

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

Yann Sweeney^{1,2,*}, Jeanette Hellgren Kotaleski², Matthias H. Hennig¹

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

* 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 connectivity. 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 neighbors 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 across cell membranes, a phenomenon often called *volume transmission* [3]. An important role of NO signaling is to regulate neural excitability through 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 responsive to changes in network dynamics. Finally, we demonstrate that these effects are preserved in networks whose recurrent synaptic inputs undergo Hebbian plasticity.

Results

We investigated the effects of diffusive homeostasis in a recurrent neural network (Fig. 1A, 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 (Fig. 1B, 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 (Fig. 2). The latter corresponds to a canonical model of HIP as investigated before [8,9].

Diffusive homeostasis enables a broad firing rate distribution

Fig. 1C 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, diffusive homeostasis 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 Fig. 1C. 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.

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 (Fig. 1D). 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 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 11 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, $\alpha = 0.8$ 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 (Fig. 1E).

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 threshold. Since setting $\alpha > 0$ (analogous to diffusive homeostasis) introduces a correlation between a neuron's threshold θ_i and the average population threshold $\bar{\theta}$, 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 $\sigma(\mu, \theta) = 0.99$ (Fig. 1F).

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. Diffusive homeostasis, on the other hand, yielded similar firing thresholds across the population (Figs. 3A-B). 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). This result was robust to changes in the rate of NO diffusion. While decreasing the rate of diffusion, D , did result

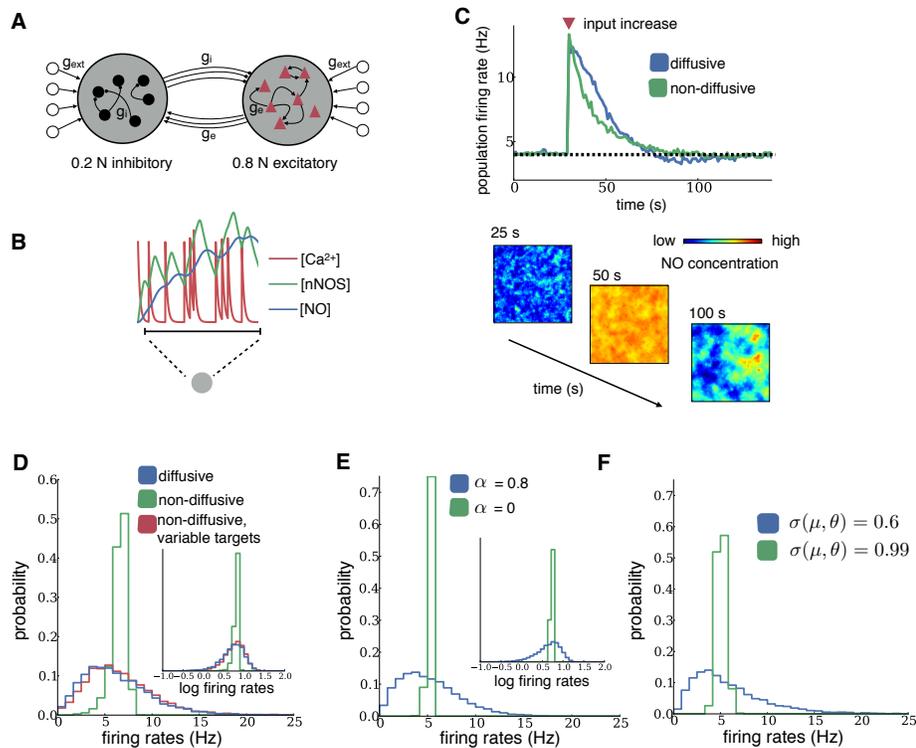


Figure 1. Steady-state behavior of diffusive and non-diffusive homeostasis. (A) Schematic of the sparsely connected recurrent network model. Neurons received homogeneous random spiking input (g_{ext}). (B) Intracellular homeostatic signals in a model neuron. Each spike triggers calcium influx, which 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.

in slightly narrower firing rate distributions, they were broader than in networks with non-diffusive homeostasis across a wide range of values (Fig. 3C). A similar trend was observed when varying the width of the external input rate distribution. While decreasing this width led to a decrease in the width of the firing rate distribution, they were consistently broader in networks with diffusive homeostasis (Fig. 3D).

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 (see Methods), we could produce a distribution of firing rates similar to that observed with diffusive homeostasis (Figs. 1D, 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.

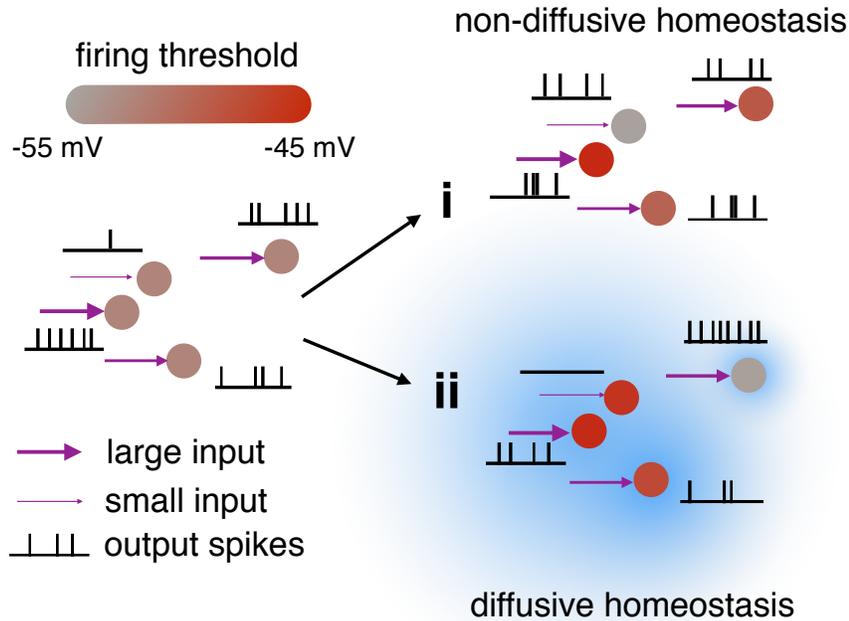


Figure 2. Illustration of the effects of non-diffusive (i) and diffusive (ii) homeostasis. Non-diffusive homeostasis adjusts each neuron's threshold (red color bar) according to its input to give identical firing rates, while diffusive homeostasis induces correlations (blue cloud) in the thresholds of neighboring neurons, thereby maintaining diverse firing rates.

