A diffusive homeostatic signal maintains neural heterogeneity and responsiveness in cortical networks

Yann Sweeney[1,2,*], Jeanette Hellgren Kotaleski[2], Matthias H. Hennig[1]

1 Institute for Adaptive and Neural Computation, School of Informatics, University of Edinburgh, Edinburgh, United Kingdom

2 Department of Computational Biology, School of Computer Science and Communication, Royal Institute of Technology, Stockholm, Sweden

∗ **E-mail: yann.sweeney@ed.ac.uk**

Abstract

Gaseous neurotransmitters such as nitric oxide (NO) provide a unique and often overlooked mechanism for neurons to communicate through diffusion within a network, independent of synaptic connectivity. NO provides homeostatic control of intrinsic excitability. Here we conduct a theoretical investigation of the distinguishing roles of NO-mediated diffusive homeostasis in comparison with canonical non-diffusive homeostasis in cortical networks. We find that both forms of homeostasis provide a robust mechanism for maintaining stable activity following perturbations. However, the resulting networks differ, with diffusive homeostasis maintaining substantial heterogeneity in activity levels of individual neurons, a feature disrupted in networks with non-diffusive homeostasis. This results in networks capable of representing input heterogeneity, and linearly responding over a broader range of inputs than those undergoing non-diffusive homeostasis. We further show that these properties are preserved when homeostatic and Hebbian plasticity are combined. These results suggest a mechanism for dynamically maintaining neural heterogeneity, and expose computational advantages of non-local homeostatic processes.

Author Summary

Neural firing rates must be maintained within a stable range in the face of ongoing fluctuations in synaptic activity. Existing cortical network models achieve this through various homeostatic mechanisms which constrain the excitability of individual neurons according to their recent activity. Here, we propose a new mechanism, *diffusive homeostasis*, in which neural excitability is modulated by nitric oxide, a gas which can flow freely across cell membranes. Information about a neurons' firing rate can be carried by nitric oxide, meaning that an individual neurons' excitability is affected by neighboring neurons' firing rates as well as its own. We find that this allows a neuron to deviate from the target population activity, as its neighbours will counteract this deviation, thus maintaining stable average activity. This form of neural heterogeneity is more flexible than assigning different target firing rates to individual neurons. Consequently, networks endowed with this diffusive mechanism have an improved representational capability compared to canonical, local homeostatic mechanisms, and allow for more efficient use of neural resources.

Introduction

Nitric oxide (NO) is a diffusive neurotransmitter which is widely synthesized in the central nervous system, from the retina to the hippocampus [1,2]. Its properties as a small nonpolar gas molecule allows rapid and unconstrained diffusion of NO across cell membrane, a phenomenon often called *volume transmission* [3]. An important role of NO signaling is to regulate neural excitability though the modulation of potassium conductances in an activity-dependent manner, effectively mediating a form of homeostatic intrinsic plasticity (HIP). Experiments characterizing this effect also demonstrated that NO-synthesizing neurons can induce changes in the excitability of neurons located up to 100 μm away [4,5]. These findings are

corroborated by a recent study demonstrating neurovascular coupling mediated through activity-dependent NO diffusion [6]. We build upon these observations, postulating a general form of HIP mediated by a diffusive neurotransmitter such as NO which we will refer to as *diffusive homeostasis*. This contrasts with canonical models of HIP, here referred to as *non-diffusive homeostasis*, which assume that each neuron has access to only its own activity [7].

Theoretical studies of HIP have generally focused on its role in maintaining stable network dynamics [8,9]. It has also been recently demonstrated that HIP can improve the computational performance of recurrent networks by increasing the complexity of network dynamics [10]. However, little is known about the effects of HIP on the heterogeneity typically observed in cortical networks; in particular, a growing body of evidence supports the finding that even neurons of the same type have a broad and heavy-tailed distribution of firing rates [11]. Rather than an epiphenomenon of biological noise, neural heterogeneity has been proposed to improve stimulus encoding by broadening the range of population responses [12, 13]

However, this form of heterogeneity is difficult to reconcile with canonical models of HIP, which generally suppress cell-to-cell variability [14]. While some degree of heterogeneity in populations of the same type of neuron may emerge naturally [15], we found that such independent sources of variability will generally limit the responsiveness of a network through neuronal saturation.

Using network models and dynamic mean field analysis, here we show that networks with HIP mediated by diffusive neurotransmission exhibit a very different and unexpected behavior. Firstly, we report that diffusive homeostasis provides a natural substrate for flexibly maintaining substantial heterogeneity across a network. Secondly, the resulting population heterogeneity enables linear network responses over a wide range of inputs. This not only improves population coding, but enables a good use of available resources by ensuring that all neurons remain functionally integrated in the network dynamics.

Results

We investigated the effects of diffusive homeostasis in a recurrent neural network (Figure 2A, see Methods) with sparse and random connectivity, based on conventional models of cortical networks giving rise to asynchronous irregular spiking activity [16]. Each neuron received external input with a rate randomly drawn from a normal distribution. Ca^{2+} influx during a spike triggered NO synthesis through nNOS activation (Figure 2B, see Methods). To simulate spatial NO signaling, neurons were randomly distributed on the surface of a torus and linear diffusion was simulated on this surface. Each neuron's firing threshold θ_i underwent modulation through negative feedback mediated by the concentration of NO (Equation 6 in Methods).

The effect of a diffusive neurotransmitter mediating HIP within the network was investigated by comparing two cases: first where NO was allowed to diffuse freely across cell membranes as observed experimentally [4], and second without diffusion such that intracellular NO concentration was affected only by a neuron's own recent activity (Figure 1). The latter corresponds to a canonical model of HIP as investigated before [8, 9].

Diffusive homeostasis enables a broad firing rate distribution

Figure 2C illustrates that both forms of homeostasis stabilized network activity following an increase in input. There was however a crucial difference in how the neurons reacted to this change. While for non-diffusive homeostasis each neuron simply returned to its target firing rate, diffusion caused each neuron to sense a mixture of its own activity level and that of the rest of the network. This can be seen in the spatial concentration profiles in Figure 2C. It is important to note that the spatial position of each neuron was random and independent of its connections, meaning that there was no explicitly defined structure in the NO concentrations.

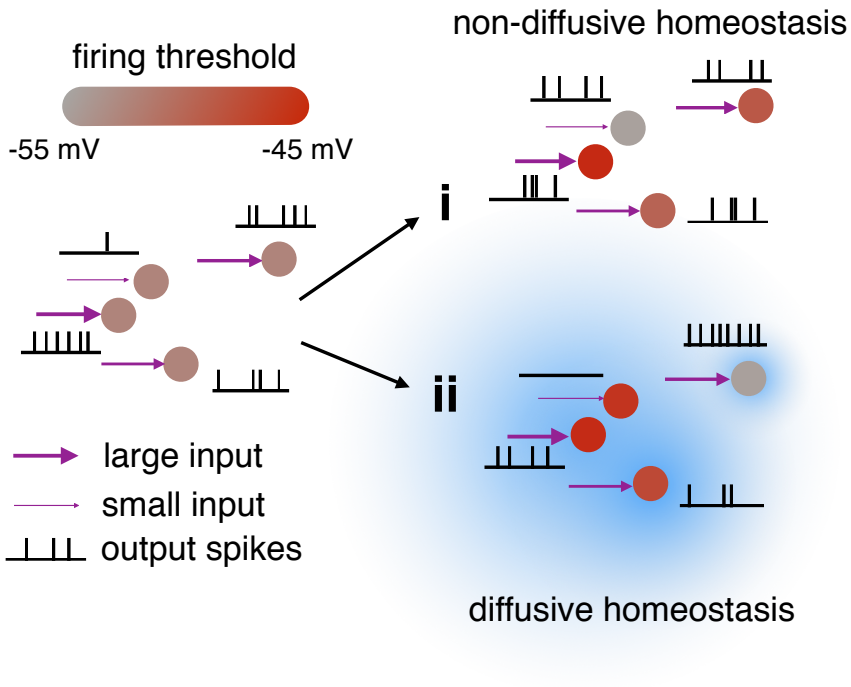

Figure 1. Illustration of the effects of non-diffusive (i) and diffusive (ii) homeostasis. Non-diffusive homeostasis adjusts each neuron's threshold (red color bar) according to its input to give identical firing rates, while diffusive homeostasis induces correlations (blue cloud) in the thresholds of neighboring neurons, thereby maintaining diverse firing rates.

