

The Role of Simplifying Models in Neuroscience: Modelling Structure and Function

Dina M. Kronhaus¹ and Stephen J. Eglan²

¹ Computer Laboratory, Cambridge University
dk323@cam.ac.uk

² Cambridge Computational Biology Institute
sje30@cam.ac.uk

Abstract. The adult human brain has around 10^{11} neurons and 10^{15} connections between these neurons, thus forming an incredibly complex network. In this article, we first describe two complementary approaches to modelling brain function, namely simplifying and realistic models. We then demonstrate, by way of two examples, the utility of building simplifying neural models. In the first example, we consider the development of neuronal positioning. In the second example, we investigate the stability of a cortical network under control and perturbed conditions.

Keywords: simplifying neural models, retinal mosaics, cortical compensation, anterior cingulate.

1 Background

This article introduces, by way of examples, the utility of theoretical modelling in understanding aspects of neural circuitry. Perhaps one of the best examples of theoretical modelling aiding our understanding in neuroscience is the seminal work of Hodgkin and Huxley concerning the ionic basis of the action potential (reviewed by Hodgkin, 1958). Since their work over fifty years ago, many researchers have built upon this framework, developing more complex models to account for new experimental findings. This work is an example of a “realistic” model such that the computational model tries to account for all known relevant details of a particular system (Sejnowski et al., 1988). The model is then evaluated by comparing model output with experimental results.

One crucial issue with this kind of modelling is deciding just how much known experimental detail to include in a model. As we now have a lot of information about the details of individual neurons, should we include all those details when we model a neuron? The answer may, perhaps, be yes if our model investigates the behaviour of just one neuron. However, if, for example, we wish to study the collective dynamics of a large population of neurons, modelling each neuron in detail may not be possible, as the system may become too large to be carefully studied: for example, the model may require the specification of too many parameters, or require prohibitively long computation resources. Furthermore,

it may be the case that for a particular study, certain details, although known, may be irrelevant for the model in question.

For example, the Hodgkin-Huxley equations capture the detailed time evolution of an action potential by modelling ion channels and voltage-dependent conductances. However, if the detailed shape of an action potential is not important, but rather we are interested in the *timing* of the action potential, then a simpler integrate and fire neuron may be sufficient for our purposes. In this case, we can refer to the integrate and fire neuron as a “simplifying model” (Sejnowski et al., 1988). It clearly does not attempt to model all the details underlying action potentials, but can be used as the basis for building larger models to study e.g. network dynamics. This distinction between realistic models and simplifying models is obviously not unique to neuroscience modelling. For example, in ecology, the corresponding terms are “mechanistic” versus “phenomenological” models (Nathan and Muller-Landau, 2000). Mechanistic models attempt to simulate mechanisms underlying key behaviours, whereas phenomenological models are more concerned with replicating the behaviours, irrespective of the actual mechanisms that may generate those behaviours.

A valid criticism against simplifying models is often that they are “biologically implausible” and, hence, do not tell us how the brain solves the task. That often is indeed the case. However, the model may still tell us interesting things about other aspects of the problem. For instance, the back-propagation learning algorithm was used to train a multiple layer perceptron to investigate how children might learn to pronounce words (Sejnowski and Rosenberg, 1987). Specifically, the network was trained to associate letters within a word to the corresponding phonetic representation. This is a difficult task to achieve (at least in English) since the context of the letter is important: consider the pronunciation of the letter *i* within the words *bite* and *bit*. The back-propagation learning algorithm was used to adapt the connection strengths within the network to learn this association. This learning algorithm is so-named because of the way that during learning, error signals “back-propagate” from the output layer back to the input layer. This back-propagation of error is unlikely to occur in neural systems, and hence this learning algorithm is validly regarded as biologically implausible. However, if one regards this learning algorithm as simply a way of training a network, we can still explore the properties of the network, and compare them with human performance. For example, during training there is a stage-like progression of behaviour, seen in young children, where an early babbling-phase can be distinguished from the later, more-refined, performance. Analysis of the structure of the internal representation formed by the networks suggests that the network autonomously learns to distinguish vowels from consonants (Sejnowski and Rosenberg, 1987). Therefore, even though the learning rule may be biologically implausible, the model as a whole gives useful insights into how the brain may acquire the ability to pronounce words.