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 Fig. 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 (Figs. 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 (Fig. 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 (Fig. 4B, colored traces). Introducing variability in homeostatic targets for the non-diffusive case, as described previously, did not maintain separability of individual groups as in the diffusive case. Instead, the different groups returned to their mean firing rates that existed before inputs were elevated (Fig. 4C).

The distribution of final firing thresholds explains these differences (Fig. 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 the final threshold of each neuron

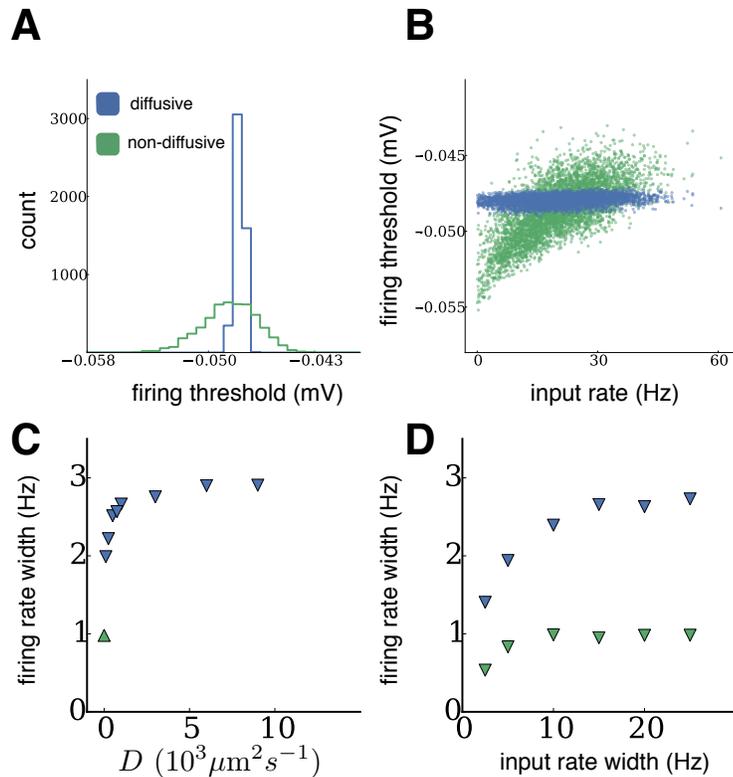


Figure 3. Steady-state firing thresholds of diffusive and non-diffusive homeostasis. (A) Distributions of firing thresholds after homeostasis from network simulations, receiving independent Poisson input drawn from a Gaussian distribution. (B) Steady-state firing thresholds plotted against external inputs received during homeostasis. (C) Width of the steady-state firing rate distribution as the diffusion coefficient D is varied ($D = 1000 \mu\text{m}^2 \text{s}^{-1}$ in panels A-B). (D) Width of the steady-state firing rate distribution as the input rate width is varied, for networks with diffusive ($D = 1000 \mu\text{m}^2 \text{s}^{-1}$) and non-diffusive homeostasis (input rate width = 10 Hz in panels A-B).

reflecting its input. Note that the distribution of firing thresholds is broader in this setup than in (Fig. 3A-B), as a broader range of inputs is given to the network. 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 distribution of final thresholds which are less segregated according to their input compared to a network with non-diffusive homeostasis. Thus, firing thresholds in neurons undergoing diffusive homeostasis were more 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 depend more strongly on their external input, similar to the non-diffusive case.

We could broadly 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 (Figs. 4G). Note that changing the input of groups of neurons in the

recurrent network also affects the activity of neurons with fixed input (Figs. 4B, red traces) due to recurrent connections, an effect that is obviously absent in the dynamic mean-field description. Increasing α led to local firing rate differences persisting for longer periods of time. However, these differences eventually decay very slowly, only remaining stable for the case where $\alpha = 1$ (Figs. 4H). The reason this occurs is, even after the population activity has quickly reached its homeostatic target, the deviations of the input groups still exert a small homeostatic force when $\alpha < 1$. For example, if $\alpha = 0.95$, there will be a relatively fast change in thresholds as the population activity reaches its target, followed by much slower changes, at $1 - \alpha = 0.05$ times the speed (Figs. 4I). This does not happen to the same extent in the spiking network simulations with diffusive homeostasis, as diffusion of NO ensure that deviations from the population activity are directly compensated for by neighboring neurons. Differences persist for 3245 ± 440 s, compared with 115 ± 6 s and 140 ± 40 s for non-diffusive homeostasis with uniform and variable homeostatic targets, respectively (Data not shown. \pm symbol denotes standard error of the mean of 6 independent network realizations in each case, see Methods). Since we have increased the speed of homeostasis in order to reduce simulation time (see Methods), a more realistic time course of 15 minutes for NO modulation would cause input differences to persist in networks with diffusive homeostasis for many hours to days [5]. 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.

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 fast changes in activity, since we assume that homeostasis is not active in these time scales. 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. 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, although this may indeed be desirable in some brain regions. 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.

Figs. 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 superlinear 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 as in the non-diffusive network with homogeneous targets, with an R^2 value of 0.38.