As a result, these networks exhibited a very different steady state behavior. The firing rate distribution was narrow as expected for non-diffusive homeostasis, but broad and heavy-tailed for diffusive homeostasis (Figure 2D). The latter is consistent with recent experimental results indicating that firing rate distributions in cortex are generally heavy-tailed, approximating log-normal distributions [11]. There were no noticeable differences in inter-spike interval statistics between networks with diffusive and non-diffusive homeostasis (not illustrated).

We explicitly investigated the difference in firing rate distributions by modeling the relation between activity read-out and homeostatic compensation in these two cases using a dynamic mean-field model (see Methods). This approach considered an unconnected population of neurons with random inputs, where each of the two scenarios was simulated by using an appropriate activity read-out. HIP was implemented as in the full spiking model, but the degree of diffusive signaling was now controlled by a single parameter, α (Equation 10 in Methods), which determined the balance between local and global activity read-out. If small, neurons used primarily their own activity to modulate their firing threshold, while increasing α caused the firing threshold to depend more strongly on the average population activity. Setting, for instance, $\alpha = .8$ led to a broad and heavy-tailed rate distribution similar to the full model, while $\alpha = 0$ yielded a narrow distribution as in the non-diffusive case (Figure 2E).

This model provides a simple and intuitive explanation for this effect. For a non-interacting population, non-diffusive homeostasis can be thought of as precisely matching a neuron's input μ_i and its threshold θ_i to maintain the target firing rate. We can imitate this by introducing a covariance $\sigma(\mu, \theta)$ between μ_i and θ_i, such that a high input rate implies a high firing threshold and a low input rate a low firing

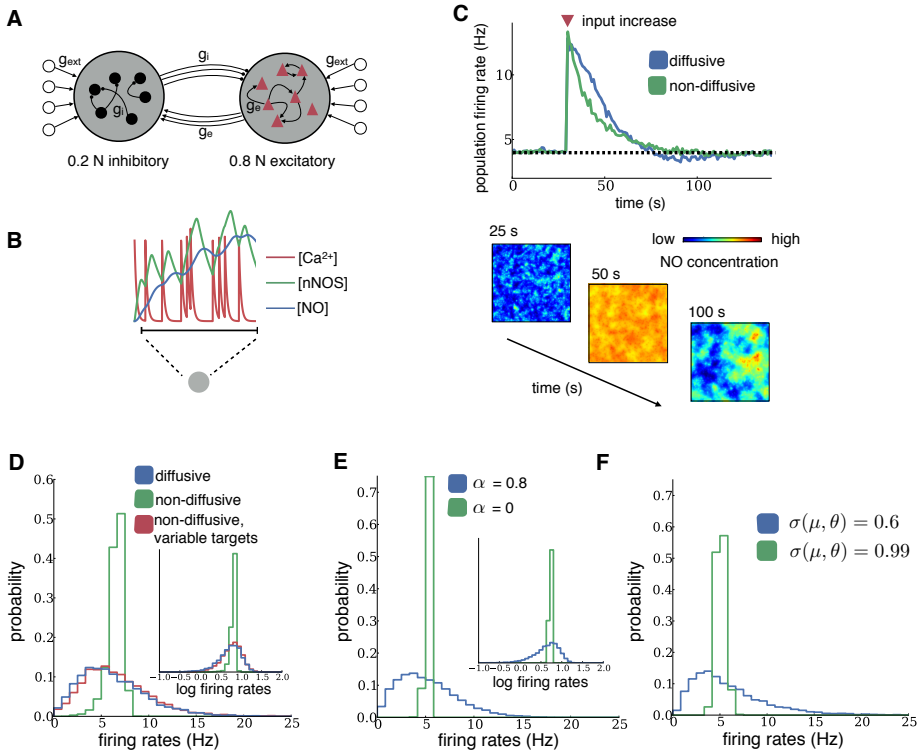

Figure 2. Steady-state behavior of diffusive and non-diffusive homeostasis. (A) Schematic of the sparsely connected recurrent network model. Neurons received homogeneous random spiking input (g_{ext}). **(B)** Intracellular homeostatic signals in a model neuron. Each spike triggers calcium influx, which leads to nNOS activation and NO synthesis. **(C)** Mean population firing rates for networks with diffusive or non-diffusive homeostasis after an increase in external input (red triangle). Spatial distribution of NO concentrations at different times across the network with diffusion are shown below. (**D**-**F**) Distributions of firing rates and log firing rates (insets) after homeostasis from network simulations (D) and mean-field analysis (E), both receiving independent Poisson inputs drawn from a Gaussian distribution. **(F)** Distributions of firing rates in the mean-field analysis for low and high covariance σ of threshold and input rate.

threshold. Since setting $\alpha > 0$ (analogous to diffusive homeostasis) introduces a correlation between a neuron's threshold θ_i and the average population threshold $\overline{\theta}$, this effectively results in a decorrelation of μ_i and θ_i in comparison with setting $\alpha = 0$ (analogous to non-diffusive homeostasis). In line with the previous results, populations with for instance $\sigma(\mu, \theta) = 0.6$ yielded a broader and more heavy-tailed distribution of firing rates than populations with $\sigma(\mu, \theta) = 0.99$ (Figure 2F). This method could serve to efficiently characterize the effects of homeostasis on population statistics, providing an alternative to computationally intensive simulations of large networks.

Since non-diffusive homeostasis directly relates the firing threshold of a neuron to its input, we observed a wider distribution of firing thresholds, which in turn ensured that all neurons assumed similar firing rates (Figure 3). Diffusive homeostasis, on the other hand, yielded similar firing thresholds across the population (Figure 3). When combined with the nonlinear input-output relation of neurons [17], this gave rise to the broad firing rate distributions we observed (see also Discussion).

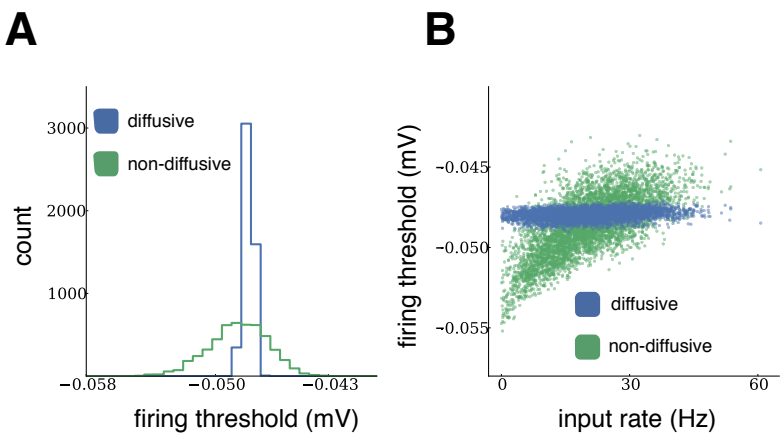

Figure 3. Steady-state firing thresholds of diffusive and non-diffusive homeostasis. (A) Distributions of firing thresholds after homeostasis from network simulations, receiving independent Poisson input drawn from a Gaussian distribution. **(B)** Steady-state firing thresholds plotted against external inputs received during homeostasis.

Since one may argue that diffusive homeostasis is merely adding variability to each neuron's homeostatic signal due to the influence of neighboring neurons' activity, we now ask whether it is possible to achieve broad firing rate distributions with non-diffusive homeostasis. Indeed, by introducing variability of homeostatic targets, we could produce a distribution of firing rates similar to that observed with diffusive homeostasis (Figures 2D, red histogram). However, as we will show next, the effect of diffusive homeostasis is quite distinct from that of activity-independent, 'quenched' heterogeneity arising from randomly distributed homeostatic targets.

Diffusive homeostasis retains input heterogeneity

To investigate the functional consequences of heterogeneity caused by a diffusive homeostatic process, we next simulated specific changes in external input. First, we stimulated small random groups of neurons at higher input rates of 5Hz and 10Hz (versus a baseline of 2Hz), as illustrated in Figure 4. Such inputs may, for instance, reflect developmental or other plastic changes that lead to a long-lasting change in network input. In these simulations, the average network firing rate was reliably brought back to the target firing rate of 2 Hz by both forms of homeostasis (Figures 4A-C, black traces). As above, in networks with non-diffusive homeostasis this was achieved by returning the rate of each neuron to the target firing rate regardless of their external input (Figure 4A, colored traces). In contrast, for networks with diffusive homeostasis, we found that the separability of firing rates of individual groups are maintained according to their input, while the firing rates of all groups were simultaneously reduced so that the average network firing rate again reached the target (Figure 4B, colored traces). Introducing variability in homeostatic targets for the non-diffusive case, as described previously, did not maintain full firing rate separability of individual groups as in the diffusive case, although smaller differences could persist (Figure 4C).