In this article, we review two recent applications of simplifying models to two different problems in neuroscience. The first model investigates an aspect of structural formation of neural circuitry, illustrating how modelling can help us

study early developmental events. The second application considers how models can help us understand functional aspects of adult circuitry.

2 Structural Development of Neuronal Positioning

The human central nervous system is an incredibly complex structure: billions of neurons connect to each other to form complex networks. A key challenge in neuroscience is to understand how such complex networks are generated during development. There are too many neurons and connections for the network to be genetically encoded in the form of “wiring diagrams” that specify which neurons connect to which other neurons. Instead, the nervous system is likely to take advantage of principles of self-organisation to create such circuitry.

One key step in the generation of neural circuitry is for neurons to be appropriately positioned within their target tissue. Once neurons have been generated, they need typically to migrate to their destination layer and move to a particular location, respecting the position of neighbouring neurons. This process is most strikingly observed in the retina (the light-sensing neural structure at the back of the eye). Figure 1 shows an example of the position of all neurons of a particular class within a rectangular field of view. From this figure it is evident that there is some spatial ordering within this population of neurons: cells are not too close to each other, nor are they too far apart from each other. This semi-regular structure is called a “retinal mosaic” due to the way that the cell bodies (and their surrounding dendritic trees, not shown here) “tile” the surface of the retina. This neuronal arrangement serves to ensure that there are cells located throughout the retinal layer (rather than leaving “holes” in the surface), and may help the subsequent wiring of neurons within different retinal layers.

Many developmental mechanisms have been implicated in the formation of retinal mosaics, reviewed by Cook and Chalupa (2000). These vary from relatively early events of neuronal differentiation through to lateral interactions mediated by dendritic interactions. In addition to many experimental approaches to understanding the formation of retinal mosaics, there has been considerable interest in using theoretical modelling approaches to investigate developmental constraints in mosaic formation. To quantitatively compare real and simulated mosaics, here we use the simplest, and most popular, measure of spatial arrangement, the regularity index (RI) (Wässle and Riemann, 1978). The RI is simply the mean of the nearest-neighbour distances divided by the standard deviation of those distances (Figure 1). Informally, the higher the RI, the more regular the spatial distribution of neurons; values less than two typically indicate a random arrangement of neurons.

2.1 The d_{\min} Model

The basic concept underlying the d_{\min} model is that each neuron has a circular exclusion zone surrounding the cell body, which prevents neighbouring cells from coming too close to it. We take a phenomenological approaching by assuming




Fig. 1. Example retinal mosaic and quantification of its spatial organisation. Left: positioning of rat cholinergic amacrine cells (data from Lucia Galli-Resta). Cells drawn assuming $10\ \mu\text{m}$ diameter; scale-bar (bottom left): $50\ \mu\text{m}$. Right: distribution of nearest-neighbour distances of the neurons on the left. The regularity index of this population, 4.2, indicates a mildly regular arrangement of neurons.

that biological mechanisms (which we do not explicitly model) can somehow enforce the exclusion zone. To model a retinal mosaic, we create a region \mathbf{A} of the same size as the real mosaic being modelled. Initially this region is empty; neurons are added to \mathbf{A} using a serial algorithm, positioning neurons one by one into the array until the number of neurons in the model region \mathbf{A} matches the real mosaic. To position a neuron into \mathbf{A} we follow the following steps:

1. Generate a trial neuron position and exclusion zone (x, y, d) . The position of the neuron (x, y) is determined by uniform sampling of \mathbf{A} . The effective diameter of the exclusion zone, d_{\min} is drawn from a truncated Normal distribution with fixed mean and standard deviation (μ, σ) . The Normal distribution is truncated at some lower bound d_{low} so that d cannot be smaller than e.g. the typical cell body diameter.
2. Find the distance of the trial neuron to all other neurons that have previously been accepted into \mathbf{A} . The smallest of those distances is labelled d .
3. If $d < d_{\min}$, the trial neuron is too close to an existing neuron. The trial neuron is thus rejected. Otherwise, if $d \geq d_{\min}$, the trial neuron is accepted.

