EJN European Journal of Neuroscience

FENS Federation of European Neuroscience Societies

European Journal of Neuroscience, Vol. 45, pp. 1057–1067, 2017 doi:10.1111/ejn.13279

COMPUTATIONAL NEUROSCIENCE SPECIAL SECTION

Emergent spatial synaptic structure from diffusive plasticity

Yann Sweeney and Claudia Clopath
Department of Bioengineering, Imperial College London, South Kensington Campus, London SW7 2AZ, UK

Keywords: diffusive neurotransmitters, neural networks, synaptic connectivity, synaptic plasticity, volume transmission

Edited by Panayiota Poirazi
Received 11 February 2016, revised 4 May 2016, accepted 13 May 2016

Abstract

Some neurotransmitters can diffuse freely across cell membranes, influencing neighbouring neurons regardless of their synaptic coupling. This provides a means of neural communication, alternative to synaptic transmission, which can influence the way in which neural networks process information. Here, we ask whether diffusive neurotransmission can also influence the structure of synaptic connectivity in a network undergoing plasticity. We propose a form of Hebbian synaptic plasticity which is mediated by a diffusive neurotransmitter. Whenever a synapse is modified at an individual neuron through our proposed mechanism, similar but smaller modifications occur in synapses connecting to neighbouring neurons. The effects of this diffusive plasticity are explored in networks of rate-based neurons. This leads to the emergence of spatial structure in the synaptic connectivity of the network. We show that this spatial structure can coexist with other forms of structure in the synaptic connectivity, such as with groups of strongly interconnected neurons that form in response to correlated external drive. Finally, we explore diffusive plasticity in a simple feedforward network model of receptive field development. We show that, as widely observed across sensory cortex, the preferred stimulus identity of neurons in our network become spatially correlated due to diffusion. Our proposed mechanism of diffusive plasticity provides an efficient mechanism for generating these spatial correlations in stimulus preference which can flexibly interact with other forms of synaptic organisation.

Introduction

A common assumption when constructing a theory of neural plasticity is that neurons require some form of synaptic connection, electrical or chemical, to communicate. There exists, however, a number of neurotransmitters whose gaseous properties allow them to diffuse freely and rapidly across cell membranes (Dawson & Snyder, 1994; Kiss & Vizi, 2001; Boehning & Snyder, 2003; Garthwaite, 2008). Experiments characterising this effect have demonstrated that neurons which synthesise nitric oxide (NO), a diffusive neurotransmitter, can alter the excitability of other neurons up to away (Steinert *et al.*, 2008; Artinian *et al.*, 2010). Diffusive neurotransmitters therefore provide an alternative means of communication among neighbouring neurons which is independent of synaptic connectivity, a phenomenon often called *volume transmission*. In a previous theoretical study we explored the consequences of this effect on network activity, finding that homeostatic plasticity mediated by a diffusive signal enables the maintenance of diverse and flexible neural responses within a network (Sweeney *et al.*, 2015). Here, we propose a new form of synaptic plasticity in which changes in the synaptic weight are mediated by such a diffusive neurotransmitter. This is achieved by modifying the Bienenstock–Cooper–Munro (BCM) learning rule (Bienenstock *et al.*, 1982) so that the variable representing the activity of the postsynaptic neuron is replaced by a variable representing the combined activity of the postsynaptic neuron and its neighbouring neurons. We name this modified learning rule *diffusive BCM* (dBCM). We explore the consequences of dBCM on synaptic connectivity following plasticity in simple rate-based network models with feedforward and recurrent architectures, finding that the resulting synaptic structure differs from that in networks with an unmodified BCM learning rule.

First, we observe emergent spatial synaptic structure in response to uncorrelated random inputs in networks with dBCM, unlike in networks with BCM. This structure reflects the underlying spatial structure within the network, with neurons that are close together more likely to be strongly connected than neurons which are further away. Second, we find that presenting groups of correlated inputs to the network leads to the development of strongly interconnected assemblies of neurons in networks with both dBCM and BCM, but that this structure is in addition to the underlying spatial structure in the synaptic connectivity of networks with dBCM. Finally, in a minimal model of receptive field development, we find that dBCM leads to spatial organisation in the input which output neurons become selective to, with neighbouring neurons more likely to share stimulus preference than in networks with BCM. These spatial

Correspondence: Claudia Clopath, as above.
E-mail: c.clopath@imperial.ac.uk

correlations are similar to those observed in sensory brain areas, such as in orientation preference and ocular dominance maps in visual cortex (Kaschube, 2014), or tonotopic maps in auditory cortex (Reale & Imig, 1980).

Methods

Neuron model

We use a simple firing rate neuron model, given by the transfer function $g(y)$ defined below, and as used previously by Rajan *et al.* (2010; Hennequin *et al.*, 2014).

$$g(y) = \begin{cases} r_0 \tanh[y/r_0] & \text{if } y < 0 \\ (r_{\max} - r_0)\tanh[y/(r_{\max} - r_0)] & \text{if } y \geq 0 \end{cases} \quad (1)$$

This leads to firing rates with a baseline of r_0 and a maximum of $r_{\max}$. Following Rajan *et al.* (2010), the firing rates y_i of neuron i driven by external input H_i in a network are described below.

$$\frac{dy_i}{dt} = -y_i + \sum_{j=i}^{N} W_{ji} g(y_j) + H_i \quad (2)$$

Representing neuron position in space

To test the effects of a diffusive mediator of synaptic plasticity, we introduce a description of physical space in our network model. A straightforward approach is to represent neuron positions along a single spatial dimension, i.e. along a line. By assuming that neurons are point sources along this line, the position of neuron i can be denoted by a single variable, p_i, where positions are bounded such that $0 < p_i < 1$.

To avoid discontinuous boundary effects we use periodic boundary conditions when calculating distance, so that a neuron located at 0.91 is a distance of 0.1 units from a neuron located at 0.01. The distance between two neurons with positions p_i and p_j is therefore given by $\min(|p_i - p_j|, |p_i + 1.0 - p_j|, |p_j + 1.0 - p_i|)$. As we only simulate synaptic plasticity in synapses between excitatory neurons, the spatial location of inhibitory neurons in the recurrent network can be ignored, and so only excitatory neurons are assigned spatial coordinates.

Modelling synaptic plasticity

We use the BCM learning rule to model synaptic plasticity in recurrent excitatory to excitatory (E-E) synapses (Bienenstock *et al.*, 1982; Blais & Cooper, 2008). I-E and E-I weights remain fixed throughout the simulation.

$$\frac{dW_{ij}^{EE}}{dt} = \alpha y_i y_j (y_j - \theta_j) \quad (3)$$

$$\frac{d\theta_i}{dt} = \frac{1}{\tau_\theta}\frac{{y_i}^2}{y_0} \quad (4)$$

The BCM learning rule has both a Hebbian component and a homeostatic component. This form of plasticity is competitive and leads to the development of stimulus selectivity, as discussed in Bienenstock *et al.* (1982). To prevent extremely strong synapses, weights are bounded so that their values lie between 0 and $w_{\max}$.

Investigating the diffusive range of NO

We use a mathematical description of NO diffusion to investigate the conditions under which a diffusive neurotransmitter can achieve significant concentrations over a volume large enough to affect the synapses of neighbouring neurons. This is based on previous work by Philippides (2001; Philippides *et al.*, 2005), later extended by Bellefontaine *et al.* (2014; Garthwaite, 2015). Given multiple spherical sources of NO, the concentration at coordinates (x, y, z) in 3D space after a time t can be written as

$$[\mathrm{NO}](x,y,z,t) = \frac{1}{10^3 N_{\mathrm{av}}} \sum_{a=0}^{n-1}\sum_{b=0}^{n-1}\sum_{c=0}^{n-1} \left[\frac{Q}{8\pi \mathrm{R} r' \sqrt{\pi \mathrm{D}}} \int_0^t \frac{1}{\sqrt{t-t'}} \right. \quad (5)$$
$$\left. \exp\left[\frac{-(R-r')^2}{4D(t-t')} - \lambda(t-t')\right] - \exp\left[\frac{-(R+r')^2}{4D(t-t')} - \lambda(t-t')\right]\right] dt'$$

where $R' = \sqrt{(x-x_a)^2 + (y-y_b)^2 + (z-z_c)^2}$. Above, N_{av} is Avagadro's number (molecules/Mole), Q is the rate of NO synthesis at a single source (molecules/s), D is the diffusion coefficient of NO (m^2/s), λ is the inactivation rate of NO (1/s), r' is the radius of the spherical source (m) and x_a, y_b, z_c are the spatial coordinates of the sources. As in Bellefontaine *et al.* (2014; Garthwaite, 2015), we use $D = 8.48 \times 10^{-10}$ m^2/s and $r' = 0.2$ μm, $Q = 40$ molecules/s for synaptic sources.