The distribution of final firing thresholds explains these differences (Figure 4D-F). For non-diffusive homeostasis, the neurons in the group receiving 10 Hz input had the highest thresholds since they needed to reduce their firing rate the most, followed by the 5 Hz and 2 Hz groups respectively. This led to a roughly trimodal distribution of final thresholds. For a diffusive signal, a neuron's firing threshold is modulated by the activity of nearby neurons. Since group membership of a neuron is independent of its

position, this effect again introduced a correlation between each neuron's threshold and the mean threshold of the entire network, resulting in a roughly unimodal and narrower distribution of final thresholds compared to a network with non-diffusive homeostasis. Thus, firing thresholds in neurons undergoing diffusive homeostasis were only weakly related to their external input. This in turn preserves local firing rate differences in input groups while maintaining constant average network activity. Introducing variable targets for non-diffusive homeostasis caused the thresholds to segregate, similar to the non-diffusive case.

We could reproduce the distinctions between diffusive and non-diffusive homeostasis in the dynamic mean-field approach by varying α. For $\alpha = 0$, modeling non-diffusive homeostasis, we obtained identical firing rates in input groups, as in the recurrent network (Supporting Information, Figure S1). Setting $\alpha = 0.8$, on the other hand, preserved different firing rates in input groups, as also observed in the recurrent network with diffusive homeostasis (Supporting Information, Figure S1). Taken together, this shows that diffusive homeostasis can retain input heterogeneity due to the influence of neighboring neurons' activity on an individual neuron's firing threshold.

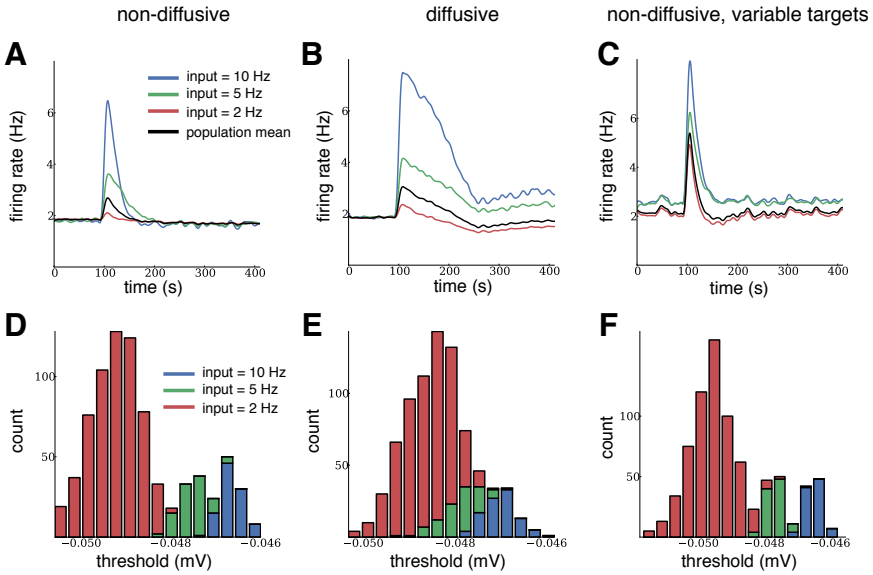

Figure 4. Relative differences between inputs are preserved during diffusive homeostasis. **(A-C)** Evolution of firing rates in a recurrent network for each input group under diffusive and non-diffusive homeostatic control, and for variable homeostatic targets. Black traces show population activity. Independent Poisson input at the target rate was given to 800 neurons (red), while two groups of 100 neurons each received elevated Poisson input. The relative rate differences of the groups were only preserved for diffusive homeostasis. **(D-F)** Final distributions of firing thresholds.

Population heterogeneity during diffusive homeostasis enables linear network responses

In the simulations shown so far, each neuron received a static input throughout since we were interested in the final network states. We now investigate how these networks respond to fast changes in input; specifically how faithfully each neuron represents its change in input. Since networks with diffusive homeostasis simultaneously maintain constant average network activity and firing rate heterogeneity, we expected that this should allow input modulations to be followed more precisely due to a greater

representational capability.

After the network reached steady state under an initial distribution of external inputs, we froze homeostasis so as to simulate changes in activity too fast to influence homeostatic processes. We then regenerated the external inputs to each neuron from the same distribution presented during homeostasis. This can be thought of as a re-configuration of inputs due to external fluctuations, or fast changes brought about by processes such as synaptic plasticity or rewiring during development. To best represent such changes in a simple population coding paradigm, each neuron should respond linearly to a change in input; non-linear transformations may lead to an information loss and hence affect neural computations. We interpreted the range of changes in input over which this response is linear, or non-saturating, as the range over which homeostasis does not interfere with the network response.

Figures 5A-C show the change in input rate versus change in output rate of each neuron. A highly nonlinear response was observed in networks with non-diffusive homeostasis, with rectification for large decreases and supralinear responses for large increases in input. This effect was quantified by an R^2 value of 0.57 from a linear regression. Conversely, networks with diffusive homeostasis exhibited a linear response across the entire range of input changes, with an R^2 value of 0.85. Population heterogeneity can also be achieved, as discussed before, by introducing target variability during non-diffusive homeostasis. This yielded a similar non-linear response function as in the non-diffusive network with homogeneous targets, with an R^2 value of 0.38.

A consequence of the asymmetry in responses to input changes for networks with non-diffusive homeostasis was that the population rate increases upon regenerating inputs, despite the fact that mean input to the network remained unchanged (Figure 5D). This did not occur for networks with diffusive homeostasis, suggesting that these networks are more adept at maintaining a target level of activity in conditions where external inputs are dynamic and fast-changing.

This difference in responses to input changes could again be reproduced in the dynamic mean-field approach. This allowed us to characterize population responses across different effective diffusive ranges, using the R^2 value from a linear regression as a measure of response linearity. Figure 5E shows normalized R^2 values across a range of different input distribution widths, δ, as α is varied to model different diffusion coefficients (see Methods). This revealed a dependency on δ: While values of $\alpha \sim 1$ exhibited the best response for small δ, hence cases where the inputs are rather narrow, the optimal α decreased as δ increased. Crucially, the benefits of a diffusive homeostatic signal can be achieved by a relatively broad range of values for α, indicating that the effects we describe should be rather robust to precise parameter choices affecting the diffusion range.

Since connection probability falls off with spatial distance in cortical networks [18], we additionally simulated recurrent networks featuring such connectivity profiles. These networks exhibited qualitatively similar behaviour under diffusive and non-diffusive homeostasis compared to networks without any spatial dependence in connectivity (Supporting Information, Figure S2).

Overall, these results suggest that networks undergoing diffusive homeostasis are better suited to represent a range of inputs. We investigated this by presenting the networks with time-varying inputs after freezing homeostasis. Groups of excitatory neurons received additional inputs which were randomly and independently generated after fixed time intervals (see Methods). Figure 6A shows the representation of such a time-varying input pattern (dotted black line) for each network (colored lines). Networks which have undergone diffusive homeostasis were capable of tracking this input significantly better than their non-diffusive counterparts, as characterized by the RMS error between the network response and input pattern (0.12 for diffusive homeostasis; 0.23 and 0.19 for non-diffusive homeostasis with uniform and variable targets, respectively; Figure 6B).

