

Cortical-like dynamics in recurrent circuits optimized for sampling-based probabilistic inference

Rodrigo Echeveste^{1,2@}, Laurence Aitchison¹, Guillaume Hennequin^{1*}, Máté Lengyel^{1,3*}

¹Computational and Biological Learning Lab, Dept. of Engineering, University of Cambridge, Cambridge, UK

²Research Institute for Signals, Systems and Computational Intelligence sinc(i) (FICH-UNL/CONICET), Santa Fe, Argentina.

³Center for Cognitive Computation, Department of Cognitive Science, Central European University, Budapest, Hungary

* equal contribution, @ corresponding author: recheveste@sinc.unl.edu.ar

Sensory cortices display a suite of ubiquitous dynamical features, such as ongoing noise variability, transient overshoots, and oscillations, that have so far escaped a common, principled theoretical account. We developed a unifying model for these phenomena by training a recurrent excitatory–inhibitory neural circuit model of a visual cortical hypercolumn to perform sampling-based probabilistic inference. The optimized network displayed several key biological properties, including divisive normalization, as well as stimulus-modulated noise variability, inhibition-dominated transients at stimulus onset, and strong gamma oscillations. These dynamical features had distinct functional roles in speeding up inferences and made predictions that we confirmed in novel analyses of awake monkey recordings. Our results suggest that the basic motifs of cortical dynamics emerge as a consequence of the efficient implementation of the same computational function—fast sampling-based inference—and predict further properties of these motifs that can be tested in future experiments.

The dynamics of sensory cortices exhibit a set of ubiquitous features across species and experimental conditions. Responses vary over time and across trials even when the same static stimulus is presented¹, and these intrinsic variations have both systematic and seemingly random components (so-called noise variability). The most prominent systematic patterns of neural activity are strong, inhibition-dominated transients at stimulus onset² (or, equivalently, strong adaptation following stimulus onset), and stimulus-dependent population oscillations in the gamma band (20–80 Hz)^{3,4}. The extent and pattern of noise variability is also stimulus-dependent: variability is quenched at stimulus onset¹, decreasing gradually with stimulus contrast in the primary visual cortex (V1)⁵, and is further modulated by the content of the stimulus, e.g. the orientation or direction of drifting gratings for cells in V1 or in the middle temporal visual area (MT)⁶.

While the mechanisms giving rise to these dynamical phenomena are increasingly well understood^{6,7}, their functional significance remains largely unknown and controversial. For example, cortical gamma oscillations have been suggested to be a substrate for binding different sources of information about a feature (known as binding by synchrony⁸), to mediate information routing (communication by synchrony⁹), or to enable a temporal code of spikes relative to the oscillation phase¹⁰. Transient overshoots have been proposed to carry novelty or prediction error sig-

nals¹¹. Noise variability, when considered to be useful at all, has been argued to bear signatures of probabilistic computations^{5,12,13}. However, it is unclear whether these explanations can be reconciled, as each of them only accounts for select aspects of the data, and has been challenged by alternative accounts^{3,14–16}.

Here, we present a unifying model in which all of these dynamical phenomena emerge as a consequence of the efficient implementation of the same computational function: probabilistic inference. Probabilistic inference provides a principled solution to forming percepts by fusing partial and noisy information from multiple sources (including multiple sensory cues, modalities, and forms of memory)¹⁷. Formally, this fusion results in a posterior distribution expressing the probability that relevant features in the environment that are not directly accessible to the brain (e.g. the three-dimensional shapes of objects) may take any particular configuration given information that is directly available to our senses (e.g. photons absorbed in the retina). Behavioral evidence in several domains, including near-optimal performance in multi-sensory integration, decision making, motor control, and learning suggests that the brain represents posterior distributions at least approximately¹⁸. There have also been several proposals for how the neural responses of sensory cortical populations may implement these probabilistic representations^{5,12,19}. While these models successfully explained important aspects

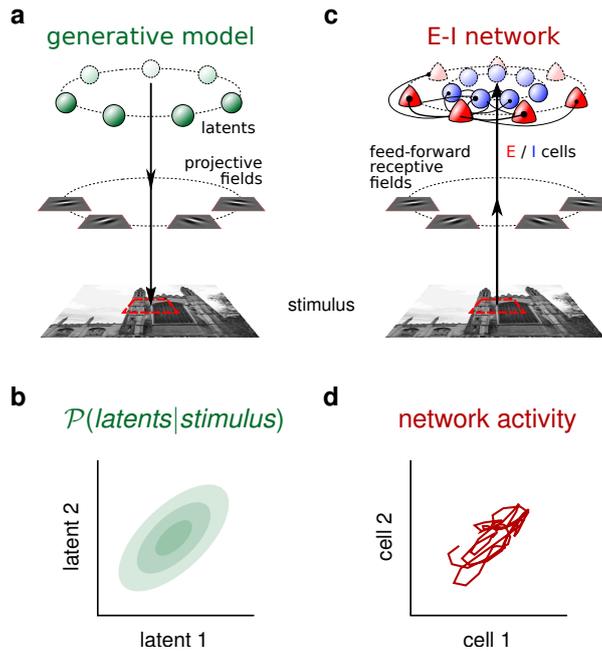


Fig. 1 | The statistical generative model, and the corresponding neural circuit implementing sampling-based probabilistic inference. **a**, Sketch of the Gaussian scale mixture (GSM) generative model. An image patch is constructed as a linear combination of a fixed set of localized, oriented, Gabor filter-like features (projective fields, differing only in their orientations, uniformly spread between -90° and 90°), with stimulus-specific feature intensities (latent variables) drawn from a multivariate Gaussian distribution. The resulting image patch is scaled by a global contrast variable and corrupted by noise (not shown). (Stimulus shown is for illustration only: the GSM model employed here was not sufficiently complex to generate photorealistic images. For a sample of generated image patches see Fig. 3c and Extended Data Fig. 1c.) **b**, 2-dimensional projection of the posterior distribution over latent variables given a visual stimulus, computed by the Bayesian ideal observer under the generative model. **c**, An excitatory–inhibitory (E–I) neural network receiving an image patch as an input, filtered by feedforward receptive fields identical to the projective fields of the generative model in **a**. The activity of each E cell represents the value of one latent variable in the generative model. As an illustration of the ring topology of the network, the outgoing connections of one E–I cell pair are shown (connection strength is indicated by line width and “synapse” size, see details in text). **d**, Responses of the two E cells corresponding to the latent variables shown in **b**. The response trajectory samples from the corresponding posterior distribution over time given the same stimulus.

of *stationary* response distributions (e.g. tuning curves, Fano factors, noise correlations), they have so far fallen short of accounting for the rich intrinsic *dynamics* of sensory cortical areas.

To bring together dynamics (cortical-like activity patterns) and function (representing posterior distributions) in a principled manner, we optimized a biologically constrained recurrent neural network for performing sampling-based probabilistic inference^{5,13,20,21}. Specifically, the network’s objective was to dynamically produce responses whose distribution matched a Bayesian ideal observer’s posterior distribution for each stimulus. The optimized neural circuit exhibited a number of appealing computational and dynamical features. Computationally, after training on a reduced stimulus set, the network exhibited strong forms of generalization by producing near-optimal response distributions to novel inputs which required qualitatively different responses. Furthermore, the network discovered out-of-equilibrium dynamics, a strategy currently employed by modern machine learning algorithms to produce samples

that become statistically independent on short timescales²².

Biologically, the optimized circuit achieved divisive normalization of its outputs and displayed marked transients at stimulus onset, as well as strong gamma oscillations. Both the magnitude of transients and the frequency of gamma oscillations scaled with stimulus contrast. Crucially, these dynamical phenomena did not emerge in control networks optimized for related objectives that did not require sampling. Further analyses of transients and oscillations in the optimized network revealed distinct functional roles for them. These analyses predicted novel properties of cortical dynamics: onset transients should be tuned to stimuli, which we confirmed by new analyses of published V1 recordings in the awake monkey²³. In addition, our model also made specific predictions about the stimulus tuning of excitatory–inhibitory lags and the distribution of gamma power across the different modes of network dynamics. Both can be readily tested in future experiments.

In summary, we constructed the first biologically constrained recurrent neural network performing sampling-based probabilistic inference that explained a plethora of electrophysiological observations in sensory cortices. Our model thus provides a unifying theoretical account of the basic motifs of sensory cortical dynamics.

Results

Optimizing a recurrent neural circuit for probabilistic inference

To study neural circuit dynamics implementing probabilistic inference, we used a novel combination of two well-established, though hitherto unrelated computational approaches: Bayesian ideal observers and the training of recurrent neural networks (Methods). First, we used a Bayesian ideal observer model to specify the computational goal of perceptual inference in a simplified visual task. Performing inference requires an internal model that encapsulates one's assumptions about how the inputs to be processed have been generated by the environment. For this, we adopted the Gaussian scale mixture (GSM) model (Fig. 1a), a generative model that has been shown to capture the statistics of natural image patches²⁴. Conversely, inference under the GSM model has been shown to account for behavioral and neural data (for stationary responses) in visual perception^{5,25,26}. The GSM model assumes that an image patch is generated as a linear combination of oriented Gabor filter-like visual features ("projective fields"), each present with a different intensity (the latent variables of the model). The image patch is further scaled by a single global "contrast" variable. Here, we focused on modeling a single hypercolumn, by choosing the projective fields of the GSM latent variables to only differ in their orientation so that they formed a ring topology (Fig. 1a, Extended Data Fig. 1a-b). The ideal observer was obtained by a Bayesian inversion of this model. Thus, for every image patch taken as sensory input, the ideal observer yielded a high-dimensional posterior distribution quantifying the probability that any particular joint combination of feature intensities may have generated the input (Fig. 1b).

Second, to model cortical circuit dynamics, we used a canonical, rate-based stochastic recurrent neural network model, the stochastic variant of the stabilized supralinear network (SSN)^{6,27}. The network was constrained to exhibit some basic biological features previously shown to be fundamental for cortical dynamics: separate but interconnected excitatory (E) and inhibitory (I) populations of neurons (Fig. 1c), supralinear (expan-

sive) input/output functions^{27,28}, and finite and stimulus-independent process noise^{6,29} incorporating intrinsic and extrinsic forms of neural variability.

We trained this network to perform sampling-based inference under the GSM. For this, we assumed a one-to-one mapping between latent variables and excitatory cells, such that the response (membrane potential) of each E cell represented the intensity of a different feature in the GSM model. Inhibitory neurons were treated as auxiliary units not explicitly constrained by the computational objective. The network was optimized to produce distributions of excitatory neural activities that matched the posteriors computed by the GSM-based ideal observer up to second-order statistics. For each stimulus in a small training set, the network was required to use its stochastic dynamics to visit different parts of state space over time with a frequency proportional to the posterior distribution corresponding to the same stimulus (Fig. 1d). Critically, as process noise in the network was stimulus-independent, the network had to use its recurrent dynamics to shape this variability appropriately for matching the target posteriors for each input. The training objective also included terms encouraging fast circuit dynamics. In summary, the network had to generate fast fluctuations with the correct stimulus-dependent patterns of across-trial mean and covariance.

Optimizing our network was challenging because modulating response variability (to match the stimulus-dependent posterior covariances of the ideal observer model) requires strong and non-linear recurrent interactions, but networks of strongly connected excitatory neurons – especially with supralinear input/output functions – are prone to becoming unstable^{6,30}. In such networks, it is non-trivial to find parameter regimes in which stability is preserved and thus optimization can proceed⁶. We therefore resorted to a reduced "ring" (hypercolumn) parametrization in which recurrent connection strength between any two cells (and the covariability of their process noise) only depended on the angular distance between their preferred stimuli and their respective cell types (E or I; Extended Data Fig. 1d). The feedforward receptive fields of the cells were fixed and identical to the projective fields of the corresponding latent variables of the ring-structured GSM (Extended Data Fig. 1a).

Inference and generalization in the optimized network

In line with neural recordings, activity in the optimized network was highly variable across time

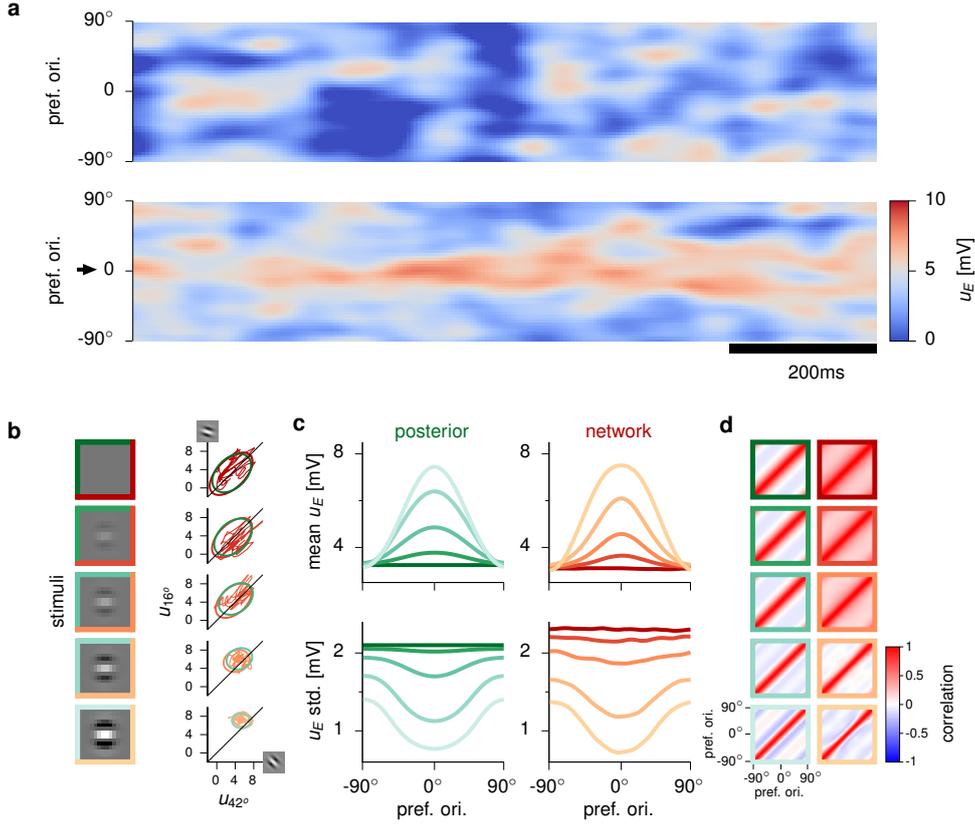


Fig. 2 | Inference and responses in the optimized network. **a**, Sample population activity of excitatory (E) cell membrane potentials u_E (color) at zero (top) and high (bottom) contrast. The high contrast stimulus has a dominant orientation at 0° (arrow). Neurons are ordered by preferred orientation. **b**, Left: stimuli in the training set (shade of frame color indicates contrast level, split green and red indicates that the same stimuli were used as input to the ideal observer and the neural network). Right: covariance ellipses (2 standard deviations) of the ideal observer's posterior distributions (green) and of the network's corresponding response distributions (red). Red trajectories show sample 500 ms-sequences of activities in the network. Projections for two representative latent variables / E cells are shown, with projective fields / receptive fields at preferred orientations 42° and 16° (insets at the end of axes). **c**, Mean (top) and standard deviation (bottom) of latent variables under the ideal observer's posterior distribution (left, green) and of E cell membrane potentials u_E under the network's stationary distribution (right, red), ordered by their preferred orientation, for each stimulus in the training set. **d**, Correlation matrices of the ideal observer's posterior distributions (left, green) and the network's stationary response distributions (right, red). Line colors in **c** and frame colors in **d** correspond to different contrast levels, same colors as stimulus frames in **b**. Response moments in **c** and **d** were estimated from $n = 20,000$ independent samples (taken 200 ms apart). Correlations in **d** are Pearson's correlations.

and trials, for both low-contrast (Fig. 2a, top) and high-contrast stimuli (Fig. 2a, bottom). Critically, the distributions of neural responses at the five training stimuli (the same image patch at five different contrast levels; Fig. 2b left) closely matched the corresponding GSM posteriors (Fig. 2b–d, compare red to green). Specifically, the mean activity of neurons increased while the variability of their responses decreased with contrast as well as with the match between stimulus orientation and their preferred orientation. This was consistent with the behavior of the moments of the GSM posterior (Fig. 2c, and circles in Fig. 3a). Thus, the network had been trained successfully to perform sampling-based inference on these stimuli.

We also tested the capacity of the network to represent the appropriate posterior distributions for novel stimuli. First, we confirmed that the mean

and variability of network responses smoothly interpolated between the corresponding target moments for intermediate contrast levels, closely following the behavior of the GSM posterior (Fig. 3a, solid curves between circles). Next, we presented the network with 500 entirely novel image patches randomly generated from the GSM (Extended Data Fig. 1c). Overall, both the means (Fig. 3b, top) and covariances of network responses (Fig. 3b, bottom) matched those of the target posteriors. This match was similarly good for test stimuli (Fig. 3b, orange) as for training stimuli (Fig. 3b, lavender). Critically, while the inputs of the training set included a single dominant orientation, many test stimuli had a more complex structure, with more than one dominant orientation (Fig. 3c, first column). Consequently, the corresponding GSM posteriors that the network was required to match became qualitatively different. Specifically, both the mean activity

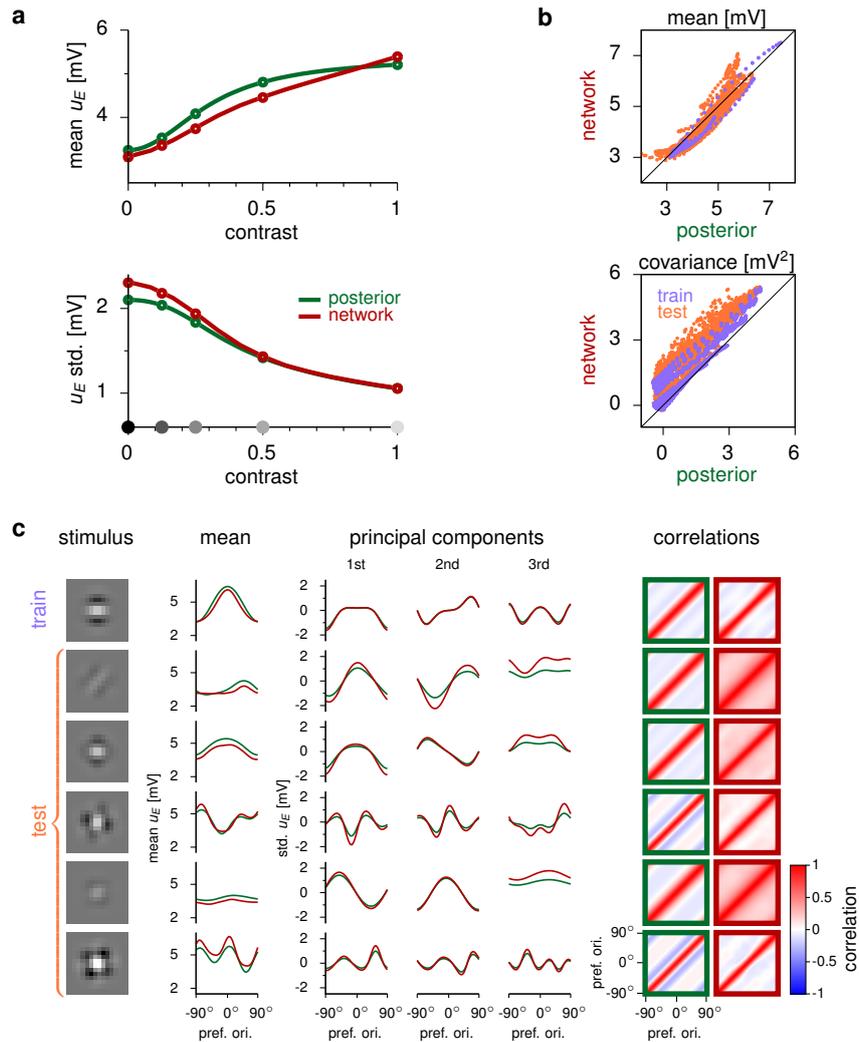


Fig. 3 | Generalization in the optimized network. **a**, Mean and standard deviation of latent variables (green) and stationary network responses (red) averaged over the population, as a function of contrast. Circles, and gray dots on x-axis indicate training contrast levels. The network correctly generalizes to untrained contrast levels (segments between circles). **b**, Stationary mean (top) and covariance (bottom) during network activity (y-axis) versus under the posterior (x-axis). Each dot corresponds to the response of an individual cell (top) or cell-pair (bottom) to one of the trained stimuli (lavender) or one of the novel, untrained stimuli in the test set (orange). **c**, Examples of generalization in the network. Each row corresponds to a different stimulus, and shows the corresponding statistical moments of latent variables under the GSM posterior (green) and stationary responses in the network (red). As a reference, the top row shows one of the training stimuli. The bottom five rows show generalization to novel test stimuli. Left: example stimuli. Middle: GSM (green) and network means (red), and the first three principal components of the GSM covariance, scaled by the square root of the variance they explain of the GSM posterior (green) and of the network covariance (red). Right: Correlation matrices of the ideal observer's posterior distributions (left, green frames) and the network's response distributions (right, red frames). Response moments in all panels were estimated from $n = 20,000$ independent samples (taken 200 ms apart). Population mean moments in **a** were further averaged across $n=50$ E cells. Correlations in **c** are Pearson's correlations.

across the population (Fig. 3c, 2nd column) and the principal components (PCs) of the noise covariances (Fig. 3c, remaining columns) became multimodal and highly dependent on the stimulus (Fig. 3c, green; compare across rows). The network was able to match the required GSM posteriors with high accuracy even in these challenging cases (Fig. 3c, red). Thus, the optimized network performed approximate Bayesian inference over a wide array of stimuli by always sampling (approximately) from the appropriate, stimulus-

dependent high-dimensional posterior distribution of the ideal observer.