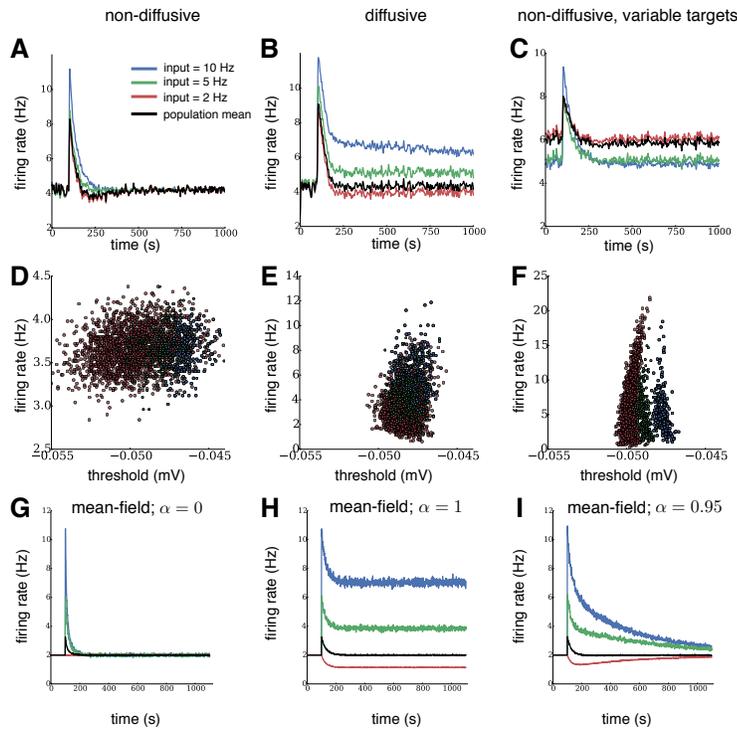


Figure 4. Relative differences between inputs are preserved during diffusive homeostasis. (A-C) Evolution of firing rates in a recurrent network for each input group under diffusive and non-diffusive homeostatic control, and for variable homeostatic targets. Black traces show population activity. Independent Poisson input at the target rate was given to 2000 neurons (red), while two groups of 250 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 rates versus firing thresholds. (G-I) Evolution of firing rates in a dynamic mean-field population for each input group for purely local ($\alpha = 0$), purely global ($\alpha = 1$), and mixed local and global activity read-out ($\alpha = 0.95$). Black traces show population activity. Independent Poisson input at the target rate is given to 2000 neurons (red), while two groups of 250 neurons each receive elevated Poisson input.

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 (Fig. 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. Crucially, the benefits of a diffusive homeostatic signal can be achieved by a relatively broad range of values for the rate of diffusion, D , indicating that the effects we describe are robust to precise parameter choices (Fig. 5E).

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. Fig. 5F shows 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 smaller δ , hence cases where the inputs are rather narrow, the optimal α decreased as δ increased, as well as the overall response linearity. This dependence on input width can be explained by considering the manner in which a population of neurons with a distribution of dynamic ranges span a range of inputs. If this range of inputs is small, then all neurons will span it regardless of their dynamic range (determined by their firing threshold), hence the high values of R^2 for $\delta = 0.1$. For an intermediate range of inputs, neurons whose dynamics range is best adapted to the average input are most responsive. This is achieved by increasing α . If the range of inputs is very large ($\delta = 1.0$), R^2 values are low since the dynamic ranges of the population cannot span the inputs. This effect is stronger at high α , as firing thresholds are more correlated, and the dynamics range of most neurons cannot capture the full input variance.

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, Fig. S1).

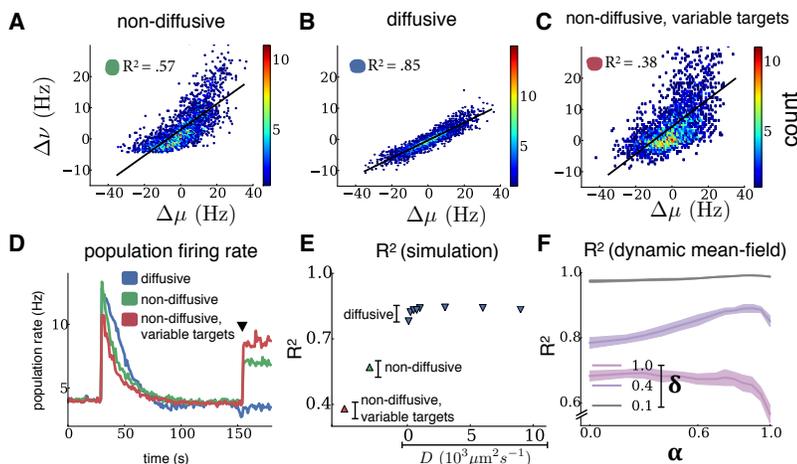


Figure 5. Diffusive homeostasis enables linear network responses. (A-C) Firing rate changes $\Delta\nu$ of all neurons following input changes $\Delta\mu$. Black lines show linear fits, with corresponding R^2 values inset. (D) Population activity before and after input change (black triangle). (E) Linearity of the population response to a change of inputs as the diffusion coefficient D is varied ($D = 1000\mu\text{m}^2\text{s}^{-1}$ in panel B). (F) 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.

Overall, these results suggest that networks undergoing diffusive homeostasis are better suited to linearly 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). Fig. 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; Fig. 6B). 244

We can explore these differences further by constructing a simplified task in which a 245
population of orientation-selective neurons respond to the orientation of a stimulus (see 246
Fig. 6C-F , Methods). This is not intended to represent circuits which perform this task 247
in the brain, but to serve purely as a demonstration of the relative merits of linear and 248
non-linear network responses. 249

Neurons in the network are randomly assigned a preferred stimulus orientation, 250
independent of their spatial position. A stimulus of a certain orientation can then be 251
presented to the network by varying the external input rates of each neuron, with 252
neurons whose preferred orientation is closest to the stimulus orientation receiving the 253
highest input rate. The stimulus orientation can be decoded from the network by taking 254
the vector average of the stimulus response across all neurons. The orientation of this 255
vector average, or population vector, is the decoded stimulus orientation. Networks with 256
linear responses perform better than those with non-linear responses in decoding 257
stimulus orientation, as measured by the standard deviation of errors in the orientation 258
of the population vector compared to the stimulus orientation (41° , 63° and 72° for 259
diffusive homeostasis, non-diffusive, and non-diffusive with variable targets respectively, 260
Fig. 6G). 261

Properties of diffusive homeostasis are conserved in networks 262 with Hebbian plasticity 263

In the networks described so far, we have used static and uniform synaptic weights for 264
recurrent connections. We next considered whether the observed properties of diffusive 265
homeostasis are altered by the presence of plastic synaptic weights, in particular when 266
Hebbian spike-timing-dependent plasticity (STDP) is introduced (see Methods). Using 267
a standard model of STDP with additive depression and potentiation for all recurrent 268
excitatory synapses, we simulated networks with both STDP and homeostasis active 269
until synaptic weight and firing rate distributions reached a steady state [19]. As before, 270
firing rate distributions were broader in networks with diffusive homeostasis (Fig. 7A). 271
Broad distributions could also be achieved by introducing variability in homeostatic 272
targets. Spiking activity remained asynchronous after STDP, as shown by the 273
distribution of inter-spike intervals (Fig. 7A, inset) [20], and the additive STDP rule led 274
to a bimodal distribution of synaptic weights (Fig. 7B), as previously reported [19]. 275
STDP amplified the differences in response linearity that was observed between 276
homeostatic cases. Inputs to each neuron were regenerated from the same distribution 277
presented during plasticity, and the corresponding change in output rate was compared 278
to the change in input rate, as in Figs. 5A-C. While the response linearity, given by the 279
mean R^2 value, for networks with diffusive homeostasis was 0.16, networks with 280
non-diffusive homeostasis exhibited much lower mean values of 0.01 and 0.02, for 281
uniform and variable homeostatic targets respectively (Fig. 7C). Networks without any 282
homeostasis had a mean value of 0.1. The lower R^2 in all cases compared to networks 283
with static synapses is due to the smaller impact that changes in external input have on 284
these networks, as STDP increases the ratio of recurrent input to external input. 285