We can explore these differences further by constructing a simplified task in which a population of orientation-selective neurons decodes the orientation of a stimulus (see Figure 6C-F , Methods). This is not intended to represent circuits which perform this task in the brain, but to serve purely as a demonstration of the relative merits of linear and non-linear network responses. Networks with linear responses perform

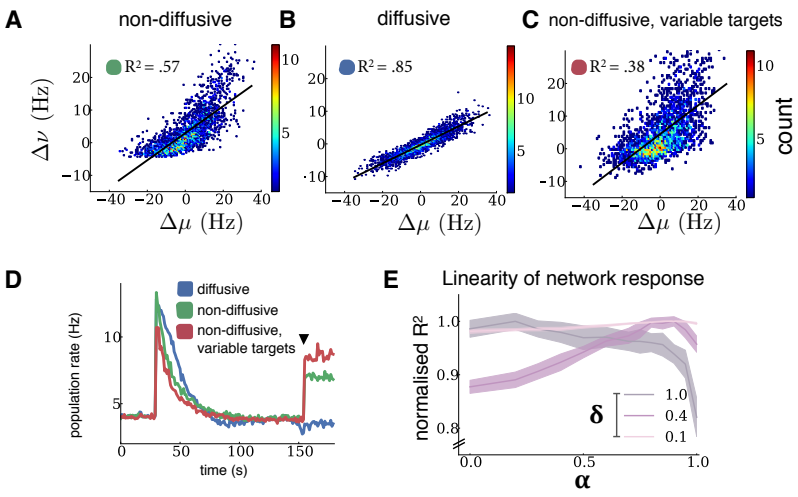

Figure 5. Diffusive homeostasis enables linear network responses. (A-C) Firing rate changes $\Delta\nu$ of all neurons following input changes $\Delta\mu$. Black lines show linear fits, with corresponding R^2 values inset. **(D)** Population activity before and after input change (black triangle). **(E)** Linearity of the population response to a change of inputs in the mean-field analysis as α is varied, shown for different input distribution widths δ. Shaded regions correspond to one standard deviation.

better than those with non-linear responses when the entire population of neurons is used to determine the stimulus orientation, as measured by the variance of error in the orientation of the population vector compared to the stimulus orientation (41°, 63° and 72° for diffusive homeostasis, non-diffusive, and non-diffusive with variable targets respectively, Figure 6G).

Properties of diffusive homeostasis are conserved in networks with Hebbian plasticity

In the networks described so far, we have used static and uniform synaptic weights for recurrent connections. We next considered whether the observed properties of diffusive homeostasis are altered by the presence of plastic synaptic weights, in particular when Hebbian spike-timing-dependent plasticity (STDP) is introduced (see Methods). Using a standard model of STDP with additive depression and potentiation for all recurrent excitatory synapses, we simulated networks with both STDP and homeostasis active until synaptic weight and firing rate distributions reached a steady state (Figure 7A) [19]. As before, firing rate distributions were broader in networks with diffusive homeostasis (Figure 7B). Broad distributions could also be achieved by introducing variability in homeostatic targets. STDP amplified the differences in response linearity that was observed between homeostatic cases; while the mean R^2 value for networks with diffusive homeostasis was 0.52, networks with non-diffusive homeostasis exhibited much lower mean R^2 values of 0.051 and 0.046, for uniform and variable homeostatic targets respectively (Figure 7D. We observed a further effect of diffusive homeostasis on the synaptic weight distributions. Additive STDP generally leads to bimodal weight distributions, as observed in networks without any homeostasis or with non-diffusive homeostasis (Figure 7C). Diffusive homeostasis, however, results in weight distributions which are roughly unimodal, again consistent with experimental results [20].

We observed qualitatively similar retention of broad firing rate distributions and response linearity with diffusive homeostasis when a weight-dependent update rule was used (not illustrated), which has been argued to lead to more realistic weight distributions [21].

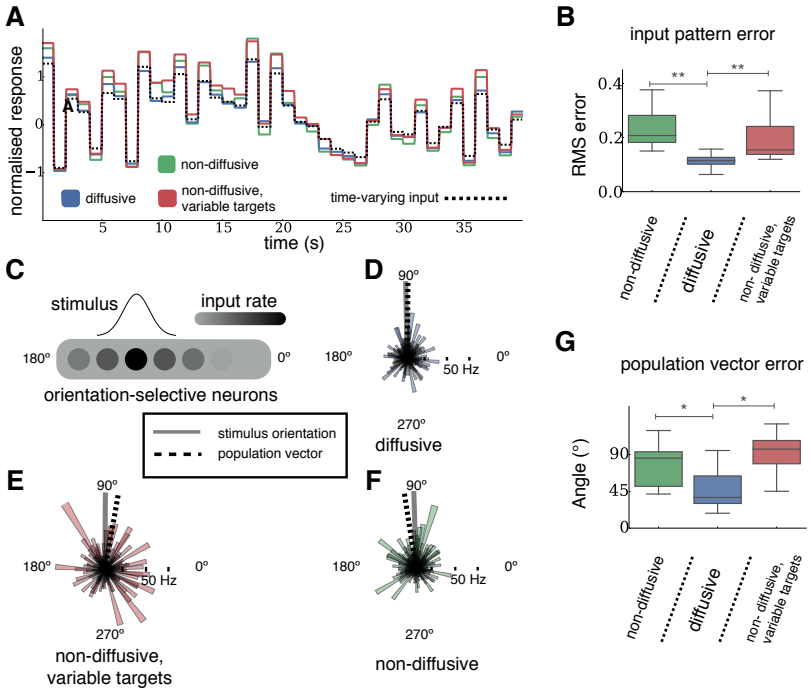

Figure 6. Performance advantages of networks with diffusive over non-diffusive homeostasis (A) Response of each network to a time-varying input pattern (dotted black line). Colored lines show normalized firing rate deviations from the mean population rate of those neurons receiving inputs. **(B)** RMS errors between the normalized input pattern and normalized response of each network (1000 s stimulation and 11 networks in each case). $^{**}p < 0.01$ from a two-sided K-S test. **(C)** Diagram showing network inputs during an orientation decoding trial. Each neuron is randomly assigned a preferred orientation, and receives external input at a rate given by the stimulus. **(D-F)** Example response of a population of orientation-selective neurons to a stimulus at an orientation of 90° (grey line), for networks with each type of homeostasis. Individual neural responses and their preferred orientation are given by the radii and direction of the coloured areas respectively. Orientation decoded using either the population vector or most responsive neuron are shown by the dot-dash and dashed lines respectively. **(G)** Variance of errors in the orientation of the population vector in response to a stimulus. $^{*}p < 0.05$ from a two-sided K-S test. The box encompasses the inter-quartile range and the whiskers extend to 1.5 times the inter-quartile range in all boxplots.

Discussion

Neural homeostasis is commonly thought of as a local process, where neurons individually sense their activity levels and respond with compensatory changes if activity changes. Here we investigated a complementary mechanism, where homeostasis is mediated by a diffusive molecule such as NO that acts as a non-local signal. Using a generic recurrent network model, we show that this form of homeostasis can have unexpected consequences. First, we found that it enables and maintains substantial population heterogeneity in firing rates, similar to that observed experimentally in intact circuits [11], and that input heterogeneities can be preserved in the population activity. Second, the specific form of neural heterogeneity brought about by diffusive homeostasis is particularly suited to support linear network responses over a broad range of inputs. It is important to note that this behavior differs from that of networks where heterogeneity is simply introduced by randomly assigning a different target to each

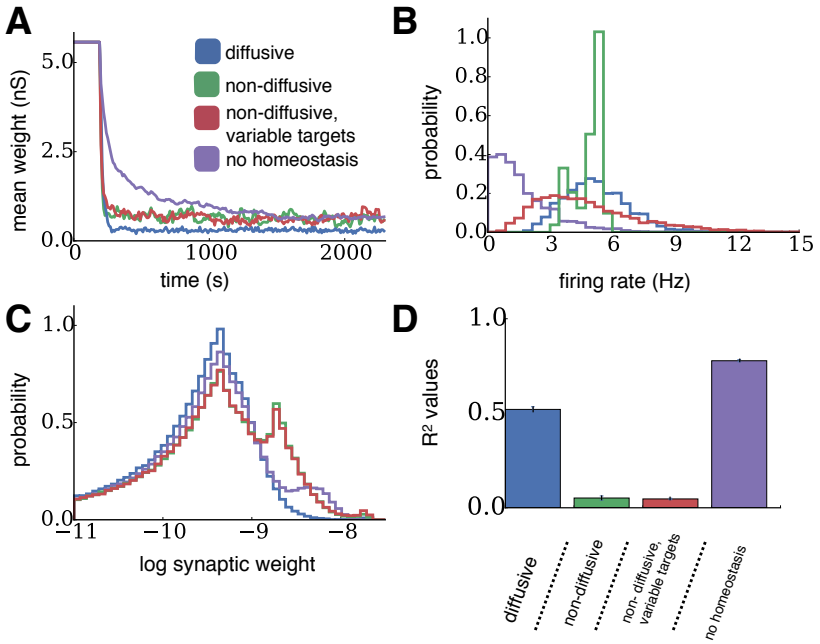

Figure 7. Diffusive homeostasis retains its properties in networks with STDP. (A) Evolution of mean synaptic weights in network with STDP and each case of homeostasis. **(B)** Distributions of steady-state firing rates after homeostasis and STDP. **(C)** Distributions of steady-state log synaptic weights after homeostasis and STDP. **(D)** Mean response linearity (measured by R^2 values) for networks with STDP and each case of homeostasis (error bars represent one standard deviation).

neuron. These results predict that disrupting neural diffusive NO signaling can affect perceptual and cognitive abilities through changes of neural population responses. While other non-diffusive homeostatic mechanisms would continue to stabilize neural activity, the lack of a signal related to the average population activity may disrupt the flexible maintenance of firing rate heterogeneity, and as a result the ability to represent network inputs.