The optimized network performs fast sampling

Under sampling-based inference, the time it takes to accurately represent a posterior distribution by collecting successive samples is directly proportional to the timescale over which these samples are correlated³¹. In our optimized network, noise variability generated new, independent samples every few tens of milliseconds across all contrast

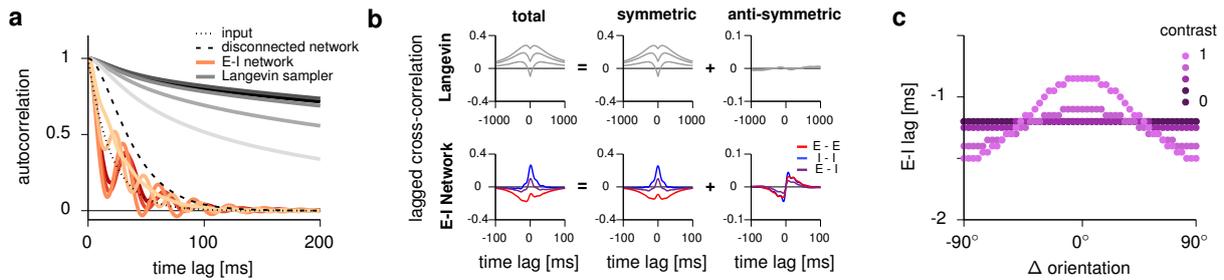


Fig. 4 | Temporal correlations in the optimized network. **a**, Membrane potential auto-correlations (population average) in the network for increasing levels of stimulus contrast (from dark to pale red; same colors as in Fig. 2b–d). The auto-correlation of a purely feedforward network (with the same process noise) is shown for comparison (dashed black line), together with those of the process noise (dotted black line), and a collection of networks implementing Langevin sampling at each contrast level (from dark to light gray). **b**, Lagged cross-correlation (left) in the Langevin sampler (top) and in the optimized E–I network (bottom), decomposed into temporally symmetric (middle) and anti-symmetric components (right). Each line corresponds to a different cell pair (3 representative pairs shown), color encodes identity of participating cells (E or I, note that there is no separation of E and I cells in the Langevin networks). **c**, Lag between total E and I inputs to each E cell, as a function of stimulus orientation (relative to preferred orientation) at different contrast levels (colors).

levels. This was evident in the fast fall-off of response autocorrelations (Fig. 4a, colored curves), similar to that observed in sensory cortices^{13,32}. In fact, our network was even faster than a disconnected network with the same membrane and input time constants (Fig. 4a, dashed curve), and close to the theoretical limit of a network of infinitely fast neurons in which sampling speed is solely limited by the input time constant (Fig. 4a, dotted curve).

To understand how the optimized network achieved fast sampling, we compared its dynamics to an algorithm known as Langevin sampling, a form of directed random walk (Methods). Sampling by Langevin dynamics was a relevant comparison for our network for two reasons. First, Langevin sampling is a popular, general-purpose algorithm in machine learning useful for benchmarking the performance of our network. Second, previous work suggested that sampling in stochastic recurrent neural networks (without respecting Dale’s principle) may be implemented by Langevin-like dynamics^{33–35}.

We found that for each input, Langevin dynamics was consistently an order of magnitude slower than our network (Fig. 4a, gray curves). The slowness of Langevin dynamics is known to arise from one of its critical features: time-reversible dynamics—i.e. that any time series of responses is as probable as its time-reversed counterpart³⁶. This was indeed reflected in temporally symmetric cross-correlograms (Fig. 4b, top). In contrast, our optimized network displayed a marked departure from time-reversibility, as evidenced by strong asymmetric components in its cross-correlograms (Fig. 4b, bottom). This dynamical irreversibility implied sequentiality in the activation of particular pairs of neurons. In particular, we found that I cells typically lagged behind E

cells. Moreover, for any cell, its total inhibitory input tended to also lag behind its overall excitatory input (Fig. 4c), consistent with known electrophysiology³⁷. Interestingly, this lag was smaller for cells that were most strongly driven by the stimulus, and this modulation became stronger with increasing contrast. These form testable predictions of our model.

Cortical-like dynamics in the optimized circuit

Having established that our network represented posterior distributions via sampling, we compared its responses with known physiological properties of V1. First, firing rates in the model had a physiologically realistic dynamic range and were tuned to stimulus orientation, similarly to neurons in macaque V1 (Fig. 5a, left-middle; Ref. 6, analysis of data recorded by Ref. 23). We also computed spike count statistics in the network from firing rates, assuming a doubly stochastic spike generation process (Methods). The quenching of membrane potential variability with increasing contrast (Fig. 2d, bottom) gave rise to quenching of spike count variability (quantified by the Fano factor), strongest at the cell’s preferred orientation (Fig. 5b, middle). These effects qualitatively matched V1 data, but were somewhat weaker (Fig. 5b, left). Moreover, stationary responses in the network exhibited clear signatures of divisive normalization (Extended Data Fig. 2a–b)³⁸. All these results were expected for a network whose stationary membrane potential response distributions represent GSM posteriors⁵.

Although our optimization procedure only constrained the network’s stationary response distributions, without prescribing any specific dynamics, we found that the optimized network exhibited realistic, cortical-like dynamics. Specifically, strong gamma oscillations emerged, with

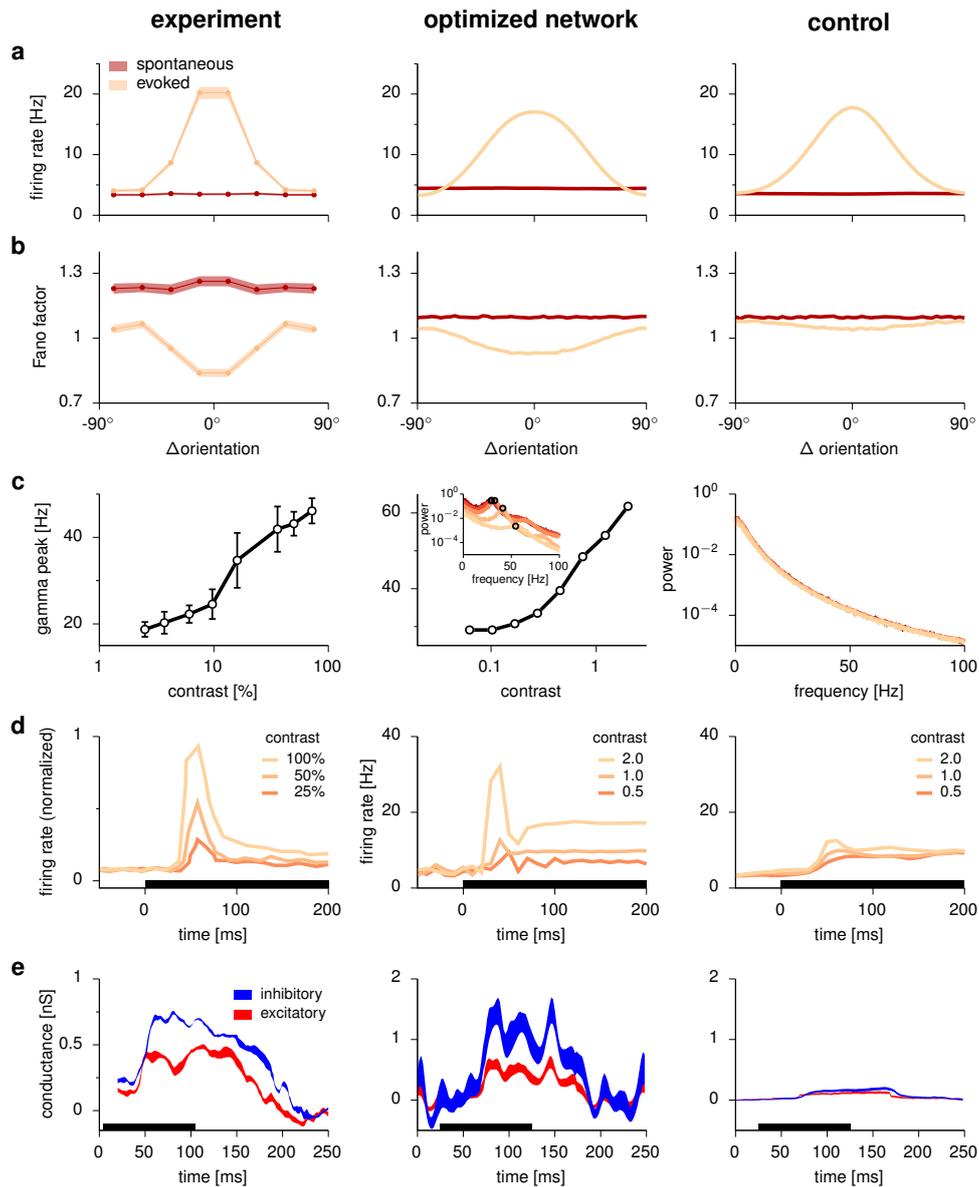


Fig. 5 | Cortical-like dynamics in the optimized network. Left: experimental data; middle: optimized network; right: control network trained to modulate its mean responses but not its variability. **a-b**, Mean firing rate (**a**) and Fano factor (**b**) of neurons as a function of stimulus orientation (relative to preferred orientation) during spontaneous (dark red) and evoked activity (light orange). Experimental results show mean \pm s.e.m. ($n=99$ cells). **c**, Peak gamma frequency in the local field potential (LFP) power spectrum as a function of contrast. Inset for the optimized network and main panel for the control network show LFP power spectra at different contrast levels (colors as in Figs. 2 and 4). Note that no dependence of gamma frequency is shown for the control network as there were no discernible gamma peaks in its power spectra. Experimental results show mean \pm s.d. ($n=14$ sessions). **d**, Average rate response around stimulus onset at different contrast levels (colors). **e**, Excitatory and inhibitory conductance (mean \pm s.e.m., relative to baseline, see Methods for details) during a transient stimulus response. Experimental results show mean \pm s.e.m. ($n=8$ trials). Numerical results show mean \pm s.e.m. ($n=20$ trials). Black bars in **d-e** show stimulus period. Panels **a-b** reproduce analyses from Ref. 6 of data from Ref. 23 (awake macaque V1). Experimental data in **c** was reproduced from Ref. 4 (awake macaque V1), in **d** reproduced from Ref. 3 (awake macaque V1), and in **e** reproduced from Ref. 2 (awake mouse V1).

a peak frequency increasing with contrast, consistent with V1 recordings in the awake monkey^{3,4} (Fig. 5c, left-middle). Moreover, selective clamping of either the E or I population abolished gamma oscillations (or stability altogether) (Extended Data Fig. 2c-d), suggesting gamma oscillations arose from dynamical interactions between E and I cells (“PING” mechanism¹⁶). The

network also showed strong transient responses such that average population rates had marked contrast-dependent overshoots at stimulus onset, consistent with recordings in V1³ (Fig. 5d, left-middle). Finally, using a conductance-based approximation of our current-based model (Methods), we found that inhibition transiently dominated over excitation during stimulus presenta-

tion, as in V1 of the awake mouse² (Fig. 5e, left-middle).

Control networks do not show cortical-like dynamics

We next sought to establish whether these dynamical properties arose simply due to the biological and architectural constraints imposed on our network – or specifically due to optimizing for sampling-based inference. For this, we used a series of ‘control’ networks in which single cell parameters (time constants and firing rate nonlinearities), overall network architecture, and receptive fields were all identical to those used in the original network. Thus, these networks only differed in the objective for which they were optimized.

First, we confirmed that the dynamics characteristic of the originally optimized network did not emerge in randomly parametrized networks without optimization, but robustly emerged after optimization starting from different random initial conditions (Extended Data Figs. 3 and 4). Next, we optimized a control network that differed from the original network in only one critical aspect: it was only required to match the means of the posterior distributions, but neither variances nor covariances. Despite clear stimulus-dependent modulations in mean responses (as required by training; Fig. 5a, right), this network exhibited only minimal modulations of both membrane potential variability (Extended Data Fig. 5a-g) and Fano factors (Fig. 5b, right; Extended Data Fig. 6). Thus, the modulations of response variability seen in the original network, which are a hallmark of sampling-based inference⁵, were not simply a by-product of non-linear E–I dynamics. Moreover, neither gamma oscillations nor marked inhibition-dominated transients emerged in the control network (Fig. 5c-e, right). In fact, matching both means and variances, but not covariances (necessary for full inference), still abolished these dynamical features (Extended Data Fig. 5h-n, and 6). Finally, oscillations were also absent in another control network specifically optimized to modulate its mean firing rates as before, while keeping its Fano factors constant (Extended Data Figs. 6 and 7), as would be required by other, non sampling-based probabilistic representations¹².

These results suggest that the dynamical features observed in the original network emerged as a consequence of the specific computation for which it was optimized. Conversely, training the network on the original cost function but without enforcing Dale’s principle resulted in substantially poorer performance and a lack of oscilla-

tions and transients (Extended Data Figs. 4, 6 and 8). Thus, achieving competent sampling performance and exhibiting realistic dynamics again appeared to be coupled.

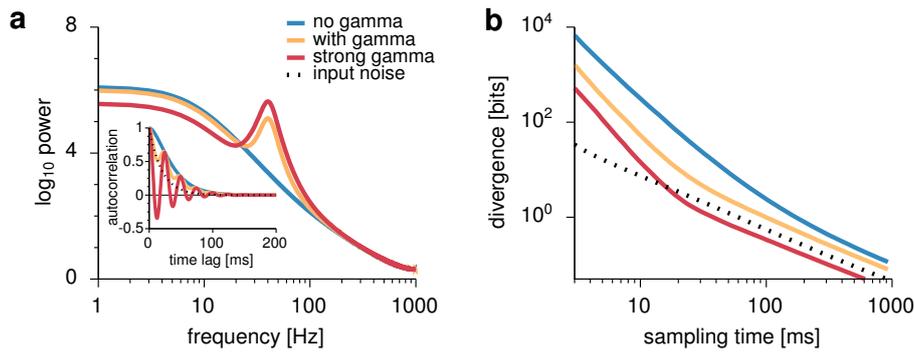
Oscillations improve mixing time

To study the potential functional benefits of oscillations, one would ideally like to “knock-out” oscillations from the network while leaving all other features of the dynamics intact. The complex and high dimensional dynamics of our network made this unfeasible. Therefore, we first studied the response of a single neuron to obtain an analytical understanding of the general role of oscillations in sampling (Fig. 6a-b). We next generalized this analysis to oscillations in network-wide activity patterns rather than single neurons, thus providing insights into the high-dimensional dynamics of the full network (Fig. 6c-d).

Assuming that the response of a neuron is statistically stationary and approximately normally distributed, it is fully characterized by its mean, variance, and autocorrelogram. As long as this neuron is part of a sampling-optimized network, the mean and variance of its response will have to match those prescribed by the target distribution sampled by the network (Fig. 2b-d). Although the autocorrelogram is not constrained by the target moments, it still contributes critically to the performance of the network. Specifically, it can be shown mathematically that the total area under the autocorrelogram directly scales “mixing time”: the time it takes for the dynamics to represent the target distribution to a given precision (Supplementary Math Note). Therefore, to understand the specific role of oscillations, we compared idealized (stationary and normally distributed) neuronal responses, constructed to have the same mean and variance as responses in our network but different autocorrelation functions (Methods).

We compared response autocorrelograms with different degrees of “oscillatoriness” (Fig. 6a, blue, orange, red), but the same envelope as that of the full network (Fig. 6a, inset, black dotted; see also Fig. 4). These oscillations substantially reduced the area under the autocorrelogram (Fig. 6a, inset) and thus accelerated the convergence of the empirical distribution of responses to the target distribution (Fig. 6b; Supplementary Math Note). Importantly, oscillations will only decrease the area under the autocorrelogram if at least one oscillation cycle fits under the envelope, i.e. if the oscillation period is sufficiently shorter than the width of the envelope (approximately 35 ms). This implied oscillation frequencies higher than 30 Hz, just as observed in the optimized network (Fig. 5c, middle).

single neuron



full network

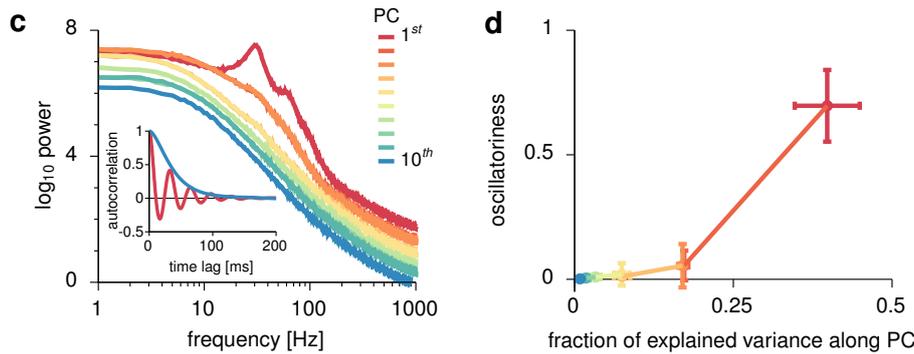


Fig. 6 | Oscillations improve mixing time. **a-b**, Analysis of oscillations in the response of a single neuron. **a**, Power spectra of three different neural responses (colored lines) with identical mean and variance but different degrees of oscillatoriness. Inset: autocorrelation functions. Black dotted line represents the autocorrelation of the process noise. **b**, Divergence between the distribution estimated from a finite sampling time (x-axis) and the true stationary distribution for the three systems (colors as in **a**). **c-d**, Analysis of oscillations in the full network. **c**, Power spectra of the network's neural activity along the directions of the principal components (PCs) of its stationary response distribution, ordered by PC rank (colors). Inset: autocorrelation of neural activity along the directions of the 1st and 10th PCs (colors as in main plot). **d**, Oscillatoriness of the autocorrelogram along each principal component (colors as in **c**) as a function of the fraction of the total variance of responses they capture. Note that our measure of oscillatoriness is based on the relative contributions of an oscillatory vs. a non-oscillatory component in a parametric fit to the autocorrelogram, and as such it is invariant to the overall magnitude of fluctuations (which is factored out by using the autocorrelation rather than the autocovariance of responses, [Methods](#)). Error bars show mean \pm s.e.m. ($n=50$ stimuli).