Experimental estimates of excitatory synapse density range from 0.25 μm^{-3} in layer 5 motor cortex, 0.72 μm^{-3} in mouse neocortex to 1.7 μm^{-3} in dentate gyrus (Schüz & Palm, 1989; Monfils *et al.*, 2005). In addition, a quantitative analysis observed nNOS expression in 8% of CA1 pyramidal neuron synapses (Burette *et al.*, 2002). Combining these gives us an estimate of ~ 0.1 μm^{-3} for the density of potential synaptic NO sources. We can investigate the effects of diffusion within a cube of length 50 μm and volume 1.25×10^5 μm^3 containing either densely activated synaptic sources, where 1% of potential sources are active (corresponding to 125 sources within the volume), or sparsely activated sources, where 0.2% of potential sources are active (corresponding to 25 sources within the volume). Assuming a neural density of roughly 9.2×10^4 mm^{-3}, this volume would be occupied by $\sim$ 10 neurons (Schüz & Palm, 1989; Rubinov *et al.*, 2015). Active synaptic sources are randomly positioned within this volume by generating their coordinates x_a, y_b, z_c from a uniform distribution. The steady-state concentration of the points on a 2D slice of this volume (at $z = 25$ μm) can by computed by solving Eqn 5 at time $t = 1$ s, summing over all active synaptic sources in the volume.

Simulating the effect of a diffusive plasticity mediator

To simulate the effect of a diffusive mediator of synaptic plasticity in our learning rule, we introduce a mechanism by which the activity of neighbouring postsynaptic neurons influence the synaptic weight changes that occurs at a particular synapse. This is achieved by replacing the postsynaptic firing rate y_j in Eqn 3 with a spatially averaged postsynaptic firing rate, denoted $\tilde{y}_j$. This alteration results in a learning rule described by the equations below and Eqn 4, which remains unchanged.

$$\frac{dW_{ij}}{dt} = \alpha y_i y_j (\tilde{y}_j - \theta_j) \quad (6)$$

$$\frac{d\tilde{y}_j}{dt} = \frac{1}{\tau_{\tilde{y}}}\left(-\tilde{y}_j + \frac{\sum_{i=1}^{N_E} D_{ij} y_i}{\sum_{i=1}^{N_E} D_{ij}}\right) \quad (7)$$

D_{ij} denotes the diffusive range between neuron i and neuron j. We assume that diffusion has a Gaussian spatial profile, such that the diffusive range between two neurons i and j with positions p_i and p_j is given by

$$D_{ij} = \frac{1}{\sigma_D\sqrt{2\pi}} e^{\frac{-p_i - p_j}{2\sigma_D{}^2}}. \quad (8)$$

The standard deviation σ_D is the width of the spatial profile, and can be thought of as the spatial range of the diffusive neurotransmitter. As the contribution of an individual neuron i to its own spatially averaged activity $\tilde{y}_l$ should be significantly larger than that of neighbouring neurons, we set the diffusive range such that $D_{ii} = 2$. This is in contrast with neighbouring neurons, where their contribution to $\tilde{y}_l$ is always <1. The diffusive range matrix, D, is shown in Fig. 3C for a network in which neurons are positioned between 0 and 1 with regular spacing, with their indices sorted according to their position. The neuron with index 1 therefore has a position $p_1 = 0.0$, and the neuron with index 40 has a position $p_{40} = 1 - \frac{1}{40} = 0.975$. Figure 3D shows D for a network in which neurons are randomly and uniformly positioned between 0 and 1, after which their indices are sorted according to their position.

Recurrent network model

The dynamics of both inhibitory (I) and excitatory (E) neurons are described by Equs 1 and 2. There is dense all-to-all synaptic connectivity in the E-E, E-I and I-E populations, and no I-I connectivity. Self-connections, or autapses, are not permitted in this network. As such, W_{ij} in Eqn 2 takes the form of a square matrix with size $(N_E + N_I)^2$, where N_E and N_I represent the number of excitatory and inhibitory neurons respectively. The strength of the inhibitory synapses are uniform, and set so that the inhibitory current roughly balances excitatory currents in the network.

Random inputs

Every 500 ms, a randomly chosen group of neurons receive a large external input with a firing rate of H_{max}, whereas the remaining neurons in the network receive external input with a lower firing rate of H_0. These firing rates are kept constant for the next 500 ms. The number of neurons in this random group is also random, drawn from a uniform distribution between 10 and 20.

Correlated input groups

As neurons are spatially ordered by index, we can introduce spatially structured groups of correlated external input by assigning four groups of 10 identical external inputs each to neurons with indices 1–10, 11–20, 21–30 and 31–40. For correlated inputs in which there is no spatial structure, four groups of 10 external inputs are randomly assigned to the 40 excitatory neurons, with no dependence on their indices and, by extension, spatial positions.

Every 500 ms, a randomly chosen group is activated, such that the 10 neurons within this group receive input at a firing rate of H_{max}, whereas the remaining neurons in the network receive external input with a lower firing rate of H_0. These firing rates are kept constant for the next 500 ms.

Feedforward network model

This feedforward network consists of excitatory neurons only. The network is composed of a input layer of five neurons and an output layer of 50 neurons. There is all-to-all synaptic connectivity between neurons in the input and output layers, and no lateral connectivity between neurons in either layer. Each neuron i in the input layer receives an external excitatory drive H_i.

Every 500 ms, a randomly chosen neuron k in the input layer receives a large external drive by setting $H_k = H_{max}$, whereas the external drive to the rest of the input neuron remains at H_0. These firing rates are kept constant for the next 500 ms.

Simulating synaptic plasticity

In all our networks simulations, we initialise the excitatory synaptic weight matrices to uniform values of 0.4 w_{max}. External input is then presented to the network as described above, with either a BCM or dBCM learning rule mediating synaptic plasticity. To ensure that synaptic weights are sufficiently stable, we run each network instantiation for 10^5 s (i.e. 2×10^5 input presentations) before analysing the final synaptic weight matrix. Simulation parameters are given in Table 1.

Network simulations and data analysis was performed with the numpy PYTHON package and plotting with the MATPLOTLIB package and IPYTHON notebook (Hunter, 2007; Perez & Granger, 2007; van der Walt *et al.*, 2011). Simulation code and IPYTHON Notebooks which perform the data analysis and plotting are available on ModelDB (http://modeldb.yale.edu) and at https://github.com/yannaodh/sweeney-2016.

Results

We explore the effect of Hebbian synaptic plasticity which is mediated by a diffusive neurotransmitter. This is implemented as a BCM learning rule in which the variable representing postsynaptic activity is modified so that it is also influenced by the activity of neighbouring neurons. This modification can be represented as follows:

$$\text{weight change} = \text{pre} \times \text{post} \times (\text{spatially averaged post} - \text{threshold}).$$

See Eqns (6) and (7) in Methods for further details. As such, whenever a synapse is potentiated or depressed, a similar but

TABLE 1. Simulation parameters

N_E	40	N_I	10	r_0	1.0 Hz	r_{max}	20.0 Hz
dt	0.05 ms	α	5×10^{-12} Hz	τ_θ	4×10^4 ms	$\tau_{\bar{y}}$	100 ms
w_{max}	0.06	H_0	2.5 Hz	H_{max}	10 Hz	y_0	5 Hz
σ_D	0.25	$w_{inhibitory}$	−0.03				
Feedforward network							
N_{input}	5	N_{output}	50	α	2.5×10^{-12} Hz	w_{max}	1.0

smaller synaptic weight change will occur in synapses connecting to neurons which neighbour the original postsynaptic neuron. This effect is illustrated in Fig. 1. The threshold term above refers to the sliding threshold that determines whether synaptic potentiation or synaptic depression occurs, and which depends on recent postsynaptic activity.