Mean-field analysis revealed that these differences are essentially due to the diffusive messenger providing each neuron with a combination of the average network activity and its own activity as the homeostatic signal. Diffusion of the signal from highly active neurons causes a reduction in the activity of their neighbors, such that firing rates of highly active neurons do not have to be completely reduced in order for the population to achieve a target rate. As a consequence, diffusive homeostasis furnishes a network with an efficient way of flexibly maintaining heterogeneity of firing rates. These effects can also be understood by considering the neural transfer functions, as illustrated in Figures 8A-B, which provides an intuitive explanation for the broader and heavy-tailed firing rate distributions observed under diffusive homeostasis [17, 22].

Narrow firing rate distributions are an obvious consequence of local homeostatic processes, as for instance shown recently with homeostasis implemented as local synaptic metaplasticity [14]. This is in apparent conflict with the growing body of experiments documenting broad and heavy-tailed distributions of firing rates in cortex [11]. One could argue that a straightforward explanation is a process, for example genetic or developmental, which randomly assigns neurons heterogenous homeostatic targets. While we show here that this can result in broader firing rate distributions, we also found that this generally leads to networks with a mismatch between the neural dynamic ranges and input statistics, which in turn limits

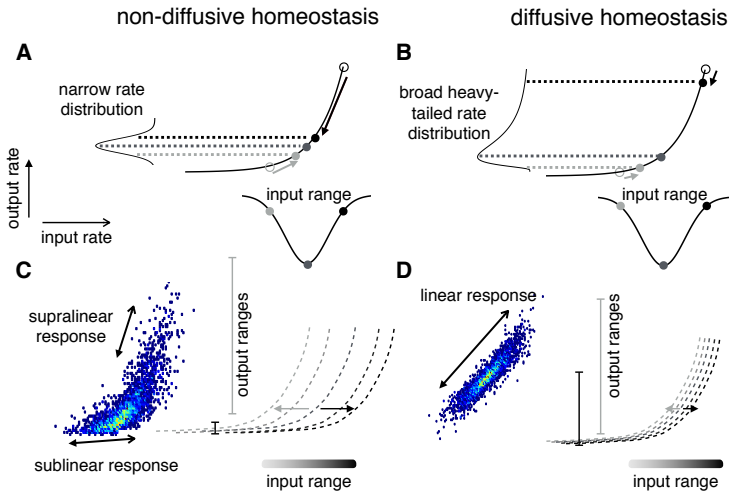

Figure 8. Comparison of the effects of diffusive and non-diffusive homeostasis on neural transfer functions. (A) For non-diffusive HIP the transfer function of each neuron is brought to center around its input in order to avoid response saturation or runaway excitation (adapted from [17]). **(B)** Diffusive homeostasis decorrelates the input and threshold of individual neurons, resulting in a population of neurons residing along the entire transfer function. The non-linear shape of the transfer function causes broad and heavy-tailed rate distributions observed with diffusive homeostasis. **(C-D)** Transfer functions following homeostasis of neurons receiving small (light grey) to large (dark grey) input are shown by the dotted lines. Following a fast input change, the output ranges across the entire range of inputs are shown as vertical grey lines, with the shade of grey corresponding to the neuron's previous input. Networks with non-diffusive homeostasis (C) exhibit response saturation (dark grey line) and supralinear responses (light grey line), while diffusive homeostasis (D) causes transfer functions to shift towards the center of the input distribution, leading to approximately linear responses.

the responsiveness of the network.

A striking feature of diffusive homeostasis is the lack of requirement for any such distribution of homeostatic targets, as the diffusive signal can be effectively exploited through providing a context for heterogeneity - neurons which maintain a significantly higher firing rate than the rest of the network also synthesize a higher level of the diffusive signal, thus ensuring that their deviation from the average firing rate is counterbalanced by lowering neighboring neurons' firing rates. This mechanism essentially allows neurons to differ in activity from the population as long as the population as a whole provides some compensation for these deviations. Moreover, this mechanism is compatible with the recent finding that a minority of cells were found to consistently be the most highly active and informative across brain states [23]. While non-diffusive homeostasis would have a disruptive effect on such a 'preserved minority' of neurons by reducing their activity towards those of the less active majority, diffusive homeostasis provides a substrate for maintaining their differentiated activity.

A significant distinction between the effects of diffusive and non-diffusive homeostasis appears when network responses to rapidly changing input are considered (Figure 5). We show that networks with diffusive homeostasis represent input changes more faithfully than those with non-diffusive homeostasis. Saturation of neurons' responses to large changes are observed in networks without diffusion - this effect is further illustrated in Figure 8C-D. Across a spatially homogeneous network, diffusing signals act to effectively shift the transfer function of each neuron towards the average network input, ensuring that neurons are responsive across the entire range of inputs presented to a network. Moreover, the asymmetric

response of networks with non-diffusive homeostasis causes the average network activity to increase after fast input changes, while it is constant for a network with diffusive homeostasis (Figure 5D). The latter case is consistent with observations that mean population firing rates are preserved across novel and familiar environments and across different episodes of slow-wave sleep [24].

Networks with diffusive homeostasis have an improved ability to accurately track time varying inputs (Figure 6A-B) as a direct consequence of their linear responses. Beneficial effects of neural heterogeneity for population coding have been suggested before [13, 25], but here we find that the broad linear response regime maintained by diffusive homeostasis further improves network performance. This improvement in network performance is also observed in a simplified stimulus orientation decoding task (Figure 6C-G). Networks with diffusive homeostasis perform better than those with non-diffusive homeostasis when the entire range of neural responses are used to construct a population vector in order to decode stimulus orientation (Figure 6G-H). Although there exist alternative methods for decoding stimuli, the population vector has been shown to exhibit performance close to the optimal maximum likelihood procedure for broad tuning, as was used in our example [26].

The distinctions between diffusive and non-diffusive homeostasis which have been described are conserved in networks with STDP (Figure 7). This demonstrates that the limitations of non-diffusive homeostasis in maintaining neural heterogeneity and responsiveness extend beyond the case of static inputs, towards more realistic situations in which neurons receive ongoing and diverse perturbations.

While NO is involved in a wide variety of neural processes throughout development and learning [27–29], these were ignored throughout for the sake of simplicity and tractability. Nonetheless, the impaired performance of nNOS knock-out mice in cognitive tasks [30] and the prevalence of epilepsy following nNOS inhibition [31] could be linked to diminished homeostatic control of neural excitability. Finally, the outcome of this study is not necessarily confined to NO, and could equally apply to other diffusive neurotransmitters observed in the brain such as hydrogen sulfide and carbon monoxide [32]. Indeed, we believe it serves to demonstrate the potential role of diffusive neurotransmitters as an economical and reliable signal of activity across a population of neurons.

Methods

Network model

We simulated a spiking network of leaky integrate-and-fire (LIF) model neurons with conductance-based synapses and injected Ornstein-Uhlenbeck noise, as described by

$$\frac{dv}{dt} = \frac{1}{\tau_m}(E_l - v) + g_e J_e(E_e - v) + g_i J_i(E_i - v) + \sigma_{OU} + J_{ext}(E_e - v)\big(\delta(t - t^{ext})\big) \tag{1}$$

$$\frac{dg_e}{dt} = -\frac{g_e}{\tau_e} + \sum_k \delta(t - t^k), \qquad \frac{dg_i}{dt} = -\frac{g_i}{\tau_i} + \sum_k \delta(t - t^k) \tag{2}$$

where v is the membrane potential, τ_m the membrane time constant, E_l the leak conductance reversal potential, and σ_{OU} the variance of the Ornstein–Uhlenbeck noise. g_e and g_i are the excitatory and inhibitory synaptic currents respectively, given by Equation (2), where t_k denotes the time of all k incoming spikes. The reversal potential of the synapses are denoted by E_e and E_i, the synaptic conductances by J_e and J_i, and the synaptic time constants by τ_e and τ_i. The external input conductance is given by J_{ext}, and t^{ext} denotes the arrival time of external input, modeled as an independent homogeneous Poisson process for each neuron i with rate μ_i. A spike is emitted whenever the membrane potential v exceeds the firing threshold θ, and the membrane potential is then reset to the resting potential value, v_r, after a refractory period, τ_{ref}.