We next studied the organization of gamma oscillations in the multidimensional responses of the full network. We were able to show that maximal sampling speed is achieved specifically when smaller response variance is associated with higher oscillation frequency (Supplementary Math Note). In turn, as we showed above, variability is quenched with increasing contrast both in our network and in the cortex (Fig. 2b-c, Fig. 3a, Fig. 5b). This explains why the frequency of gamma oscillations increased with contrast in our network after optimization. These results suggest that contrast-modulated gamma oscillations observed in the cortex^{3,4} may reflect a speed-optimized sampling strategy (Fig. 5c).

Our mathematical analyses also predicted that oscillations in an efficiently sampling network should be predominantly expressed where they matter most: in the (stimulus-dependent) network-wide activity patterns capturing most of

the overall response variability. Namely, for each stimulus, we expected the strongest oscillations along the top PCs of the corresponding stationary covariance (Supplementary Math Note). This was indeed apparent in the power spectra of our network associated with the top 10 PCs (Fig. 6c), and the corresponding autocorrelograms that even showed negative-going lobes (Fig. 6c, inset). Specifically, there was a positive relationship between oscillatoriness along successive PCs and the fraction of variance explained. This meant that the network oscillated more in the directions along which its responses had the largest variance (Fig. 6d). (Note that our measure of oscillatoriness was based on autocorrelograms, and therefore had no *a priori* dependence on response variance; [Methods](#).) In sum, the network used non-trivial temporal dynamics, in the form of contrast-dependent, pattern-selective gamma oscillations, to ensure that even short segments of its activity were sufficiently rep-

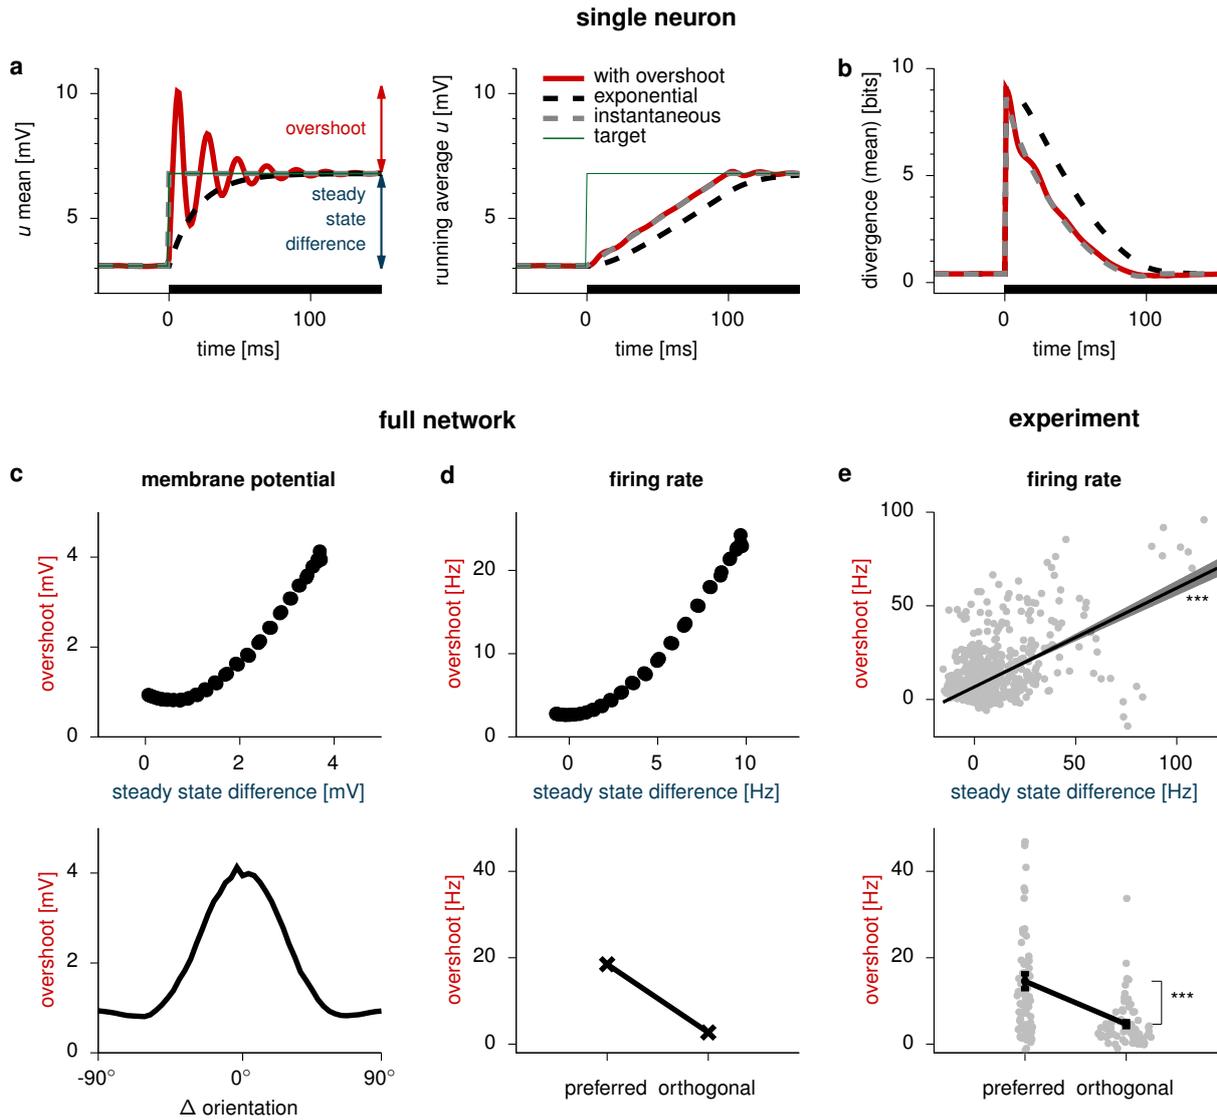


Fig. 7 | Transients support continual inference. **a-b**, Analysis of transients in the response of a single neuron. **a**, Temporal evolution of the mean (left) membrane potential (u_E), and its running average (right), in three different neural responses (thick lines) with identical autocorrelations (matched to neural autocorrelations in the full network, [Extended Data Fig. 9a](#); cf. [Fig. 6](#)) but different time-dependent means (shown here) and variances ([Extended Data Fig. 9a](#)). Thin green line shows the time-varying target mean. **b**, Divergence between the target distribution at a given point in time and the distribution represented by the neural activity sampled in the preceding 100 ms, for each of the three responses (colors as in **a**). The mean-dependent term of the divergence is shown here, which depends on the difference between the target mean and the running average of samples (shown in **a**, right; see [Extended Data Fig. 9b](#) for the full divergence). **a-b** Black bars show stimulus period. Mean and divergence computed as an average over multiple trials ($n = 10,000$). **c**, Top: overshoot magnitude versus steady state difference in membrane potentials (see **a** for legend). Each dot corresponds to the response of one cell to one particular stimulus. Bottom: overshoot magnitude as a function of stimulus orientation (relative to preferred orientation). **d**, Top: same as **c**, top, for firing rates. Bottom: average rate overshoot across stimuli whose orientation is aligned with the cells preferred orientation ($0 \pm 30^\circ$), or near-orthogonal to it ($90 \pm 30^\circ$). **e**, Analysis of experimental recordings from awake macaque V1²³. Top: overshoot magnitude versus steady state difference, as in **d**, top. Black line shows linear regression ($\pm 95\%$ confidence bands); ***: two-sided Wald test $p = 3 \times 10^{-114}$ ($n = 1280$ cell-stimulus pairs, $R^2 \simeq 0.33$; see also [Extended Data Fig. 9d-e](#)). Bottom: each gray dot represents one cell's average rate overshoot across stimuli whose orientation is aligned with the cells preferred orientation ($0 \pm 30^\circ$), or near-orthogonal to it ($90 \pm 30^\circ$), the mean of each group (\pm s.e.m.) is presented in black, as in **d**; ***: two-sided paired t-test $p = 6 \times 10^{-9}$ ($n = 80$ cells, $DF = 79$). (For a better comparison with **d**, the y-range is truncated at 50 Hz, clipping 2 data points. Statistical analyses used all data.)

representative of the posterior distribution it represented for each stimulus.

Transients support continual inference

The foregoing results showed that oscillations increase mixing speed in the stationary regime, i.e.

once network responses have become representative of the target distribution. Complementing this, we found that transients in our network mitigate the other main temporal constraint of sampling: the “burn-in” time it takes for responses to become representative in the first place³¹. We

observed that, in line with experimental data, during stimulus onset, neural responses tended to overshoot the corresponding stationary response levels (Fig. 5d, Fig. 7a). One might naively expect such transients to be detrimental for representing a distribution, as they clearly deviate from the target (represented by the steady-state responses). However, in a realistic setting with a changing environment, distributions need to be represented continually, without waiting for the system to achieve steady state.

In order to understand the role of transients in continual inference, we considered how a moving decoder of neural responses over a finite trailing time window approximated the target. As with oscillations, we performed this analysis in two steps. First, to isolate the potential functional benefits of transients we once again considered the response of just a single idealized neuron that is part of a sampling-optimized network (Methods). For this idealized neuron, we fixed the autocorrelation (thus controlling for oscillations) as well as the before- and after-stimulus onset steady-state means and variances to those of an actual, representative neuron in our network (Extended Data Fig. 9a). We then compared three ways in which this neuron could transition between these two steady states (Fig. 7a): 1. as an upper bound on performance, instantaneously switching between the two steady-states (Fig. 7a, gray dashed); 2. exponentially approaching the new steady state with the characteristic time constant of the cells in the network, thus lacking overshoots (Fig. 7a, black dashed); and 3. undergoing overshoots as seen in our optimized network (Fig. 7a, red).

We found that overshoots performed close to the upper bound, provided by instantaneous switching. In particular, they generated samples that allowed a substantially more accurate estimate of the target mean than that afforded by approaching the new steady-state exponentially without overshoots (Fig. 7a–b). (These results extended qualitatively to the case when the match in the full distributions was considered, Extended Data Fig. 9b.) This was because without overshoots at stimulus onset, responses were still sampling from the distribution corresponding to the baseline input. Thus, including them in the estimation of the new stimulus-related mean inevitably biased the estimate downwards. The overshoot largely compensated for this bias. Indeed, we were able to show analytically that optimal compensation requires transient overshoots at stimulus onset (Supplementary Math Note). This is because continual averaging of responses formally corresponds to a temporal convolution, and so the optimal response is the deconvolution of the

target with the averaging kernel. Under basic smoothness constraints, the deconvolution of a step function with such an averaging kernel yields transients like those we observed in the network (Extended Data Fig. 9c).

The hypothesis of increased sampling accuracy by transient compensation made a distinct prediction (Supplementary Math Note): transient overshoots should scale with the change in steady state responses. Indeed, our network exhibited this effect in both membrane potentials and firing rates (Fig. 7c–d, top). Importantly, transient overshoots thus inherited the orientation tuning of stationary responses (Fig. 7c–d, bottom). While stimulus-onset transients have been widely observed^{3,4}, their stimulus-tuning has not been analyzed. Therefore, we analyzed a previously published dataset of V1 responses in the awake monkey²³. In line with the predictions of the model, overshoot sizes were orientation tuned (Fig. 7e, bottom) and, more generally, they scaled with the change in stationary responses (Fig. 7e, top; these results were robust to excluding the outliers with high firing rates, e.g. above 60 Hz, Extended Data Fig. 9d–e).

Discussion

We have shown that a canonical neural network model^{6,27} produces cortical-like dynamics when optimized for sampling-based inference, but not when optimized for non-probabilistic objectives, or non sampling-based probabilistic objectives. Further controls demonstrated that these dynamics were not mere side products of the particular biological constraints or optimization approach we adopted. Instead, they played well-defined functional roles in performing inference.

The Gaussian scale mixture model and the stochastic stabilized supralinear network

We used a canonical model of neural network dynamics (the stochastic SSN) to embody a set of biologically relevant constraints for cortical circuits. It was not trivial *a priori* that this model would be able to modulate its responses as necessary for successful sampling-based inference under a canonical generative model of visual image patches (the GSM). A hint that this might indeed be possible came from previous studies showing that both in the SSN^{6,27} and the GSM^{5,39}, a range of parameters exists for which the average response or posterior mean monotonically increases while the variance decreases with increasing stimulus strength. Empirically, we found a good quantitative match that went beyond this coarse, qualitative trend, with the SSN also capturing much of the detailed structure of

the GSM posteriors. However, this match was not perfect: for example, the GSM posteriors systematically showed negative correlations of larger magnitude than what the network was able to express (Fig. 2d and Fig. 3b-c). It might be possible to achieve a more accurate match by allowing negative input correlations, and in general a more flexible parameterization of the SSN. Indeed, once the optimization of larger-scale, more flexibly parameterized SSNs becomes feasible, we also expect them to be able to sample from richer, deeper generative models.

Function-optimized neural networks

Our approach is complementary to classical approaches for training neural network models. Previous work showed how various steady-state properties of cortical responses (such as receptive fields, or trial-averaged activities) emerge from optimizing neural networks for some computationally well-defined objective^{30,40–44}. Notably, our sampling-based computational objective required our network to modulate not only the mean but also the variability of its responses in a stimulus-dependent manner. This made the training of networks significantly more challenging than conventional approaches training networks for deterministic targets without explicitly requiring them to modulate their variability^{40,43,45}. In return, the dynamics of our network exhibited rich, stimulus-modulated patterns of variability. These responses captured a variety of ubiquitous features of the trial-by-trial behavior of cortical responses (noise variability, transients, and oscillations) beyond the steady-state or trial-average properties that could be addressed by previous work.

Typically, previous network optimization approaches aimed to determine the types of dynamics that arise when a task is executed under minimal mechanistic constraints, using a neural network as a universal function approximator. As a result, they yielded fundamental insights about the macroscopic organization of network dynamics (e.g. the presence of line attractors⁴⁵) but did not attempt to incorporate some of the most salient constraints on the detailed organization of cortical circuits. Specifically, they used networks that were either purely feedforward⁴⁰, utilized neuronal transfer functions that lacked the expansive nonlinearities characteristic of cortical neurons^{40,42,44,45}, had no separation of E and I cells^{40,44,45}, or had noiseless dynamics⁴³.

In contrast, our goal was to study the emergence of (probabilistic) computations through dynamics and to connect these dynamics to experimental data at (or near) single cell resolution

(e.g. the neuron- and stimulus-specific reduction of variability, or the lag between total inhibitory and excitatory inputs in individual cells). This required respecting all the aforementioned biological constraints. Nevertheless, this additional realism came at the cost of having to limit the number of optimized parameters to be far lower than standard approaches with feedforward networks or recurrent networks for which dynamical stability is more easily achieved. While this reduced parametrization made it easier to find stable solutions, it was still sufficiently expressive. In particular, we found that our results could not have been obtained without optimization (Extended Data Figs. 3 and 4), or with the optimization of other objective functions (Fig. 5 and Extended Data Figs. 5–7). Indeed, this parametrization still included networks that were unstable, or showed a decrease in mean responses and/or increase in variability with increasing stimulus strength (i.e. the opposite of what was required for matching the GSM), or were modulated in a non-monotonic way or only minimally altogether (Extended Data Fig. 3).

Neural representations of uncertainty

Our approach markedly differed from previous work on the neural bases of probabilistic inference. Previous models were typically derived using a *top-down* approach (but see Ref. 43), using hand-designed network dynamics that explicitly mimicked specific existing approximate inference algorithms from machine learning based on sampling^{33–36,46} or other representations^{12,19,33,47}. As a result, these models came with strong theoretical guarantees for their performance but often offered only a mostly phenomenological match to neural circuit dynamics. In particular, they did not respect some basic biological constraints (e.g. Dale’s principle^{33,35,47}), or had to assume an unrealistically rapid and direct influence of stimuli on network parameters (e.g. synaptic weights^{35,46}). In contrast, we used a more *bottom-up* approach, starting from known constraints of cortical circuit organization, and then optimizing the parameters of networks under such constraints to achieve efficient sampling-based probabilistic inference – without pre-specifying the details of the dynamics that needed to be implemented. While this approach cannot provide formal guarantees on performance, our optimized network “discovered” novel algorithmic motifs (oscillations and transients) for speeding up probabilistic inference. Although some of these motifs have been observed in previous work⁴⁶, their function remained unclear as they were built-in by design rather than obtained as a result of optimization, or appeared purely epiphenomenal. In contrast, these motifs

served computationally well-defined functions in our network.

The dynamics of our network may also provide useful clues for constructing novel machine learning algorithms. In general, the kind of time-irreversible, out-of-equilibrium dynamics we demonstrate for our network have only recently been appreciated in machine learning^{22,36}. At the same time, sampling-based inference algorithms using second-order dynamics with so-called “momentum” variables, such as Hamiltonian Monte Carlo, have long been known to improve sampling speed³¹. Indeed, it might be interesting to explore how much the dynamics of our network can be interpreted as a neural implementation of Hamiltonian Monte Carlo⁴⁶. Nevertheless, despite such second-order dynamical systems often exhibiting oscillations and transient overshoots, their sampling efficiency has usually been analyzed only in the more generic terms of the suppression of random walk-like behavior. In contrast, our analyses revealed specific roles for oscillations and transients. In fact, the setting of continual inference that we used to demonstrate the benefits of transients has not been considered in machine learning applications so far, although we expect it to be highly relevant for both biological and artificial cognition.

Cortical variability, transients, and oscillations

Our work suggests a novel unifying function for three ubiquitous properties of sensory cortical responses: stimulus-modulated variability, transient overshoots, and gamma oscillations. In previous work, these phenomena have traditionally been studied in isolation and ascribed separate functional roles that have been difficult to reconcile. In particular, they have not been derived normatively, i.e. by starting from some functional objective and then optimizing that objective in a principled manner (but see e.g. Ref. 47). For example, cortical variability has most often been considered a nuisance, diminishing the accuracy of neural codes²³. Theories postulating a functional role of variability in probabilistic computations have only considered the steady-state distribution of responses without making specific predictions about their dynamical features^{5,12}. Conversely, transient responses prominently feature as central ingredients of models of predictive coding, where they signal novelty or deviations between predicted and observed states⁴⁷. However, these theories did not address response variability.

Our work accounts for both transients and variability starting from a single principle, using only

the equivalent of “internal representation neurons”⁴⁸ of predictive coding but without invoking specific prediction error-coding neurons. In particular, our model correctly predicted a specific scaling relationship between transients and steady-state responses which we tested by novel analyses of experimental data (Fig. 7). Furthermore, our mathematical analysis suggested that prediction-error-like signals (more formally, responses that scale with the magnitude of change in the target distribution; Extended Data Fig. 9c) are a generic signature of continual inference using sampling-based dynamics, and will thus not only appear at stimulus onsets but in any situation when predictions change temporally. A conclusive test of whether prediction-error-like responses in the cortex are due to this mechanism or classical predictive coding mechanisms will require more specific manipulations of prior expectations.