This proposed form of synaptic plasticity, which we call dBCM, is explored in a simple network model of rate-based neurons. Networks with both feedforward and recurrent architectures are investigated, as illustrated in Figs 3A and 5A.

Diffusive range of multiple synaptic sources

We first investigate the conditions necessary for the transmission of a diffusive signal over a range large enough to affect the synapses of neighbouring neurons. We estimate the density of potential excitatory synaptic sources of NO, a neurotransmitter whose diffusive properties have been widely characterised, as $\sim$ 0.1 μm^{-3} (Kiss & Vizi, 2001; Burette *et al.*, 2002; Garthwaite, 2015). From this we calculate the steady-state concentrations of NO for a 2D slice within a cubic volume of radius 50 μm, estimated to contain $\sim$ 10 neurons (Fig. 2, see Methods) (Schüz & Palm, 1989; Rubinov *et al.*, 2015). As the inactivation rate of NO *in vivo* is not well constrained by experiment, we investigate diffusion with both a low inactivation rate (λ = 5 s^{-1}, corresponding to a half-life of 139 ms), and a high inactivation rate (λ = 138 s^{-1}, corresponding to a half-life of 5 ms).

Figure 2A,B show the steady-state concentrations for λ = 5 s^{-1}, with either 0.2% (sparse) or 1.0% (dense) of potential sources actively synthesising NO. Figure 2C shows the concentration of NO for points on the slice shown in Fig. 2A,B, depending on the distance from each point to its nearest active synaptic source. For sparsely activated synaptic sources, concentrations are around 50 pM within 5 μm of a synaptic source, falling to around 10 pM for points over 10 μm away from a synaptic source. For densely activated synaptic sources, concentrations decrease from 150 pM to around 50 pM over 10 μm away (note that there are no points more than 12 μm from a synaptic source because of the higher density of these sources). These concentrations lie within the range for which signal transduction of NO occurs, which can be as low as 10 pM (Batchelor *et al.*, 2010).

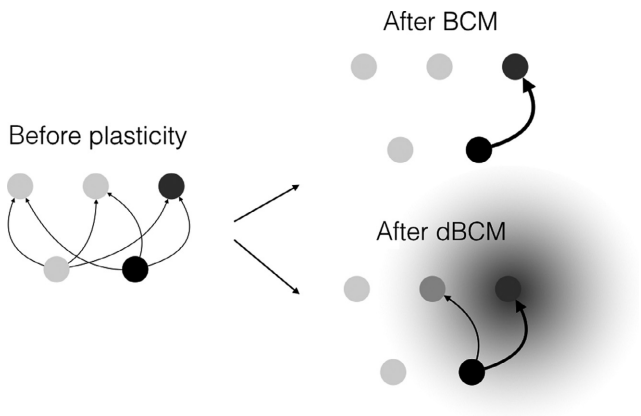

FIG. 1. Illustration of the effect of synaptic plasticity mediated by a diffusive neurotransmitter (dBCM, bottom) in comparison with synaptic plasticity determined by a standard Bienenstock–Cooper–Munro (BCM) learning rule (top). The grey cloud surrounding the active postsynaptic neuron represents the spatial range of the diffusive neurotransmitter which mediates synaptic plasticity in the dBCM rule.

Figure 2D,F shows the same as Fig. 2A–C, but for λ = 138 s^{-1}. The steady-state concentrations are much lower in this case, with significant concentrations only occurring <5 μm from synaptic sources. Increasing the density of synaptic sources does not have a large effect on the effective diffusive range for λ = 138 s^{-1}. A low inactivation rate is important to maintain significant concentrations, whereas the density of synaptic sources of NO will also increase the diffusive range. Although the diffusive range would not extend over the entire dendritic arbour of neocortical neurons, these results demonstrate that the volume transmission of NO is possible over distances such that a portion of synapses at neighbouring neurons, though not all, may be affected. This is especially plausible when we consider the high densities of axonal-dendritic contacts observed in ultrastructural EM reconstructions (Mishchenko *et al.*, 2010). To remain agnostic about the properties of the diffusive neurotransmitter we are modelling, we will explore the effects of dBCM over a broad range of effective diffusive ranges, assuming that the concentration decreases with distance from the synaptic source, as observed in Fig. 2C,F.

dBCM leads to spatial structure in the synaptic weights

To test whether dBCM can form spatial structure in the synaptic weights, we simulate synaptic plasticity in a recurrent network model. The network is composed of an excitatory neuron population with dense, all-to-all connectivity, and an inhibitory neuron population with no recurrent connections. The excitatory and inhibitory neurons are reciprocally connected in an all-to-all manner. Only excitatory-excitatory synapses undergo synaptic plasticity, whereas the weights of other synapse types remain fixed.

We test the effects of plasticity under both BCM and dBCM by presenting each excitatory neuron in the network with random, uncorrelated external inputs, as described in Methods. In networks with BCM, without a diffusive neurotransmitter and where synaptic updates are purely local, the synaptic weight matrix after a long period of synaptic plasticity exhibits no apparent spatial structure (Fig. 3E). Synaptic weight gradually increase or decrease from their initial uniform value, and tend towards either their minimum or maximum values, where they remain relatively stable (Fig. 3B, grey traces).

Similar synaptic weight dynamics are observed in networks with dBCM (Fig. 3B, black traces), although differences emerge in the spatial structure of the synaptic weight matrix after plasticity (Fig. 3F). Neurons are much more likely to have strong synaptic connections with neighbouring neurons, and to have very weak connections with neurons that are further away. This is a direct consequence of the effect of diffusion, as illustrated in Fig. 1, and the competitive nature of BCM, which ensures that a small proportion of incoming synapses reach their maximum weight, whereas the remaining synapses tend towards their minimum weight.

The resulting spatial structure in the synaptic weight matrix reflects the underlying spatial structure described by the pairwise diffusive range between neurons. This is demonstrated by the blocks of strongly interconnected neurons which emerge following plasticity in Fig. 3G. These blocks reflect the blocks in the diffusive range matrix which emerge due to the randomly generated positions of excitatory neurons in Fig. 3D. Indeed, the competitive nature of the learning rule enhances the spatial structure, as synapses between closely neighbouring neurons approach the maximum weight w_{max}, whereas the remaining synapses approach the minimum synaptic weight of 0. Simulations of multiple networks confirmed this effect, by measuring the mean synaptic weight difference between a

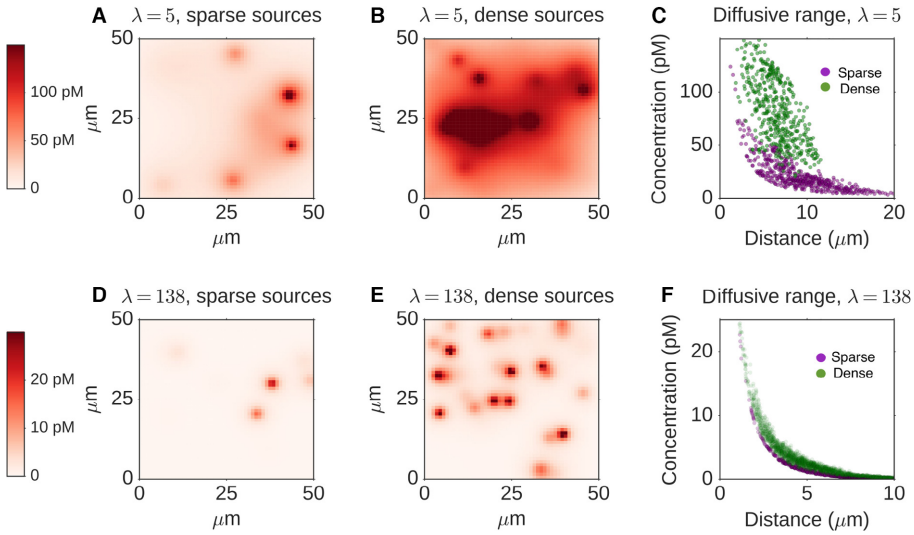

FIG. 2. Diffusive range of multiple synaptic sources within a volume. (A) Steady-state concentration of nitric oxide (NO) for a 2D slice of a cube with length 50 μm, containing 25 synaptic sources of NO. The inactivation rate of NO is set at $\lambda = 5\ s^{-1}$. (B) Same as A, but for 125 synaptic sources of NO. (C) Steady-state concentration of NO for points on the slice shown in Fig. 2A, B, depending on the distance from each point to its nearest active synaptic source in the volume. (D–F), same as A–C, but for $\lambda = 138\ s^{-1}$.

neurons five nearest neighbours and five furthest neighbours. For networks with BCM, there was a mean weight difference of -0.003 ± 0.006, whereas for networks with dBCM there was a significant mean weight difference of 0.04 ± 0.01 ($P = 0.01$, unpaired t-test, $\pm$ represents the standard deviation).