The network was made up of N neurons; $0.8N$ excitatory and $0.2N$ inhibitory, with excitatory and inhibitory synaptic conductances scaled so that the network was in a balanced state [16]. Recurrent

connections were random and sparse, with connection probability $\epsilon = \frac{C}{N}$ independent of neuron type, where C defines the mean number of synapses per neuron. The balanced state was achieved in the network through scaling the inhibitory synaptic conductances by a factor of g, such that $J_i = gJ_e$.

NO synthesis and diffusion

We assumed that neuronal NO synthase (nNOS) is activated by Ca^{2+} influx following a spike (3) and describe the relationship through the Hill equation (4), which results in a sigmoidal concentration dependence, where n and K are parameters of the Hill equation [33] and $\tau_{Ca^{2+}}$ and τ_{nNOS} are the timescales of Ca^{2+} decay and nNOS activation respectively.

$$\frac{d[\mathrm{Ca}^{2+}]}{dt} = -\frac{[\mathrm{Ca}^{2+}]}{\tau_{\mathrm{Ca}^{2+}}} + [\mathrm{Ca}^{2+}_{\mathrm{spike}}]\delta(t_{\mathrm{spike}}) \tag{3}$$

$$\frac{d[\mathrm{nNOS}]}{dt} = \frac{1}{\tau_{\mathrm{nNOS}}} \frac{[\mathrm{Ca}^{2+}]^n}{[\mathrm{Ca}^{2+}]^n + K^n} \tag{4}$$

Throughout this paper we considered the case where all neurons, inhibitory and excitatory, express nNOS. The 2D diffusion equation (5) was solved numerically using a spatial resolution ds, with periodic boundary conditions defined by the torus, diffusion coefficient D and a decay term λ [34]. Neurons were represented by a point source according to their activated nNOS concentration.

$$\frac{d[\mathrm{NO}]}{dt} - D\nabla^2[\mathrm{NO}] = [\mathrm{nNOS}] - \lambda[\mathrm{NO}] \tag{5}$$

The homeostatic effect of NO was represented in neuron i by an increase in θ_i, the firing threshold, according to the relative difference in intracellular NO concentration [NO] and a target concentration $[\mathrm{NO}]_0$;

$$\frac{d\theta_i}{dt} = \frac{1}{\tau_{\mathrm{HIP}}} \frac{[\mathrm{NO}] - [\mathrm{NO}]_0}{[\mathrm{NO}]}, \tag{6}$$

where τ_{HIP} is the timescale of homeostasis.

Dynamic mean-field analysis

For a detailed derivation of equations used in our dynamic mean-field analysis, see [16] and [17]. Briefly, under the assumptions that the network is in an asynchronous regime and that a single EPSP is sufficiently small compared to the voltage required to elicit a spike from resting membrane potential, we can extract the mean firing rate of an LIF neuron in a recurrent network by solving a pair of equations under the condition of self-consistency. The synaptic current for a neuron i in a time interval τ can be described by its mean μ_i and standard deviation σ_i as follows:

$$\mu_i = JC\nu\tau, \qquad \sigma_i = J\sqrt{C\nu\tau}, \tag{7}$$

where J is the synaptic efficacy, C the number of synapses per neuron and ν the average population firing rate. The expected mean firing rate $\phi_i(\mu_i, \sigma_i)$ of an LIF neuron with this synaptic current is given by

$$\phi_i(\mu_i, \sigma_i) = \left[\sqrt{\pi}\tau_m \int_{\frac{v_r-\mu_i}{\sigma_i}}^{\frac{\theta_i-\mu_i}{\sigma_i}} dv e^{v^2} \mathrm{erfc}(-v)\right]^{-1}, \tag{8}$$

where erfc is the complementary error function. Since the firing rate described by (8) is determined by the synaptic current parameters μ_i and σ_i, which are in turn determined by the population firing rate ν,

self-consistency requires that the rate which determines the synaptic current parameters must also be equal to the rate which is produced by these parameters, that is:

$$\nu = \phi_i(\mu_i(\nu), \sigma_i(\nu)). \tag{9}$$

We simulated a non-interacting population of neurons described by the mean-field theory, in which all neurons are identical except for their threshold θ_i. Although there is no recurrent excitation within the population, the synaptic current statistics are comparable to that which a neuron within a recurrent network would receive. This enabled us to consider the firing rate distributions arising from presenting single neurons with distributions of synaptic currents, similar to the approach by [17].

We assumed that a neuron embedded in a homogeneous network receiving a diffusive homeostatic signal is analogous to a neuron using a combination of its own firing rate and the average population firing rate as a signal.

The network can then be reduced to a population in which the firing threshold θ_i of each neuron i is modulated according to

$$\frac{d\theta_i}{dt} = \frac{1}{\tau_{\text{HIP}}}\left((1-\alpha)\frac{\phi_i - \phi_0}{\phi_i} + \alpha\frac{\bar{\phi} - \phi_0}{\bar{\phi}}\right), \tag{10}$$

where ϕ_0 is the target firing rate and ϕ_i and $\bar{\phi}$ are the firing rate of the neuron i and the population respectively. α was varied between 0 and 1 and can be thought of as the proportion of NO which a neuron receives due to diffusion from other neurons, with $\alpha = 0$ indicating that each neuron senses only its own activity and $\alpha = 1$ indicating that all neurons share an identical population-wide signal.

In order to implement homeostasis in this setup, we iterated through (7)-(8) until (9) is satisfied to a precision of 10^{-4} Hz, where (8) returns ϕ_i for each neuron in the population, and $\bar{\phi} = \frac{\sum \phi_i}{N}$ is used as ν in (7). At each timestep the thresholds θ_i of each neuron were modulated according to (10), and rates ϕ_i were subsequently recalculated from (8).

Procedure for Figure 2.

External input rates μ_i for each neuron i were randomly drawn from a Gaussian distribution such that $\mu_i \sim \mathcal{N}(25, 10)$ Hz. Since the mean NO concentration takes time to reach a steady state in the recurrent network simulations, we ran the network for 100 s without homeostasis and with all neurons receiving 5 Hz input, defining the target NO concentration $[\text{NO}]_0$ to be the mean NO concentration across all neurons at 100 s.

For the dynamic mean field analysis, we chose parameters which roughly match the rate statistics of the recurrent network simulations. Inputs to each neuron were drawn from a Gaussian distribution such that $\mu_i \sim \mathcal{N}(5.7, \delta)$, $\sigma_i = \sqrt{\mu_i}$. $\delta = 0.4$ is the width of the input distribution to the population.

Adding target variability.

In order to match the distribution of effective targets observed during diffusive homeostasis for networks with non-diffusive homeostasis, we assigned each neuron a target $[\text{NO}_0]_i$, which is its steady-state intracellular NO concentration after the network has been presented with inputs which led to a firing rate distribution matching that of the network with diffusive homeostasis. A similar approach was adopted in the dynamic mean-field analysis, with each neuron assigned a target firing rate $\phi_{0,i}$ from the steady-state firing rate distribution of a network with $\alpha = 0.8$.

Procedure for Figure 4.

External input for each neuron i was $\mu_i = 2$ Hz ($N = 1000$). NO_0 was set as described previously, although with an input rate of 2 Hz. 2 groups of 100 excitatory neurons each were randomly chosen,

independent of neuron position, and stimulated with $\mu_{\text{green}} = 5$ Hz and $\mu_{\text{blue}} = 10$ Hz, keeping the inputs to the remaining neurons unchanged. Firing rates plotted in Figures 4A-B were smoothed with a uniform time window of 20 s.

Procedure for Figure 5.

