Nitric Oxide Neuromodulation

Michael O'Shea, Phil Husbands, Andy Philippides
University of Sussex

Synonyms
Nitric oxide volume signalling

Definition
Neuromodulators are a class of neurotransmitter that diffuse into the region surrounding an emitting neuron and affect potentially large numbers of other neurons by modulating their responses, irrespective of whether or not they are electrically connected to the modulating neuron. Nitric Oxide (NO) is a particularly interesting example of a neuromodulator because of its very small size and gaseous state. The type of modulatory signalling NO is involved in, sometimes known as volume signalling, is in sharp contrast to the connectionist point-to-point electrical transmission picture that dominated thinking about the nervous system for many decades, whereby neural signalling could only occur between electrically connected neurons.

Detailed Description

Background

The discovery that the toxic gas nitric oxide (NO) is synthesised in biological systems and functions as a physiological signalling molecule earned Robert Furchgott, Louis Ignarro and Ferid Murad the Nobel Prize in 1998. By showing that physiological functions could be regulated by information carried from cell to cell by a gas, their research suggested an entirely new way of thinking about how cells communicate with one another.

Increasing blood flow was the first physiological function attributed to NO. It was known that the cells of the endothelial lining of blood vessels released a factor which caused the relaxation of circular muscle leading to vessel dilation and increased blood flow. The endothelial factor was unstable and thus elusive. But the combined work of the three Nobel laureates showed conclusively, if unexpectedly, that this factor was NO.

If this were the whole story, there would be of little or no relevance for our understanding of brain function. When however Solomon Snyder and colleagues isolated an enzyme from the mammalian brain that synthesised NO and showed that it was present in sub-populations of neurons (Bredt and Snyder 1990), interest in NO's signalling functions switched from endothelial cells to neurons.

Three forms of the NO synthesising enzyme, known as NOS, have been identified; it is the neuronal form (nNOS) which attracted most interest among neuroscientists (Bredt and Snyder 1990, Bredt et al. 1990, Mungrue and Bredt 2004). Since NO is both a signalling molecule and a toxic free radical gas, its generation by nNOS must be strictly controlled. The primary regulator of nNOS activity is the concentration of calcium inside the nNOS expressing neurons (Bredt and Snyder 1990). Intracellular calcium concentration in resting neurons is very low – below that required for the synthesis of NO. However, as a

consequence of electrical activity calcium enters the neuron via voltage gated calcium channels and the concentration rises transiently from micro-molar to milli-molar levels. At these elevated levels the synthesis of NO is triggered. Calcium therefore neatly links the synthesis of NO in the brain to the electrical activity of neurons. Because NO is so small (the 10th smallest molecule of the Universe) and non-polar it can diffuse more or less freely away from its source, unimpeded by the lipid bi-layer of the cell membrane. Thus for NO, release is an inevitable consequence of NO synthesis.

Calcium influx in electrically excited neurons both triggers NO release and also is the trigger for the release of conventional neurotransmitters from synaptic vesicles. The two modes of release, diffusive and vesicular, are not mutually exclusive alternatives at the neuronal level. Neurons capable of NO synthesis also store conventional neurotransmitters in synaptic vesicles. NO is therefore best regarded as a co-transmitter, acting in conjunction with classical inhibitory and excitatory transmitters.

Clearly NO signalling violates some of the classical tenets of synaptic transmission and there has been an understandable reluctance to regard NO as a genuine neurotransmitter. Perhaps the term neuromodulator is more appropriate. Neuromodulators are a class of neurotransmitter that, rather than causing excitation or inhibition directly, modulate the responses of neurons to direct excitatory and inhibitory transmitters. A modulator may for example enhance or facilitate the inhibitory effect of an inhibitory neurotransmitter. In this way modulators are synaptic gain-control agents. Whereas direct inhibition and excitation is mediated by transmitter-gated ion channels, modulators bind to receptors that regulate the synthesis of so-called second messengers and it is these that can change the sensitivity of a neuron to other neurotransmitters.

Because they act indirectly via a neuron's metabolic machinery, the temporal dynamics of neuromodulator effects are delayed and more long-lasting than the quick and transient effects of the classical transmitters. The spatial dynamics of neuromodulation are also noteworthy with respect to signalling by NO. As noted above, the physical properties of NO allow it to diffuse into the volume surrounding an emitting neuron. Therefore, depending on the size of the affected volume, a NO emitting neuron may modulate large numbers of other neurons within the affected volume. Importantly then, a nNOS expressing neuron may communicate with other neurons irrespective of whether or not they are connected by synapses. In this way NO is a particularly interesting example of a neuromodulator because it can act as a non-synaptic volume signal. This clearly is in sharp contrast to the connectionist point-to-point synaptic transmission picture that dominated thinking about the nervous system for many decades.

In the traditional connectionist model of neuron-to-neuron communication, neurons generate brief electrical signals (action potentials), which propagate along wire-like axons terminating at highly localized junctions (synapses) on other neurons, where the release of a neurotransmitter is triggered. The neurotransmitter is confined to the region of the synapse, and here the receiving neuron is equipped with receptors that directly translate the chemical signal into a brief electrical signal, either excitatory or inhibitory (Purves 1997, Brazier 1961). Hence in standard artificial neural network (ANN) models based on this incomplete picture, chemical signalling can be safely factored out, leaving only the idea of electrical signals flowing between nodes in a network.

However, the discovery of non-synaptic chemical signalling by NO has potentially greatly extended the spatial scale over which neurons can communicate. The most important feature of this derives from the ability of NO to diffuse away from its site of release and to occupy a volume of the nervous system perhaps containing many other neurons and synapses (Edelman and Gally 1992). Crucial to understanding the functional significance to volume signalling by NO (and perhaps other gases) is a measure of the spatial and temporal dynamics of NO diffusion. We need to know how large is the affected volume and how long does it take for NO to occupy it? As these questions have proven to be difficult to answer empirically (Philippides et al. 2003), an important alternative approach has been to develop realistic computational models of diffusion to illuminate the parameters of NO mediated volume signalling.

The following section looks at how NO diffusion in the nervous can be modelled computationally and how this can be used to illuminate neuromodulatory mechanisms. We will also explore how volume signalling can now be added to the growing list of phenomena in the nervous system that might be a source of inspiration for new and perhaps improved styles of ANNs. This is probably especially true for ANNs intended for use as artificial nervous systems in robots or simulated autonomous agents, an area where biomimetic techniques are often particularly fruitful. Classes of ANN directly inspired by NO signalling, which can be used in robotic applications or to shed light on biological questions, will be discussed later in this article.