Gamma oscillations have also been proposed as a substrate for a number of functional roles in the past, related to how information is encoded, combined, or routed in the brain^{8–10,49}. These putative functions need not be mutually exclusive to that played in our network. Nevertheless, some of these functions seem difficult to reconcile with specific experimental findings^{3,14,15,50}. More generally, theories of gamma oscillations do not typically address transients.

Predictive coding models naturally account for transients, and can also account for gamma oscillations¹¹. However, it is unclear whether these theories would also account for properties beyond the mere existence of gamma oscillations. These would include the frequency modulation by contrast^{3,4} that our model reproduced (Fig. 5), or indeed any aspect of the ubiquitous variability of cortical responses, and its modulation by stimuli, which our model also reproduced as a core feature (Figs. 2, 3 and 5). In contrast, our results show that variability, transients, and gamma oscillations can all emerge from the same functional objective: that neural circuits use an efficient sampling-based representation of uncertainty under time constraints.

The mechanism by which gamma oscillations are generated in the brain, particularly whether it involves interactions between E and I cells (‘PING’ mechanism) or among I cells only (‘ING’ mechanism), is a subject of current debate¹⁶. In our model, voltage-clamping of E cells eliminated gamma oscillations (Extended Data Fig. 2c-d), pointing to the ‘PING’ mechanism. However, our network only included a single inhibitory cell type, and heavily constrained connectivity, therefore it remains for future work to study how the pre-

cise mechanism of gamma generation depends on such architectural constraints. Studying more hierarchical or spatially extended versions of our model may also allow us to study longer-range aspects of gamma oscillations, such as gamma synchronization⁴⁹.

Acknowledgements

This work was supported by the Wellcome Trust (New Investigator Award 095621/Z/11/Z and Investigator Award in Science 212262/Z/18/Z to M.L., and Seed Award 202111/Z/16/Z to G.H.), and the Human Frontiers Science Programme (Research Grant RGP0044/2018 to M.L.). We are grateful to A. Ecker, P. Berens, M. Bethge, and A. Tolias for making their data publicly available, to G. Orbán, A. Bernacchia, R. Haefner, and Y. Ahmadian for useful discussions, and to J.P. Stroud for detailed comments on the manuscript.

Author Contributions

R.E., G.H. and M.L. designed the study. R.E. and G.H. developed the optimization approach, R.E. ran all numerical simulations, R.E. and G.H. analyzed experimental data, all authors performed analytical derivations, R.E. G.H. and M.L. interpreted results and wrote the paper, with comments from L.A.

Competing Interests statement

The authors declare no competing interests.

References

1. Churchland, M. *et al.* Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature Neuroscience* **13**, 369 (2010).
2. Haider, B., Häusser, M. & Carandini, M. Inhibition dominates sensory responses in the awake cortex. *Nature* **493**, 97–100 (2013).
3. Ray, S. & Maunsell, J. H. Differences in gamma frequencies across visual cortex restrict their possible use in computation. *Neuron* **67**, 885–896 (2010).
4. Roberts, M. *et al.* Robust gamma coherence between macaque V1 and V2 by dynamic frequency matching. *Neuron* **78**, 523–536 (2013).
5. Orbán, G., Berkes, P., Fiser, J. & Lengyel, M. Neural variability and sampling-based probabilistic representations in the visual cortex. *Neuron* **92**, 530–543 (2016).
6. Hennequin, G., Ahmadian, Y., Rubin, D., Lengyel, M. & Miller, K. The dynamical regime of sensory cortex: stable dynamics around a single stimulus-tuned attractor ac-

count for patterns of noise variability. *Neuron* **98**, 846–860 (2018).

7. Buzsáki, G. & Wang, X. Mechanisms of gamma oscillations. *Annual review of neuroscience* **35**, 203–225 (2012).
8. Gray, C., König, P., Engel, A. & Singer, W. Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* **338**, 334–337 (1989).
9. Akam, T. & Kullmann, D. Oscillations and filtering networks support flexible routing of information. *Neuron* **67**, 308–320 (2010).
10. Masquelier, T., Hugues, E., Deco, G. & Thorpe, S. Oscillations, phase-of-firing coding, and spike timing-dependent plasticity: an efficient learning scheme. *Journal of Neuroscience* **29**, 13484–13493 (2009).
11. Bastos, A. *et al.* Canonical microcircuits for predictive coding. *Neuron* **76**, 695–711 (2012).
12. Ma, W., Beck, J., Latham, P. & Pouget, A. Bayesian inference with probabilistic population codes. *Nature Neuroscience* **9**, 1432–1438 (2006).
13. Berkes, P., Orbán, G., Lengyel, M. & Fiser, J. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* **331**, 83–87 (2011).
14. Shadlen, M. & Movshon, J. Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron* **24**, 67–77 (1999).
15. Thiele, A. & Stoner, G. Neuronal synchrony does not correlate with motion coherence in cortical area MT. *Nature* **421**, 366–370 (2003).
16. Tiesinga, P. & Sejnowski, T. Cortical enlightenment: are attentional gamma oscillations driven by ING or PING? *Neuron* **63**, 727–732 (2009).
17. Knill, D. & Richards, W. *Perception as Bayesian inference* (Cambridge University Press, 1996).
18. Fiser, J., Berkes, P., Orbán, G. & Lengyel, M. Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences* **14**, 119–130 (2010).
19. Deneve, S., Latham, P. & Pouget, A. Efficient computation and cue integration with noisy population codes. *Nature Neuroscience* **4**, 826 (2001).
20. Haefner, R., Berkes, P. & Fiser, J. Perceptual decision-making as probabilistic inference by neural sampling. *Neuron* **90**, 649–660 (2016).

21. Bányai, M. *et al.* Stimulus complexity shapes response correlations in primary visual cortex. *Proceedings of the National Academy of Sciences* **116**, 2723–2732 (2019).
22. Sohl-Dickstein, J., Mudigonda, M. & DeWeese, M. Hamiltonian Monte Carlo without detailed balance. *arXiv preprint arXiv:1409.5191* (2014).
23. Ecker, A. *et al.* Decorrelated neuronal firing in cortical microcircuits. *Science* **327**, 584–587 (2010).
24. Wainwright, M. & Simoncelli, E. Scale mixtures of Gaussians and the statistics of natural images. In *Advances in Neural Information Processing Systems*, 855–861 (2000).
25. Coen-Cagli, R., Kohn, A. & Schwartz, O. Flexible gating of contextual influences in natural vision. *Nature Neuroscience* (2015).
26. Schwartz, O., Sejnowski, T. & Dayan, P. Perceptual organization in the tilt illusion. *Journal of Vision* **9**, 19–19 (2009).
27. Ahmadian, Y., Rubin, D. & Miller, K. Analysis of the stabilized supralinear network. *Neural Computation* **25**, 1994–2037 (2013).
28. Priebe, N. & Ferster, D. Inhibition, spike threshold, and stimulus selectivity in primary visual cortex. *Neuron* **57**, 482–497 (2008).
29. van Vreeswijk, C. & Sompolinsky, H. Chaotic balanced state in a model of cortical circuits. *Neural Computation* **10**, 1321–1371 (1998).
30. Hennequin, G., Vogels, T. & Gerstner, W. Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron* **82**, 1394–1406 (2014).
31. MacKay, D. *Information theory, inference and learning algorithms* (Cambridge university press, 2003).
32. Murray, J. *et al.* A hierarchy of intrinsic timescales across primate cortex. *Nature Neuroscience* **17**, 1661 (2014).
33. Grabska-Barwinska, A., Beck, J., Pouget, A. & Latham, P. Demixing odors-fast inference in olfaction. In *Advances in Neural Information Processing Systems*, 1968–1976 (2013).
34. Buesing, L., Bill, J., Nessler, B. & Maass, W. Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons. *PLoS computational biology* **7**, e1002211 (2011).
35. Savin, C. & Deneve, S. Spatio-temporal representations of uncertainty in spiking neural networks. In *Advances in Neural Information Processing Systems*, 2024–2032 (2014).
36. Hennequin, G., Aitchison, L. & Lengyel, M. Fast sampling-based inference in balanced neuronal networks. In *Advances in neural information processing systems*, 2240–2248 (2014).
37. Okun, M. & Lampl, I. Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nature Neuroscience* **11**, 535 (2008).
38. Carandini, M. & Heeger, D. Normalization as a canonical neural computation. *Nature Reviews Neuroscience* **13**, 51 (2012).
39. Echeveste, R., Hennequin, G. & Lengyel, M. Asymptotic scaling properties of the posterior mean and variance in the Gaussian scale mixture model. *arXiv preprint arXiv:1706.00925* (2017).
40. Yamins, D. *et al.* Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences* **111**, 8619–8624 (2014).
41. Festa, D., Hennequin, G. & Lengyel, M. Analog memories in a balanced rate-based network of EI neurons. In *Advances in Neural Information Processing Systems*, 2231–2239 (2014).
42. Song, H., Yang, G. & Wang, X. Training excitatory-inhibitory recurrent neural networks for cognitive tasks: A simple and flexible framework. *PLoS computational biology* **12**, e1004792 (2016).
43. Orhan, A. & Ma, W. Efficient probabilistic inference in generic neural networks trained with non-probabilistic feedback. *Nature communications* **8**, 138 (2017).
44. Remington, E., Narain, D., Hosseini, E. & Jazayeri, M. Flexible sensorimotor computations through rapid reconfiguration of cortical dynamics. *Neuron* **98**, 1005–1019 (2018).
45. Mante, V., Sussillo, D., Shenoy, K. & Newsome, W. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78 (2013).
46. Aitchison, L. & Lengyel, M. The Hamiltonian brain: efficient probabilistic inference with excitatory-inhibitory neural circuit dynamics. *PLoS computational biology* **12**, e1005186 (2016).
47. Rao, R. & Ballard, D. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience* **2**, 79 (1999).
48. Keller, G. & Mrsic-Flogel, T. Predictive processing: A canonical cortical computation. *Neuron* **100**, 424–435 (2018).
49. Vinck, M. & Bosman, C. More gamma more predictions: gamma-synchronization as a key mechanism for efficient integration of classical receptive field inputs with surround predictions. *Frontiers in systems neu-*

rosience **10**, 35 (2016).

50. Roelfsema, P., Lamme, V. & Spekreijse, H. Synchrony and covariation of firing rates in the primary visual cortex during contour grouping. *Nature Neuroscience* **7**, 982 (2004).

Methods

Ideal observer model

Following Refs. 5,25, we adopted the Gaussian scale mixture model (GSM)²⁴ as the generative model of natural image patches under which the primary visual cortex (V1) performs inference. Thus, an image patch $\mathbf{x} \in \mathbb{R}^{N_x}$ was assumed to be constructed by linearly combining a set of local features, the columns of $\mathbf{A} \in \mathbb{R}^{N_x \times N_y}$, weighted by a set of image patch-specific feature coefficients, $\mathbf{y} \in \mathbb{R}^{N_y}$, and scaled by a single global (at the scale of the image patch) contrast variable, $z \in \mathbb{R}$, plus additive white Gaussian noise, resulting in the following likelihood for the feature coefficients \mathbf{y} :

$$\mathbf{x}|\mathbf{y}, z \sim \mathcal{N}(z \mathbf{A} \mathbf{y}, \sigma_x^2 \mathbf{I}) \quad (1)$$

where the feature coefficients were assumed to be drawn from a multivariate Gaussian prior distribution:

$$\mathbf{y} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}) \quad (2)$$

and z was assumed to be drawn from a Gamma prior: $z \sim \Gamma(K, \vartheta)$ (Table S1, see also Ref. 39).

To model inferences in a V1 hypercolumn, we chose the columns of \mathbf{A} (the so-called projective fields of the latent variables) to be oriented Gabor filters that only differed by their orientation (evenly spaced between -90° and 90° , four examples are shown in Fig. 1a, see also Extended Data Fig. 1a). The prior covariance matrix \mathbf{C} was a circulant matrix whose elements varied smoothly as a function of the angular distance between the orientations of the projective fields of the corresponding latent variables, from positive (for similarly oriented projective fields) to negative (for orthogonally oriented projective fields) (Extended Data Fig. 1b).

The ideal observer's posterior over latent feature coefficients \mathbf{y} under the GSM for a given image patch, \mathbf{x} , and a known contrast z , can be written as⁵:

$$\mathcal{P}_{\text{GSM}}(\mathbf{y}|\mathbf{x}, z) = \mathcal{N}(\mathbf{y}; \boldsymbol{\mu}_{\text{GSM}}, \boldsymbol{\Sigma}_{\text{GSM}}) \quad (3)$$

$$\text{with } \boldsymbol{\mu}_{\text{GSM}} = \frac{z}{\sigma_x^2} \boldsymbol{\Sigma}_{\text{GSM}} \mathbf{A}^T \mathbf{x} \quad (4)$$

$$\text{and } \boldsymbol{\Sigma}_{\text{GSM}} = \left(\mathbf{C}^{-1} + \frac{z^2}{\sigma_x^2} \mathbf{A}^T \mathbf{A} \right)^{-1} \quad (5)$$

In general, z would also need to be inferred. However, as z is just a single scalar of which the inference pools information across all pixels in the input, we approximated the posterior over z with a delta distribution at z^* , the true value of z that was

used to generate the input³⁹. Thus, the final posterior over \mathbf{y} , after marginalizing out the unknown z , was approximated by substituting z^* into Eq. 3:

$$\mathcal{P}_{\text{GSM}}(\mathbf{y}|\mathbf{x}) \simeq \mathcal{P}_{\text{GSM}}(\mathbf{y}|\mathbf{x}, z^*) \quad (6)$$

Following Ref. 5, membrane potentials, \mathbf{u} , were taken to represent a weakly non-linear function of visual feature activations \mathbf{y} (Supplementary Math Note):

$$u_i(y_i) = \alpha_{\text{nl}} \lfloor y_i + \beta_{\text{nl}} \rfloor^{\gamma_{\text{nl}}} \quad (7)$$

where $\lfloor \cdot \rfloor$ is the threshold-linear function, and α_{nl} , β_{nl} , and γ_{nl} are respectively the scaling, baseline, and power of the transformation (Table S1, Fig. S1a).

Network dynamics and architecture

Our nonlinear, stochastic E/I network consisted of N_E excitatory and N_I inhibitory neurons. Following Ref. 6, we modeled the dynamics of each neuron i as:

$$\tau_i \frac{du_i}{dt} = -u_i(t) + h_i(t) + \sum_j W_{ij} r_j(t) + \eta_i(t) \quad (8)$$

where u_i represented the membrane potential of neuron i , τ_i was its membrane time constant, h_i its feedforward input, η_i was process noise (incorporating intrinsic and extrinsic forms of neural variability), and W_{ij} was the weight of the synapse connecting neuron j to neuron i . Firing rates r_i were given by a supralinear transformation of membrane potentials:

$$r_i(t) = k \lfloor u_i(t) \rfloor^n \quad (9)$$

where k and n were respectively the scale and exponent of the firing rate nonlinearity (Table S1).

We reasoned that any network performing accurate sampling-based inference under our ring-structured GSM would need to exhibit the same circular symmetry. We therefore parametrized the recurrent connectivity of the network to be rotationally symmetric, such that neurons were arranged into pairs of E and I cells around a "ring" according to their preferred orientations (Fig. 1c) and the connectivity of the network (as well as the process noise covariance, see below) was a smoothly decaying (circular Gaussian) function of the tuning difference between two cells. Specifically, each quadrant of the weight matrix (E \rightarrow E, E \rightarrow I, I \rightarrow E, and I \rightarrow I) was defined as:

$$\mathbf{W}_{XY}(\theta_i, \theta_j) = a_{XY} \exp \left(\frac{\cos[2(\theta_i - \theta_j)] - 1}{d_{XY}^2} \right) \quad (10)$$

where $X, Y \in \{E, I\}$ and $\theta_i = \pi i / N_{E/I}$ was the orientation represented by the i^{th} E/I neuron. Thus, we did not optimize all elements of the weight matrix, but only the eight free parameters a_{XY} and d_{XY} . We also constrained the a_{XY} amplitudes to

be positive for $Y = E$ and negative for $Y = I$, such that the network obeyed Dale's principle. This circulant parametrization implied that training the network on one particular stimulus-posterior pair in effect trained the network on all possible rotations of this pair. This reduced the size of the training set necessary to achieve good generalization, and therefore sped up training.

The stimulus-independent process noise (last term in Eq. 8) was spatially and temporally correlated zero-mean Gaussian (e.g. modeling inputs from other brain areas, or intrinsic variability in the network):

$$\langle \boldsymbol{\eta}(t) \rangle = 0, \quad \langle \boldsymbol{\eta}(t) \boldsymbol{\eta}(t+s)^\top \rangle = \boldsymbol{\Sigma}^\eta \exp(-s/\tau_\eta) \quad (11)$$

where τ_η was the timescale of the process noise (Table S1) and $\boldsymbol{\Sigma}^\eta$ was the stationary (zero-lag) covariance matrix parametrized block-wise as:

$$\boldsymbol{\Sigma}_{EE/II}^\eta(\theta_i, \theta_j) = \sigma_{E/I}^2 \exp\left(\frac{\cos[2(\theta_i - \theta_j)] - 1}{d_\sigma^2}\right) \quad (12)$$

$$\boldsymbol{\Sigma}_{EI}^\eta(\theta_i, \theta_j) = \rho \sigma_E \sigma_I \exp\left(\frac{\cos[2(\theta_i - \theta_j)] - 1}{d_\sigma^2}\right) \quad (13)$$

which introduced four additional free parameters: $\sigma_E > 0$, $\sigma_I > 0$, ρ , d_σ .

As in standard models of V1 simple cells⁵¹, the stimulus-dependent input to each neuron was obtained by applying a linear filter \mathbf{W}^{ff} to the stimulus followed by a static nonlinearity:

$$h_i(t) = \alpha_h \left[\beta_h + \sum_j \mathbf{W}_{ij}^{\text{ff}} x_j(t) \right]^{\gamma_h} \quad (14)$$

where $\mathbf{x}(t)$ was the stimulus (input image patch) received at time t , and α_h , β_h , and γ_h were respectively the scale, baseline, and exponent of the input nonlinearity (Table S1, Extended Data Fig. 1e). Given the one-to-one correspondence between the latent variables of the GSM and excitatory-inhibitory neuron pairs of the network model (Fig. 1a and c), we determined the external input to each neuron via an input receptive field that was identical (up to a constant factor) to the projective field of the corresponding GSM latent variable, as this was suggested to be optimal for sampling by previous work³⁰: $\mathbf{W}^{\text{ff}} = [\mathbf{A} \ \mathbf{A}]^\top / 15$, where \mathbf{A} was the same matrix as in the generative model (Eq. 1), and $[\mathbf{A} \ \mathbf{A}]$ denotes concatenating \mathbf{A} with itself column-wise.

In summary, we optimized a total of 15 parameters: 8 describing the weight matrix \mathbf{W} (Eq. 10), 4 describing $\boldsymbol{\Sigma}^\eta$ (Eqs. 12 and 13), and 3 specifying the mapping from stimuli to network inputs (Eq. 14).

Computing the moments of neural responses

For every time t relative to stimulus onset, we denote the across-trial moments of neural responses by

$$\boldsymbol{\mu}(t) = \langle \mathbf{u}(t) \rangle \quad (15)$$

$$\boldsymbol{\Sigma}(t, \Delta t) = \langle (\mathbf{u}(t) - \boldsymbol{\mu}(t)) (\mathbf{u}(t + \Delta t) - \boldsymbol{\mu}(t + \Delta t))^\top \rangle \quad (16)$$

where $\langle \cdot \rangle$ denotes trial-averaging. To compute these moments, we employed two different methods. The first approach, which we refer to as the "stochastic method", consisted of approximating the averages via sampling, i.e. simulating stochastic network dynamics in a set of trials using the same stimulus (Eqs. 8, 9 and 14) and computing the across-trial sample mean and sample covariance at each time step. We used this approach in the first phase of network training and for obtaining results from the network once it was trained (see below).