Figures 5A-D were generated using the same simulation setup as described previously. After the network has reached the homeostatic target firing rate, we froze homeostasis. Input rates to each neuron were then regenerated from the same original input distribution, such that $\mu_i^{\text{after}} \sim \mathcal{N}(25, 10)$ Hz. $\Delta\mu_i = \mu_i^{\text{after}} - \mu_i^{\text{before}}$ is the difference in input rate each neuron experiences upon this change, and $\Delta\nu_i = \nu_i^{\text{after}} - \nu_i^{\text{before}}$ is the corresponding change in output rate for each neuron. The black lines in Figures 5A-C are from least-squares linear regression, and the R^2 values given were derived from this fit. A similar approach was used in the dynamic mean-field analysis, while varying δ, the width of the input distribution. R^2 value in Figure 5E were normalised so that the maximal value across each δ is 1. Other parameters remained unchanged.

Time-varying input.

In addition to the external input μ_i previously described, the network was randomly separated into groups of 250 neurons. Each group j was stimulated with external Poisson input with a rate given by $\mu_{j,t} \sim \mathcal{N}(0, 25)$ Hz. These inputs were regenerated at each timestep t of length 1 s. The time-varying input $\mu_{j,t}$ was also presented during homeostasis. The dotted black line in Figure 6A shows the normalised $\mu_{j,t}$, while coloured lines show the normalised rate deviation of a randomly chosen group j from the mean population firing rate.

Decoding stimulus orientation.

Each excitatory neuron was randomly assigned a preferred orientation. After the network reached a steady state, homeostasis was frozen. For each trial, each neuron i with preferred orientation θ_i was stimulated with external Poisson input at a rate given by $\mu_b + \mathcal{N}(\theta_s, \sigma_s, \theta_i)$, where $\mu_b = 20$ Hz is the base input rate and $\mathcal{N}(\theta_s, \sigma_s, \theta_i)$ is the amplitude at θ_i of a Gaussian tuning curve centered around the stimulus orientation θ_s, with a width given by $\sigma_s = 90°$ and a peak amplitude of 2.5 Hz. The angle decoded using the population vector method is the angle of the vector sum of all neural responses.

Spike-timing-dependent plasticity.

A spike-timing dependent plasticity rule, as described in [19], was implemented in each recurrent excitatory-excitatory synapse. Both potentiation and depression are additive in this rule, with no weight dependence. For each pair of pre- and post-synaptic spikes separated by a time Δt, the synaptic weight is updated by a value Δw given by

$$\Delta w = \begin{cases} A_+ \exp(\Delta t/\tau_+) g_{\max}, & \text{if } \Delta t < 0. \\ -A_- \exp(\Delta t/\tau_-) g_{\max}, & \text{if } \Delta t \geq 0. \end{cases} \tag{11}$$

Above, τ_+ and τ_- denote the timecourse over which potentiation and depression occur respectively, while A_+ and A_- denote the relative strengths of potentiation and depression. $g_{\max}$ is the maximum synaptic weight. $A_+ < A_-$ such that irregular firing is maintained within a reasonable range of rates.

Model parameters.

Unless explicitly defined, the parameters used throughout the paper were as follows: $N = 5000$, $\tau_{\text{HIP}} = 2500$ms, $J_{\text{ext}} = 80.0$ nS, $J_e = 4.0$ nS, $J_i = 64.0$ nS, $C = 100$, $\theta_0 = -50$ mV, $\tau_m = 20$ ms, $\tau_e = 3$ ms,

$\tau_i = 7$ ms, $\tau_{\text{ref}} = 5$ ms, $c_m = 0.2$ nF, $E_l = -80$ mV, $E_e = 0$ mV, $E_i = -70$ mV, $v_r = -60$ mV, $[\text{Ca}^{2+}_{\text{spike}}] = 1$, $\tau_{\text{Ca}^{2+}} = 10$ ms, $\tau_{\text{nNOS}} = 100$ ms, $n = 3$, $K = 3$, $D = 0.1$, $\lambda = 10$, $ds = 5\text{x}10^{-4}$, $\tau_{OU} = 1$ ms, $\sigma_{OU} = 1$ mV. For the dynamic mean-field model $\theta_0 = 10$ and $v_r = 0$. For the STDP rule, $g_{\max} = 20$nS, $\tau_+ = \tau_- = 20$ms, $A_+ = .025$, $A_- = .0625$.

For synthesis, diffusion, and decay of NO we have attempted to match data when available [34], although the dearth of experimental measurements does not permit for great precision [35, 36]. Additionally, parameters were chosen such that the timescale of homeostasis is separated from that of firing rate fluctuations. This is a reasonable assumption, given that activity-dependent NO modulation likely acts within 10 minutes or slower [5]. τ_{HIP} was chosen to be long enough so as to avoid oscillations but short enough so as to allow feasible large scale simulations. This is a common assumption in computational studies [10]. Larger simulations, up to $N = 25000$, were run with no discernible difference in results.

All numerical simulations were implemented using the Brian simulator [37], and the mean-field analysis was implemented using IPython Notebook [38]. Data analysis was performed with the numpy Python package and plotting with the matplotlib package and seaborn library [39, 40]. Code and IPython Notebooks which perform the data analysis and plotting will be available on ModelDB and a public github repository following peer review. In the meantime a minimal example of diffusion of a neurotransmitter over a 2D surface is available at http://tinyurl.com/sweeney-diffusion, which may easily be incorporated into existing Brian models.

Acknowledgments

We thank Mark van Rossum, David Sterratt, and Alex Cayco Gajic for their careful reading of the manuscript and helpful advice. This work was supported by MRC CDA Fellowship G0900425 (MHH), and the EuroSPIN Erasmus Mundus doctoral programme, UK Engineering and Physical Sciences Research Council, UK Biotechnology and Biological Sciences Research Council and the UK Medical Research Council (grant numbers EP/F500385/1 and BB/F529254/1 for the University of Edinburgh School of Informatics Doctoral Training Centre in Neuroinformatics and Computational Neuroscience) (YS).

References

1. Garthwaite J (2008) Concepts of neural nitric oxide-mediated transmission. Eur J Neurosci 27: 2783-2802.

2. Pape HC, Mager R (1992) Nitric oxide controls oscillatory activity in thalamocortical neurons. Neuron 9: 441–448.

3. Gally J, Montague P, Reeke G, Edelman G (1990) The NO hypothesis: possible effects of a short-lived, rapidly diffusible signal in the development and function of the nervous system. P Natl Acad Sci USA 87: 3547-3551.

4. Steinert J, Cornelia K, Baker C, Challiss R, Mistry R, et al. (2008) Nitric oxide is a volume transmitter regulating postsynaptic excitability at a glutamatergic synapse. Neuron 60: 642-656.

5. Steinert J, Robinson S, Tong H, Haustein M, Cornelia K, et al. (2011) Nitric oxide is an activity-dependent regulator of target neuron intrinsic excitability. Neuron 71: 291-305.

6. Lourenço CF, Santos RM, Barbosa RM, Cadenas E, Radi R, et al. (2014) Neurovascular coupling in hippocampus is mediated via diffusion by neuronal-derived nitric oxide. Free Radic Biol Med 73: 421-429.

7. LeMasson G, Marder E, Abbott L (1993) Activity-dependent regulation of conductances in model neurons. Science 259: 1915-1917.

8. Lazar A, Pipa G, Triesch J (2009) SORN: a self-organizing recurrent neural network. Front Comput Neurosci 3: 23.

9. Olypher A, Prinz A (2010) Geometry and dynamics of activity-dependent homeostatic regulation in neurons. J Comput Neurosci 28: 361-374.

10. Naudé J, Cessac B, Berry H, Delord B (2013) Effects of cellular homeostatic intrinsic plasticity on dynamical and computational properties of biological recurrent neural networks. J Neurosci 33: 15032-15043.

11. Wohrer A, Humphries M, Machens C (2013) Population-wide distributions of neural activity during perceptual decision-making. Prog Neurobiol 103: 156-193.

12. Marsat G, Maler L (2010) Neural heterogeneity and efficient population codes for communication signals. J Neurophysiol 104: 2543–2555.

13. Tripathy SJ, Padmanabhan K, Gerkin RC, Urban NN (2013) Intermediate intrinsic diversity enhances neural population coding. P Natl Acad Sci USA 110: 8248-53.

14. Zenke F, Hennequin G, Gerstner W (2013) Synaptic plasticity in neural networks needs homeostasis with a fast rate detector. PLoS Comput Biol 9: e1003330.

15. O'Leary T, Williams AH, Caplan JS, Marder E (2013) Correlations in ion channel expression emerge from homeostatic tuning rules. P Natl Acad Sci USA 110: E2645–E2654.