Modelling NO diffusion in the nervous system

The features that make NO different to a standard neurotransmitter – it's free diffusion through neural structure with no discrete target receptor and thus inactivation – make the basis of a model relatively straightforward. Movement is governed by Fick's second law of diffusion; essentially molecules move from high concentration to low concentration. NO does not have a specific inactivating mechanism, and is lost through reaction with oxygen species and metals as well as with heme-containing proteins (Lancaster 1996, Vaughn et al. 1998a). This means that the movement of other molecules and receptors and their interactions need not be modeled and instead a more general loss function can be used, typically as concentration-dependent decay throughout the volume being modeled. These features were the basis of the first wave of models of diffusion (Gally et al., 1990, Edelman and Gally, 1991, Montague et al., 1991, Lancaster, 1994, 1996, 1997, Wood and Garthwaite 1994) which prompted key insights into the action of this kind of messenger, for instance, noting it as a 'four-dimensional volume signal' (Edelman and Gally, 1991), and highlighted the potential range of action of a diffusible signal (e.g. Wood and Garthwaite, 1994).

To make these general points it was not important to model in detail the structure of either the volume within which the neuron was diffusing (meaning that space and time were typically handled coarsely) or the source neuron itself which was typically taken to be a point in space. This meant they were unable to capture two very important features of NO signalling.

The first regards the production of NO and follows from the fact that nitric oxide synthase (NOS, responsible for the production of NO) is soluble and thus highly likely to be distributed throughout a neuron's cytoplasm. This means that the volume containing NOS is the source and the whole surface of the neuron is a potential release site for NO, in marked contrast to conventional transmitter release. These properties suggest that the 3D structure of

the NO source will have a profound influence on the dynamics of NO spread. Hence an accurate structure-based model of neuronal NO diffusion is an indispensable tool in gaining deeper insight into the signalling capacity of the molecule. The second is that depletion of NO might also be spatially heterogeneous. While a general decay rate can be assumed for most tissue, reflecting background oxidization and binding events, certain structures (importantly including blood vessels containing very high concentrations of NO-binding hemes) can have a much higher decay rate and effectively act as NO sinks. The kinetics of these reactions are not understood perfectly (Wood and Garthwaite 1994), but empirical data indicates either first or second order decay (Laurent et al. 1996,Lancaster 1996,Vaughn et al. 1998b, Liu et al. 1998,Thomas et al. 2001).

These considerations gave rise to a second set of more detailed NO diffusion models (Philippides et al., 1998, 2000, 2005a, Vaughn et al., 1998a, 1998b, Ott et al., 2007) based on the modified diffusion equation shown in Equation 1. The terms on the left-hand side describe the general diffusion process and those on the right-hand side model NO production and depletion mechanisms.

$$\frac{\partial C(\vec{x},t)}{\partial t} - D\nabla^2 C(\vec{x},t) = P(\vec{x},t) - \lambda C(\vec{x},t) - S(\vec{x})\,(C(\vec{x},t))^n \quad \text{(Eq. 1)}$$

where $\boldsymbol{x}$ is a point in space, $C(\boldsymbol{x}, t)$ is the concentration at point $\boldsymbol{x}$, and time t, D is the diffusion coefficient, typically 3300μm^2 s^{-1} and $P(\boldsymbol{x}, t)$ is the concentration of NO produced per second at point $\boldsymbol{x}$, and time t. $\lambda\, C(\boldsymbol{x}, t)$ models background inactivation of NO at all locations in the modelled volume as first-order decay with an inactivation rate (half-life) of λ. In contrast, $S(\boldsymbol{x})\,(C(\boldsymbol{x}, t))^n$ models spatially localised NO sinks such as blood vessels where the decay-rate is greatly elevated within a discrete region. Decay can be modelled as first-order ($n=1$) or second-order ($n=2$), though the latter is mainly used for very strong blood flow such as when modelling diffusion of NO from endothelial cells surrounding large blood vessels (Vaughn et al., 1998a, 1998b).

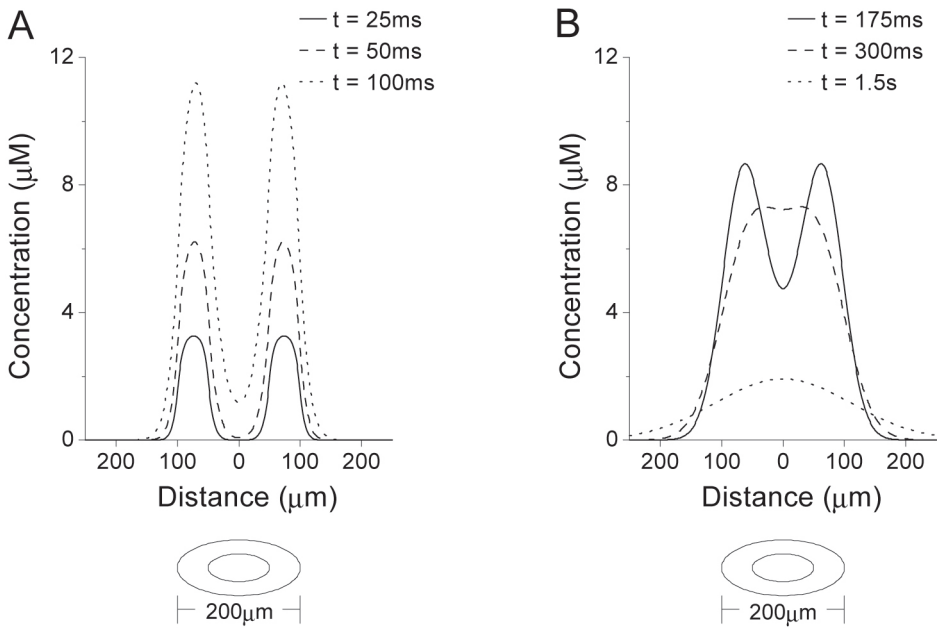

Figure 1: Concentration of NO plotted against distance from the centre of a hollow spherical source neuron of inner radius 50μm and outer radius 100μm for a 100ms burst of synthesis

starting at time t = 0. The graphics underneath each plot depict the source structure. A. Concentration of NO at times t = 25, 50 and 100ms, two time points during and one at the end of synthesis. B. Concentration of NO after synthesis at times t = 175, 300 and 1.5s. The reservoir effect following the end of synthesis is clearly seen as the centrally accumulated NO is trapped by the higher surrounding concentrations. Figure used with permission from Philippides et al., 2000.