To model a given retinal mosaic, the only parameters required by this model are the mean and standard deviation of the exclusion zone. These parameters can be determined by trial and error, or by systematic searching over a range of suitable values. Furthermore, the model can also be used to generate randomly distributed neurons within a layer, subject only to the constraint that their cell bodies do not overlap, by setting the mean and s.d. of the d_{\min} model to match the mean and s.d. of observed cell body diameters. Figure 2 compares a real mosaic against two d_{\min} simulations: one where the exclusion zone has been selected to generate patterns similar to those observed, and one where the exclusion zone simply reflects non-overlap of cell bodies. Although in this case the regularity index of the simulated mosaic (6.3) is higher than the regularity index of the retinal mosaic (4.7), as will be discussed later, multiple simulations




Fig. 2. Comparison of a retinal mosaic with two d_{\min} simulations. Left: positions of off-centre beta retinal ganglion cells from cat; data taken from (Wässle et al., 1981). Cell bodies drawn to size, assuming $15 \mu\text{m}$ diameter. Regularity index: 4.7. Middle: matching d_{\min} simulation ($\mu = 130 \mu\text{m}$; $\sigma = 25 \mu\text{m}$). Regularity index: 6.3. Right: d_{\min} simulation with exclusion zone set to only reflect non-overlap of cell bodies ($\mu = 15 \mu\text{m}$; $\sigma = 0.01 \mu\text{m}$). Regularity index: 2.4. (Each sample field $1090 \times 750 \mu\text{m}$).

generate a range of RIs similar to the observed value. By contrast, when the exclusion zone merely prevents overlap of cell bodies (Figure 2 right), the pattern is typically quite disorganised, and both visually and quantitatively distinct from the retinal mosaic: many areas of the sample field are devoid of neurons.

So far, all the d_{\min} model shows us is that some local-acting mechanism is sufficient for generating patterns similar to those observed experimentally. It therefore does not constrain the underlying biological mechanism that may generate such a local exclusion zone (Galli-Resta et al., 1997). This is a limitation of this simplifying model, and we must rely on either experimental results or more detailed theoretical models to inform us. (In this case, both experimental and theoretical evidence suggest the exclusion zone could be the product of lateral migration mediated by dendritic interactions (Eglen et al., 2000).) However, we can now use this d_{\min} model to allow us to ask other questions, outlined next.

2.2 Bivariate Patterning of Retinal Mosaics

Figure 2 shows a sample of off-centre beta retinal ganglion cells (RGCs); these cells respond to the offset of light stimulation. By contrast, there is a complementary group of neurons, the on-centre beta RGCs that respond to the onset of light stimulation. The cell bodies of these neurons occupy the same layer of the retina, generating a bivariate mosaic pattern, as shown in Figure 3, left (data from Wässle et al., 1981). A striking feature of this pattern is that neighbouring neurons are usually of the opposite type, which has led to the question of whether there are interactions between the two cell types during development that generate such a pattern (Eglen and Willshaw, 2002).

To address this question, the d_{\min} model can be extended to generate bivariate patterns. The null hypothesis of the model is that there are no functional interactions between neurons of opposite type. In this case, we generate N_1 type 1 neurons (here the on-centre neurons), and N_2 type 2 neurons (off-centre). When positioning a neuron in the array, e.g. a type 1 neuron, the type 2 exclusion




Fig. 3. Bivariate d_{\min} simulation. Left: bivariate pattern of on- and off-centre beta retinal ganglion cells (on-centre coloured green and off-centre coloured red). Sample field as in Figure 2. Middle: typical output from bivariate simulation. Parameters used: ($\mu_1 = 116 \mu\text{m}$, $\sigma_1 = 20 \mu\text{m}$, $\mu_2 = 130 \mu\text{m}$, $\sigma_2 = 25 \mu\text{m}$, $d_{12} = 9 \mu\text{m}$). Right: quantitative comparison of the regularity index of the retinal mosaic (horizontal red lines) with each of 99 simulations (black dots; dotted black line indicates median). In each case, the regularity index of the real mosaic is within the range generated by the simulations.