The spatial structure in the synaptic weights which emerge due to dBCM also leads to spatial correlations between the firing rates of neurons in the network. This is shown in Fig. 3I, where pairwise correlations in the firing rates between neurons in response to random inputs also reflect the spatial structure of the diffusion range matrix. Neurons that are located close together exhibit highly correlated firing rates, due to their strong synaptic coupling. These simulations demonstrate that plasticity which is mediated by a diffusive neurotransmitter can lead to spatial structure in the synaptic connectivity and activity of recurrent networks.

Having demonstrated that dBCM leads to spatial structure in the synaptic weight matrix, and as Fig. 2 demonstrated how the spatial extent of diffusion varies with properties of the diffusive neurotransmitter, we next wondered whether the resulting structure depends on the width of the diffusive range, σ_D.

To test this, we decrease the width to $\sigma_D = 0.025$. Figure 3H shows the final synaptic weight matrix after this change. The number of neighbouring neurons that are highly interconnected is reduced, and strong synapses which do not follow a particular spatial structure begin to emerge, similar to Fig. 3E. The impact of decreasing σ_D is explored further in Fig. 3J, which shows the relationship between the firing rate correlations of pairs of neurons and their distance apart, for multiple values of the diffusive range. As the diffusive range is reduced, the dependence of firing rate correlations on distance becomes weaker. The resulting spatial structure does not depend significantly on the rate at which the spatial activity of neighbouring neurons is averaged. Qualitatively similar results emerge if the time constant $\tau_{\bar{y}}$ determining the spatial average of the postsynaptic firing rate is slower, or indeed if instantaneous averaging is simulated, where $\tau_{\bar{y}} \to 0$ (data not shown). These results show that dBCM leads to spatial structure in the synaptic weights, and that this structure depends on the diffusive range.

A dBCM learning rule develops specific connectivity biased by spatial organisation

As we have shown differences in the manner which synaptic weight changes occur between BCM and dBCM, we now ask whether the development of specific connectivity, which is a feature of the canonical BCM learning rule, remains present in our dBCM learning rule. To test this, we assign correlated external inputs to different groups of neurons within the network, as described in Methods. Each of these four groups of 10 correlated inputs may be thought of as representing a certain stimulus, such as the orientation of moving bars presented as visual stimuli.

We first test the effects of both BCM and dBCM in response to groups of correlated external inputs when these groups have a spatial structure, as shown in Fig. 4A. Figure 4B shows the final synaptic weight matrix after a period of plasticity mediated by the standard BCM learning rule, where weight updates are local. Assemblies of strongly interconnected neurons clearly emerge, due to neurons within these assemblies receiving external inputs from the same correlated input group (Ko *et al.*, 2013). This results in synapses within these assemblies potentiating in response to high co-activity, whereas synapses between assemblies undergo depression as the sliding threshold, θ_i (which determines the postsynaptic activity level necessary for potentiation), gradually increases in response to increased recurrent excitation. The emergence of strongly interconnected assemblies of neurons in response to correlated inputs are well described in the literature (Hopfield, 1982; Mongillo *et al.*, 2005; Clopath *et al.*, 2010; Sadeh *et al.*, 2015; Zenke *et al.*, 2015).

For networks with dBCM, similar assemblies of strongly interconnected neurons which share correlated inputs develop during synaptic plasticity (Fig. 4C,D). However, as demonstrated in Fig. 3, dBCM also leads to spatial structure in the synaptic weight matrix which reflect the diffusive range matrix. The final synaptic weight matrices shown in Fig. 4C,D are therefore influenced by both the spatial structure of correlated external inputs and the underlying spatial structure of the neurons, given by the diffusive range matrix. This leads to neuronal assemblies which, while primarily

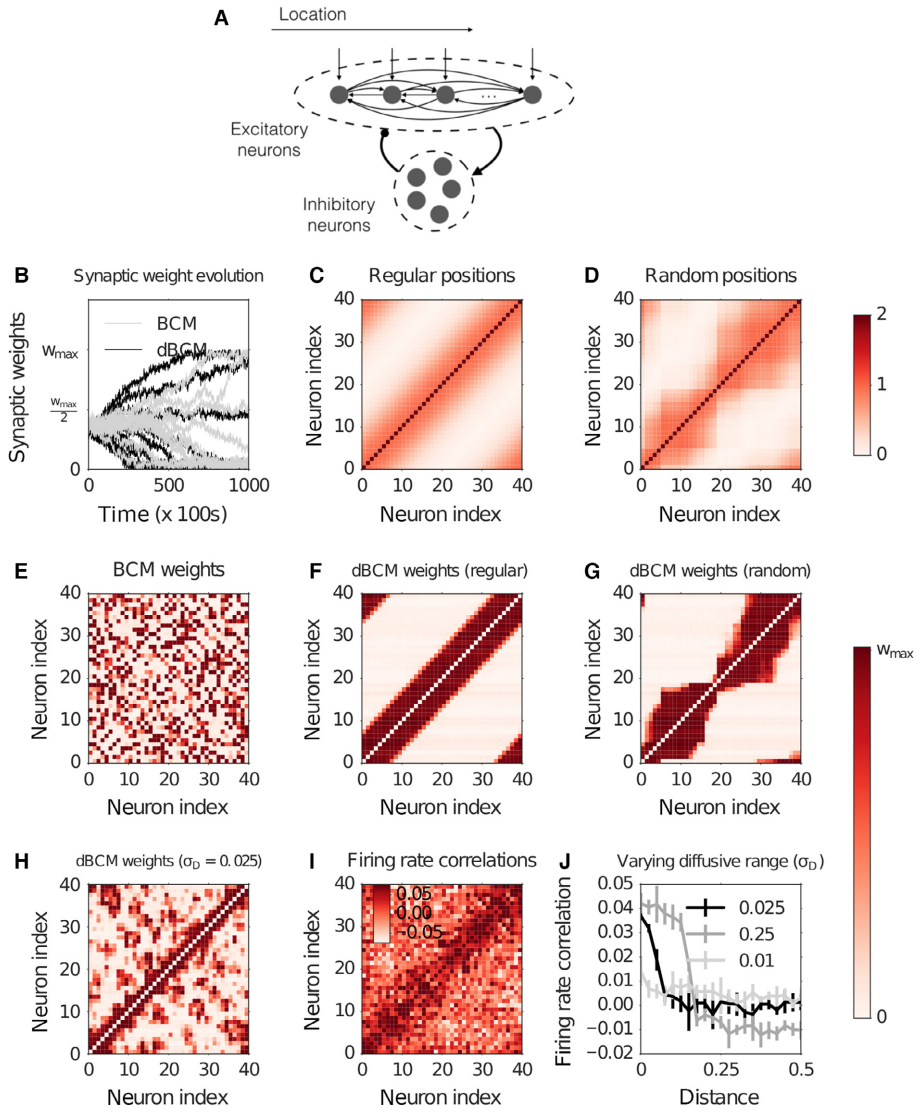

FIG. 3. Diffusive Bienenstock–Cooper–Munro (BCM) plasticity leads to spatial structure in the synaptic weights. (A) Architecture of the recurrent network model. (B) Evolution of sample synaptic weights during plasticity mediated by either the BCM (grey traces) or diffusive BCM (dBCM, black traces) learning rule. Synaptic weights are sampled every 100 s. (C) Diffusive range matrix for excitatory neurons within the recurrent network, for the case where neurons locations are regularly distributed between 0 and 1. Note that periodic boundary conditions are used to calculate distance. (D) Diffusive range matrix for the case where neuron locations are randomly distributed between 0 and 1, then ordered by index. (E) Final excitatory synaptic weight matrix of a network receiving random inputs with synaptic plasticity mediated by a BCM learning rule. (F) Same as E, but for a network with synaptic plasticity mediated by a dBCM learning rule, and with regularly distributed neuron locations. (G) Same as F, but for a network with randomly distributed neuron locations. (H) Same as F, but when the diffusive range width is decreased to $\sigma_D = 0.025$. (I) Pairwise Pearson correlation coefficients between the firing rates of neurons in the network shown in F. (J) The dependence of the mean Pearson correlation coefficient of the firing rates of two neuron on their distance apart, as the diffusive range width σ_D is varied between 0.01 and 0.25. Correlation coefficients are averaged over five independent network instantiations for each value of σ_D. Errorbars denote the standard deviation.