Armed with these equations, the time-course of NO diffusion from neurons of different shapes and sizes was modeled in the presence of different levels and types of decay. Figures 1 and 2 shows the results generated by the first accurate model of NO diffusion from a continuous biologically realistic structure (Philippides et al., 1998, 2000, 2003). The source was a large neuron-like structure modelled as a hollow sphere as NO is synthesized throughout the cytoplasm but not in the nucleus. In this case, there are no sinks and thus $S(\boldsymbol{x})=0$ in equation 1 and decay rate λ is set so the NO half-life is 5s and $P(\boldsymbol{x}, t)$ is a constant value for all points inside the hollow spherical source during synthesis (the amount of NO produced per second by a single NO-producing unit multiplied by the density of these units) and 0 elsewhere and thus the solution is radially symmetric. (For full details of modeling methods and biological justification for parameter values, see Philippides et al., 2000, 2003, 2005a.)

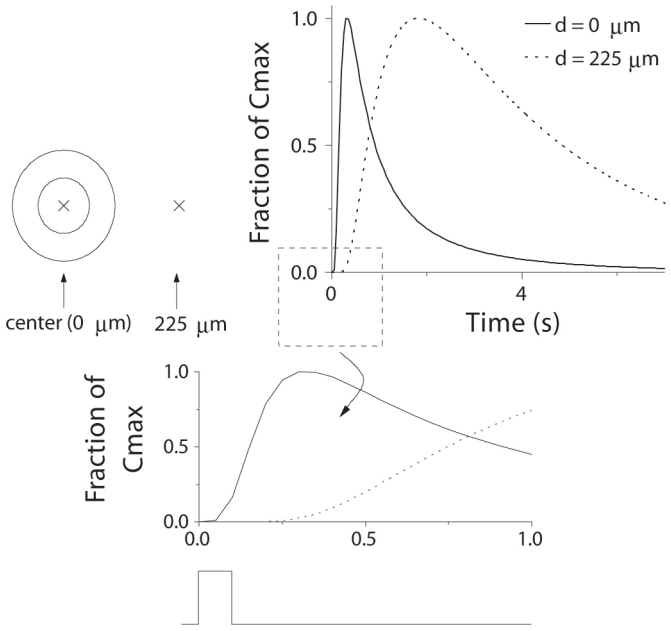

Figure 2: Concentration of NO plotted against time after synthesis for a hollow spherical source of inner radius 50µm and outer radius 100µm for a 100ms burst of synthesis. Here the solid line depicts the concentration at the centre of the cell (0µm), whilst the dotted line shows the concentration at 225µm from the centre. Because the absolute values attained at the two positions differ from one another markedly, the concentration is given as a fraction of the peak concentration attained. These peak values are 7.25µM (centre) and 0.25µM at 225µm. The cell and the points at which the concentration is measured are depicted to the left of the main figure. Note the high central concentration, which persists for a long time (above 1µM for about 2s. Also, there is a significant delay to a rise in concentration at distant points which is more clearly illustrated in the expanded inset figure. The square-wave beneath the inset figure represents NO synthesis. Figure used with permission from Philippides et al., 2000.

Examination of the time-course of the evolution of NO concentration during and after 100ms of NO synthesis shows two very interesting observations. Firstly, the NO concentration remains elevated within and throughout the cell for significant lengths of time after the end of synthesis; a 'reservoir' effect (Fig 1). Second, when viewed over time, there is a delay before the concentration peaks at points where NO is not synthesised, both within the neuron itself and outside it (Fig 2). Further, the delay increases, while peak concentration decreases, with distance from the NO source neuron. To determine the volume affected by NO, we would need to know the physiologically effective NO concentration, which is likely to vary depending on a number of factors. However, taking a threshold of 0.1 μM (close to the equilibrium dissociation constant for the NO receptor soluble guanylyl cyclase, Stone and Marletta, 1996), with 100ms of NO synthesis within the walls of hollow spherical neurons with radii over 10μm, a volume approximately 3 times the source radius would be affected. Although, importantly, the concentration time-course would depend on distance from the source. While spatially localised NO sinks such as small blood vessels do affect diffusion, they do not alter these overall properties (Philippides et al., 2000).