We observed qualitatively similar retention of broad firing rate distributions and 286
response linearity with diffusive homeostasis when a weight-dependent update rule was 287
used (not shown), which has been argued to lead to more realistic weight 288
distributions [21]. 289

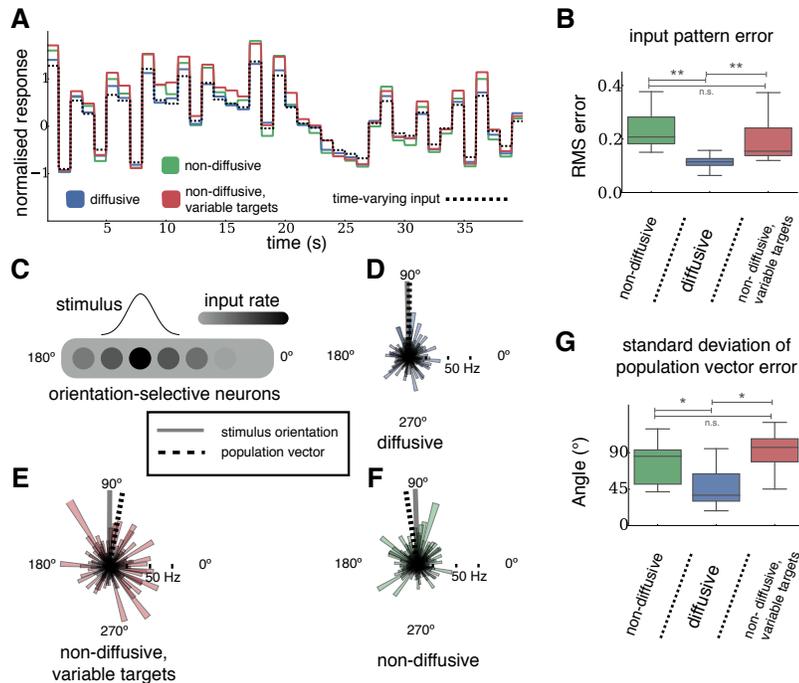


Figure 6. Performance advantages of networks with diffusive over non-diffusive homeostasis (A) Response of each network to a time-varying input pattern (dotted black line). Colored lines show normalized firing rate deviations from the mean population rate of those neurons receiving inputs. **(B)** RMS 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 Kolmogorov-Smirnov 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 the population vector is shown by the dashed line. **(G)** Standard deviation of errors in the orientation of the population vector in response to a stimulus (100 trials for 24 networks in each case). * $p < 0.05$ from a two-sided Kolmogorov-Smirnov 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

290

291

292

293

294

295

296

297

298

299

300

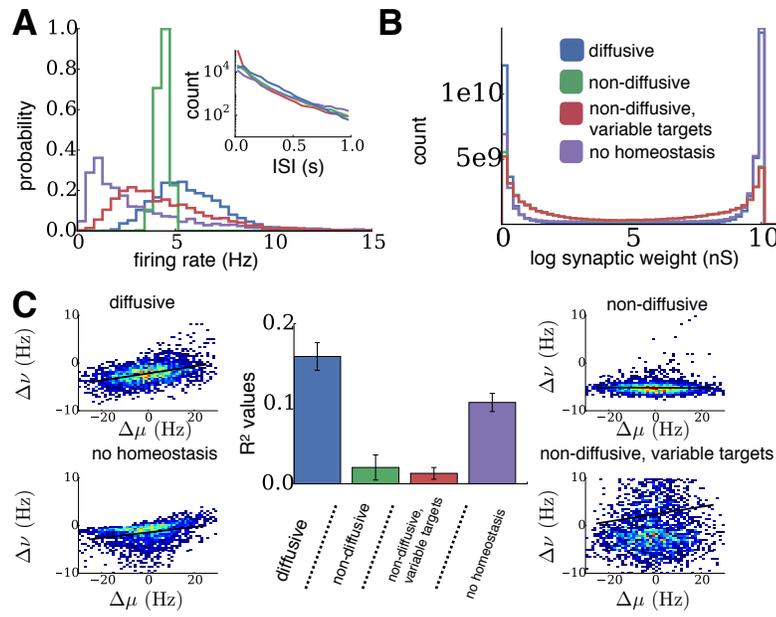


Figure 7. Diffusive homeostasis retains its properties in networks with STDP. (A) Distributions of steady-state firing rates after homeostasis and STDP. (Inset) Distribution of inter-spike intervals at the steady-state (log scale). (B) Distributions of steady-state synaptic weights after homeostasis and STDP. (C) Mean response linearity (measured by R^2 values) for networks with STDP and each case of homeostasis ($n=9$ networks, error bars represent one standard deviation). Surrounding plots show firing rate changes $\Delta\nu$ of all neurons following input changes $\Delta\mu$ for an example network in each case (black line shows linear fit).

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 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 Figs. 8A-B, which provides an intuitive explanation for the differences in firing rate distributions observed under diffusive homeostasis [17, 22]. For non-diffusive homeostasis the transfer function of each neuron is brought to center around its input, leading to a narrowing of the firing rate distribution. Diffusive homeostasis decorrelates the input and threshold of individual neurons, resulting in a population of neurons residing along the entire transfer function. This preserves the non-linear shape of the transfer function, causing broad and heavy-tailed firing rate distributions.

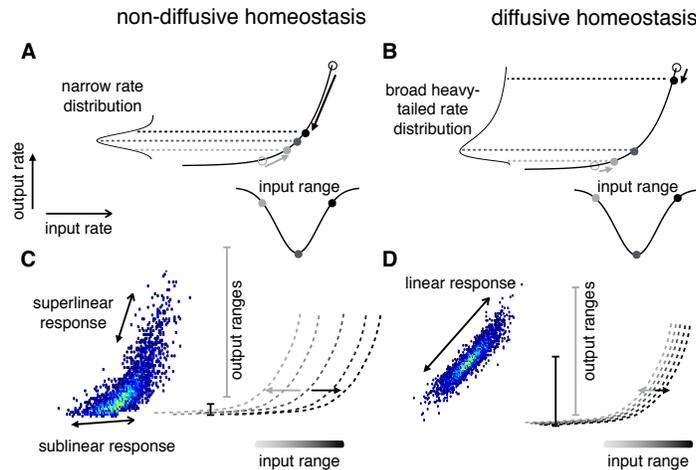


Figure 8. Comparison of the effects of diffusive and non-diffusive homeostasis on neural transfer functions. (A) For non-diffusive HIP the transfer function of each neuron is brought to center around its input in order to avoid response saturation or runaway excitation. Arrows show the action of HIP on neurons (circles) with a given input (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 an instantaneous 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 superlinear 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.