interconnected with neurons sharing correlated inputs, also exhibit a significant degree of overlapping connectivity with neighbouring neuronal assemblies. Having tested the network response to spatially structured correlated inputs, we will now consider the response to spatially unstructured correlated inputs. This is achieved by randomly assigning four groups of 10 external inputs to the neurons independent of spatial position, as shown in Fig. 4A and as described in Methods. Assemblies of strongly interconnected neurons which share inputs from the same group can be more easily visualised if neuron indices in the final synaptic weight matrix are reordered so that neurons are ordered by input group identity, and not by spatial position. This reordered weight matrix is shown in Fig. 4F. Although a significant portion of strong synapses within the network can be attributed to these assemblies formed by shared correlated inputs, there still remains some underlying spatial structure within the weight matrix due to the diffusive effect of dBCM. Figure 4G shows the weight matrix, ordered again by spatial location as in Fig. 4E, but where synapses between neurons that share correlated external inputs are coloured blue. This demonstrates that, even when correlated input leads to the development of strongly interconnected neuron assemblies which are not spatially structured, a small amount of spatial structure still remains in these networks due to the influence that dBCM has on neighbouring neurons.

These results suggest that, together with the emergence of spatial synaptic structure, specific connectivity driven by correlated stimuli can develop in networks with dBCM.

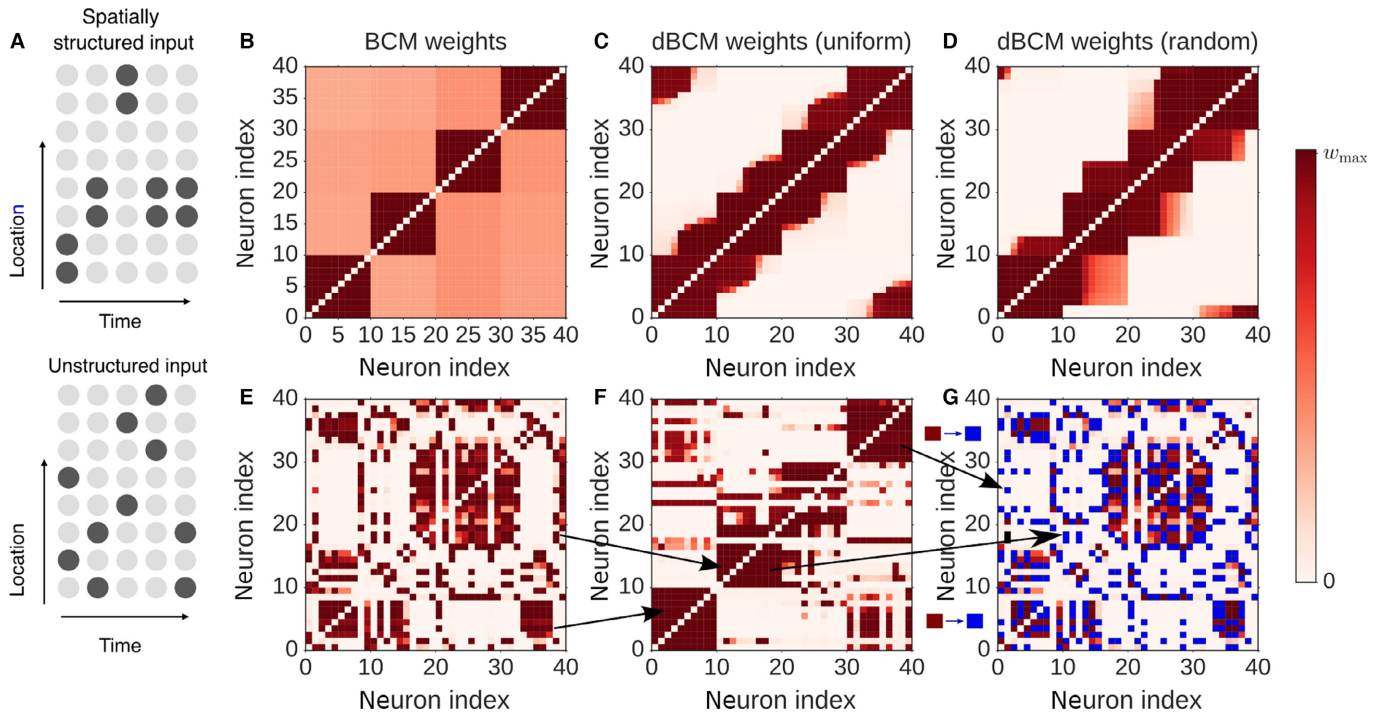

FIG. 4. Specific connectivity in response to spatially structured or unstructured groups of correlated inputs. (A) Illustration of the inputs to a network, when groups of correlated inputs are either spatially structured (top) or have no spatial structure (bottom). (B) Final excitatory synaptic weight matrix of a network receiving correlated inputs from spatially structured groups with synaptic plasticity mediated by a Bienenstock–Cooper–Munro (BCM) learning rule. (C) Same as B, but for a network with synaptic plasticity mediated by a diffusive BCM (dBCM) learning rule, and with regularly distributed neuron locations. (D) Same as C, but for a network with randomly distributed neuron locations. (E) Final excitatory synaptic weight matrix of a network receiving correlated inputs from groups with no spatial structure, and with synaptic plasticity mediated by a BCM learning rule. (F) The same synaptic weight matrix as in E, with neuron indices reordered (see black arrows) such that they correspond to input group and no longer to spatial location. (G) The same synaptic weight matrix as in E, but with synapses that form the strongly interconnected assemblies of neurons, shown in F, coloured blue to highlight the underlying spatial structure which remains.

Receptive fields developed with dBCM are spatially organised

We have shown that dBCM leads to spatial structure in recurrent networks. We hypothesise that diffusive plasticity can also form spatially organised receptive fields, such that neighbouring neurons will develop similar stimulus tuning. To test this, we compare the compare the effects of a BCM and a dBCM learning rule in feedforward networks, with architecture as described in Methods and illustrated in Fig. 5A. We simulate the development of receptive fields to specific stimuli by sequentially presenting a randomly chosen input out of the five possible inputs to the network as synaptic plasticity occurs. This is further described in Methods. Over time, the competitive nature of BCM will lead to the strengthening of a single incoming synapse for each output neuron, which determines the input that particular output neurons becomes responsive to, whereas the remaining incoming synaptic weights tend towards 0. This is shown in Fig. 5B, where we plot the evolution of synaptic weights from a single example input neuron to the 50 output neurons. In this case, five output neurons have become responsive to this input, whereas the remaining output neurons become unresponsive.

A different outcome to the standard BCM learning rule emerges in networks with a dBCM learning rule. Figure 5C shows the evolution of postsynaptic weights for an example input neuron, as in Fig. 5A. Although output neurons become selective to only a single input as before, some spatial structure now emerges regarding which output neurons respond to a particular input. This is clearly in demonstrated in Fig. 5C, where the nine outputs responsive to the input are all located within the same spatial region. This can again be explained by the effect that changes in synaptic weights have on incoming synapses of neighbouring postsynaptic neurons, illustrated in Fig. 1. Figure 5D,E further illustrates the difference between dBCM and BCM, where we plot the final incoming synaptic weights for the 50 output neurons. For standard BCM, there is no significant correlation in the input which a particular output neuron is selective to compared with its neighbouring output neurons (Fig. 5D). For dBCM, output neurons are highly likely to develop selectivity to the same input as their neighbouring neurons (Fig. 5E).