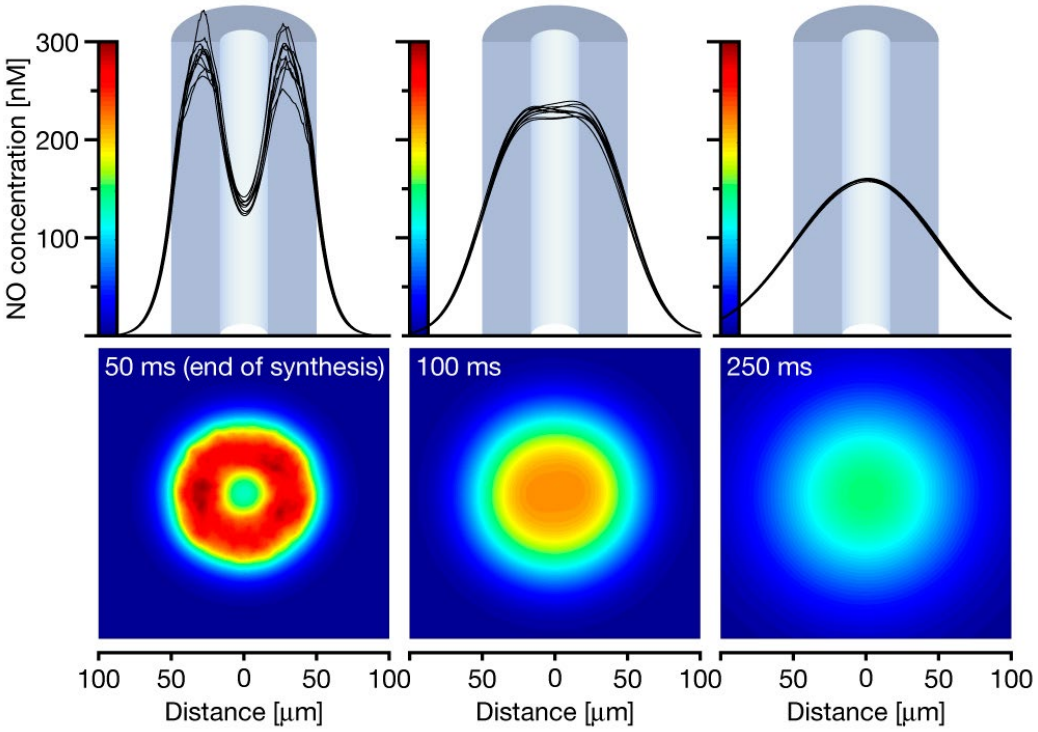

Figure 3: Sparse activity in the nitric oxide (NO) synthase (NOS)-positive outer region of the mushroom body (MB) yields a strong volume signal in the NOS-negative core. The MB is made up of a parallel array of 30,525 Kenyon cell axons represented as 0.5 μm diameter fibres within a tubular region of 99 μm diameter. Axons within the inner 33μm diameter tubular core of the MB are NOS-negative and designated as targets. The remaining 27, 212 NOS-positive axons form an annular 'wall' (shown as a dark gray cut-out in the upper part) around the core and are potential NO sources. In 10 independent runs of the diffusion model, a random 5% of these peripheral NOS-expressing Kenyon cell axons in the tube wall synthesize NO simultaneously for 50 ms. NO half-life is 5 s though results are qualitatively similar for other values. The resultant NO concentrations are shown along a line through the centre of the stalk as 10 superimposed plots (upper part) and, for a single run, in a cross-section through the MB (lower part) at the end of synthesis, and 50 and 200 ms later. Figure used with permission from Ott et al., 2007.

In many instances, however, NO is not produced in one source but in multiple distinct sources. For instance, the mushroom bodies (MB) of the locust (Schistocerca gregaria), are composed of parallel axons of intrinsic neurones (Kenyon cells, KCs) arranged in a tubular structure in which peripheral NOS-positive KC axons form NO-producing zones, the tube 'walls', surrounding central cores of NO-receptive KC axons, which do not produce NO. This segregated architecture requires NO to spread at physiological concentrations up to 60 µm from the outer regions of the tube where it is synthesised to the target axons in the inner NO-receptive cores (Fig 3). Despite NO being produced by small numbers of distinct sets of KCs in response to odours, and that single KCs in the tube walls cannot produce enough NO to generate effective signals in the central targets, they act co-operatively and a reservoir of NO builds up in the NOS-negative cores. During synthesis, NO is highest in the tube walls containing the NOS-positive axons and the resulting concentration gradient drives NO into the centre where it accumulates. About 50 ms after the end of synthesis, the peak concentration is in NOS-negative cores where NO remains locally elevated for several hundred milliseconds. Further, NO signals in the core of the MB are highly invariant for a given percentage of randomly selected active sources. Importantly, the general features of the diffusion dynamics illustrated here are unaffected by the strength or duration of the synthesis pulse, showing that an effective volume signal can be produced by segregated small sources.