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 heterogeneous 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 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 (Fig. 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 Fig. 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. This is in contrast to networks with non-diffusive homeostasis, in which individual neurons are only responsive in a range around their current input. 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 (Fig. 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, 25]

Networks with diffusive homeostasis have an improved ability to accurately track time varying inputs (Fig. 6A-B) as a direct consequence of their linear responses. Beneficial effects of neural heterogeneity for population coding have been suggested before [13, 26], 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 (Fig. 6C-G). Networks with diffusive homeostasis perform better than those with non-diffusive homeostasis when a population vector is constructed from the neural responses 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 [27].

These distinctions between diffusive and non-diffusive homeostasis are conserved in networks with STDP (Fig. 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. Indeed, networks with non-diffusive homeostasis lost almost all sensitivity to external inputs after STDP, while networks with diffusive homeostasis retained this sensitivity (Fig. 7C).

The consequences of diffusive and non-diffusive homeostasis coexisting were also explored, by implementing these mechanisms simultaneously in a single network (not shown). Stable homeostatic activity could be robustly maintained, with the resulting network behaviour depending on the relative timescales of the non-diffusive and diffusive mechanisms. If non-diffusive homeostasis acted faster than diffusive homeostasis, the network exhibited a narrow rate distribution and a low responsiveness to input changes. Conversely, if diffusive homeostasis acted faster than non-diffusive homeostasis, the network exhibited broad firing rate distributions and linear responses to input changes. This is a plausible scenario, as NO modulation of ion channels through occurs in a timescale of 15 minutes [5], while other homeostatic processes which require transcriptional changes occur in a timescale of hours to days [28]. These results reflect what is observed as α is varied in the dynamic mean-field analysis, as local and global homeostatic mechanisms are simultaneously active for values in the range $0 < \alpha < 1$.

It is important to note that modelling HIP as a force acting on the threshold of an

integrate-and-fire neuron in order to achieve a target firing rate is a significant simplification. More physiologically realistic descriptions of homeostatic processes reveal the complex relationship between ion channel concentrations and the regulation of a wide range of neural activity [28]. Moreover, a number of previous studies have explored the effects of volume transmission on network dynamics, including its potential in implementing a winner-take-all function [29], the ability of a diffusive signal to reflect the average activity of a group of neurons [30], and the role of another diffusive neurotransmitter, $\text{TNF}\alpha$, in epileptogenesis [31]. Here, we add a functional interpretation by exploring its effects on neural heterogeneity and responsiveness within a network.

While NO is involved in a wide variety of neural processes throughout development and learning [32–34], these were ignored throughout for the sake of simplicity and tractability. Nonetheless, the impaired performance of nNOS knock-out mice in cognitive tasks [35] and the prevalence of epilepsy following nNOS inhibition [36] 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 [37]. Indeed, we conclude that it demonstrates 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

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

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

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 injected noise. $\eta(t)$ is an Ornstein-Uhlenbeck process with zero mean, unit variance, and correlation time $\tau_{OU} = 1$ ms [38]. 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

439

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 [39] and $\tau_{\text{Ca}^{2+}}$ and τ_{nNOS} are the timescales of Ca^{2+} decay and nNOS activation respectively.

440

441

442

443

444

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

$$\frac{d[\text{nNOS}]}{dt} = \frac{1}{\tau_{\text{nNOS}}} \frac{[\text{Ca}^{2+}]^n}{[\text{Ca}^{2+}]^n + K^n} \quad (4)$$

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 [40]. Neurons were represented by a point source according to their activated nNOS concentration. Periodic boundary conditions were used, as we assume we are simulating a subsection of a cortical network embedded in a larger cortical network with similar network activity.

445

446

447

448

449

450

451

$$\frac{d[\text{NO}]}{dt} - D\nabla^2[\text{NO}] = [\text{nNOS}] - \lambda[\text{NO}] \quad (5)$$

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 $[\text{NO}]$ and a target concentration $[\text{NO}]_0$;

452

453

454

$$\frac{d\theta_i}{dt} = \frac{1}{\tau_{\text{HIP}}} \frac{[\text{NO}] - [\text{NO}]_0}{[\text{NO}]}, \quad (6)$$

where τ_{HIP} is the timescale of homeostasis.

455

Non-diffusive homeostasis

456

For simplicity, the implementation of non-diffusive homeostasis is almost identical to that of diffusive homeostasis, in that the putative non-diffusive neuromodulator $[\text{NO}_{\text{non-diffusive}}]$ is synthesized through equations (3) and (4), and modulates firing thresholds through equation (6). The only difference is that the diffusion term in equation (5) is removed, so that by $[\text{NO}_{\text{non-diffusive}}]$ is entirely determined by the rate of intracellular synthesis and decay;

457

458

459

460

461

462

$$\frac{d[\text{NO}_{\text{non-diffusive}}]}{dt} = [\text{nNOS}] - \lambda[\text{NO}_{\text{non-diffusive}}] \quad (7)$$

Dynamic mean-field analysis

463

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:

464

465

466

467

468

469

470

$$\mu_i = JC\nu\tau, \quad \sigma_i = J\sqrt{C\nu\tau}, \quad (8)$$

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

$$\phi_i(\mu_i, \sigma_i) = \left[\sqrt{\pi} \tau_m \int_{\frac{\nu \tau - \mu_i}{\sigma_i}}^{\frac{\theta_i - \mu_i}{\sigma_i}} dv e^{v^2} \operatorname{erfc}(-v) \right]^{-1}, \quad (9)$$

where erfc is the complementary error function. Since the firing rate described by (9) 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:

$$\nu = \phi_i(\mu_i(\nu), \sigma_i(\nu)). \quad (10)$$

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 /d according to

$$\frac{d\theta_i}{dt} = \frac{1}{\tau_{\text{HIP}}} \left((1 - \alpha) \frac{\phi_i - \phi_0}{\phi_i} + \alpha \frac{\bar{\phi} - \phi_0}{\bar{\phi}} \right), \quad (11)$$

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 (8)-(9) until (10) is satisfied to a precision of 10^{-4} Hz, where (9) returns ϕ_i for each neuron in the population, and $\bar{\phi} = \frac{\sum \phi_i}{N}$ is used as ν in (8). At each timestep the thresholds θ_i of each neuron were modulated according to (11), and rates ϕ_i were subsequently recalculated from (9).