The second approach, which we refer to as "assumed density filtering" (ADF), used deterministic equations of motion for computing the across-trial moments. This approach was only used in the last phase of network training (see below). Based on Ref. 52, the following exact differential equations were used to describe the evolution of $\boldsymbol{\mu}(t)$ and $\boldsymbol{\Sigma}(t, \cdot)$:

$$\frac{d\boldsymbol{\mu}(t)}{dt} = \mathbf{T}^{-1} [-\boldsymbol{\mu}(t) + \mathbf{h}(t) + \mathbf{W} \boldsymbol{\nu}(t)] \quad (17)$$

$$\begin{aligned} \frac{d\boldsymbol{\Sigma}(t, 0)}{dt} &= [\mathbf{T}^{-1} \boldsymbol{\Sigma}^*(t)] + [\mathbf{T}^{-1} \boldsymbol{\Sigma}^*(t)]^\top + \\ &+ \mathcal{J}(t) \boldsymbol{\Sigma}(t, 0) + \boldsymbol{\Sigma}(t, 0) \mathcal{J}(t)^\top \end{aligned} \quad (18)$$

$$\frac{d\boldsymbol{\Sigma}^*(t)}{dt} = -\frac{1}{\tau_\eta} \boldsymbol{\Sigma}^*(t) + \boldsymbol{\Sigma}^\eta \mathbf{T}^{-1} + \boldsymbol{\Sigma}^*(t) \mathcal{J}(t)^\top \quad (19)$$

$$\begin{aligned} \frac{d\boldsymbol{\Sigma}(t, \Delta t)}{d\Delta t} &= e^{-\Delta t/\tau_\eta} [\mathbf{T}^{-1} \boldsymbol{\Sigma}^*(t)]^\top + \\ &+ \boldsymbol{\Sigma}(t, \Delta t) \mathcal{J}(t + \Delta t)^\top \quad \forall \Delta t > 0 \end{aligned} \quad (20)$$

where \mathbf{T} is the diagonal matrix of membrane time constants (Table S1), $\boldsymbol{\nu} = \langle \mathbf{r} \rangle$ is the average firing rate of neurons, $\boldsymbol{\Sigma}(t, -\Delta t) = \boldsymbol{\Sigma}^\top(t, \Delta t)$ is the time-lagged cross-covariance of membrane potentials in the network, $\boldsymbol{\Sigma}^* = \langle \boldsymbol{\eta} (\mathbf{u} - \boldsymbol{\mu})^\top \rangle$ is the instantaneous cross-covariance between membrane potentials and temporally correlated process noise with instantaneous covariance $\boldsymbol{\Sigma}^\eta$ (Eq. 11), and

$$\mathcal{J} = \mathbf{T}^{-1} \left[-\mathbf{I} + \mathbf{W} \text{diag} \left(\frac{\partial \boldsymbol{\nu}}{\partial \boldsymbol{\mu}} \right) \right] \quad (21)$$

is the Jacobian of Eq. 17 w.r.t. $\boldsymbol{\mu}$. Integrating Eqs. 17-20 in turn required evaluating some nonlinear moments of \mathbf{u} , namely covariances between membrane potentials, \mathbf{u} , and firing rates,

r. For the SSN, these moments can be obtained in closed form, assuming the full joint (space-time) distribution of membrane potentials is Gaussian⁵². Thus, in contrast to the first (stochastic) method, which leads to unbiased, but potentially high-variance, estimation of the moments, the ADF method leads to zero-variance, but potentially biased estimates. To train the network, as described below, we combined the strengths of these two approaches.

Training and test stimuli and target moments

The training set (Fig. 2b) consisted of five image patches:

$$\mathbf{x}^\alpha = z_\alpha \mathbf{A} \bar{\mathbf{y}} \quad (22)$$

with $\alpha = 1, \dots, 5$ and $z_\alpha \in \{0, 0.125, 0.25, 0.5, 1.0\}$. Therefore, these stimuli had the same content $\bar{\mathbf{y}}$ – a 27° -wide Gaussian function centered around 0° (i.e. a single dominant orientation) –, and differed only in their contrast (Fig. 2b). As the parametrization of our network was rotationally invariant (see above), such a stimulus was in fact representative of all image patches that could be obtained by rotating this patch around the center. For each training stimulus \mathbf{x}^α , we computed the corresponding posterior distributions over \mathbf{u} under the GSM (Eqs. 6 and 7). We called these distributions the “target distributions”, and their corresponding means $\boldsymbol{\mu}_{\text{tgt}}^\alpha$ and covariances $\boldsymbol{\Sigma}_{\text{tgt}}^\alpha$, the “target moments” (Fig. 2c,d):

$$\boldsymbol{\mu}_{\text{tgt}}^\alpha = \int \mathbf{u}(\mathbf{y}) \mathcal{P}_{\text{GSM}}(\mathbf{y}|\mathbf{x}^\alpha) d\mathbf{y} \quad (23)$$

$$\boldsymbol{\Sigma}_{\text{tgt}}^\alpha = \int \mathbf{u}(\mathbf{y}) \mathbf{u}^\top(\mathbf{y}) \mathcal{P}_{\text{GSM}}(\mathbf{y}|\mathbf{x}^\alpha) d\mathbf{y} - \boldsymbol{\mu}_{\text{tgt}}^\alpha \boldsymbol{\mu}_{\text{tgt}}^{\alpha \top} \quad (24)$$

To test generalization in the network, we generated a set of 500 novel image patches with the GSM, which were thus not constrained to have a single dominant orientation (as the prior allowed multiple elements of \mathbf{y} with different projective fields to be non-zero, Eq. 2). To be consistent with the training set, we did not include additive noise in \mathbf{x} , and added a contrast-dependent baseline to \mathbf{y} so that its mean was modulated by contrast in the same way as in the training set. For each image patch in the test set, we also computed the corresponding posterior moments (Eqs. 23 and 24) to evaluate the network’s test performance.

Network training

The cost function \mathcal{F} which we minimized during network training consisted of four terms for each

input stimulus α in the training set:

$$\mathcal{F} = \sum_{\alpha} (\epsilon_{\text{mean}} \phi_{\text{mean}}^\alpha + \epsilon_{\text{var}} \phi_{\text{var}}^\alpha + \epsilon_{\text{cov}} \phi_{\text{cov}}^\alpha + \epsilon_{\text{slow}} \phi_{\text{slow}}^\alpha) \quad (25)$$

The first three terms of Eq. 25 penalized differences between the (across trial) moments of the network’s response distribution (Eqs. 15 and 16) averaged over a finite time window ending at $T_{\text{max}} = 500$ ms after stimulus onset, and the respective target moments of the ideal observer’s corresponding posterior distributions (Eqs. 23 and 24):

$$\phi_{\text{mean}}^\alpha = \int_{T_{\text{min}}}^{T_{\text{max}}} \|\boldsymbol{\mu}^\alpha(t) - \boldsymbol{\mu}_{\text{tgt}}^\alpha\|_F^2 dt \quad (26)$$

$$\phi_{\text{var}}^\alpha = \int_{T_{\text{min}}}^{T_{\text{max}}} \|\boldsymbol{\sigma}^\alpha(t) - \boldsymbol{\sigma}_{\text{tgt}}^\alpha\|_F^2 dt \quad (27)$$

$$\phi_{\text{cov}}^\alpha = \int_{T_{\text{min}}}^{T_{\text{max}}} \|\boldsymbol{\Sigma}^\alpha(t, 0) - \boldsymbol{\Sigma}_{\text{tgt}}^\alpha\|_F^2 dt \quad (28)$$

where $\boldsymbol{\sigma}^\alpha(t) = \text{diag}(\boldsymbol{\Sigma}^\alpha(t, 0))$ and $\boldsymbol{\sigma}_{\text{tgt}}^\alpha = \text{diag}(\boldsymbol{\Sigma}_{\text{tgt}}^\alpha)$ are the response and target variances, respectively. The last term of Eq. 25 was an additional slowness cost, penalizing the total lagged neural response autocorrelation, given by the diagonal of $\mathbf{C}(\tau) = \text{corr}(\mathbf{u}(t), \mathbf{u}(t + \tau))$, within a $\tau_{\text{max}} = 100$ ms time window:

$$\phi_{\text{slow}}^\alpha = \int_0^{\tau_{\text{max}}} \|\text{diag}(\mathbf{C}^\alpha(\tau))\|_F^2 d\tau \quad (29)$$

The coefficients ϵ controlled the relative importance of these terms (Table S1). In the first control network (Fig. 5, right column; Extended Data Fig. 5a-g, and 6), we set $\epsilon_{\text{var}} = \epsilon_{\text{cov}} = \epsilon_{\text{slow}} = 0$, but kept all other meta-parameters and target means the same. In the second control network (Extended Data Fig. 5h-n, and 6), we set only $\epsilon_{\text{cov}} = \epsilon_{\text{slow}} = 0$, but left ϵ_{var} and other meta-parameters the same as in the original network. In the third control network (Extended Data Figs. 6 and 7, right), all ϵ_{\dots} parameters were the same as for the optimization of the original network, but the target covariances were modified to induce contrast-independent Fano factors (see below).

Optimization involved back-propagation through time⁵³, for which we used automatic differentiation. We trained the network in two stages. During the first stage, we employed a stochastic gradient method using $N_{\text{trial}} = 50$ trials for each training stimulus to estimate the corresponding moments of network responses (see above), and performed 250 iterations of the ADAM optimizer⁵⁴. Both the network’s initial conditions and the process noise were re-sampled for each trial and iteration. Initial conditions were drawn from a Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ (Table S1). Moreover,

across iterations, the beginning of the averaging time window, T_{\min} in Eqs. 26-28, was systematically changed (“annealed”) from $T_{\min} = 0$ ms (stimulus onset) to $T_{\max} - 50$ ms. The finite length of the averaging window, and in particular including samples immediately or shortly following stimulus onset, encouraged fast sampling. Thus, setting the explicit slowness cost $\epsilon_{\text{slow}} = 0$ did not qualitatively affect our results (Extended Data Figs. 6 and 10).

In the second stage, we continued optimization using the L-BFGS-B optimizer⁵⁵, now using the ADF method to (deterministically) compute the moments of the network’s response distribution (see above). We kept the cost-integration time window at its minimum ($T_{\max} - T_{\min} = 50$ ms, as reached by the end of the first phase). The slowness penalty cost in Eq. 29 was only applied during ADF-based optimization and, for simplicity, it was approximated using stationary lagged correlations predicted by the ADF-method (Eq. 20) in the limit of temporally white process noise⁵².

As the cost function that we used (Eq. 25) was non-convex, we checked the robustness of our findings by performing 10 further optimization attempts from random initial conditions. No solutions achieved substantially lower costs, and those whose final cost was at least approximately as low as the network presented in the main text behaved qualitatively similarly (in particular, they showed contrast-dependent oscillations and transients, Extended Data Fig. 4). Nevertheless, our results should not be taken to represent a global optimum of our cost function.

Langevin sampling

As a comparison (Fig. 4), we also implemented Langevin dynamics³¹ to sample from the same target posteriors as those used to train the optimized network. As the GSM target posteriors were Gaussian (Eqs. 3 and 6), the resultant Langevin dynamics was isomorphic to that of a generic stochastic linear recurrent neural network:

$$\tau_E \dot{\mathbf{u}} = \mathbf{W}_L \mathbf{u} + \mathbf{h} + \boldsymbol{\eta} \quad (30)$$

where, for a fair comparison, τ_E and $\boldsymbol{\eta}$ were respectively the same time constant and process noise as in our optimized network (Eqs. 8 and 11-13). Without loss of generality, we set the input to the network $\mathbf{h} = 0$, as it would be completely determined by the requirement to match the response mean to the target mean, $\boldsymbol{\mu}_{\text{tgt}}$, but would not affect the autocorrelogram of the system, which was the focus of our investigation here. As variability in a linear network does not

depend on the input (unlike in our nonlinear circuit model), we used a different \mathbf{W}_L to match each target covariance, $\boldsymbol{\Sigma}_{\text{tgt}}$:

$$\mathbf{W}_L = \frac{1}{\gamma_L} \left(\mathbf{I} - \sqrt{\mathbf{I} + \gamma_L^2 \boldsymbol{\Sigma}^\eta \boldsymbol{\Sigma}_{\text{tgt}}^{-1}} \right) \quad (31)$$

where $\gamma_L = 2\tau_\eta/\tau_E$.

Numerical experiments after training

To obtain a reliable estimate of the stationary moments of neural responses to a fixed input (Figs. 2 and 3), a total of 20,000 independent samples (taken 200 ms apart) were drawn from the network, not including transients, as neural activity evolved according to Eq. 8. Neural activities in Fig. 2a show 1 s of simulated network activity, convolved with a 20 ms sliding window to match the effects of spike binning to compute average rates in experiments. Neural trajectories in Fig. 2b correspond to the neural activity of two cells in the network with preferred orientations 42° (u_i) and 16° (u_j), over a post-transient period of 500 ms. To illustrate both the degree of modulation of the posterior covariances and the match between posterior and network covariances in Fig. 3c, the top three PCs of each posterior covariance were computed. Neural activity was then projected onto each PC, and the amount of variance along each direction was computed. The middle plots of Fig. 3c present these posterior PCs scaled by either the square root of the total variance along that direction in the GSM (in green) or in the network (in red).

Autocorrelograms in Fig. 4a were computed in 500 non-overlapping windows of 2 s of simulated neural activity each (subsampling at 0.4 ms) after stimulus onset (excluding transients), and then averaged across these windows. Autocorrelograms were first computed for individual cells’ membrane potentials and then averaged across all cells. Cross-correlograms and E-I lags in Fig. 4b–c were computed from a single 400 s-long simulation after stimulus onset, excluding transients (without subsampling). The E-I lag for each cell was determined as the location of the maximum in the anti-symmetric component of the cross-correlogram between its total E and I input. Langevin samplers in Fig. 4a–b corresponded to neural networks with linear, time-reversible dynamics, not respecting Dale’s principle, as defined by Eqs. 30 and 31. Autocorrelograms and cross-correlograms for the Langevin sampler were computed as for the original network.

Average firing rates in Fig. 5a were computed from the same neural traces used in Fig. 2 to compute \mathbf{u} moments (here taking the average

of \mathbf{r} instead of \mathbf{u}). To compute Fano factors in Fig. 5b, we used an inhomogeneous Gamma process with the time-varying rate given by $r_i(t)$ for each neuron i , and the shape of the interspike interval distribution controlled by an additional parameter, K_{ISI} (Table S1). We computed spike counts in a 100 ms window over 500,000 independent trials. Our results were qualitatively robust to the choice of K_{ISI} , which primarily determined the overall magnitude of Fano factors – in particular $K_{\text{ISI}} > 1$ was needed to achieve Fano factors < 1 at high contrast – but not their modulation by stimuli.

Power spectra in Fig. 5c were based on simulated local field potentials (LFPs), computed as the (across-cell) average neural activity (membrane potentials), following standard approaches⁶, using the same samples as the autocorrelograms of Fig. 4a (see above). Gamma peak frequency was identified as the location of the local maximum (within the gamma band, 20–80 Hz) of the power spectrum. Transients in Fig. 5d were computed from average firing rates across E cells and trials ($n = 100$), further averaged over a sliding 10-ms time window to mimic the resolution of experimental data. To account for the response delays observed in experimental data, we used a random delay time (truncated Gaussian, with 45 ms mean and 5 ms s.d.) for the feedforward input of each E–I cell pair in the network.

To estimate input conductances (Fig. 5e), we equated the total (excitatory or inhibitory) input current in our model to the (excitatory or inhibitory) current in a canonical conductance-based model⁵¹. This gave the following expression:

$$\mathbf{g}_i^{\text{E/I}}(t) \approx \frac{C_m \sum_{j \in \text{E/I}} \mathbf{W}_{ij} \mathbf{r}_j(t)}{\tau_i [V^{\text{E/I}} - (\mathbf{u}_i(t) + V_{\text{rest}})]} \quad (32)$$

where C_m is the membrane capacitance, $V^{\text{E/I}}$ denote the reversal potentials for E/I currents, and V_{rest} is a baseline (resting) potential added to the membrane potentials of our model. We chose $C_m = 20.0$ pF, $V^{\text{E}} = 0$ mV, $V^{\text{I}} = -80$ mV, and $V_{\text{rest}} = -65$ mV. Conductances in Fig. 5e are shown relative to their steady-state values during spontaneous activity, and averaged across 20 trials for a single neuron with preferred orientation aligned to that of the stimulus.

Autocorrelograms and power spectra of Fig. 6c were computed as in Fig. 4a and Fig. 5c, but for the directions in the space of neural responses that corresponded to the first ten principal components (PCs) of neural variability. To quantify oscillatoriness in neural responses along some direction in state space (PC, Fig. 6d; or LFP, Extended

Data Fig. 4b), we computed the corresponding projection of neural responses, and then fitted the following parametric function to its autocorrelogram:

$$C_\chi(\Delta t) = [(1 - \chi) + \chi \cos(2\pi f \Delta t)] \cdot \left[\frac{\hat{\tau}}{\hat{\tau} - \hat{\tau}_\eta} e^{-\frac{|\Delta t|}{\hat{\tau}}} - \frac{\hat{\tau}_\eta}{\hat{\tau} - \hat{\tau}_\eta} e^{-\frac{|\Delta t|}{\hat{\tau}_\eta}} \right] \quad (33)$$

where $\hat{\tau}$ and $\hat{\tau}_\eta$ are two time constant parameters, $\chi \in [0, 1]$ quantifies the degree of oscillatoriness, and f represents the dominant oscillation frequency. Fits of Eq. 33 to simulated network activity were performed using Tensorflow. The form of Eq. 33 was motivated by noting that, for $\chi = 0$, it reduces to

$$C_0(\Delta t) = \frac{\hat{\tau}}{\hat{\tau} - \hat{\tau}_\eta} e^{-\frac{|\Delta t|}{\hat{\tau}}} - \frac{\hat{\tau}_\eta}{\hat{\tau} - \hat{\tau}_\eta} e^{-\frac{|\Delta t|}{\hat{\tau}_\eta}} \quad (34)$$

which is the autocorrelation function of the fluctuations in a single (isolated, and thus not oscillating) neuron with membrane time constant $\hat{\tau}$ receiving noisy inputs with correlation time $\hat{\tau}_\eta$ (this can be seen by integrating Eqs. 18–20). More generally, for $\chi > 0$, $C_0(\Delta t)$ (Eq. 34) determines the envelope of $C_\chi(\Delta t)$ (Eq. 33):

$$(1 - 2\chi) C_0(\Delta t) \leq C_\chi(\Delta t) \leq C_0(\Delta t) \quad (35)$$

Overshoots in Fig. 7c and d were obtained using the same stimulus that was used to train the network at 0.7 contrast, and computed as the maximal across-trial average ($n = 100$) response of each E cell (membrane potential for c, firing rate for d), minus its stationary mean response, further averaged over 1000 delay configurations in our network (as for Fig. 5d, see above). Steady-state differences denote the magnitude of mean evoked responses of each cell with respect to its mean pre-stimulus response. Results in the bottom plot of Fig. 7d were computed by averaging stimuli presented at each neuron’s preferred orientation ($\pm 30^\circ$) or orthogonal to its preferred orientation ($\pm 30^\circ$).