zone is ignored, and the trial neuron is rejected if it falls within the exclusion zone of an existing type 1 neuron. Each exclusion zone is again described by a Normal distribution with given (μ, σ) . However, even under the assumption that there are no functional interactions, cell bodies of two neurons of opposite type still cannot overlap, and so a trial neuron is also rejected if the distance to the nearest neuron of opposite type is less than some small value, d_{12} , which typically is around 9–15 μm , matching cell body diameter.

Typical results from this bivariate d_{\min} simulation are shown in Figure 3. The regularity index of either solely the on-centre cells, the off-centre cells, or both types of cells fall within the range observed from 99 runs of the simulation, leading us to accept the null hypothesis that there are no functional interactions between neurons of opposite type, contrary to results from earlier modelling work (Eglén and Willshaw, 2002). A fuller treatment of this problem, using a more general style of model for simulating point patterns (the pairwise interaction point process model) is given elsewhere (Eglén et al., 2005).

3 Investigating Compensation in Cortical Networks

In our second example, we consider a higher level problem: inferring the putative activity dynamics from connectivity maps that describe effective connections between a network of cortical areas performing a cognitive task (Kronhaus and Willshaw, 2006). Again, we take the approach of using a simplified model that allows us to investigate subsequent ideas. In this case, we study the stability of network performance in the presence of global and local perturbations, as well as compensation in altered networks.

3.1 Cortical Interactions Implicated in the Delayed Match to Sample Task

Positron emission tomography (PET) was used to image brain activation during performance of a recognition task, the Delayed Match to Sample (DMTS) task. In this task, subjects were asked to identify a stimulus that was presented earlier in the experiment (Haxby et al., 1995). Here, two conditions were considered. In the first condition, perceptual matching, there were no distracting stimuli between the two presentations of the stimulus to be identified, and just one second delay between the two stimuli. In the second condition, long-delay, four distracting stimuli appeared over an interval of 21 seconds. The stimuli used in this study were pictures of male and female faces. The imaging data under the two conditions were analysed, generating activity maps for each condition. Structural equation models were then applied to estimate (effective) connectivity between key brain regions thought to be involved in these tasks (McIntosh et al., 1996). The path coefficients estimated by McIntosh et al. (1996) were used to guide the generation of the networks studied here (Figure 4, Table 1).




Fig. 4. Human perceptual matching (left) and long-delay (right) networks inferred from PET imaging data. Areas in each hemisphere are identified using the index described in Table 1, which can be used to identify the corresponding Brodmann area. The areas are arranged topographically, respecting their order within the brain. The colour and width of arrows indicates the value and type (excitatory/inhibitory) of connection: thick red (+0.65), thin red (+0.35), thin blue (-0.35) and thick blue (-0.65). Connectivity data derived from McIntosh et al. (1996).

3.2 Modelling Cortical Dynamics

We assume that N brain regions are being modelled. (For the McIntosh data, $N = 22$, eleven of the same brain regions from each hemisphere.) Each brain region i is summarised simply by a real value a_i denoting the overall activity of that area; the vector $\mathbf{a}(t)$ represents the activity of all brain regions at given time t . Cortical interactions (for each of the two conditions: perceptual matching and long-delay) were encoded into an $N \times N$ connectivity matrix \mathbf{W} , with

Table 1. Brodmann Areas used in this study. Left and Right hemisphere (LH/RH) index denote the region number assigned to this area in the model.