Procedure for Fig. 1.

External input rates μ_i for each neuron i were randomly drawn from a Gaussian distribution such that $\mu_i \sim \mathcal{N}(25, 10^2)$ Hz. Since the mean NO concentration takes time to reach a steady state in the recurrent network simulations, we ran the network for 100s 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 100s.

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 $\mu_i \sim \mathcal{N}(5.7, \delta^2)$, $\sigma_i = \sqrt{\mu_i}$. $\delta = 0.4$ is the width of the distribution of mean inputs to the population. Note that the parameter δ referred to here differs from the δ in (1)-(3).

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 in the non-diffusive network a different homeostatic target, $[\text{NO}_0]_i$. A network without any homeostasis is presented with input statistics ($\mu_i \sim \mathcal{N}(2, 5^2)$ Hz), tuned such that the firing rate distribution match that of the network with diffusive homeostasis. $[\text{NO}_0]_i$ for each neuron i can then be drawn from the distribution of steady-state intracellular concentrations of NO for this network. This results in a broad and heavy-tailed distribution of homeostatic targets, as opposed to the single homeostatic target which is used in networks with diffusive homeostasis and the unmodified non-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 Fig. 4.

External input for each neuron i was $\mu_i = 2$ Hz ($N = 2500$). NO_0 was set as described previously, although with an input rate of 2 Hz. 2 groups of 250 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 Figs. 4A-C were smoothed with a uniform time window of 20 s. Persistence of input differences were calculated by measuring the length of time it took for the signal-to-noise ratio between the two groups receiving elevated inputs to fall below 0. The signal-to-noise ratio is defined as $(\mu_1 - \mu_2)/(\sigma_1 + \sigma_2)$, where μ_i and σ_i correspond to the mean and standard deviation of the firings rate of group i .

Procedure for Fig. 5.

Figs. 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 Figs. 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 Fig. 5E were normalized 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 Fig. 6A shows the normalized $\mu_{j,t}$, while coloured lines show the normalized rate deviation of a randomly chosen group j from the mean population firing rate.

Decoding stimulus orientation. 554

Each excitatory neuron was randomly assigned a preferred orientation. After the 555
network reached a steady state, homeostasis was frozen. For each trial, each neuron i 556
with preferred orientation θ_i was stimulated with external Poisson input at a rate given 557
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 558
amplitude at θ_i of a Gaussian tuning curve centered around the stimulus orientation θ_s , 559
with a width given by $\sigma_s = 90^\circ$ and a peak amplitude of 2.5 Hz. The angle decoded 560
using the population vector method is the angle of the vector sum of all neural 561
responses. 562

Spike-timing-dependent plasticity. 563

A spike-timing dependent plasticity rule, as described in [19], was implemented in each 564
recurrent excitatory-excitatory synapse. Both potentiation and depression are additive 565
in this rule, with no weight dependence. For each pair of pre- and post-synaptic spikes 566
separated by a time Δt , the synaptic weight is updated by a value Δw given by 567

$$\Delta w = \begin{cases} A_+ \exp(\Delta t / \tau_+) g_{\max}, & \text{if } \Delta t < 0. \\ -A_- \exp(\Delta t / \tau_-) g_{\max}, & \text{if } \Delta t \geq 0. \end{cases} \quad (12)$$

Above, τ_+ and τ_- denote the timecourse over which potentiation and depression occur 568
respectively, while A_+ and A_- denote the relative strengths of potentiation and 569
depression. g_{\max} is the maximum synaptic weight. $A_+ < A_-$ such that irregular firing 570
is maintained within a reasonable range of rates. The external input in this case was 571
given by $\mu_i \sim \mathcal{N}(10, 10^2)$ Hz, and $J_{\text{ext}} = 40$ nS, $C = 250$. 572

Model parameters. 573

Unless explicitly defined, the parameters used throughout the paper are given in Table 574
1. For synthesis, diffusion, and decay of NO we have attempted to match data when 575
available [40, 41], although the dearth of experimental measurements does not permit for 576
great precision [42, 43]. Additionally, parameters were chosen such that the timescale of 577
homeostasis is separated from that of firing rate fluctuations. This is a reasonable 578
assumption, given that activity-dependent NO modulation likely acts within 10 minutes 579
or slower [5], although NO diffusion occurs in the order of 10 seconds. τ_{HIP} was chosen 580
to be long enough so as to avoid oscillations but short enough so as to allow feasible 581
large scale simulations. This is a common assumption in computational studies [10]. 582
Larger simulations, up to $N = 25000$, were run with no discernible difference in results. 583

All numerical simulations were implemented using the Brian simulator, v1.4.1 [44], 584
and the mean-field analysis was implemented using IPython Notebook [45]. Data 585
analysis was performed with the numpy Python package and plotting with the 586
matplotlib package and seaborn library [46, 47]. Code and IPython Notebooks which 587
perform the data analysis and plotting will be available on ModelDB and a public github 588
repository following peer review. In the meantime a minimal example of diffusion of a 589
neurotransmitter over a 2D surface is available at <http://tinyurl.com/sweeney-diffusion>, 590
which may easily be incorporated into existing Brian models. 591

Table 1. Simulation Parameters.

Neuron		NO synthesis and diffusion		Network		STDP	
E_l	-80 mV	$[Ca_{spike}^{2+}]$	1	N	5000	τ_+	20 ms
v_r	-60 mV	$\tau_{Ca^{2+}}$	10 ms	N_e	$0.8 N$	τ_-	20 ms
c_m	0.2 nF	τ_{nNOS}	100 ms	J_e	4.0 nS	A_+	0.025
τ_m	20 ms	n	3	J_i	64.0 nS	A_-	0.0275
τ_{ref}	5 ms	K	3	C	100	g_{max}	10 nS
θ_0	-50 mV	D	$1000 \mu m^2 s^{-1}$	J_{ext}	80.0 nS		
τ_e	3 ms	λ	$0.1 s^{-1}$	τ_{OU}	1 ms		
τ_i	7 ms	Grid size	$1 mm^2$	σ_{OU}	1 mV		
E_e	0 mV	ds	$2 \mu m$	Dynamic mean-field			
E_i	-70 mV	dt (diffusion)	1 ms	θ_0		10	
dt	0.1 ms	τ_{HIP}	2500 ms	v_r		0	

Supporting Information

Procedure for Fig. S1.