Given these differences in receptive field development, we hypothesised that functional differences would emerge between these networks, such as in their response to changes in input statistics. We test this directly by turning off one input to each network and allowing them to run with plasticity for a further. Figure 5F,G show the incoming synaptic weights after this period of plasticity for the same networks as in Fig. 5D,E. In both networks, the synaptic weights decrease from the input we have turned off (labelled with index 0), and the corresponding output neurons compensate by increasing their synaptic weights from inputs which have remained active. As such, neurons which are selective to inactive stimuli undergo plasticity to become selective to active stimuli. This rewiring occurs significantly faster in networks with dBCM than in networks with BCM (Fig. 5F,G. Mean absolute weight across five independent network instantiations is 0.03 ± 0.01 for BCM and 0.06 ± 0.01 for dBCM. $P = 0.01$, unpaired t-test, $\pm$ represents the standard deviation). In addition, output neurons in the network with

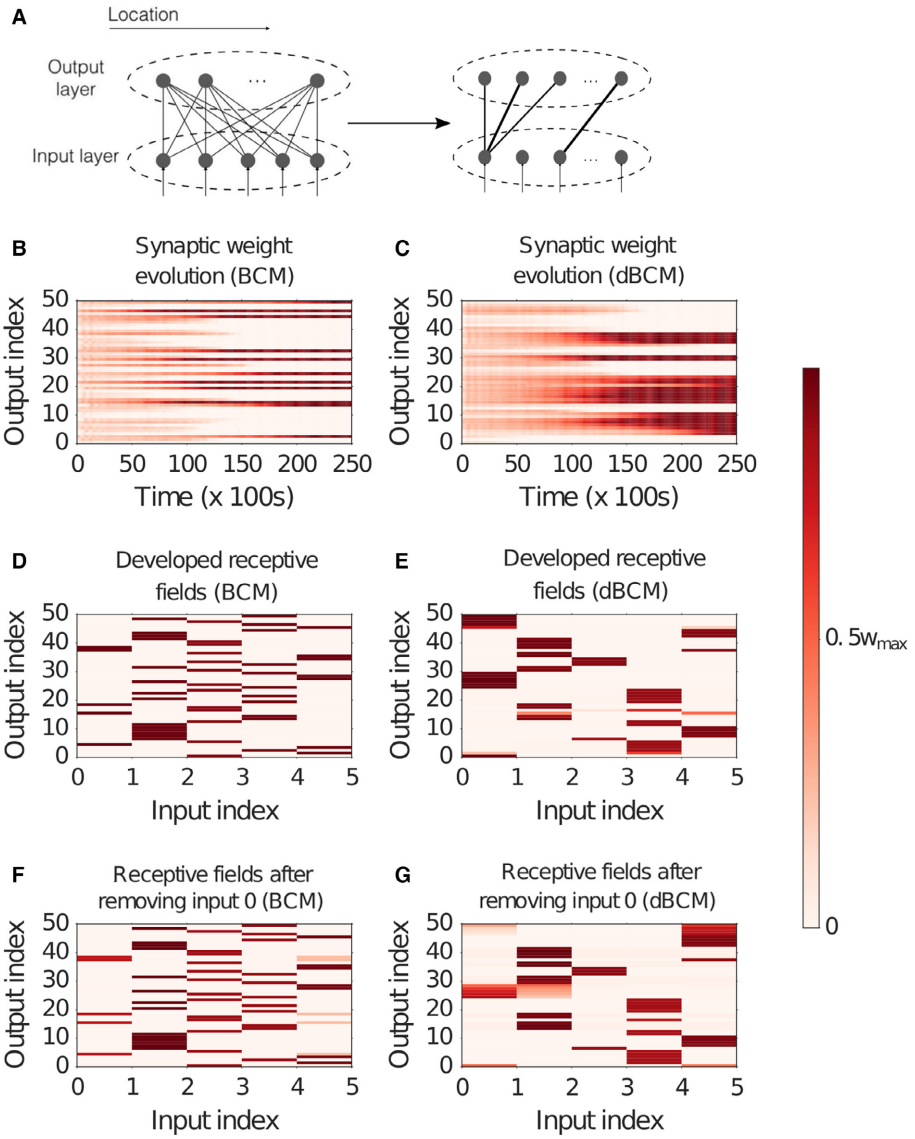

FIG. 5. Emergence of spatially organised receptive fields in a feedforward network with diffusive Bienenstock–Cooper–Munro (BCM). (A) Architecture of the feedforward network model (left), and connectivity after diffusive BCM (right). (B) Evolution of the synaptic weights from an example input neuron onto the 50 output neurons in a feedforward network with a BCM learning rule, as a randomly chosen input out of the five possible inputs are sequentially presented. Synaptic weights are sampled every 100 s. (C) Same as in B, but for a network with a diffusive BCM (dBCM) learning rule. (D) Final synaptic weights between all five input neurons and 50 output neurons in a network with a BCM learning rule. (E) Same as in D, but for a network with a dBCM learning rule. (F) Synaptic weights between all five inputs neurons and 50 outputs neurons in a network with BCM, after a period of plasticity following the removal of input 0. (G) Same as in F, but for a network with a dBCM learning rule.

dBCM rewire so that their new selective inputs are the same as those of neighbouring neurons (Fig. 5G), suggesting that the spatial influence of synaptic weight updates facilitates the recruitment of inactive neurons.

These simulations demonstrate that the influence of neighbouring neurons on synaptic plasticity in a network with dBCM leads to the development of receptive fields which are spatially organised, and which have a more flexible representation of stimulus statistics.

Discussion

This study proposes a new form of Hebbian plasticity which is mediated by a diffusive neurotransmitter, such that synaptic weight updates are influenced by the activity of neighbouring postsynaptic neurons (Fig. 1). This is in contrast with canonical models of neural plasticity, which generally only allow for communication between neurons with existing synaptic connections. We establish the conditions necessary for effective diffusive neurotransmission to synapses of neighbouring neurons (Fig. 2), and explore the consequences of this mechanism in a simple network model.

The primary consequence of our proposed dBCM learning rule is the emergence of spatial structure in the synaptic weight matrix after a period of plasticity. We first demonstrate this in a simple recurrent network model receiving random uncorrelated inputs (Fig. 3), which results in neighbouring neurons developing strong synaptic connections, whereas neurons that are located further away from each other develop weak synaptic connections. This differs from the standard BCM learning rule with a purely local synaptic update, which results in a spatially uncorrelated synaptic weight matrix (Fig. 3E). Indeed, the extent of spatial correlations in the firing rates of

neurons after dBCM increases with the range of the diffusive neurotransmitter (Fig. 3J).

We next test how the synaptic weights evolve in response to groups of correlated inputs. As expected from previous studies (Mongillo *et al.*, 2005; Sadeh *et al.*, 2015), correlated input leads to the development of strongly interconnected neuron assemblies driven by common input, in both networks with a BCM or dBCM learning rule (Fig. 4). However, in networks with dBCM, there still remains some underlying spatial structure within the synaptic weight matrix due to the influence of the activity of neighbouring neurons on synaptic plasticity. This occurs both when the groups of correlated inputs have some spatial structure (Fig. 4C,D), or when they are spatially unstructured (Fig. 4E–G).

The spatial structure which emerges due to dBCM in addition to the spatial structure imposed by correlated inputs may be of functional relevance. For brain regions in which similar stimuli are topographically organised such as the primate early visual cortex or cat auditory cortex (Reale & Imig, 1980; Kaschube, 2014), assemblies developed from dBCM would be more likely to coactivate with similar stimuli compared with assemblies developed from BCM. Such coactivation could ensure more flexible stimuli representations in environments with dynamic stimulus statistics, as explored in Fig. 5. It is also possible that the spatial correlations in stimulus preferences introduced by dBCM could have a detrimental impact on neural coding by introducing redundancy of stimulus representations across neighbouring neurons.