Lobe	Brodman Area	Region	LH index	RH index
Frontal	BA46	Middle frontal	1	12
	BA10	Inferior frontal	2	13
	BA47	Ventral inferior frontal	3	14
Temporal	BA21	Middle temporal gyrus	4	15
	BA37	Inferior temporal gyrus	5	16
Visual	BA18v	Fusiform	6	17
	BA19d	Cuneus	7	18
	BA17/18	Cuneus	8	19
Limbic	GH	Hippocampus	9	20
	BA24	Anterior cingulate gyrus	10	21
	BA23	Posterior cingulate gyrus	11	22

positive/negative values denoting excitatory/inhibitory connections, respectively. Given a pattern of activity at time t , the activity at time $t + 1$ is given by:

$$\mathbf{a}(t + 1) = f(\mathbf{W}\mathbf{a}(t)) \quad (1)$$

$$\text{where } f(x) = \frac{1}{1 + \exp(-k(x - \theta))} \quad (2)$$

The sigmoidal function $f(\cdot)$ is applied elementwise to \mathbf{a} to ensure that the activity of each unit stays bounded within $[0, 1]$ (typically $k=10$, $\theta=0.5$.) Given an initial pattern of activity at time 0 (described in the results below), network activity was typically updated for 100 iterations by which time the network dynamics usually converged to a stable pattern.

3.3 Characteristic Behaviour

The basic network behaviour was examined by initialising the network with a small amount of activity, and seeing how the activity propagated. In particular, one brain region (i) was selected to be initially active, whilst the remaining brain regions were silent. Figure 5 shows typical results of the spread and stabilisation of brain activity. Brain activity was normally initialised in BA18 (left hemisphere; $i=6$); this area was chosen as it is part of the visual cortex, and so likely to be activated by visual stimulation. For the perceptual matching network, transient visual activity led to sustained activity in several left hemisphere regions; some activity propagated transiently to the right hemisphere, but this did not persist. By contrast, activation of only BA18 did not lead to persistent activity in the long-delay network (data not shown). Instead, to generate persistent activity, initial input was needed in both BA18 and another area, such as BA37 (Figure 5 right). Visual comparison of the activity patterns in the perceptual matching and long-delay networks shows that although persistent activity

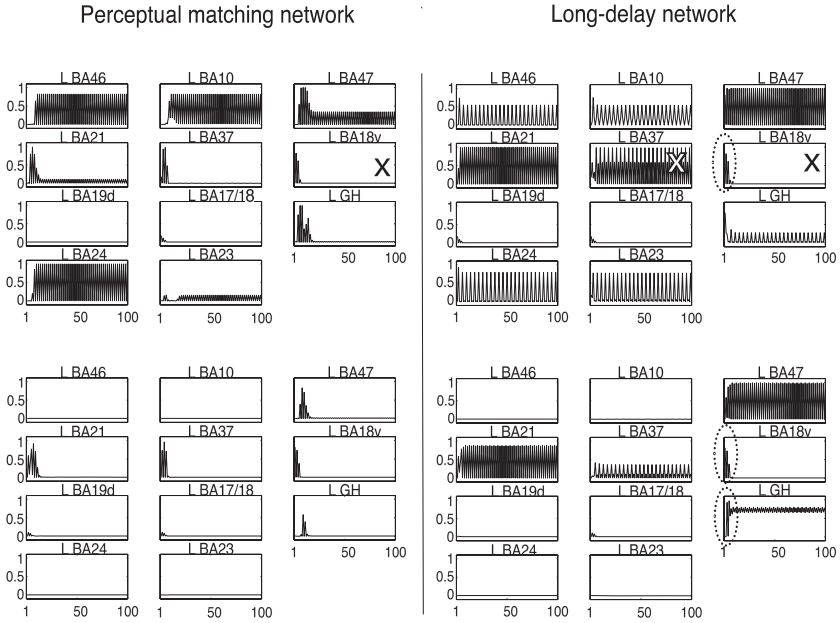


Fig. 5. Characteristic activation in the perceptual matching (left) and long-delay (right) networks. Each graph shows the activity of one brain region for 100 time steps. The legend above each graph denotes the brain region being plotted and its hemisphere (L/R). X within a graph (e.g. L BA18v on left) denotes the activation of this area was set to 1 at time 0; remaining areas were set to 0. This figure adapted from (Kronhaus and Willshaw, 2006), with permission from Oxford University Press.

is generated by both networks, in the long-delay network more activity persists in the right hemisphere. This result could not be predicted by visual comparison of the two networks in Figure 4.

