer Ströme in körperlichen Leitern, mit Anwendung auf die thierisch-elektrischen Versuche." *Annalen Der Physik* 165 (6): 211–33. <https://doi.org/10.1002/andp.18531650603>.
- Isett, Philip. 2018. "A Proof of Onsager's Conjecture." *Annals of Mathematics* 188 (3). <https://doi.org/10.4007/annals.2018.188.3.4>.
- Jefferys, J. G. 1995. "Nonsynaptic Modulation of Neuronal Activity in the Brain: Electric Currents and Extracellular Ions." *Physiological Reviews* 75 (4): 689–723. <https://doi.org/10.1152/physrev.1995.75.4.689>.
- Kaufman, Matthew T, Mark M Churchland, Stephen I Ryu, and Krishna V Shenoy. 2014. "Cortical Activity in the Null Space: Permitting Preparation Without Movement." *Nature Neuroscience* 17 (3): 440–48. <https://doi.org/10.1038/nn.3643>.
- Lindén, Henrik, Tom Tetzlaff, Tobias C. Potjans, et al. 2011. "Modeling the Spatial Reach of the LFP." *Neuron* 72 (5): 859–72. <https://doi.org/10.1016/j.neuron.2011.11.006>.
- Logothetis, Nikos K., Christoph Kayser, and Axel Oeltermann. 2007. "In Vivo Measurement of Cortical Impedance Spectrum in Monkeys: Implications for Signal Propagation." *Neuron* 55 (5): 809–23. <https://doi.org/10.1016/j.neuron.2007.07.027>.
- Lundqvist, Mikael, Jonas Rose, Pawel Herman, Scott L. Brincat, Timothy J. Buschman, and Earl K. Miller. 2016. "Gamma and Beta Bursts Underlie Working Memory." *Neuron* 90 (1): 152–64. <https://doi.org/10.1016/j.neuron.2016.02.028>.
- Marder, Eve, and Jean-Marc Goaillard. 2006. "Variability, Compensation and Homeostasis in Neuron and Network Function." *Nature Reviews Neuroscience* 7 (7): 563–74. <https://doi.org/10.1038/nrn1949>.
- Miller, Earl K., Mikael Lundqvist, and André M. Bastos. 2018. "Working Memory 2.0." *Neuron* 100 (2): 463–75. <https://doi.org/10.1016/j.neuron.2018.09.023>.
- Müller, Stefan. 1999. "Variational Models for Microstructure and Phase Transitions." In *Calculus of Variations and Geometric Evolution Problems*, edited by S. Hildebrandt and M. Struwe, vol. 1713. Lecture Notes in Mathematics. Springer. <https://doi.org/10.1007/bfb0092670>.
- Müller, Stefan, and Vladimír Šverák. 2003. "Convex Integration for Lipschitz Mappings and Counterexamples to Regularity." *Annals of Mathematics* 157 (3): 715–42. <https://doi.org/10.4007/annals.2003.157.715>.
- Murray, John D., Alberto Bernacchia, Nicholas A. Roy, Christos Constantinidis, Ranulfo Romo, and Xiao-Jing Wang. 2017. "Stable Population Coding for Working Memory Coexists with

- Heterogeneous Neural Dynamics in Prefrontal Cortex." *Proceedings of the National Academy of Sciences* 114 (2): 394–99. <https://doi.org/10.1073/pnas.1619449114>.
- Nash, John. 1954. "C¹ Isometric Imbeddings." *The Annals of Mathematics* 60 (3): 383. <https://doi.org/10.2307/1969840>.
- Nunez, Paul L., and Ramesh Srinivasan. 2006. *Electric Fields of the Brain: The Neurophysics of EEG*. 2nd ed. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195050387.001.0001>.
- Pinotsis, Dimitris A., Gene Fridman, and Earl K. Miller. 2023. "Cytoelectric Coupling: Electric Fields Sculpt Neural Activity and "Tune" the Brain's Infrastructure." *Progress in Neurobiology* 226: 102465. <https://doi.org/10.1016/j.pneurobio.2023.102465>.