Finally, we explore the effect of dBCM in a simplified model of receptive field development in a feedforward network (Fig. 5). The emergence of spatially correlated receptive fields suggests a potential mechanism for the development of orientation preference maps in visual cortex, in which the extent of spatial correlation in these maps would be determined by the range of the diffusive neurotransmitter. The spatial organisation resulting from networks with a very small diffusive range, or with only BCM, would resemble the salt-and-pepper organisation of orientation preferences observed in mouse visual cortex (Kaschube, 2014). There are a range of mechanisms which have been proposed to explain the development of topographic organisation in visual cortex, extensively reviewed by Erwin *et al.* (1995; Vidyasagar & Eysel, 2015). Notably, there is much debate on the role of spatially structured feedforward thalamic input which may already be biased to a certain orientation selectivity, compared with the emergence of orientation selectivity through intracortical interactions (Somers *et al.*, 1995; Kuhlmann & Vidyasagar, 2011). Although dBCM provides an energetically efficient alternative mechanism to these theories as it acts independently of synaptic connectivity, it is not to our knowledge in contradiction with them. Indeed, it is plausible that dBCM may act in concert with other mechanisms to enhance the topographic organisation of stimulus preferences. This hypothesis can be tested in experiment by blocking the action of diffusive neurotransmission during development, and testing whether the development of spatially structured stimulus selectivity is impaired. Animal models which are deficient in NO synthase, a candidate mediator of dBCM, have been developed (Huang, 2000). In addition, we demonstrate a functional advantage of dBCM over BCM. Networks with dBCM exhibit a more flexible response to a loss of input by rewiring, compared to networks with BCM (Fig. 5F, G). These functional properties may also be investigated in an animal model deficient in NO synthase by testing the plasticity of responses to chronic changes in environmental stimulus statistics.

A significant assumption of our model is that neurons occupy a single point in space. As the spatial extent of dendritic branches can be very large compared with the distance between neighbouring neurons (London & Häusser, 2005), it is not necessarily the case that synapses are spatially distributed in an ordered manner. However, a recent overview discussing the plausibility of NO volume transmission from synaptic sources found that the most likely scenario was one in which multiple synaptic sources combined to create a cloud of gas which would affect all local targets (Garthwaite, 2015), which we have explored in Fig. 2. As such, our assumption is a reasonable approximation of this scenario. Future research could investigate the emergent spatial structure in a more biophysically realistic model of a cortical network which includes a 3D descriptions of space (Bauer *et al.*, 2014).

Note that the y_j and y_i term outside the bracket in Eqn 6 ensures that some amount of activity is still required in both pre and postsynaptic neurons for synaptic plasticity to occur regardless of the activity of neighbouring neurons, as a value of 0 for either leads to no change in the synaptic weight. As the BCM learning rule is equivalent to a biologically realistic triplet STDP rule for rate-based patterns (Gjorgjieva *et al.*, 2011), we would expect qualitatively similar results in a network of spiking neurons. The consequences of diffusive neurotransmitters mediating other forms of Hebbian or homeostatic plasticity have been explored previously by Kohonen (1993; Smith *et al.*, 2002; Bhaumik & Mathur, 2003; Yin *et al.*, 2005; Savin *et al.*, 2009; Gupta & Markan, 2014; Sweeney *et al.*, 2015). Although we remain agnostic to the identity of the diffusive neurotransmitter which mediates synaptic plasticity in our proposed dBCM learning rule, there are a number of candidates which have been previously identified. Nitric oxide, carbon monoxide and hydrogen sulphide are all gaseous neurotransmitters which are activity-dependent and have been implicated in synaptic plasticity signalling pathways (Dawson & Snyder, 1994; Boehning & Snyder, 2003). Of these, the most studied is NO, whose diffusive properties are consistent with our putative neurotransmitter (Philippides *et al.*, 2000; Garthwaite, 2015; Sweeney *et al.*, 2015), and which is implicated in signalling pathways related to both long-term potentiation and long-term depression (Kiss & Vizi, 2001; Gallo & Iadecola, 2011; Hardingham *et al.*, 2013). Indeed, there is some experimental demonstration which implicate NO in long-term potentiation that is locally distributed across neighbouring neurons (Bonhoeffer *et al.*, 1989; Schuman & Madison, 1994; Madison & Schuman, 1995).

It is also worth noting that the consequences of NO in mediating a form of homeostatic intrinsic plasticity have been investigated, in which diffusion was modelled in detail in a spiking neural network model (Sweeney *et al.*, 2015). Although the effect of this proposed diffusive homeostasis is increased heterogeneity in neural firing rates, here we propose an unrelated mechanism which acts through synaptic as opposed to intrinsic plasticity.

The structure of synaptic connectivity within sensory cortical networks has been of keen interest in recent years, as experiments elucidate the relationship between the sensory features that neurons encode and their functional connectivity (Harris & Mrsic-Flogel, 2013; Cossell *et al.*, 2015). Our proposed dBCM learning rule demonstrates that diffusive neurotransmission can introduce spatial structure in the synaptic connectivity of a network, providing possibilities for the interaction between functional organisation driven by common sensory features and functional organisation driven by the underlying spatial distribution of neurons within a network.

Acknowledgements

CC and YS were supported by the Engineering and Physical Sciences Research Council (EPSRC), the Leverhulme Trust and Google Faculty Award.

Abbreviations

BCM, Bienenstock–Cooper–Munro; dBCM, diffusive Bienenstock–Cooper–Munro; NO, Nitric oxide.

References

Artinian, L., Tornieri, K., Zhong, L., Baro, D. & Rehder, V. (2010) Nitric oxide acts as a volume transmitter to modulate electrical properties of spontaneously firing neurons via apamin-sensitive potassium channels. *J. Neurosci.*, **30**, 1699–1711.

Batchelor, A.M., Bartus, K., Reynell, C., Constantinou, S., Halvey, E.J., Held, K.F., Dostmann, W.R., Vernon, J. *et al.* (2010) Exquisite sensitivity to subsecond, picomolar nitric oxide transients conferred on cells by guanylyl cyclase-coupled receptors. *Proc. Natl. Acad. Sci. USA*, **107**, 22060–22065.

Bauer, R., Zubler, F., Pfister, S., Hauri, A., Pfeiffer, M., Muir, D.R. & Douglas, R.J. (2014) Developmental self-construction and -configuration of functional neocortical neuronal networks. *PLoS Comput. Biol.*, **10**, e1003994.

Bellefontaine, N., Chachlaki, K., Parkash, J., Vanacker, C., Colledge, W., D'Anglemont De Tassigny, X., Garthwaite, J., Bouret, S.G. *et al.* (2014) Leptin-dependent neuronal no signaling in the preoptic hypothalamus facilitates reproduction. *J. Clin. Invest.*, **124**, 2550–2559.

Bhaumik, B. & Mathur, M. (2003) A cooperation and competition based simple cell receptive field model and study of feed-forward linear and nonlinear contributions to orientation selectivity. *J. Comput. Neurosci.*, **14**, 211–227.

Bienenstock, E., Cooper, L. & Munro, P. (1982) Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J. Neurosci.*, **2**, 32–48.

Blais, B. & Cooper, L. (2008) BCM theory. *Scholarpedia*, **3**, 1570.

Boehning, D. & Snyder, S.H. (2003) Novel neural modulators. *Annu. Rev. Neurosci.*, **26**, 105–131.

Bonhoeffer, T., Staiger, V. & Aertsen, A. (1989) Synaptic plasticity in rat hippocampal slice cultures: local "Hebbian" conjunction of pre- and postsynaptic stimulation leads to distributed synaptic enhancement. *Proc. Natl. Acad. Sci. USA*, **86**, 8113–8117.

Burette, A., Zabel, U., Weinberg, R.J., Schmidt, H.H.H.W. & Valtschanoff, J.G. (2002) Synaptic localization of nitric oxide synthase and soluble guanylyl cyclase in the Hippocampus. *J. Neurosci.*, **22**, 8961–8970.

Clopath, C., Buesing, L., Vasilaki, E. & Gerstner, W. (2010) Connectivity reflects coding: a model of voltage-based STDP with homeostasis. *Nat. Neurosci.*, **13**, 344–352.