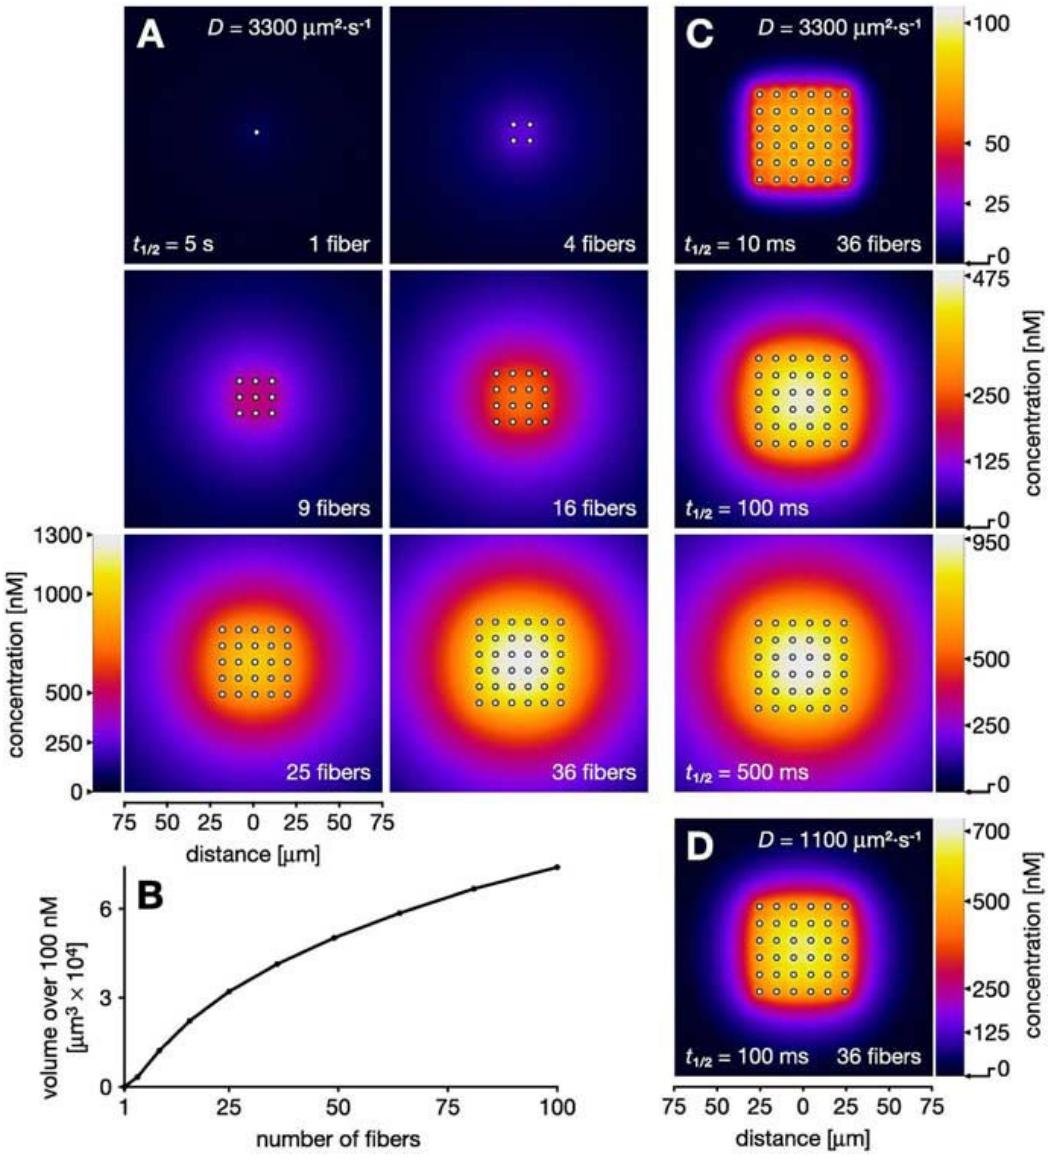

Figure 4: Cooperative volume signals produced by ordered arrays of parallel NO-synthesizing fibers at the end of a 1s burst of NO synthesis. Fiber diameters (2 μm) and spacing (10 μm) approximate that in the locust optic lobe (Elphick et al., 1996). A: NO concentration in slices across increasing numbers of active fibers (white dots). A single fiber is a relatively ineffectual NO source (1 fiber in A), but increasing numbers (4-36 fibers) result in cumulative build-up of NO. B: As the number of fibers increases, so does the volume of tissue over a particular concentration, here 100 nM. NO half-life $t_{1/2}$ is 5s in A and B. C: The cooperative volume signal is robust to variation in NO half-life $t_{1/2}$ within the limits reported in the literature, from top to bottom $t_{1/2}$ *=10, 100 and 500 ms*. D: Strong cooperation is still observed when $t_{1/2}$ is relatively short (100 ms) and D is reduced to 1100μm^2 s^{-1}, one-third of its standard value. Figure used with permission from Philippides et al., 2005a.

Such co-operative volume signalling can again be seen in a model of NO production in the optic lobe of the locust (Fig 4). Here NO is again produced in multiple distinct fine fibers of approximately 2 μm diameter. As with the locust mushroom body, despite a single fiber producing an ineffective concentration, NO from multiple fibers summate to produce effective concentrations over large regions (Fig. 4A, B). This is also the case if sources are arranged irregularly instead of in parallel arrays and if their firing is asynchronous. Finally, the volume signal persists despite large changes to the parameters governing diffusion namely the half-life of NO, $t_{1/2}$, and the diffusion coefficient, D. When half-life is reduced to 100 ms (Griffiths and Garthwaite, 2001), 1⁄50 of the value used in Fig 4A, the concentrations in the target volume fall to only approximately one-third (250–475 vs 700 –1300 nM). Further reduction of $t_{1/2}$ to 10 ms yields a more spiky concentration distribution with 100 nM peaks at the fibers. However, at least 50% or more of this peak concentration is still reached everywhere throughout the target volume. Similarly, while reducing the speed of diffusion dramatically by decreasing D to a third of its standard value and using a relatively short half-life, $t_{1/2}$ *= 100 ms*, means that the spread of the volume signal outside of the source array is limited (compare Fig 4D with bottom middle of 4A), a strong cooperative volume signal is still observed within the array.

The examples in Figures 3 and 4 show that small discrete NO sources can act to produce a co-operative volume signal, despite asynchronous activity and irregular arrangement and spacing. In these contexts, the diffusion process can be seen as a spatio-temporal integrator of neural activity in discrete sources. Interestingly, while the two features seen in large continuous sources still exist: concentrations within the source neurons persist after the end of synthesis and central concentrations can peak after synthesis has ended; there is one key difference. With these dispersed sources, the dynamics of diffusion mean that the concentration quickly becomes quite uniform within and close to the volume containing the NO sources. This means that if the goal of the signalling system is to produce an even concentration within a volume, this is best achieved with a network or plexus of fine NO producing fibers, perhaps giving rise to a different type of volume signal (explored in abstract form in the plexus GasNet model described below and in Philippides et al., 2005b). In turn this work led to further theoretical and empirical investigations that have provided experimental evidence for the kind of cooperative signalling predicted by the models described above (Steinert et al. 2008).

GasNets: NO inspired ANNs

Detailed models of long-range neurotransmitter diffusion in the nervous system are necessarily computationally expensive so the use of such models in the study of whole behaviour generating circuits has so far been beyond the state-of-the-art. Hence to further investigate the functional roles of volume signals and related mechanisms in the generation of behavior, a number of authors have advocated the study of more abstract artificial robot nervous systems incorporating simplified models of volume signalling and related processes (Grand 1997, Husbands et al. 1998, Kondo et al. 1999, Eggenberger et al. 2000, Kondo 2007, Buckley 2008). These systems are computationally tractable and can generate sensorimotor behaviors in real time in simulations or in the real world.

This section describes a style of artificial neural network directly inspired by NO neuromodulation, making use of an analogue of volume signalling. The class of artificial neural networks developed to explore artificial volume signaling are known as GasNets (Husbands et al. 1998). These are essentially standard neural networks augmented by a chemical signaling system comprising a diffusing *virtual* gas which can modulate the response of other neurons. As outlined below, a number of GasNet variants, inspired by different aspects of real nervous systems, have been explored in an evolutionary robotics (Floreano et al. 2008) context as artificial nervous systems for mobile autonomous robots. Evolutionary robotics involves populations of artificial genomes (lists of characters and numbers, the members of which act as 'genes') which encode the structure and other properties of artificial neural networks that are used to control autonomous mobile robots required to carry out a particular task or to exhibit some set of behaviors. Other properties of the robot, such as sensor layout or body morphology, may also be under 'genetic' control. The genomes are mutated and interbred creating new generations of robots according to a Darwinian scheme (i.e. an evolutionary search algorithm) in which the fittest individuals are most likely to produce offspring. Fitness is measured in terms of how well a robot behaves according to some evaluation criteria; this is usually automatically measured. Evolutionary robotics can operate with fewer assumptions about neural architectures and behaviour generating mechanisms than other methods; this means that whole general classes of designs and processes can be explored (Floreano et al. 2008, Vargas et al. 2014). This makes it a particularly attractive technique for synthesizing models in neurobiology, allowing explorations of possible processes and mechanisms as in the work described here (Floreano et al. 2008).

GasNets have been shown to be significantly more evolvable, in terms of speed to a good solution, than other forms of neural networks for a variety of robot tasks and behaviors (Husbands et al. 1998, McHale and Husbands 2004, Smith et al. 2003, Philippides et al. 2005b). They are being investigated as potentially useful engineering tools and as a way of gaining helpful insights into biological systems (Philippides et al. 2000, 2003, 2005b, Husbands et al. 2010).

By analogy with biological neuronal networks, GasNets incorporate two distinct signalling mechanisms, one 'electrical' and one 'chemical'. The underlying 'electrical' network is a discrete time step, recurrent neural network with a variable number of nodes. These nodes are connected by either excitatory or inhibitory links with the output, O_i^t, of node i at time step t determined by the following equation.

$$O_i^t = \tanh\left[k_i^t\left(\sum_{j\in\Gamma_i} w_{ji}O_j^{t-1} + I_i^t\right) + b_i\right]$$

where Γ_i is the set of nodes with connections to node i and $w_{ji} = \pm 1$ is a connection weight. I_i^t is the external (sensory) input to node i at time t, and b_i is a genetically set bias. Each node has a genetically set default transfer function gain parameter, k_i^0, which can be altered at each time-step according to the concentration of diffusing 'gas' at node i to give k_i^t (as described later).