- Pinotsis, Dimitris A., and Earl K. Miller. 2022. "Beyond Dimension Reduction: Stable Electric Fields Emerge from and Allow Representational Drift." *NeuroImage* 253: 119058. <https://doi.org/10.1016/j.neuroimage.2022.119058>.
- Pinotsis, Dimitris A., and Earl K. Miller. 2026. "Ephaptic Coupling Can Explain Variability in Neural Activity." *Cerebral Cortex* 36 (6): bhag098. <https://doi.org/10.1093/cercor/bhag098>.
- Pinotsis, Dimitris A, and Earl K Miller. 2023. "In Vivo Ephaptic Coupling Allows Memory Network Formation." *Cerebral Cortex* 33 (17): 9877–95. <https://doi.org/10.1093/cercor/bhad251>.
- Plonsey, Robert. 1963. "Reciprocity Applied to Volume Conductors and the ECG." *IRE Transactions on Bio-Medical Electronics* 10 (1): 9–12. <https://doi.org/10.1109/tbmel.1963.4322775>.
- Potthast, Roland, and Peter Beim Graben. 2010. "Existence and Properties of Solutions for Neural Field Equations." *Mathematical Methods in the Applied Sciences* 33 (8): 935–49. <https://doi.org/10.1002/mma.1199>.
- Prinz, Astrid A, Dirk Bucher, and Eve Marder. 2004. "Similar Network Activity from Disparate Circuit Parameters." *Nature Neuroscience* 7 (12): 1345–52. <https://doi.org/10.1038/nn1352>.
- Qin, Shanshan, Shiva Farashahi, David Lipshutz, Anirvan M. Sengupta, Dmitri B. Chklovskii, and Cengiz Pehlevan. 2023. "Coordinated Drift of Receptive Fields in Hebbian / Anti-Hebbian Network Models During Noisy Representation Learning." *Nature Neuroscience* 26 (2): 339–49. <https://doi.org/10.1038/s41593-022-01225-z>.
- Rule, Michael E, Adrianna R Loback, Dhruva V Raman, Laura N Driscoll, Christopher D Harvey, and Timothy O'Leary. 2020. "Stable Task Information from an Unstable Neural Population." *eLife* 9. <https://doi.org/10.7554/elife.51121>.
- Rule, Michael E, Timothy O'Leary, and Christopher D Harvey. 2019. "Causes and Consequences of Representational Drift." *Current Opinion in Neurobiology* 58: 141–47. <https://doi.org/10.1016/j.conb.2019.08.005>.

- Scholz, J. P., and Gregor Schöner. 1999. "The Uncontrolled Manifold Concept: Identifying Control Variables for a Functional Task." *Experimental Brain Research* 126 (3): 289–306. <https://doi.org/10.1007/s002210050738>.
- Schoonover, Carl E., Sarah N. Ohashi, Richard Axel, and Andrew J. P. Fink. 2021. "Representational Drift in Primary Olfactory Cortex." *Nature* 594 (7864): 541–46. <https://doi.org/10.1038/s41586-021-03628-7>.
- Todorov, Emanuel, and Michael I. Jordan. 2002. "Optimal Feedback Control as a Theory of Motor Coordination." *Nature Neuroscience* 5 (11): 1226–35. <https://doi.org/10.1038/nn963>.
- Wilson, Hugh R., and Jack D. Cowan. 1972. "Excitatory and Inhibitory Interactions in Localized Populations of Model Neurons." *Biophysical Journal* 12 (1): 1–24. [https://doi.org/10.1016/s0006-3495\(72\)86068-5](https://doi.org/10.1016/s0006-3495(72)86068-5).
- Ziv, Yaniv, Laurie D Burns, Eric D Cocker, et al. 2013. "Long-Term Dynamics of CA1 Hippocampal Place Codes." *Nature Neuroscience* 16 (3): 264–66. <https://doi.org/10.1038/nn.3329>.