Analyses with Gaussian processes

For Fig. 6a-b, we considered a stationary, moment-matched Gaussian process (GP⁵⁶) with an autocorrelation function that was in the same parametric form as that fitted to network responses, $C_\chi(\Delta t)$ (Eq. 33), by setting both time constants to their values characterizing the network ($\hat{\tau}_\eta = \hat{\tau} = \tau_{\text{E}} = \tau_\eta = 20$ ms). This yielded the following form:

$$C_{\text{GP}}(\Delta t) = [(1 - \chi) + \chi \cos(2\pi f |\Delta t|)] \cdot \left(1 + \frac{|\Delta t|}{\tau_{\text{E}}} \right) e^{-\frac{|\Delta t|}{\tau_{\text{E}}}} \quad (36)$$

We used $\chi \in \{0, 0.2, 0.7\}$ (where $\chi = 0$ corresponds to an isolated single neuron) at $f = 40$ Hz.

We set the mean and variance of the GP to $\mu_{\text{GP}} = 3$ and $\sigma_{\text{GP}}^2 = 4$, respectively, without loss of generality (Supplementary Math Note).

For Fig. 7a and b, and Extended Data Fig. 9, we used GPs whose stationary distributions all matched the same target distribution (see below) before stimulus onset (time $t < 0$), and converged to the same new target distribution after stimulus onset (for $t \gg 0$). These GPs differed only in their transient behavior. The GP that most faithfully captured our optimized network was constructed by taking the temporal evolution of the mean and variance of an actual neuron in the full network around stimulus onset, including overshoots (“with overshoots” in Fig. 7a, and Extended Data Fig. 9a, red). For this, we chose the neuron whose preferred orientation matched that of the presented stimulus. Moreover, the autocorrelogram of this GP was also set to match that neuron’s autocorrelogram (inset of Extended Data Fig. 9a, right). Two additional GPs were constructed with the same autocorrelogram as the first GP, but different mean and variance time courses. In one, the mean and variance converged exponentially with a time constant of τ_E (“exponential” in Fig. 7a, and Extended Data Fig. 9a, dashed black). In the other, the mean and variance immediately jumped at stimulus onset to their new stationary values (“instantaneous” in Fig. 7a, and Extended Data Fig. 9a, dashed gray). While such an instantaneous process is not realisable by any continuous dynamical system, it provides a useful lower bound on sampling error.

In Fig. 6b and Extended Data Fig. 9b, we used the average symmetrized Kullback-Leibler divergence between the GSM target distribution, \mathcal{P} , and the distribution sampled by the GP, \mathcal{Q}_T over some finite time window T to measure the performance of these systems:

$$\bar{D}_{\text{SKL}}[\mathcal{P}||\mathcal{Q}_T] = \frac{1}{4} [\langle \mathcal{E}^2(T) (\sigma_{\mathcal{P}}^{-2} + \sigma_{\mathcal{Q}}^{-2}(T)) \rangle + \langle \sigma_{\mathcal{Q}}^2(T) \rangle \sigma_{\mathcal{P}}^{-2} + \langle \sigma_{\mathcal{Q}}^{-2}(T) \rangle \sigma_{\mathcal{P}}^2 - 2] \quad (37)$$

where $\langle \cdot \rangle$ denotes trial-averaging (as above), $\mathcal{E}(T) = \mu_{\mathcal{Q}}(T) - \mu_{\mathcal{P}}$, and $\mu_{\mathcal{P}}$, $\mu_{\mathcal{Q}}(T)$, $\sigma_{\mathcal{P}}^2$, and $\sigma_{\mathcal{Q}}^2(T)$ are the means and variances of \mathcal{P} and \mathcal{Q}_T , respectively (Supplementary Math Note). Fig. 7b only shows the term of \bar{D}_{SKL} that depends on the sample mean, $\mu_{\mathcal{Q}}(T)$ (via $\mathcal{E}^2(T)$):

$$\bar{D}_{\text{SKL}}^{\mu} = \frac{1}{4} \langle \mathcal{E}^2(T) (\sigma_{\mathcal{P}}^{-2} + \sigma_{\mathcal{Q}}^{-2}(T)) \rangle \quad (38)$$

In Fig. 6b, we used a single target distribution (matching the stationary moments of the GP, see above) and averaged over $n=100,000$ trials of sampling from the GP for an increasing amount of time, T (x-axis), as we also varied χ (lines). GPs

were sampled at a $dt=0.5$ ms time resolution. In Fig. 7b and Extended Data Fig. 9b, the target distribution was changed instantaneously at stimulus onset ($t=0$) from a response distribution corresponding to a 0-contrast stimulus (spontaneous activity) to that corresponding to a high contrast stimulus (training image at contrast level 0.7; Fig. 7a and Extended Data Fig. 9b, green lines). Response distributions were obtained from the E cell in the original network that was tuned to the dominant orientation of the stimulus. At each time point t (x-axis), \bar{D}_{SKL} (or $\bar{D}_{\text{SKL}}^{\mu}$) was computed between the momentary target distribution and the distribution of GP samples collected in a 100 ms-long sliding time window ending at t , averaging over $n=10,000$ trials, for each of the three transition profiles (lines). GPs were sampled at a $dt=1$ ms time resolution.

Experimental data analysis

Fig. 7e shows novel analyses of experimental recordings from awake macaque V1 during the presentation of moving gratings of different orientations²³. Following the same procedure as in Ref. 6, only cells that were significantly tuned (orientation tuning index greater than 0.75) and had an average evoked rate above 1 spike per second were included in the analysis. For each cell and each stimulus, a time-dependent firing rate trace was first obtained by averaging spikes across trials in a 50 ms sliding square window. From these traces, the steady-state difference and overshoot size (dots in Fig. 7e top) were then computed respectively as the average evoked response excluding transients ($t > 160$ ms after stimulus onset) and the maximum of the response trace during the transient ($t < 160$ ms after stimulus onset), minus the average baseline response (computed from the 300 ms prior to stimulus presentation). Results in the bottom plot of Fig. 7e were computed in the same way as for model neurons (see above).

Statistics

Sample sizes (number of trials) were chosen with the following criteria. First, when studying properties of the network, a large enough n was selected ($> 10,000$ independent samples) such that standard errors would be smaller than line widths in the corresponding plots. Second, when comparing the network’s behaviour to experimental results, the same order of magnitude for n was selected as in the original experiments.

Linear regressions in Fig. 7e (top) and Extended Data Fig. 9d were performed using SciPy’s ‘linregress’ function, which reports a two-sided p-

value using a Wald test with a t-distribution of the test statistic.

In Fig. 7e (bottom) we tested for significance in overshoot tuning using a two-sided paired t-test using SciPy's 'ttest_rel' function. Data distribution was assumed to be normal but this was not formally tested.

Reporting Summary. Further information on research design is available in the Life Sciences Reporting Summary linked to this article.

Randomization. No new experimental data was gathered for this paper.

Blinding. As data collection had been performed by other labs well before our study, it was effectively blind to the purposes of our study. Data analysis was not performed blind to the conditions of the experiments.

Data exclusion. In Fig. 7e only cells with a tuning index greater than 0.75 were selected. No other data were excluded in the main text. In Extended Data Fig. 9 the same analysis of Fig. 7e was repeated excluding outliers to test for robustness.

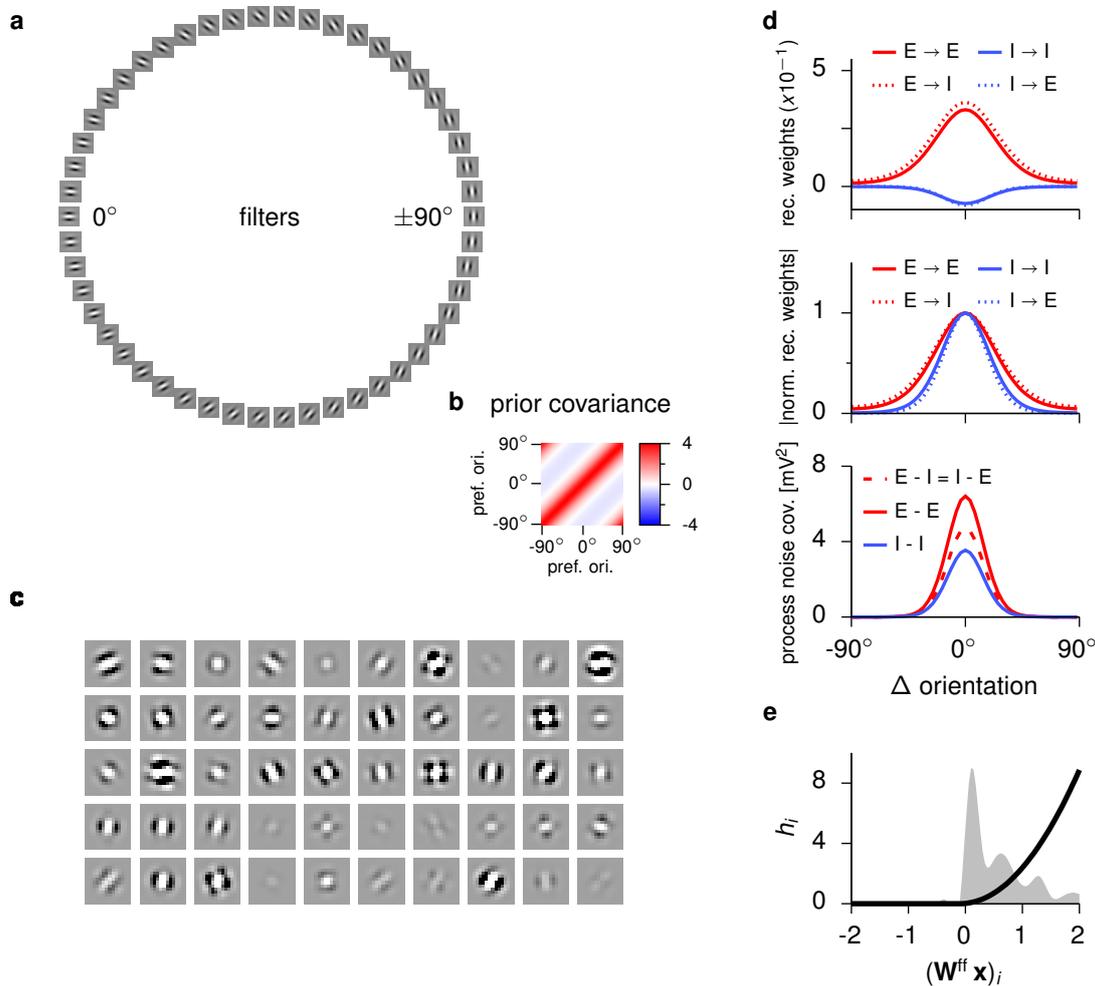
Data availability. All experimental data reported here has been collected by others and published previously. Experimental data in Fig. 5a-b reproduce analyses from Ref. 6 of data from Ref. 23 (data released in the repository of Ref. 57). Experimental data in Fig. 5c-d was captured directly from the plots in the original papers Refs. 2–4. Experimental data in Fig. 7 is a novel analysis of the data from Ref. 23, with permission from the authors. The code and parameters to generate the data corresponding to numerical experi-

ments in the paper has been made publicly available (see Code availability statement below).

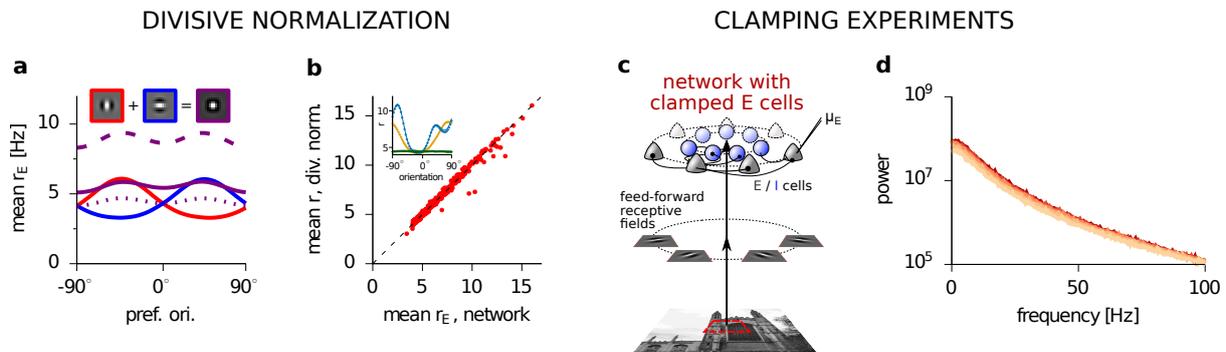
Code availability. The (Python) code and parameters for the numerical experiments are available at: bitbucket.org/RSE_1987/ssn_inference_numerical_experiments. The (OCaml) code for the optimization procedure can be found at: bitbucket.org/RSE_1987/ssn_inference_optimizer.

References

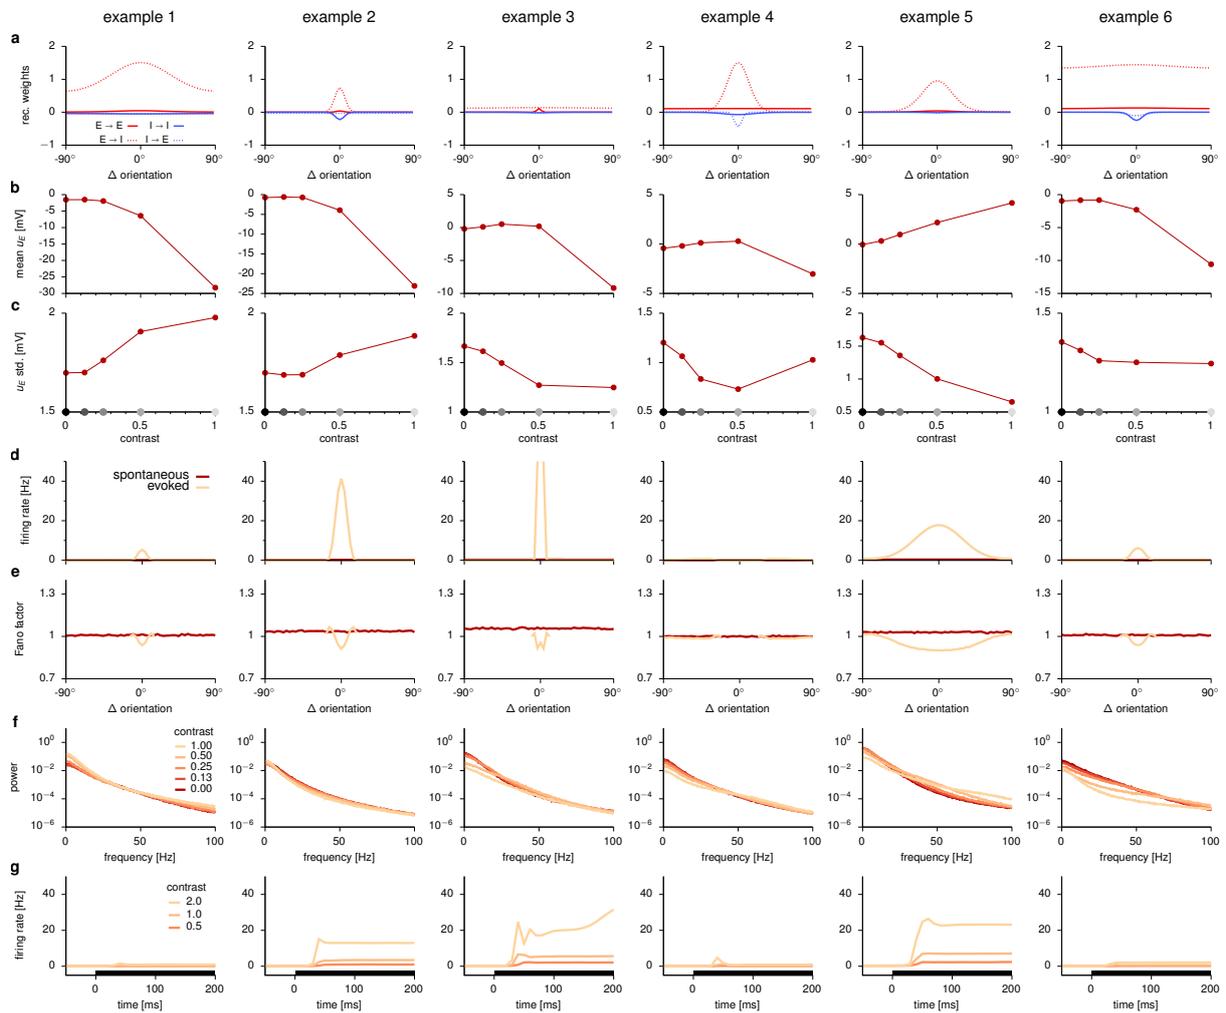
51. Dayan, P. & Abbott, L. *Theoretical neuroscience*, vol. 806 (Cambridge, MA: MIT Press, 2001).
52. Hennequin, G. & Lengyel, M. Characterizing variability in nonlinear recurrent neuronal networks. *arXiv preprint arXiv:1610.03110* (2016).
53. Werbos, P. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE* **78**, 1550–1560 (1990).
54. Kingma, D. & Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
55. Zhu, C., Byrd, R., Lu, P. & Nocedal, J. Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Transactions on Mathematical Software (TOMS)* **23**, 550–560 (1997).
56. Williams, C. & Rasmussen, C. *Gaussian processes for machine learning*, vol. 2 (MIT Press Cambridge, MA, 2006).
57. Ecker, A. *et al.* State dependence of noise correlations in macaque primary visual cortex. *Neuron* **82**, 235–248 (2014).



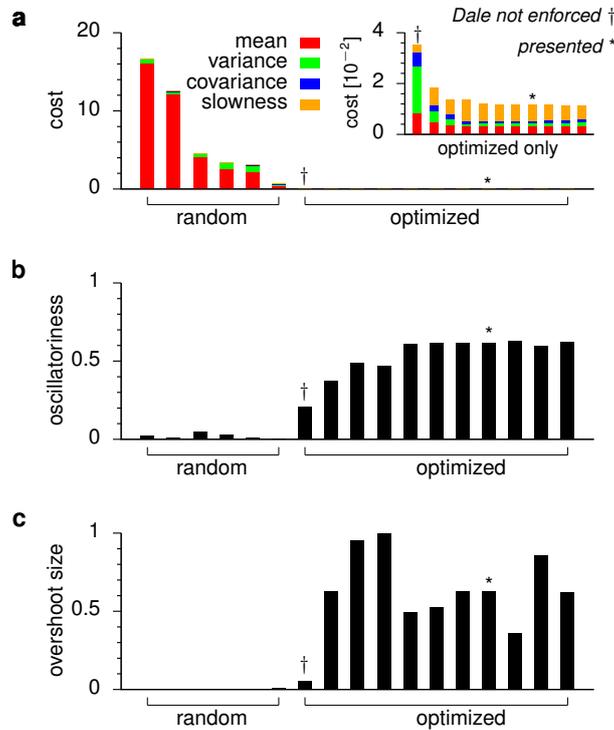
Extended Data Fig. 1 | GSM and network parameters. **a**, Filters: projective fields of the GSM and receptive fields of the network. Each filter image shows the projective field of a latent variable (columns of A ; Eq. 1), which was the same as the receptive field of the corresponding E-I cell pair in the network (rows of $\mathbf{W}^{\text{ff}} = [AA]^T / 15$; Eq. 14; cf. Fig. 1a and c). **b**, Prior covariance in the GSM (C in Eq. 2). **c**, Sample stimuli generated by the GSM, also used for testing the network's generalization in Fig. 3b-c. **d-e**, Parameters of the optimized network. **d**, Recurrent weights (top: raw weights; middle: normalized absolute values) and process noise covariance (bottom) after training. Weights and covariances are shown for only one row (of each quadrant) of \mathbf{W} (Eq. 8) and Σ^{η} (Eq. 11), respectively, as they are circulant. Thus, each line shows the weights connecting, or the covariance between, cells of different types (see legend) as a function of the difference in their preferred stimuli. As the figure shows, the connectivity profile of either E or I cells in the optimized network was largely independent of whether the postsynaptic cell was excitatory or inhibitory (top). Overall, recurrent E and I connections had similar tuning widths, with E connections being slightly more broadly tuned than I ones (middle). Nevertheless, the net E input to any one cell in the network was still more narrowly tuned than the net I input, due to the responses of presynaptic E cells being more narrowly tuned than those of I cells (not shown). The optimized network also retained a substantial amount of process noise that was larger in E than in I cells, and highly correlated both between the E and I cell of a pair and between cells with different tuning (up to a $\sim 30^\circ$ tuning difference; bottom). **e**, Input nonlinearity (Eq. 14), converting feedforward receptive field activations $(\mathbf{W}^{\text{ff}} \mathbf{x})_i$ into network inputs h_i (black). For comparison, the distribution of inputs across all cells for the training set is presented in gray. As the figure shows, the optimized input transformation, capturing the nonlinear effects of upstream preprocessing of visual stimuli, had a threshold that was just below the distribution of receptive field outputs, ensuring that all stimulus-related information was transmitted in the input signal, and an exponent close to two, remarkably similar to that used by the cells of the network (cf. Eq. 9 and table S1).