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

$$P_c(d) = \epsilon e^{-\frac{d^2}{2s^2}}, \quad (13)$$

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 Fig. 1 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.

Acknowledgments

We thank Mark van Rossum, David Sterratt, Jochen Treisch 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. Concepts of neural nitric oxide-mediated transmission. *Eur J Neurosci.* 2008;27(11):2783–2802. Available from: <http://dx.doi.org/10.1111/j.1460-9568.2008.06285.x>.
2. Pape HC, Mager R. Nitric oxide controls oscillatory activity in thalamocortical neurons. *Neuron.* 1992;9(3):441–448.
3. Gally J, Montague P, Reeke G, Edelman G. 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.* 1990;87(9):3547–3551. Available from: <http://dx.doi.org/10.1073/pnas.87.9.3547>.

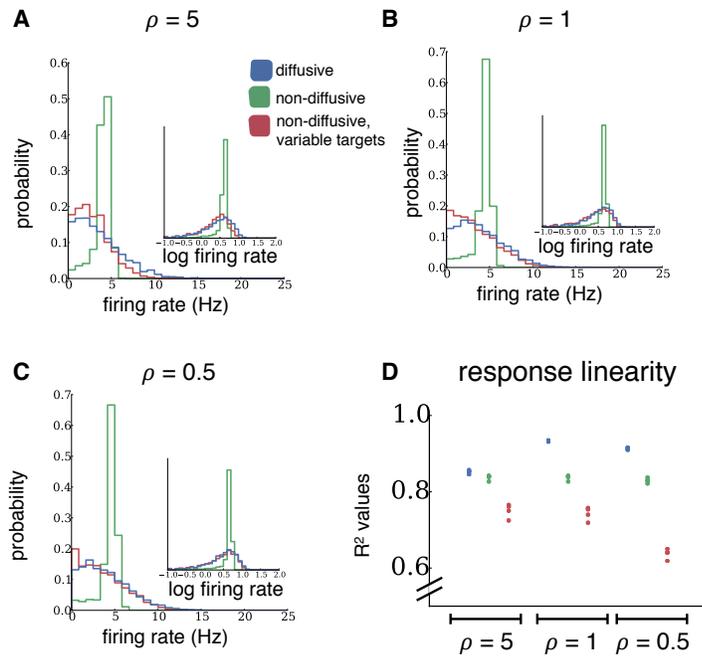


Figure S1. Linearity of network responses in networks with spatially restricted connection probabilities. These results qualitatively agree with those for the random networks used throughout the study. The simulations shown here were identical to those in Figs. 2D and 5 of the main text, but the connection 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 ($\rho=0.5, 1.0, 5.0$). (D) Each point represents the R^2 value of a linear fit as in Figs. 5A-C, for one network.

4. Steinert J, Cornelia K, Baker C, Challiss R, Mistry R, Haustein M, et al. Nitric oxide is a volume transmitter regulating postsynaptic excitability at a glutamatergic synapse. *Neuron*. 2008;60(4):642–656. Available from: <http://dx.doi.org/10.1016/j.neuron.2008.08.025>.
5. Steinert J, Robinson S, Tong H, Haustein M, Cornelia K, Forsythe I. Nitric oxide is an activity-dependent regulator of target neuron intrinsic excitability. *Neuron*. 2011;71(2):291–305. Available from: <http://dx.doi.org/10.1016/j.neuron.2011.05.037>.
6. Lourenço CF, Santos RM, Barbosa RM, Cadenas E, Radi R, Laranjinha J. Neurovascular coupling in hippocampus is mediated via diffusion by neuronal-derived nitric oxide. *Free Radic Biol Med*. 2014;73(0):421–429.
7. LeMasson G, Marder E, Abbott L. Activity-dependent regulation of conductances in model neurons. *Science*. 1993;259(5103):1915–1917.
8. Lazar A, Pipa G, Triesch J. SORN: a self-organizing recurrent neural network. *Front Comput Neurosci*. 2009;3:23. Available from: <http://dx.doi.org/10.3389/neuro.10.023.2009>.
9. Olypher A, Prinz A. Geometry and dynamics of activity-dependent homeostatic regulation in neurons. *J Comput Neurosci*. 2010;28(3):361–374. Available from: <http://dx.doi.org/10.1007/s10827-010-0213-z>.

10. Naudé J, Cessac B, Berry H, Delord B. Effects of cellular homeostatic intrinsic plasticity on dynamical and computational properties of biological recurrent neural networks. *J Neurosci*. 2013;33(38):15032–15043. Available from: <http://dx.doi.org/10.1523/JNEUROSCI.0870-13.2013>.
11. Wohrer A, Humphries M, Machens C. Population-wide distributions of neural activity during perceptual decision-making. *Prog Neurobiol*. 2013;103:156–193. Available from: <http://dx.doi.org/10.1093/cercor/bhg100>.
12. Marsat G, Maler L. Neural heterogeneity and efficient population codes for communication signals. *J Neurophysiol*. 2010;104(5):2543–2555.
13. Tripathy SJ, Padmanabhan K, Gerkin RC, Urban NN. Intermediate intrinsic diversity enhances neural population coding. *P Natl Acad Sci USA*. 2013 May;110(20):8248–53.
14. Zenke F, Hennequin G, Gerstner W. Synaptic plasticity in neural networks needs homeostasis with a fast rate detector. *PLoS Comput Biol*. 2013 Nov;9(11):e1003330.
15. O’Leary T, Williams AH, Caplan JS, Marder E. Correlations in ion channel expression emerge from homeostatic tuning rules. *P Natl Acad Sci USA*. 2013;110(28):E2645–E2654.
16. Brunel N. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J Comput Neurosci*. 2000;8(3):183–208. Available from: [http://dx.doi.org/10.1016/S0925-2312\(00\)00179-X](http://dx.doi.org/10.1016/S0925-2312(00)00179-X).
17. Roxin A, Brunel N, Hansel D, Mongillo G, van Vreeswijk C. On the distribution of firing rates in networks of cortical neurons. *J Neurosci*. 2011;31(45):16217–16226. Available from: <http://dx.doi.org/10.1523/JNEUROSCI.1677-11.2011>.
18. Holmgren C, Harkany T, Svennenfors B, Zilberter Y. Pyramidal cell communication within local networks in layer 2/3 of rat neocortex. *J Physiol*. 2003;551(1):139–153.
19. Song S, Miller KD, Abbott LF. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat Neurosci*. 2000;3(9):919–926.
20. Vogels TP, Abbott LF. Signal propagation and logic gating in networks of integrate-and-fire neurons. *The Journal of neuroscience*. 2005;25(46):10786–10795.
21. Van Rossum MC, Bi GQ, Turrigiano GG. Stable Hebbian learning from spike timing-dependent plasticity. *J Neurosci*. 2000;20(23):8812–8821.
22. Ermentrout B. Linearization of FI curves by adaptation. *Neural Comput*. 1998;10(7):1721–1729.
23. Mizuseki K, Buzsáki G. Preconfigured, Skewed Distribution of Firing Rates in the Hippocampus and Entorhinal Cortex. *Cell Reports*. 2013;4(5):1010–1021. Available from: <http://dx.doi.org/10.1016/j.celrep.2013.07.039>.
24. Hirase H, Leinekugel X, Czurkó A, Csicsvari J, Buzsáki G. Firing rates of hippocampal neurons are preserved during subsequent sleep episodes and modified by novel awake experience. *P Natl Acad Sci USA*. 2001;98(16):9386–9390.