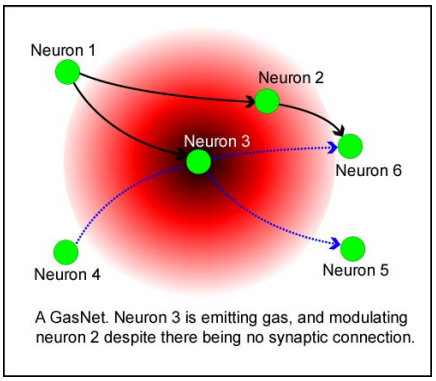

Figure 5: A basic GasNet showing excitatory (solid) and inhibitory (dashed) 'electrical' connections and a diffusing virtual gas creating a 'chemical' gradient.

In addition to this underlying network in which positive and negative 'signals' flow between units, an abstract process loosely analogous to the diffusion of gaseous modulators is at play. Some units can emit virtual 'gases' which diffuse and are capable of modulating the behaviour of other units by changing their transfer functions. The networks occupy a 2D space; the diffusion processes mean that the relative positioning of nodes is crucial to the functioning of the network. Spatially, the gas concentration varies as an inverse exponential of the distance from the emitting node with spread governed by a parameter, r, genetically set for each node, which governs the radius of influence of the virtual gas from the node as described by the equations below and illustrated in Figure 5. The maximum concentration at the emitting node is 1.0 and the concentration builds up and decays linearly as dictated by the time course function, $T(t)$, defined below.

$$C(d,t) = \begin{cases} e^{-2d/r} \times T(t) & d < r \\ 0 & \text{else} \end{cases}$$

$$T(t) = \begin{cases} 0 & t = 0 \\ min(1, (T(t-1) + \frac{1}{s})) & \text{emitting} \\ max(0, (T(t-1) - \frac{1}{s})) & \text{not emitting} \end{cases}$$

where *C(d,t)* is the concentration at a distance *d* from the emitting node at time *t* and *s* (controlling the slope of the function *T*) is genetically determined for each node. The range of *s* is such that the gas diffusion timescale can vary from 0.5 to 0.09 of the timescale of

'electrical' transmission (i.e. a little slower to much slower). The total concentration at a node is then determined by summing the contributions from all other emitting nodes (nodes are not affected by their own emitted gases to avoid runaway positive feedback). The diffusion process is modelled in this simple way to provide extreme computational efficiency, allowing arbitrarily large networks to be run very fast.

For mathematical convenience, in the original basic GasNet there are two 'gases', one whose modulatory effect is to increase the transfer function gain parameter (k_i^t) and one whose effect is to decrease it. It is genetically determined whether or not any given node will emit one of these two gases (gas 1 and gas 2), and under what circumstances emission will occur (either when the 'electrical' activation of the node exceeds a threshold, or the concentration of a genetically determined gas in the vicinity of the node exceeds a threshold; note these emission processes provide a coupling between the electrical and chemical mechanisms). The concentration-dependent modulation is described by the following equation, with transfer function parameters updated on every time step as the network runs:

$$k_i^t = k_i^0 + \alpha C_1^t - \beta C_2^t$$

where k_i^0 is the genetically set default value for k_i, C_1^t and C_2^t are the concentrations of gas 1 and gas 2 respectively at node *i* on time step *t*, and α and β are constants such that $k_i^t \in [-4, 4]$. Thus the gas does not alter the electrical activity in the network directly but rather acts by continuously changing the mapping between input and output for individual nodes, either directly or by stimulating the production of further virtual gas. The general form of diffusion is based on the properties of a (real) single source neuron as modelled in detail in Philippides et al. (2000; 2003). The modulation chosen is motivated by what is known of NO modulatory effects at synapses (Baranano et al. 2001). For further details see (Husbands et al. 1998; Philippides et al. 2005b, Husbands et al. 2010).

When they were first introduced, GasNets were demonstrated to be significantly more evolvable than a variety of standard ANNs on some noisy visually guided evolutionary robotics tasks (Husbands et al., 1998). Typically the increase in evolvability, in terms of number of fitness evaluations to a reliable good solution, was an order of magnitude or more. The solutions found were often very lean with few nodes and connections, typically far fewer than were needed for other forms of ANN (Husbands et al. 1998, 2010). But the action of the modulatory gases imbued such networks with intricate dynamics: they could not be described as simple. Oscillatory sub-networks based on interacting 'electrical' and 'gas' feedback mechanisms acting on different timescales were found to be very easy to evolve and cropped up in many forms, from CPG circuits for locomotion (McHale and Husbands, 2004) to noise filters and timing mechanisms for visual processing (Husbands et al., 1998, Smith et al., 2002). GasNets appeared to be particularly suited to noisy sensorimotor behaviours which could not be solved by simple reactive feedforward systems, and to rhythmical behaviours.

Two recent extensions of the basic GasNet, the receptor and the plexus models, incorporated further influence from neuroscience (Philippides et al., 2005b). In the receptor model, modulation of a node is now a function of gas concentration and the quantity and type of receptors (if any) at the node. This allows a range of site specific modulations within the same network. In the plexus model, inspired by a type of NO signalling seen in the mammalian cerebral cortex (Philippides et al, 2005a), as described earlier in this article, the

emitted gas ‘cloud’, which now has a flat concentration, is no longer centred on the node controlling it but at a distance from it. Both these extended forms proved to be significantly more evolvable again than the basic GasNet. Other varieties include non-spatial GasNets where the diffusion process is replaced by explicit gas connections with complex dynamics (Vargas et al. 2007) and version with other forms of modulation and diffusion (Husbands et al. 2010). In order to gain insight into the enhanced evolvability of GasNets, detailed comparative studies of these variants with each other, and with other forms of ANN, were performed using the robot task illustrated in Figure 6 (Philippides et al. 2005b, Husbands et al. 2010).