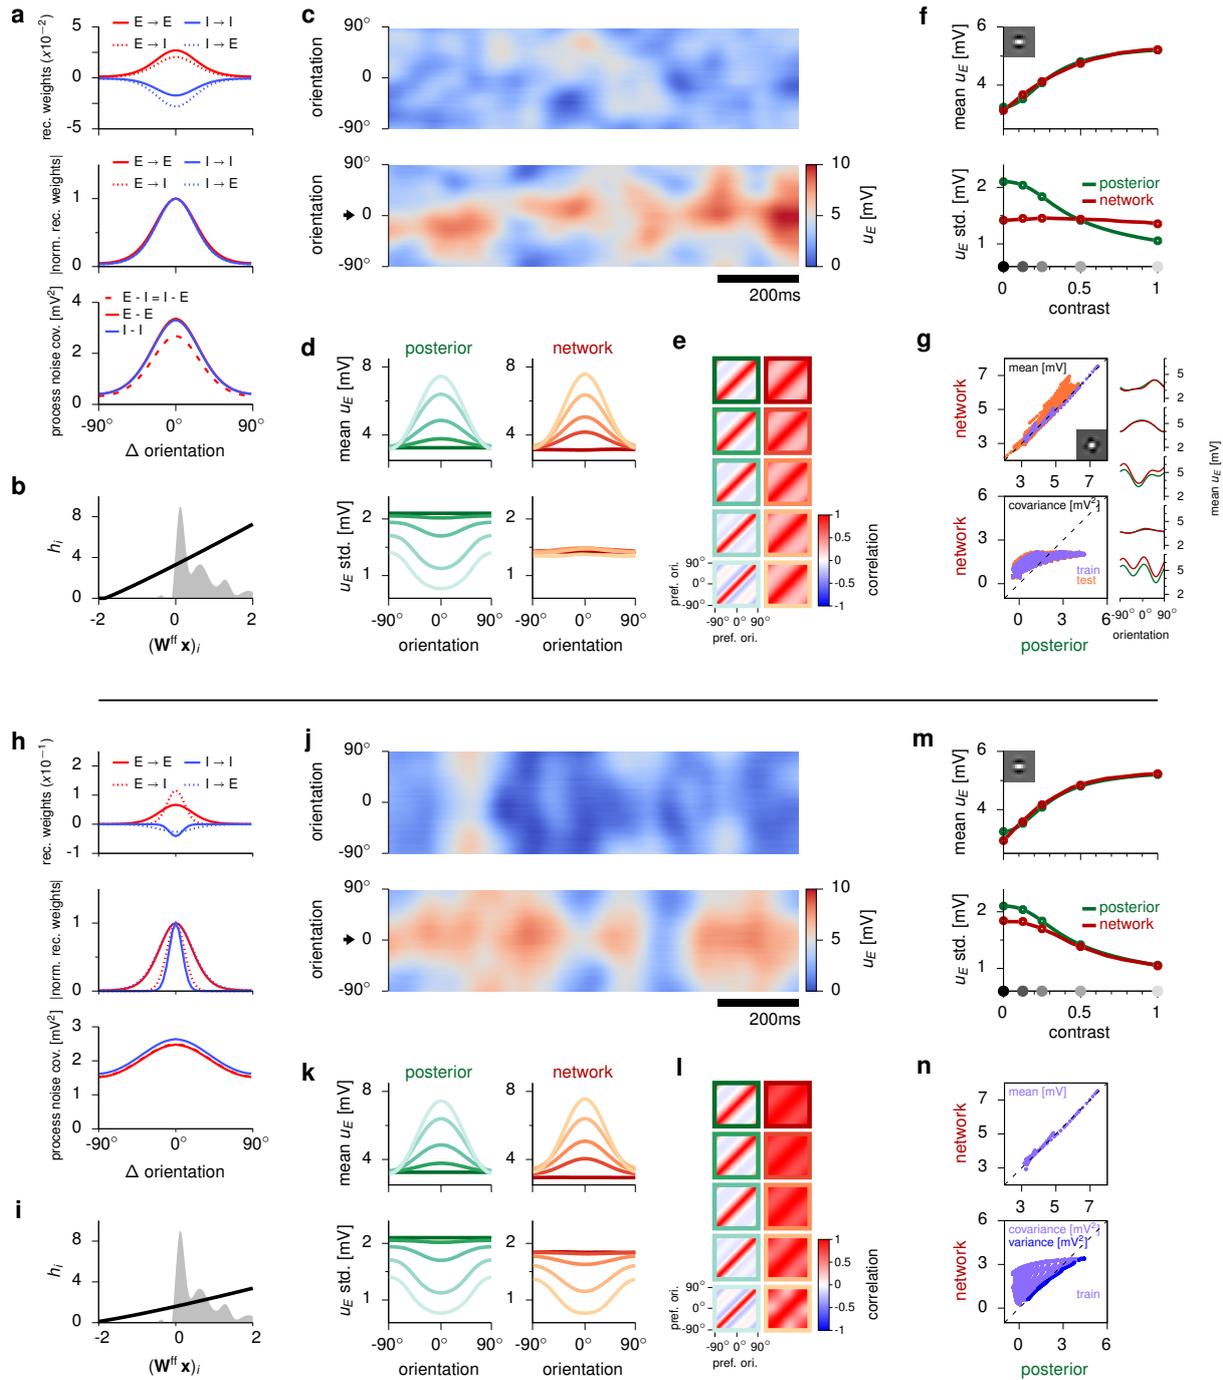
Extended Data Fig. 2 | Divisive normalization and the mechanism underlying oscillations in the optimized network. **a-b**, Divisive normalization, or sublinear summation of neural responses, has been proposed as a canonical computation in cortical circuits³⁸. In turn, the stabilized supralinear network (SSN), which formed the substrate of our optimized network, has been proposed to provide the dynamical mechanism underlying divisive normalization²⁷. We thus wondered whether our optimized network also exhibited it. **a**, In accordance with divisive normalization, the network's response to the sum of two stimuli (solid purple) was smaller than the sum of its responses to the individual stimuli (solid red/blue), and lay between the average (dotted purple) and the sum (dashed purple). Inset shows stimuli used in this example. **b**, Generic divisive normalization in the optimized network. We fitted a standard phenomenological model of divisive normalization (adapted from Ref. 38) to the (across-trial) mean firing-rate responses of E cells in the optimized network, $\langle r_E \rangle$, as a function of the feedforward input \mathbf{h} to the network (i.e. without regard to its recurrent dynamics; Eq. 14): $\langle r \rangle_i^{(\beta)} = b_2 + \left(h_i^{(\beta)} + b_1 \right) \left[(\mathbf{M} \mathbf{h}^{(\beta)})_i + s^2 \right]^{-1}$, where $h_i^{(\beta)}$ and $\langle r \rangle_i^{(\beta)}$ were the feedforward input and the average firing rate of cell i in response to stimulus β , respectively, and b_1 , b_2 and s were constant parameters. The parameter matrix \mathbf{M} was responsible for normalization, by dividing the input h_i by a mixture of competing inputs to other neurons. \mathbf{M} was parameterized as a symmetric circulant matrix to respect the rotational symmetry of the trained network (Extended Data Fig. 1). In total, our model of divisive normalization had $3 + (N_E/2) + 1 = 29$ free parameters. Model fitting was performed via minimization of the average squared difference between network and model rates, plus an elastic energy regularizer for neighbouring elements of \mathbf{M} . Shown here is a scatter plot of neural responses to a set of 500 random stimuli (generated as the generalization dataset, Methods), predicted by the phenomenological model vs. produced by the actual network. Each dot corresponds to a stimulus-neuron pair. Inset shows three representative average response profiles across the network (dots) and the phenomenological model's fit (lines). Note the near perfect overlap between the network and the phenomenological model in all three cases. The divisive normalization model also outperformed both a linear model and a model of subtractive inhibition (not shown). These results show comprehensively that, in line with empirical data, our trained network performed divisive normalization of its inputs under general conditions. **c-d**, Using voltage-clamp to study the mechanisms underlying oscillations in the optimized network. To determine whether oscillations in our network resulted from the interaction of E and I cells, or whether they arose within either of these populations alone, we conducted two simulated experiments. In each simulation, either of the two populations (E or I) was voltage-clamped to its (temporal) mean, as calculated from the original network, for each input in the training set. Thus, recurrent input from the clamped population was effectively held constant to its normal mean, but did not react to changes in the other population. As expected for a network in the inhibition-stabilized regime, clamping of the I cells resulted in unstable runaway dynamics, precluding further analysis of oscillations (not shown). **c**, Illustration of the E-clamping experiment in the optimized network (cf. Fig. 1b). For each stimulus, each E cell's voltage was clamped to its mean voltage, μ_E , obtained when the network was presented with the same stimulus without voltage clamp. **d**, LFP power spectra in the network after voltage-clamping of the E population at different contrast levels (colors as in Fig. 2). The network remained stable, but the peak in its LFP power spectrum characteristic of gamma oscillations was no longer present (cf. Fig. 5c, inset). This shows that gamma oscillations in the original network required interactions between E and I cells (i.e. they were generated by the so-called "PING" mechanism¹⁶).



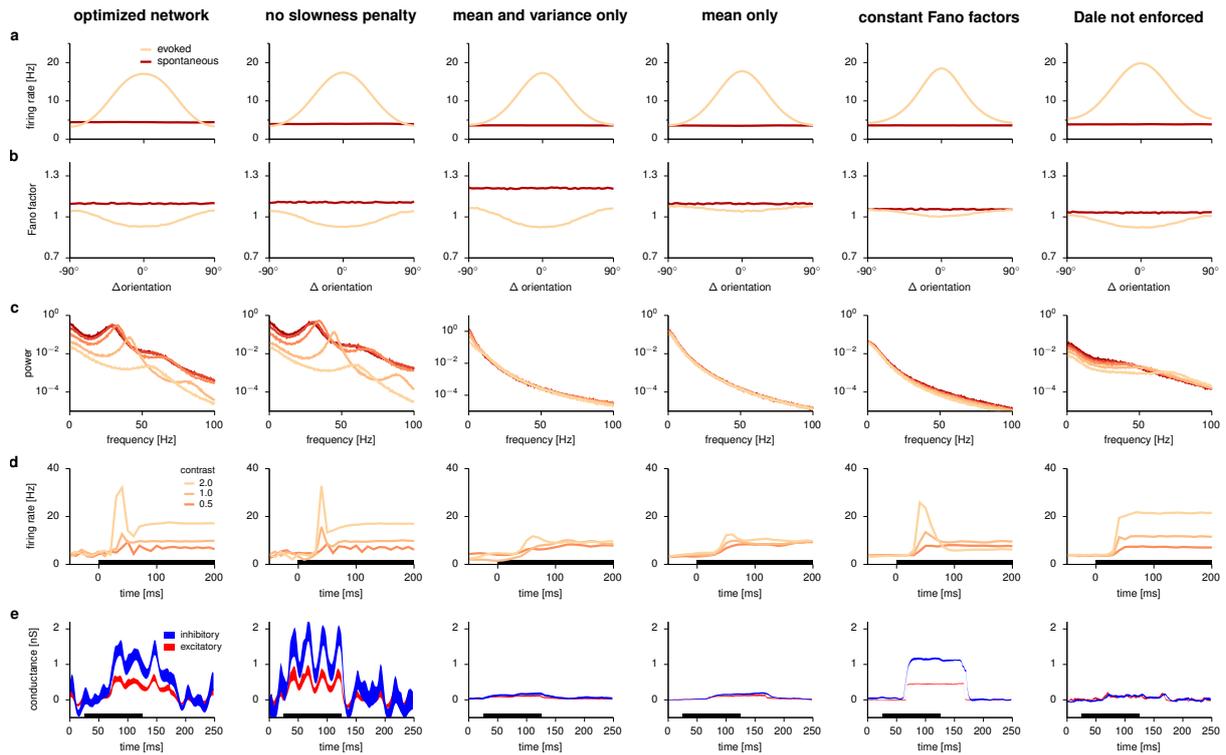
Extended Data Fig. 3 | Random networks. The parametrization of our network was highly constrained (e.g. ring topology with circulant and symmetric weight and process noise covariance structure, 1:1 E:I ratio, fixed receptive fields). To test whether these constraints alone, without any optimization, were sufficient to generate the results we obtained in the optimized network, we sampled weight matrices at random, by drawing each of the 8 hyperparameters of the weight matrix (Eq. 10) from an exponential distribution truncated between 0.1 and 10 times the values originally found by optimization. We discarded matrices that were either unstable or converged to a trivial solution (all mean rates equal 0). Less than 20% of the generated matrices satisfied these criteria, further confirming that optimization was non-trivial. Results for six such example random networks are shown (columns). **a**, Recurrent weights as a function of the difference in the preferred stimuli of two cells (cf. Extended Data Fig. 1d, top). Different lines are for weights connecting cells of different types (legend). **b-c**, Mean (**b**) and standard deviation of membrane potential responses (**c**), averaged over the population, as a function of contrast (cf. Fig. 3a). Gray dots on x-axis indicate training contrast levels. Note the wide array of behaviors displayed by these networks. For example, the standard deviation of responses could go up, down, or even be non-monotonic with contrast, while the range of mean rates also varied wildly. **d-e**, Mean firing rate (**d**) and Fano factor (**e**) of neurons as a function of stimulus orientation (relative to their preferred orientation) during spontaneous (dark red) and evoked activity (light orange; cf. Fig. 5a-b). The peak mean rate of example 3 exceeded 50 Hz and is thus shown as clipped in this figure. Note that mean rate tuning curves (during evoked activity) were very narrow for most networks (all but example 5), resulting in 0 Hz rates and thus undefined Fano factors for stimuli further away from the preferred orientation. **f**, LFP power spectra at different contrast levels (colors as in Fig. 2). Note the absence of gamma peaks (cf. Fig. 5c, inset). **g**, Average rate response around stimulus onset at different contrast levels (colors as in Fig. 2). Black bars show stimulus period. Note the absence of transients (cf. Fig. 5d). To estimate network moments in **b-d**, $n = 20,000$ independent samples (taken 200 ms apart) were used. Population averages ($n=50$ cells) were computed for **b** and **c**. Mean firing rates in panel **g** were computed over $n = 100$ trials.



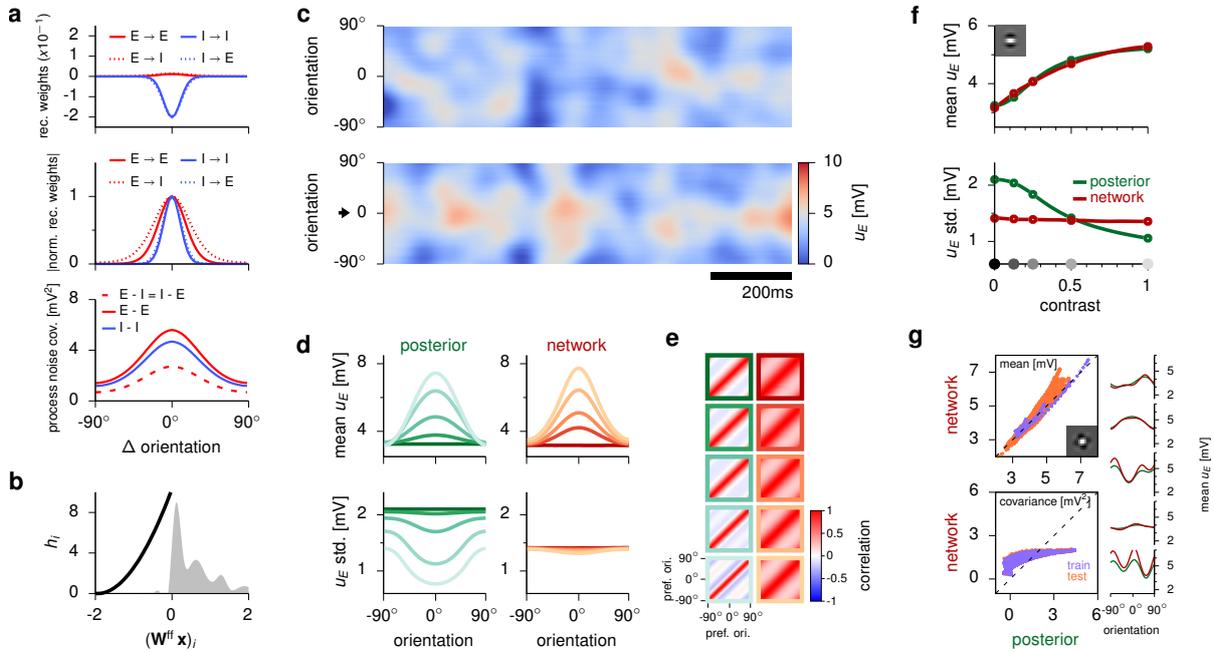
Extended Data Fig. 4 | Comparison of random and optimized networks: cost achieved and dynamical features. Networks are ranked in all panels in order of decreasing total cost achieved by them (shown in **a**). Random networks are those shown in [Extended Data Fig. 3](#). The originally optimized network presented in the main text is indicated with †, and the network optimized without enforcing Dale's principle ([Extended Data Fig. 8](#)) is marked with *. Other optimized networks were studied to confirm that well-optimized networks reliably showed similar behavior. This was important because our cost function was highly non-convex. Therefore, any minimum our optimizer found, such as that corresponding to the originally optimized network, had no guarantee of being the global minimum. Therefore, we trained 10 further networks on the original cost function ([Methods](#)), starting from random initial conditions, and show here those whose final cost was at least approximately as low as that of the original network (9 out of 10). **a**, Total cost ([Eq. 25](#)) computed for each of the random networks and (left) for networks that were optimized for the original cost (right). Colors indicate different components of the cost function (legend, see [Eqs. 26-29](#) for mathematical definitions). The inset shows the optimized networks only (note different y-scale). Note that the cost achieved by the random networks was 1-3 orders of magnitude higher than that achieved by the optimized networks. Furthermore, none of the optimized networks achieved substantially lower costs than the one we presented in the main text. **b-c**, Oscillatoriness (**b**) and transient overshoot size (**c**) for each network in **a**. Oscillatoriness was computed by numerical fits of [Eq. 33](#) to the autocorrelogram of the LFP generated by the network ([Methods](#)). Transients in population-average firing rates were quantified as the size of the overshoot normalized by the change in the steady state mean (see also [Fig. 7a](#)). Note that compared to the optimized networks – including the one presented in the main text –, oscillations and transients were almost entirely absent from random networks. Furthermore, all optimized networks had substantial oscillatoriness and transient overshoots ([Extended Data Fig. 4b-c](#)). This suggests that the results we obtained for the originally optimized network were representative of the best achievable minima of the cost function.