Cossell, L., Iacaruso, M.F., Muir, D.R., Houlton, R., Sader, E.N., Ko, H., Hofer, S.B. & Mrsic-flogel, T.D. (2015) Functional organization of excitatory synaptic strength in primary visual cortex. *Nature*, **518**, 1–5.

Dawson, T. & Snyder, S. (1994) Gases as biological messengers: nitric oxide and carbon monoxide in the brain. *J. Neurosci.*, **14**, 5147–5159.

Erwin, E., Obermayer, K. & Schulten, K. (1995) Models of orientation and ocular dominance columns in the visual cortex: a critical comparison. *Neural Comput.* **7**, 425–468.

Gallo, E.F. & Iadecola, C. (2011) Neuronal nitric oxide contributes to neuroplasticity-associated protein expression through cGMP, protein kinase G, and extracellular signal-regulated kinase. *J. Neurosci.*, **31**, 6947–6955.

Garthwaite, J. (2008) Concepts of neural nitric oxide-mediated transmission. *Eur. J. Neurosci.*, **27**, 2783–2802.

Garthwaite, J. (2015) From synaptically localized to volume transmission by nitric oxide. *J. Physiol.*, **594**, 9–18.

Gjorgjieva, J., Clopath, C., Audet, J. & Pfister, J.-P. (2011) A triplet spike-timing-dependent plasticity model generalizes the Bienenstock-Cooper-Munro rule to higher-order spatiotemporal correlations. *Proc. Natl. Acad. Sci. USA*, **108**, 19383–19388.

Gupta, P. & Markan, C.M. (2014) An adaptable neuromorphic model of orientation selectivity based on floating gate dynamics. *Front. Neurosci.*, **8**, 54.

Hardingham, N., Dachtler, J. & Fox, K. (2013) The role of nitric oxide in pre-synaptic plasticity and homeostasis. *Front. Cell. Neurosci.*, **7**, 190.

Harris, K.D. & Mrsic-Flogel, T.D. (2013) Cortical connectivity and sensory coding. *Nature*, **503**, 51–58.

Hennequin, G., Vogels, T.P. & Gerstner, W. (2014) Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron*, **82**, 1394–1406.

Hopfield, J.J. (1982) Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. USA*, **79**, 2554–2558.

Huang, P.L. (2000) Mouse models of nitric oxide synthase deficiency. *J. Am. Soc. Nephrol.*, **11**(Suppl 1), S120–S123.

Hunter, J.D. (2007) Matplotlib: a 2D graphics environment. *Comput. Sci. Eng.*, **9**, 90–95.

Kaschube, M. (2014) Neural maps versus salt-and-pepper organization in visual cortex. *Curr. Opin. Neurobiol.*, **24**, 95–102.

Kiss, J.P. & Vizi, E.S. (2001) Nitric oxide: A novel link between synaptic and nonsynaptic transmission. *Trends Neurosci.*, **24**, 211–215.

Ko, H., Cossell, L., Baragli, C., Antolik, J., Clopath, C., Hofer, S.B. & Mrsic-Flogel, T.D. (2013) The emergence of functional microcircuits in visual cortex. *Nature*, **496**, 96–100.

Kohonen, T. (1993) Physiological interpretation of the self-organizing map algorithm. *Neural Networks*, **6**, 895–905.

Kuhlmann, L. & Vidyasagar, T.R. (2011) A computational study of how orientation bias in the lateral geniculate nucleus can give rise to orientation selectivity in primary visual cortex. *Front. Syst. Neurosci.*, **5**, 81.

London, M. & Häusser, M. (2005) Dendritic computation. *Annu. Rev. Neurosci.*, **28**, 503–532.

Madison, D. & Schuman, E. (1995) Diffusible messengers and intercellular signaling: locally distributed synaptic potentiation in the hippocampus. In Koprowski, H. & Hiroshi, M. (Eds), *The Role of Nitric Oxide in Physiology and Pathophysiology*. Springer, Berlin, pp. 5–6.

Mishchenko, Y., Hu, T., Spacek, J., Mendenhall, J., Harris, K.M. & Chklovskii, D.B. (2010) Ultrastructural analysis of hippocampal neuropil from the connectomics perspective. *Neuron*, **67**, 1009–1020.

Monfils, M.-H., Bray, D.F., Driscoll, I., Kleim, J.A. & Kolb, B. (2005) A quantitative comparison of synaptic density following perfusion versus immersion fixation in the rat cerebral cortex. *Microsc. Res. Techniq.*, **67**, 300–304.

Mongillo, G., Curti, E., Romani, S. & Amit, D.J. (2005) Learning in realistic networks of spiking neurons and spike-driven plastic synapses. *Eur. J. Neurosci.*, **21**, 3143–3160.

Perez, F. & Granger, B.E. (2007) IPython: a system for interactive scientific computing. *Comput. Sci. Eng.*, **9**, 21–29.

Philippides, A. (2001) Modelling the diffusion of nitric oxide in brains. Doctoral dissertation, University of Sussex, Brighton, UK.

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

Philippides, A., Ott, S.R., Husbands, P., Lovick, T.A. & O'Shea, M. (2005) Modeling cooperative volume signaling in a plexus of nitric oxide synthase-expressing neurons. *J. Neurosci.*, **25**, 6520–6532.

Rajan, K., Abbott, L. & Sompolinsky, H. (2010) Stimulus-dependent suppression of chaos in recurrent neural networks. *Phys. Rev. E*, **82**, 011903.

Reale, R.A. & Imig, T.J. (1980) Tonotopic organization in auditory cortex of the cat. *J. Comp. Neurol.*, **192**, 265–291.

Rubinov, M., Ypma, R.J., Watson, C. & Bullmore, E.T. (2015) Wiring cost and topological participation of the mouse brain connectome. *Proc. Natl. Acad. Sci. USA*, **112**, 10032–10037.

Sadeh, S., Clopath, C. & Rotter, S. (2015) Emergence of functional specificity in balanced networks with synaptic plasticity. *PLoS Comput. Biol.*, **11**, e1004307.

Savin, C., Triesch, J. & Meyer-Hermann, M. (2009) Epileptogenesis due to glia-mediated synaptic scaling. *J. R. Soc. Interface*, **6**, 655–668.

Schuman, E.M. & Madison, D.V. (1994) Locally distributed synaptic potentiation in the hippocampus. *Science (New York, N.Y.)*, **263**, 532–536.

Schüz, A. & Palm, G. (1989) Density of neurons and synapses in the cerebral cortex of the mouse. *J. Comp. Neurol.*, **286**, 442–455.

Smith, T., Husbands, P., Philippides, A. & O'Shea, M. (2002) Neuronal plasticity and temporal adaptivity: GasNet robot control networks. *Adapt. Behav.*, **10**, 161–183.

Somers, D.C., Nelson, S.B. & Sur, M. (1995) An emergent model of orientation selectivity in cat visual cortical simple cells. *J. Neurosci.*, **15**, 5448–5465.

Steinert, J.R., Kopp-Scheinpflug, C., Baker, C., Challiss, R.A., Mistry, R., Haustein, M.D., Griffin, S.J., Tong, H. *et al.* (2008) Nitric oxide is a volume transmitter regulating postsynaptic excitability at a glutamatergic synapse. *Neuron*, **60**, 642–656.

Sweeney, Y., Hellgren Kotaleski, J. & Hennig, M.H. (2015) A diffusive homeostatic signal maintains neural heterogeneity and responsiveness in cortical networks. *PLoS Comput. Biol.*, **11**, e1004389.

Vidyasagar, T.R. & Eysel, U.T. (2015) Origins of feature selectivities and maps in the mammalian primary visual cortex. *Trends Neurosci.*, **38**, 475–485.
van der Walt, S., Colbert, S.C. & Varoquaux, G. (2011) The NumPy Array: a structure for efficient numerical computation. *Comput. Sci. Eng.*, **13**, 22–30.
Yin, J., Hu, D., Chen, S. & Zhou, Z. (2005) DSOM: a novel self-organizing model based on NO dynamic diffusing mechanism. *Sci. China Ser. F*, **48**, 247–262.
Zenke, F., Agnes, E.J. & Gerstner, W. (2015) Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks. *Nat. Commun.*, **6**, 6922.