Starting from an arbitrary position and orientation in a black-walled arena, a robot equipped with a forward facing camera must navigate under extremely variable lighting conditions to one shape (a white triangle) while ignoring the second shape (a white rectangle). The robot must successfully complete the task over a series of trials in which the relative position and size of the shapes varies. Both the robot control network and the robot sensor input morphology, i.e. the number and positions of the camera pixels used as input and how they were connected into the network, were under evolutionary control as illustrated in Figure 6. The network architecture (including number of nodes) and all properties of the nodes and connections and gas diffusion parameters were set by an evolutionary search algorithm. Because of the noise and variation, and limited sensory capabilities (only very few pixels are used), this task is challenging, requiring robust, general solutions. The gantry robot shown in the figure was used. The evolutionary search algorithm employed a special validated simulation of the robot and its environment to calculate fitness (by measuring behavioural performance) (Husbands et al. 1998, 2010). Neural controllers evolved in this manner generate the same behaviour when downloaded onto the real robot.

The comparative studies revealed that the rich dynamics and additional timescales introduced by the gas played an important part in enhanced evolvability, but were not the whole story (Philippides et al. 2005b; Husbands et al. 2010). The particular form of modulation was also important – multiplicative or exponential modulation (in the form of changes to the transfer function) were found to be effective, but additive modulations were not. The former kind of modulations may well confer evolutionary advantages by allowing nodes to be sensitive to different ranges of input (internal and sensory) in different contexts. The spatial embedding of the networks also appears to play a role in producing the most effective coupling between the two distinct signalling processes (‘electrical’ and ‘chemical’). By exploiting a loose, flexible coupling between the two processes, it is possible to significantly reduce destructive interference between them, allowing one to be ‘tuned’ against the other while searching for good solutions. It has been suggested that similar forces may be at play in spiking networks, where sub-threshold and spiking dynamics interact with each other, which have been evolved to drive vision-based robot behaviours (Floreano et al. 2006; Floreano et al. 2008). In the most successful varieties of GasNet, dynamics, modulation and spatial embedding act in concert to produce highly evolvable degenerate (Tononi et al. 1999) networks.

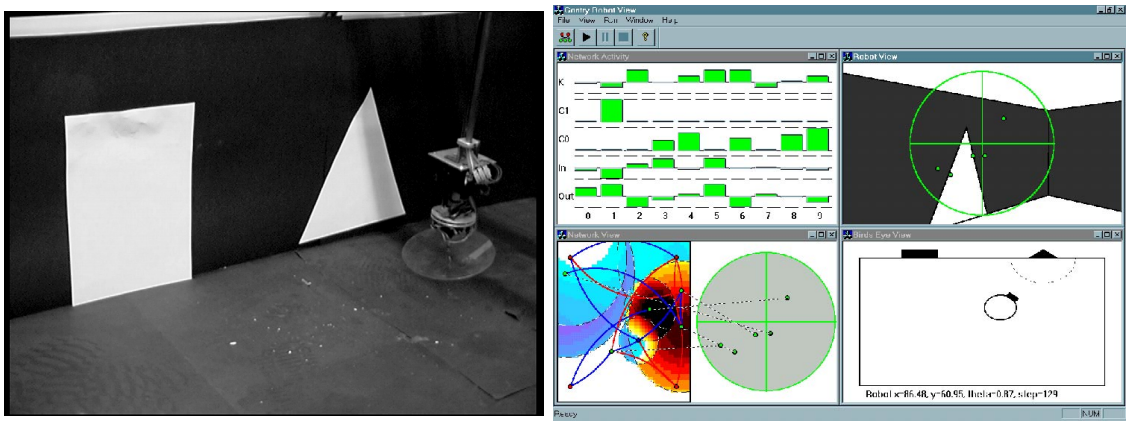

Figure 6: LEFT: The gantry robot. A CCD camera head moves at the end of a gantry arm allowing full 3D movement. In the study referred to in the text, 2D movement was used, equivalent to a wheeled robot with a fixed forward pointing camera. A validated simulation was used: controllers developed in the simulation work at least as well on the real robot. RIGHT: The simulated arena and robot. The bottom right view shows the robot position in the arena with the triangle and rectangle. Fitness is evaluated on how close the robot approaches the triangle. The top right view shows what the robot 'sees', along with the pixel positions selected by evolution for visual input. Only the values of these pixels from the camera image are fed into the neural network as visual input signals. All other pixel values are discarded. In this way the visual sensor morphology is effectively evolved. The bottom left view shows how the genetically determined pixels are connected into the control network whose gas levels are illustrated. The top left view shows current activity of nodes in the GasNet.

Conclusions

Nitric Oxide neuromodulation is an important example of a non-classical neural signalling mechanism which reveals subtle new dimensions to the functioning of the nervous system. Simple connectionist point-to-point notions of information transfer are no longer valid: new understandings and new metaphors are needed.

Detailed models of neural processes incorporating diffusing chemicals are computationally very expensive and to date have been restricted to small numbers of neurons rather than whole behavior-generating neuronal circuits. Computational studies of NO neuromodulation in functioning circuits have therefore been of a more abstract nature and have resulted in new styles of ANN that have direct engineering applications. However, recent advances in parallel computing, including the advent of general purpose GPU computing, mean that we are about to enter an era where larger-scale detailed models will be feasible. This will allow important advances in understanding the functional roles of NO neuromodulation through the careful integration of computational and experimental studies of (biological) neural circuits operating under the influence of NO modulation.

Cross-References/Related Terms

References

Baranano, D., Ferris, C. and Snyder, S. (2001) Atypical neural messengers. *Trends in Neuroscience* **24**(2):99-106.

Brazier, M.A.B. (1961) *A History of the Electrical Activity of the Brain*. Pitman: London.

Bredt, D. S. and Snyder, S. H. (1990). Isolation of nitric oxide synthetase, a calmodulin-requiring enzyme. *Proc. Natl. Acad. Sci. USA* **87**: 682-685.

Bredt, D. S., Hwang, P. M. and Snyder, S. H. (1990) Localization of nitric oxide synthase indicating a neural role for nitric oxide. *Nature* **347**:768-770.

Buckley, C. (2008) A systemic analysis of the ideas immanent in neuromodulation, PhD Thesis, University of Southampton.

Edelman, G.M. and Gally, J.A. (1992) Nitric oxide: Linking space and time in the brain. *Proceedings of the National Academy of Sciences, USA* **89**:11651-11652.

Eggenberger, P., Ishiguro, A., Tokura, S., Kondo, T. and Uchikawa, Y. (2000) Toward Seamless Transfer from Simulated to Real Worlds: A Dynamically—Rearranging Neural Network Approach, *Advances in Robot Learning*, 44-60.

Elphick M, Williams L, O'Shea M (1996) New features of the locust optic lobe: evidence of a role for nitric oxide in insect vision. *J Exp Biol*, 199:2395–2407.