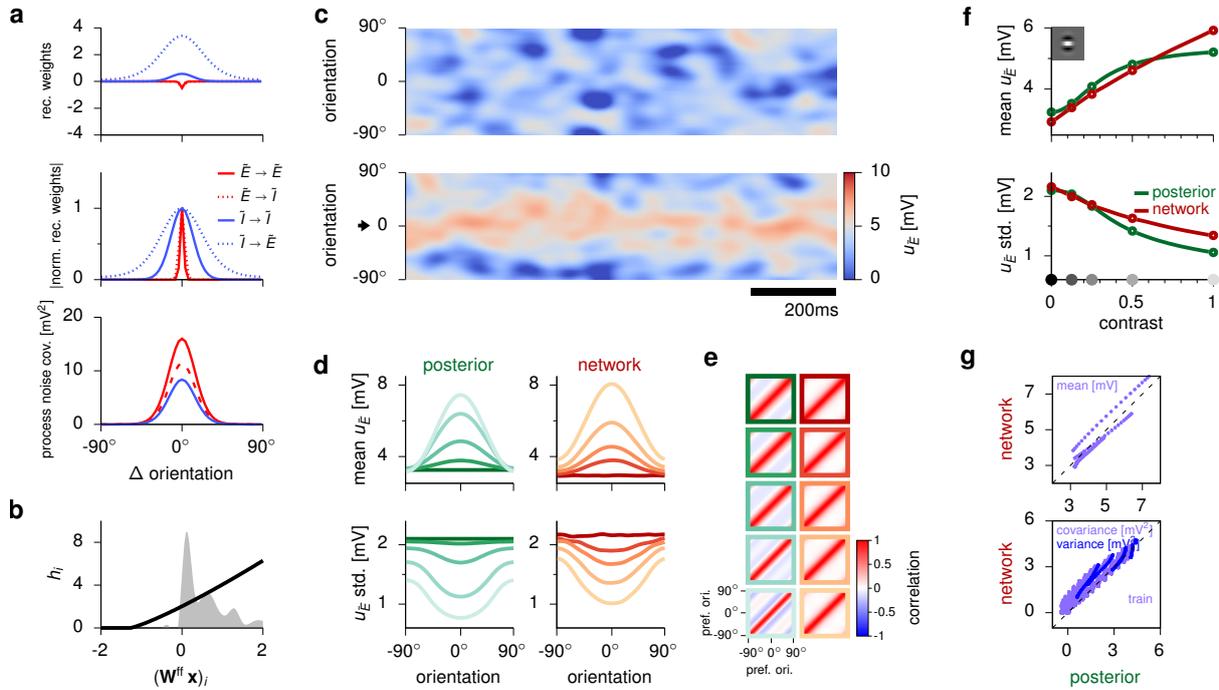
Extended Data Fig. 5 | Control networks without full variability modulation. Variability modulations are a hallmark of sampling-based inference⁵. To see whether they were also critical for our results, we optimized networks with modified cost functions, either setting both $\epsilon_{\text{var}} = 0$ and $\epsilon_{\text{cov}} = 0$ (Eq. 25), requiring only response means to be matched (a-g; see also Fig. 5, right column), or setting only $\epsilon_{\text{cov}} = 0$, requiring the matching of response means and variances but not of covariances (h-n; see also Methods). **a-b, h-i**, Network parameters as in Extended Data Fig. 1d-e. Both networks developed weak connection weights (a and h, top), with near-identical widths for E and I inputs onto both E and I cells (a and h, middle), and an almost linear input transformation (b and i). **c, j**, Sample population activity as in Fig. 2a. **d-e, k-l**, Matching moments between the ideal observer and the network for training stimuli as in Fig. 2c-d. Extremely weak coupling in the first network (a, top) meant an essentially feed-forward architecture. Thus, its response covariance simply reflected its process noise covariance (compare e and a, bottom). High input correlations in the second network (h, bottom) resulted in a single, global mode of output fluctuations (l). **f-g, m-n**, Generalization to test stimuli as in Fig. 3a-b. Insets in g show GSM posterior and network response means for example test stimuli as in Fig. 3c. Response moments in n are shown only for training stimuli, not for test stimuli, but by distinguishing variances (blue, bottom) and covariances (lavender, bottom). Both networks completely failed to fit moments that they were not explicitly required to match. Thus, firing rate tuning curves were preserved in both networks, but Fano factors were barely modulated in the first network (Extended Data Fig. 6a-b). Critically, neither of these networks showed discernible oscillations or transient overshoots Extended Data Fig. 6c-e). Response moments in d-g and k-n were estimated from $n = 20,000$ independent samples (taken 200 ms apart). Population mean moments in f and m were further averaged across $n=50$ E cells. Correlations in e and l are Pearson's correlations.



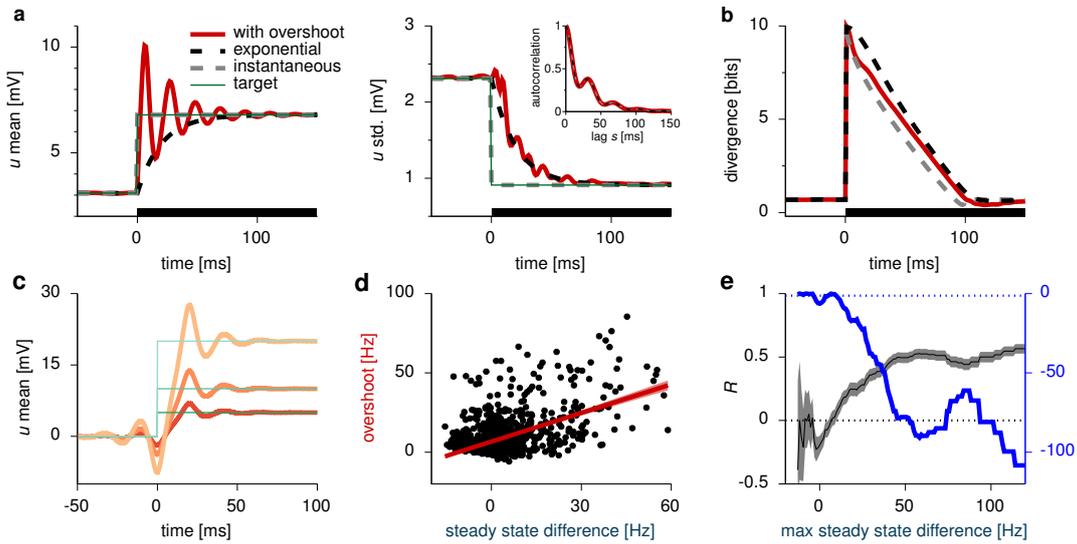
Extended Data Fig. 6 | Comparison of neural dynamics between the originally optimized network and the control networks. **a-e**, Dynamics of optimized networks as in Fig. 5a-e. The originally optimized network of the main paper (left) is compared to various control networks, from left to right: without slowness penalty (Extended Data Fig. 10), without covariance modulation (matching means and variances; Extended Data Fig. 5h-n), without covariance and variance modulation (matching means only; Extended Data Fig. 5a-g), enforcing constant Fano factors (Extended Data Fig. 7), and with Dale's principle not enforced (Extended Data Fig. 8). For ease of comparison, only stimulus-dependent power spectra are shown for the optimized network of the main paper, without showing the dependence of gamma peak frequency on contrast (cf. Fig. 5c, middle), as most control networks had no discernible gamma peaks. For more details on control networks, see the captions of the corresponding figures (Extended Data Figs. 5, 7, 8 and 10). Response moments in **a** were estimated from $n = 20,000$ independent samples (taken 200 ms apart). Mean firing rates in **d** were computed over $n = 100$ trials. Panel **e** shows mean \pm s.e.m. ($n = 20$ trials).



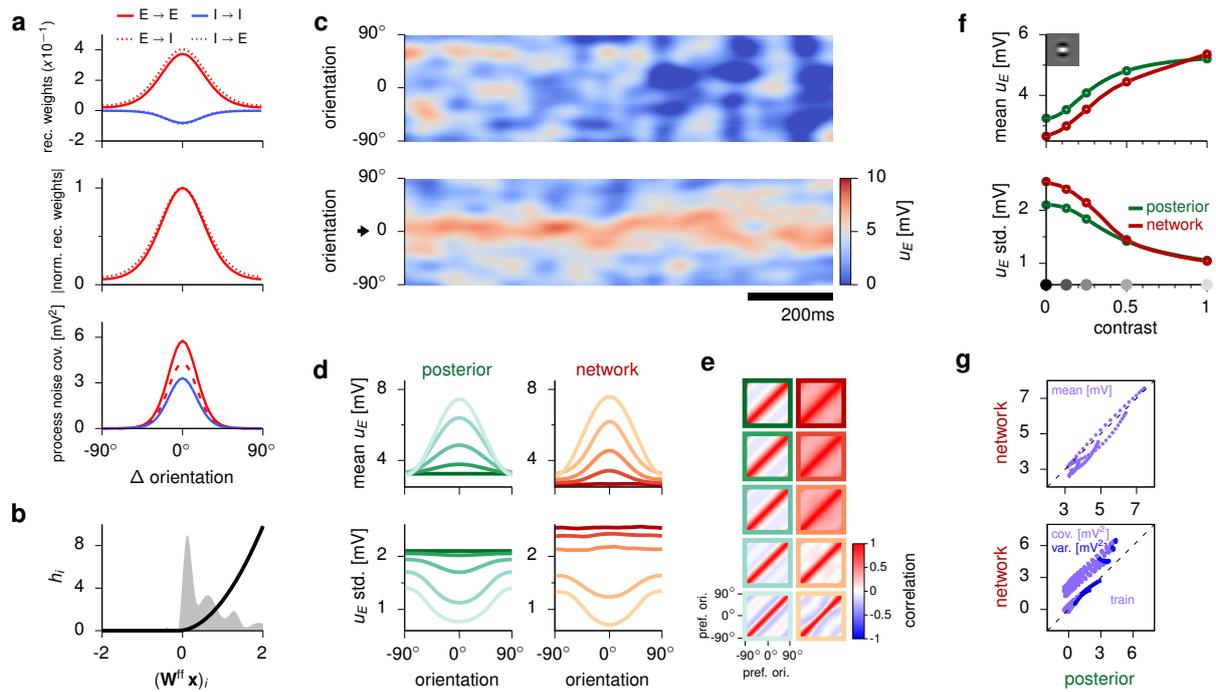
Extended Data Fig. 7 | Control network: enforcing constant Fano factors. Fano factors need to be specifically stimulus-independent for a class of models, (linear) probabilistic population codes (PPCs¹²), that provide a conceptually very different link between neural variability and the representation of uncertainty than that provided by sampling^{5,18}, which we pursue here. Therefore, we used our optimization-based approach to directly compare the circuit dynamics required by PPCs to those of our originally optimized network implementing sampling. For this, we trained a further control network whose goal was to match the mean modulation of the control network (resulting in realistic tuning curves), while keeping Fano factors constant. We achieved this by devising a set of target covariances that would result (together with the target mean responses used by all other networks) in constant Fano factors – assuming Poisson spiking and an exponentially decaying autocorrelation function (using analytic results in Ref. 52). Training then proceeded exactly as for the other networks, with the same ϵ parameters in the cost function as for the originally optimized network, only employing the new covariance targets. **a-b**, Network parameters as in [Extended Data Fig. 1d-e](#). The network made use of strong inhibitory connections (**a**, top), large shared process noise (**a**, bottom), and strongly modulated inputs (**b**). **c**, Sample population activity as in [Fig. 2a](#). **d-e**, Matching moments between the ideal observer and the network for training stimuli as in [Fig. 2c-d](#). **f-g**, Generalization to test stimuli as in [Fig. 3a-b](#). Insets in **g** show GSM posterior and network response means for example test stimuli as in [Fig. 3c](#). The network was able to match mean responses in the training set and to generalize to novel stimuli (**d**, **f-g**), while keeping Fano factors relatively constant as required ([Extended Data Fig. 6b](#)). For consistency with previous results, we obtained Fano factors by numerically simulating the same type of inhomogeneous Gamma process as in the other networks ([Methods](#)), thus violating the Poisson spiking assumptions under which we computed the target covariances of the network (see above) – hence the remaining small modulations of Fano factors. Critically, although the training procedure was identical to that used for the original network, only differing in the required variability modulation provided by the targets (see above), this control network displayed no gamma-band oscillations ([Extended Data Fig. 6c](#)). Inhibition-dominated transients did emerge, but were weaker than in the original network ([Extended Data Fig. 6d](#) and [e](#)). Response moments in **d-g** were estimated from $n = 20,000$ independent samples (taken 200 ms apart). Population mean moments in **f** were further averaged across $n=50$ E cells. Correlations in **e** are Pearson's correlations.



Extended Data Fig. 8 | Control network: Dale’s principle not enforced. In order to see how much the biological constraints we used for the optimized network, and in particular enforcing Dale’s principle, were necessary to achieve the performance and dynamical behavior of the original network, we optimized a network with the same cost function as for the original network (Eq. 25) but without enforcing Dale’s principle. This meant that the signs of synaptic weights in each quadrant of the weight matrix (the α_{XY} coefficients in Eq. 10) were not constrained. Otherwise, optimization proceeded in the same way as before (Methods). The training of this network proved to be much more difficult and prone to result in unstable networks, which we avoided by early stopping. **a-b**, Network parameters as in Extended Data Fig. 1d-e. As Dale’s principle was not enforced, only notional cell types can be shown (legend). Nevertheless, interestingly, the network still obeyed Dale’s principle after optimization (top): all outgoing synapses of any one cell had the same sign. Note that the outgoing weights of the cells whose moments were constrained (\bar{E} cells, whose activity is shown and analysed in **c-g**) were actually negative. Therefore, in effect, these cells became inhibitory during training. **c**, Sample population activity as in Fig. 2a. **d-e**, Matching moments between the ideal observer and the network for training stimuli as in Fig. 2c-d. **f**, Generalization to test stimuli as in Fig. 3a. **g**, Matching moments between the ideal observer and the network for training stimuli as in Fig. 3b (lavender). Note that here, unlike in Fig. 3, response moments are shown only for training stimuli, not for test stimuli, but by distinguishing variances (blue, bottom) and covariances (lavender, bottom). Overall, the *stationary* behavior of this network was broadly similar to that of the originally optimized network (cf. Figs. 2 and 3). Therefore, it achieved a performance that was far better than the random networks’ (Extended Data Fig. 4a). However, it still performed substantially worse than networks optimized with Dale’s principle enforced (Extended Data Fig. 4a), and its *dynamics* were also qualitatively different: oscillations and transient overshoots were largely absent from it (Extended Data Fig. 4b-c, Extended Data Fig. 6). Response moments in **d-g** were estimated from $n = 20,000$ independent samples (taken 200 ms apart). Population mean moments in **f** were further averaged across $n=50$ \bar{E} cells. Correlations in **e** are Pearson’s correlations.



Extended Data Fig. 9 | Further analyses of the role of transients in supporting continual inference. **a-b**, Analysis of transients in the response of a single neuron. (Left panel in **a** is reproduced from Fig. 7a.) **a**, Temporal evolution of the mean (left) membrane potential (u_E), and the membrane potential standard deviation (right) in three different neural responses (thick lines) with identical autocorrelations (matched to neural autocorrelations in the full network, inset, cf. Fig. 4a) but different time-dependent means (left) and standard deviations (right). Thin green line shows the time-varying target mean (left) and standard deviation (right). **b**, Total divergence (Eq. 37, Methods) between the target distribution at a given point in time and the distribution represented by the neural activity sampled in the preceding 100 ms, for each of the three responses (colors as in **a**). In comparison, note that Fig. 7b only shows the mean-dependent term of the divergence (Eq. 38, Methods). Black bars in **a-b** show stimulus period. Mean and divergence computed as an average over multiple trials ($n = 10,000$). **c**, Optimal response trajectories (red lines) for continual estimation of the mean of a target distribution (thin green lines). Different shades of red indicate optimal trajectories corresponding to three different target levels (5, 10, and 20 mV, emulating different contrast levels). We optimized neural response trajectories (Eq. S39, Supplementary Math Note) so that the distance between their temporal average, computed using a prospective box-car filter ($k(t) = 1/T$ for $t \in [0, T]$ and 0 otherwise, with $T = 20$ ms), and the corresponding Heaviside step target signal would be minimized under a smoothness constraint ($\epsilon_{\text{smooth}} = 1$; Eq. S37, Supplementary Math Note). Note the transient overshoot in the optimal response very closely resembling those observed in the optimized network (Fig. 5d, Fig. 7c-d): its magnitude scales with the value of the target mean, and it is followed by damped oscillations. Similar results (with less ringing following the overshoot) were obtained also for an exponentially decaying, “leaky” kernel (not shown). **d-e**, Relationship between overshoot magnitude and steady state difference: analysis of experimental recordings from awake macaque V1²³. **d**, Overshoot magnitude versus steady state difference: same as Fig. 7e, but restricting the analysis to steady state differences below 60 Hz (to exclude outliers). Red line shows linear regression ($\pm 95\%$ confidence bands). The correlation between overshoot size and steady state difference is still significant: two-sided Wald test $p = 1 \times 10^{-89}$ ($n = 1263$ cell-stimulus pairs $R^2 \simeq 0.27$). **e**, Systematically changing the maximal steady state difference (x-axis) used for restricting the analysis of the correlation between overshoot size and steady state difference (**d** and Fig. 7e) reveals that the correlation is robust (black line $\pm 95\%$ confidence intervals) and remains highly significant (blue line, showing corresponding p-values; note logarithmic scale) for all but the smallest threshold (and thus smallest sample size). Horizontal dotted lines show $R = 0$ correlation (black) and $p = 0.05$ significance level (blue) for reference. As the maximal steady state difference increases, the number of points n considered also increases from $n = 4$ to $n = 1279$. Correlations in **a** are Pearson’s correlations. Pearson’s R values and corresponding p-values in **e** were obtained by linear regression performed as in **d**.



Extended Data Fig. 10 | Control network: no explicit slowness penalty. To test whether the explicit slowness penalty in our cost function (Eq. 29) was necessary for obtaining the dynamical behaviour exhibited by the optimized network of the main text, we set $\epsilon_{\text{slow}} = 0$ in Eq. 25. **a-b**, Network parameters as in Extended Data Fig. 1d-e. **c**, Sample population activity as in Fig. 2a. **d-e**, Matching moments between the ideal observer and the network for training stimuli as in Fig. 2c-d. **f**, Generalization to test stimuli as in Fig. 3a. **g**, Matching moments between the ideal observer and the network for training stimuli as in Fig. 3b (lavender). Note that here, unlike in Fig. 3, response moments are shown only for training stimuli, not for test stimuli, but by distinguishing variances (blue, bottom) and covariances (lavender, bottom). Note that this network behaved largely identically to the originally optimized network (see also Extended Data Fig. 6). This could be attributed to the fact that our optimization implicitly encouraged fast sampling by default, simply by using a finite averaging window for computing average moments of network responses, and in particular by including samples immediately or shortly following stimulus onset (Methods). Response moments in **d-g** were estimated from $n = 20,000$ independent samples (taken 200 ms apart). Population mean moments in **f** were further averaged across $n=50$ E cells. Correlations in **e** are Pearson's correlations.