16. Brunel N (2000) Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. J Comput Neurosci 8: 183-208.

17. Roxin A, Brunel N, Hansel D, Mongillo G, van Vreeswijk C (2011) On the distribution of firing rates in networks of cortical neurons. J Neurosci 31: 16217-16226.

18. Holmgren C, Harkany T, Svennenfors B, Zilberter Y (2003) Pyramidal cell communication within local networks in layer 2/3 of rat neocortex. J Physiol 551: 139–153.

19. Song S, Miller KD, Abbott LF (2000) Competitive hebbian learning through spike-timing-dependent synaptic plasticity. Nat Neurosci 3: 919–926.

20. Sjöström PJ, Turrigiano GG, Nelson SB (2001) Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. Neuron 32: 1149–1164.

21. Van Rossum MC, Bi GQ, Turrigiano GG (2000) Stable hebbian learning from spike timing-dependent plasticity. J Neurosci 20: 8812–8821.

22. Ermentrout B (1998) Linearization of FI curves by adaptation. Neural Comput 10: 1721–1729.

23. Mizuseki K, Buzsáki G (2013) Preconfigured, skewed distribution of firing rates in the hippocampus and entorhinal cortex. Cell Reports 4: 1010–1021.

24. Hirase H, Leinekugel X, Czurkó A, Csicsvari J, Buzsáki G (2001) Firing rates of hippocampal neurons are preserved during subsequent sleep episodes and modified by novel awake experience. P Natl Acad Sci USA 98: 9386–9390.

25. Shamir M, Sompolinsky H (2006) Implications of neuronal diversity on population coding. Neural Comput 18: 1951–1986.

26. Seung HS, Sompolinsky H (1993) Simple models for reading neuronal population codes. P Natl Acad Sci USA 90: 10749–10753.

27. Nikonenko I, Nikonenko A, Mendez P, Michurina T, Enikolopov G, et al. (2013) Nitric oxide mediates local activity-dependent excitatory synapse development. P Natl Acad Sci USA 110: E4142–E4151.

28. Tamagnini F, Barker G, Warburton C, Burattini C, Aicardi G, et al. (2013) Nitric oxide-dependent LTD but not endocannabinoid-LTP is crucial for visual recognition memory. J Physiol 591: 3963–3979.

29. Son H, Hawkins RD, Martin K, Kiebler M, Huang PL, et al. (1996) Long-term potentiation is reduced in mice that are doubly mutant in endothelial and neuronal nitric oxide synthase. Cell 87: 1015-23.

30. Weitzdoerfer R, Hoeger H, Engidawork E, Engelmann M, Singewald N, et al. (2004) Neuronal nitric oxide synthase knock-out mice show impaired cognitive performance. Nitric Oxide 10: 130-140.

31. Del-Bel E, Oliveira P, Oliveira J, Mishra P, Jobe P, et al. (1997) Anticonvulsant and proconvulsant roles of nitric oxide in experimental epilepsy models. Braz J Med Biol Res 30: 971–979.

32. Wang R (2002) Two's company, three's a crowd: can H2S be the third endogenous gaseous transmitter? FASEB J 16: 1792-1798.

33. Hill AV (1910) The possible effects of the aggregation of the molecules of haemoglobin on its dissociation curves. J Physiol 40: 4–7.

34. Philippides A, Husbands P, O'Shea M (2000) Four-dimensional neuronal signaling by nitric oxide: a computational analysis. J Neurosci 20: 1199-1207.

35. Hall C, Garthwaite J (2009) What is the real physiological NO concentration in vivo? Nitric Oxide 21: 92-9103.

36. Batchelor AM, Bartus K, Reynell C, Constantinou S, Halvey EJ, et al. (2010) Exquisite sensitivity to subsecond, picomolar nitric oxide transients conferred on cells by guanylyl cyclase-coupled receptors. P Natl Acad Sci USA 107: 22060-5.

37. Goodman D, Brette R (2008) Brian: a simulator for spiking neural networks in python. Front Neuroinform 2: 192–7.

38. Pérez F, Granger BE (2007) IPython: a system for interactive scientific computing. Comput Sci Eng 9: 21–29.

39. Van Der Walt S, Colbert SC, Varoquaux G (2011) The numpy array: a structure for efficient numerical computation. Comput Sci Eng 13: 22–30.

40. Hunter JD (2007) Matplotlib: A 2d graphics environment. Comput Sci Eng 9: 0090–95.

Supporting Information

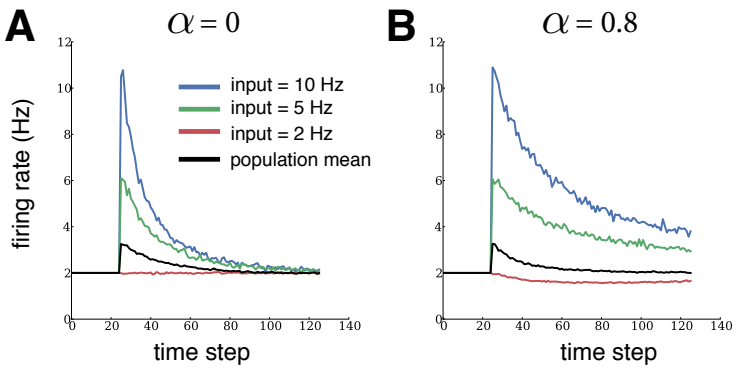

Figure S1 Input group separation can be achieved in the mean-field model by including global activity read-out. (A-B) Evolution of firing rates in a dynamic mean-field population for each input group for purely local ($\alpha = 0$) and mixed local and global activity read-out ($\alpha = 0.8$). Black traces show population activity. Independent Poisson input at the target rate is given to 800 neurons (red), while two groups of 100 neurons each receive elevated Poisson input. Note that changing the input of groups of neurons in the recurrent network also affects the activity of neurons with fixed input (Figures 4B, red traces) due to recurrent connections, an effect that is obviously absent in the mean field description. Thresholds are modulated by homeostasis at each time step.

Procedure for Figure S2.

Spatial dependence in the connection probability between two neurons was introduced as follows:

$$P_c(d) = \epsilon \, \mathrm{e}^{-\frac{d^2}{2s^2}}, \tag{12}$$

where d is the Euclidean distance between the neurons and s is a constant defining the connectivity range of the network. 2D positions on the torus are bounded such that $x, y \in (0, 1)$. Given a diffusive range of 0.1, values for s in Figure 2 were therefore set as 0.05, 0.1, and 0.5. The ratio of s and the diffusive range was defined as ρ, which had values of 0.5, 1.0 and 5.0.

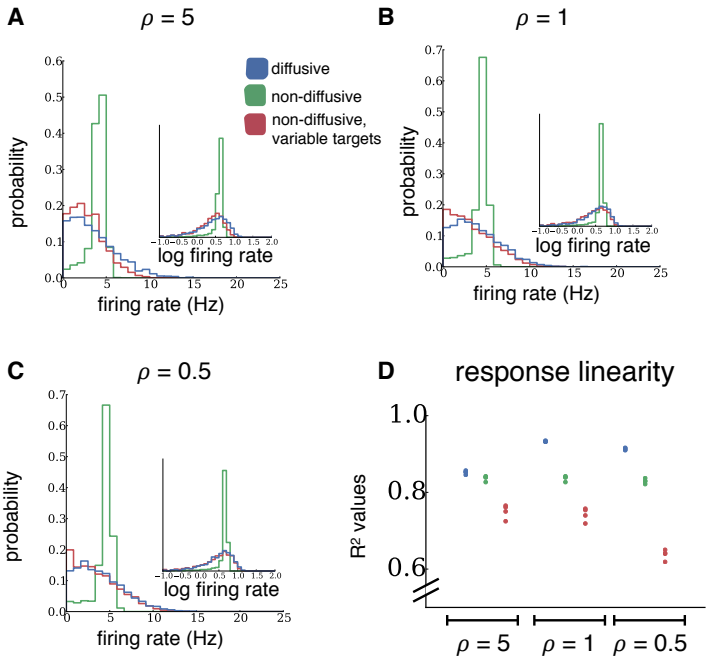

Figure S2 Linearity of network responses in networks with spatially restricted connection probabilities. These results qualitatively agree with those for the random networks used throughout the study. The simulations shown here were identical to those in Figures 2D and 5 of the main text, but the connection probability between neurons had a Gaussian shape. ρ, the ratio of the connectivity range and the diffusive range, is varied across a wide range of values (ρ=0.5,1.0,5.0). **(D)** Each point represents the R^2 value of a linear fit as in Figures 5A-C, for one network.