25. Slomowitz E, Styr B, Vertkin I, Milshtein-Parush H, Nelken I, Slutsky M, et al. Interplay between population firing stability and single neuron dynamics in hippocampal networks. *eLife*. 2015;4.
26. Shamir M, Sompolinsky H. Implications of neuronal diversity on population coding. *Neural Comput*. 2006;18(8):1951–1986.
27. Seung HS, Sompolinsky H. Simple models for reading neuronal population codes. *P Natl Acad Sci USA*. 1993;90(22):10749–10753.
28. O’Leary T, Williams AH, Franci A, Marder E. Cell types, network homeostasis, and pathological compensation from a biologically plausible ion channel expression model. *Neuron*. 2014;82(4):809–821.
29. Kohonen T. Physiological interpretation of the self-organizing map algorithm. *Neural Networks*. 1993;6(6):895–905.
30. Ott SR, Philippides A, Elphick MR, O’Shea M. Enhanced fidelity of diffusive nitric oxide signalling by the spatial segregation of source and target neurones in the memory centre of an insect brain. *European Journal of Neuroscience*. 2007;25(1):181–190.
31. Savin C, Triesch J, Meyer-Hermann M. Epileptogenesis due to glia-mediated synaptic scaling. *Journal of The Royal Society Interface*. 2009;6(37):655–668.
32. Nikonenko I, Nikonenko A, Mendez P, Michurina T, Enikolopov G, Muller D. Nitric oxide mediates local activity-dependent excitatory synapse development. *P Natl Acad Sci USA*. 2013;110(44):E4142–E4151. Available from: <http://dx.doi.org/10.1073/pnas.1311927110>.
33. Tamagnini F, Barker G, Warburton C, Burattini C, Aicardi G, Bashir Z. Nitric oxide-dependent LTD but not endocannabinoid-LTP is crucial for visual recognition memory. *J Physiol*. 2013;591(16):3963–3979. Available from: <http://dx.doi.org/10.1113/jphysiol.2013.254862>.
34. Son H, Hawkins RD, Martin K, Kiebler M, Huang PL, Fishman MC, et al. Long-term potentiation is reduced in mice that are doubly mutant in endothelial and neuronal nitric oxide synthase. *Cell*. 1996;87(6):1015–23.
35. Weitzdoerfer R, Hoeger H, Engidawork E, Engelmann M, Singewald N, Lubec G, et al. Neuronal nitric oxide synthase knock-out mice show impaired cognitive performance. *Nitric Oxide*. 2004;10(3):130–140. Available from: <http://dx.doi.org/10.1016/j.niox.2004.03.007>.
36. Del-Bel E, Oliveira P, Oliveira J, Mishra P, Jobe P, Garcia-Cairasco N. Anticonvulsant and proconvulsant roles of nitric oxide in experimental epilepsy models. *Braz J Med Biol Res*. 1997;30(8):971–979.
37. Wang R. Two’s company, three’s a crowd: can H₂S be the third endogenous gaseous transmitter? *FASEB J*. 2002;16(13):1792–1798. Available from: <http://dx.doi.org/10.1096/fj.02-0211hyp>.
38. Uhlenbeck GE, Ornstein LS. On the theory of the Brownian motion. *Phys Rev*. 1930;36(5):823.
39. Hill AV. The possible effects of the aggregation of the molecules of haemoglobin on its dissociation curves. *J Physiol*. 1910;40:4–7.

40. Philippides A, Husbands P, O'Shea M. Four-dimensional neuronal signaling by nitric oxide: a computational analysis. *J Neurosci*. 2000;20(3):1199–1207.
41. Salerno JC, Ghosh DK. Space, time and nitric oxide–neuronal nitric oxide synthase generates signal pulses. *FEBS journal*. 2009;276(22):6677–6688.
42. Hall C, Garthwaite J. What is the real physiological NO concentration in vivo? *Nitric Oxide*. 2009;21(2):92–9103. Available from: <http://dx.doi.org/10.1016/j.niox.2009.07.002>.
43. Batchelor AM, Bartus K, Reynell C, Constantinou S, Halvey EJ, Held KF, et al. Exquisite sensitivity to subsecond, picomolar nitric oxide transients conferred on cells by guanylyl cyclase-coupled receptors. *P Natl Acad Sci USA*. 2010 Dec;107(51):22060–5.
44. Goodman D, Brette R. Brian: a simulator for spiking neural networks in python. *Front Neuroinform*. 2008;2:192–7. Available from: <http://dx.doi.org/10.3389/neuro.11.005.2008>.
45. Pérez F, Granger BE. IPython: a System for Interactive Scientific Computing. *Comput Sci Eng*. 2007 May;9(3):21–29. Available from: <http://ipython.org>.
46. Van Der Walt S, Colbert SC, Varoquaux G. The NumPy array: a structure for efficient numerical computation. *Comput Sci Eng*. 2011;13(2):22–30.
47. Hunter JD. Matplotlib: A 2D graphics environment. *Comput Sci Eng*. 2007;9(3):0090–95.