Floreano, D., Epars, Y., Zufferey, J.C. and Mattiussi, C. (206) Evolution of spiking neural circuits in autonomous mobile robots, *Int. J. Int. Syst*. **21**(9), 1005 -1024.

Floreano, D., Husbands, P. and Nolfi, S. (2008). Evolutionary Robotics. In Siciliano, B., & Khatib, O. (Eds.), *Springer Handbook of Robotics* (pp.1423‑1451). Berlin: Springer.

Gally JA, Montague PR, Reeke Jnr GN and Edelman GM (1990) The NO hypothesis: possible effects of a short-lived, rapidly diffusible signal in the development and function of the nervous system. *Proc Natl Acad Sci USA*, 87:3547-3551.

Grand, S. (1997) Creatures: An exercise in creation. *IEEE Intelligent Systems Magazine* **12**(4):4-6.

Griffiths C, Garthwaite J (2001) The shaping of nitric oxide signals by a cellular sink. *J Physiol (Lond)* 536:855– 862.

Husbands, P., Smith, T., Jakobi, N. and O'Shea, M.. (1998) Better living through chemistry: Evolving GasNets for robot control. *Connection Science* **10**(4):185-210.

Husbands, P., Philippides, A., Vargas, P., Buckley,C., Fine, P., Di Paolo, E., and OShea, M. (2010). Spatial, temporal and modulatory factors affecting GasNet evolvability in a visually guided robotics task. *Complexity* **16**(2):35-44

Kondo, T., Ishiguro, A., Uchikawa, Y. and Eggenberger, P. (1999) Autonomous robot control by a neural network with dynamically-rearranging function. In Fourth International Symposium on Artificial Life and Robotics: AROB99, pages 324-329. Oita, Japan.

Kondo, T. (2007) Evolutionary design and behavior analysis of neuromodulatory neural networks for mobile robots control, *Applied Soft Computing* **7**(1):189-202.

J. Lancaster (1994). Simulation of the diffusion and reaction of endogenously produced nitric oxide. *Proc Natl Acad Sci USA*, 91:8137-8141.

J. Lancaster (1996). Diffusion of free nitric oxide. *Methods in Enzymology*, 268:31-50.

J. Lancaster (1997). A tutorial on the diffusibility and reactivity of free nitric oxide. *Nitric Oxide*, 1:18-30.

M. Laurent, M. Lepoivre and J.-P. Tenu (1996). Kinetic modelling of the nitric oxide gradient generated in vitro by adherent cells expressing inducible nitric oxide synthase. *Biochem J*, 314:109-113.

Liu, X., Miller, M., Joshi, M., Sadowska-Krowicka, H., Clark, D., and Lancaster, J. (1998). Diffusion-limited reaction of free nitric oxide with erythrocytes. *Journal of Biological Chemistry*, 273(30):18709–18713.

McHale, G. and Husbands, P. (2004) Quadrupedal locomotion: GasNets, CTRNNs and hybrid CTRNN/PNNS compared. In J. Pollack, M. Bedau, P. Husbands, T. Ikegami, and R. Watson, editors, Proc. 9th Int. Conference on the Simulation and Synthesis of Living Systems (Alife IX), page 106-112. MIT Press.

Montague, P., Gally, J., and Edelman, G. (1991). Spatial signaling in the development and function of neural connections. *Cerebral Cortex*, 1(3):1047–3211.

Mungrue, I. N. and Bredt, D.S. (2004) nNOS at a glance: implications for brain and brawn. *Journal of Cell Science* **117**: 2627-2629.

Ott, S.R., Philippides, A., Elphick, M.R. and O'Shea, M. (2007). Enhanced fidelity of diffusive NO signalling by the spatial segregation of source and target neurones in the memory center of an insect brain. *European Journal of Neuroscience*, 25: 181-190

Philippides, A.O., Husbands, P. and O'Shea, M. (2000). Four-dimensional neuronal signaling by nitric oxide: A computational analysis. *Journal of Neuroscience* **20**(3):1199-1207.

Philippides, A.O., Husbands, P., Smith, T. and O'Shea, M. (2003) Structure based models of NO diffusion in the nervous system. In J. Feng, editor, Computational Neuroscience: a Comprehensive Approach, 97-130. Chapman and Hall/CRC Press.

Philippides, A., Ott, S., Husbands, P., Lovick, T. and OShea, M. (2005a) Modeling co-operative volume signaling in a plexus of nitric oxide synthase-expressing neurons. *Journal of Neuroscience* ***25***(28): 6520‑6532.

Philippides, A. and Husbands, P. and Smith, T. and O'Shea, M. (2005b) Flexible couplings: diffusing neuromodulators and adaptive robotics. *Artificial Life* **11**(1&2):139-160.

Purves, D. (1997) *Neuroscience*. Sinauer.

Smith, T., Husbands, P. and O'Shea, M. (2003) Local evolvability, neutrality, and search difficulty in evolutionary robotics. *Biosystems* **69**:223-243.

Smith, T.M.C., Husbands, P., Philippides, A. O'Shea, M. (2002) Neuronal plasticity and temporal adaptivity: Gasnet robot control networks. *Adaptive Behavior* **10**(3&4):161-184.

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

Stone, J. and Marletta, M. (1996). Spectral and kinetic studies on the activation of soluble guanylyl cyclase by nitric oxide. *Biochemistry*, 35:1093–1099.

Tononi, G., Sporns, O., & Edelman, G. (1999). Measures of degeneracy and redundancy in biological networks. *Proceedings of the National Academy of Sciences of the USA* ***96*:** 3257.

Vargas, P., DiPaolo, E., and Husbands, P. (2007) Preliminary investigations on the evolvability of a non-spatial gasnet model. In e Costa, F. A. et al. (Eds) *Proceedings ECAL'07*, Lecture Notes in Computer Science **4648**, Springer-Verlag: Berlin, pp. 966–975.

Vargas, P., DiPaolo, E., Harvey, I. and Husbands, P. (2014) *The Horizons of Evolutionary Robotics*, Cambridge, MA: MIT Press.

M. Vaughn, L. Kuo and J. Laio (1998a). Effective diffusion distance of nitric oxide in the microcirculation. *Am J Physiol* 1705-1714.

Vaughn, M., Kuo, L., and Liao, J. (1998b) Estimation of nitric oxide production and reaction rates in tissue by use of a mathematical model. *Am J Physiol* 2163-2176.

J. Wood and J. Garthwaite (1994). Model of the diffusional spread of nitric oxide – implications for neural nitric oxide signaling and its pharmacological properties. Neuropharmacology, 33: 1235-1244.

Further Reading