As in Section 2, so far, this model can be regarded as a simplifying model; clearly simplifying the neural activity within each brain region down to a single number is not biologically plausible. However, we believe that we can use the characteristic activity patterns (e.g. from Figure 5) as a signature of activity that we can use as a reference when comparing these networks under different perturbations to represent various experimental conditions, as shown next.

3.4 Compensation for Localised Dysfunction

In this section we show how our network approach can be applied to investigate how cortical networks might adapt to either global or local impairment in a clinical population of interest, namely depressed patients. One key experimental finding suggests that in depressed patients there is an decrease in cortical excitability throughout the entire brain (Shajahan et al., 1999). We have therefore

investigated how characteristic network behaviour is shaped by a reduction in cortical excitation. Surprisingly, characteristic network behaviour is unaffected by global reduction in all excitatory weight connections (inhibitory connections were held constant) by at least 20% (Kronhaus, 2004).

Furthermore, we systematically tested the role of each brain region in sustaining this characteristic behaviour in the presence of reduced global excitation. Specifically, all output connections from a chosen area were reduced by 10%. Characteristic activity was maintained in each case except for reducing connectivity from L BA24 (anterior cingulate). In this case, this local perturbation of the anterior cingulate had the dramatic effect of silencing all activity within around 20 iterations. Since network architecture was constrained by McIntosh et al. (1996), this prominent role of the anterior cingulate was an unexpected, emergent, finding. This novel result from the model echoed findings from the experimental literature reporting reduced neuronal and glial cell density in the anterior cingulate of depressed patients (Drevets et al., 1997).




Fig. 6. Compensation in the long-delay network. Left: recovery by mostly visual areas; Right: compensation by BA37. Arrows indicate the direction in which an area’s outgoing connections were modified by the given percentage. X indicates the areas which were initially stimulated. This figure adapted from (Kronhaus and Willshaw, 2006), with permission from Oxford University Press.

How might the brain compensate for such changes in network connectivity? One possibility is that activity in other areas may increase to compensate for these changes. To test this idea, we focused on alterations to the long-delay network, i.e. overall reduction of all excitatory connections by 20%, together with reducing all outgoing connections from L BA24 by a further 10%; in such a network, activity does not persist. However, it is possible to recover sustained

activity, as shown in Figure 6, if we increase the outgoing connections from other areas. In the first case (Figure 6 left), an increase of outgoing connections from three visually-related areas led to a recovery in activity in both hemispheres. This activity is far from the characteristic pattern (Figure 5) observed in the long-delay network, as many areas that were previously active (e.g. both L and R BA47) were inactive. Further, selective increase of outgoing connections from just one area that plays a prominent role in generating the characteristic activity pattern (10% increase to L BA37) produced sustained activity in many areas, similar to that observed in the default network. Comparing Figure 6 right with Figure 5 right, activation patterns are broadly similar, although L BA46 and L BA10 are noticeably different. Thus, our model predicts that network dynamics in brains where both global and local activity are compromised (e.g. depressed patients) may be restored by intensifying activity elsewhere in the network.

4 Concluding Remarks

This paper has reviewed two recent examples of how simplified models of brain structure and function can be useful to investigate problems in neuroscience. Even though aspects of each model (e.g. the exclusion zone in the d_{\min} model, and the characterisation of brain activity within a region by a single scalar value) do not have a direct biological interpretation, we believe that making such simplifying assumptions in models allows us to use these components within a bigger framework to test particular hypotheses and predict the outcome under novel situations. Furthermore, even though both models exclude fine-grained neurobiological details, they are consistent with the neural underpinnings. Thus, changes in efficacy (excitation or inhibition) are not simply ascribed to excitatory or inhibitory neurotransmitters. Instead, we argue that *dynamic* changes in the network are more appropriate to emulate behavioural phenomena such as those observed in clinical conditions such as depression.

In the context of this Volume, it is interesting to speculate on how these neuroscience models may inspire the next generation of wireless networks. One key feature in both our neuroscience models is robustness, in both design and function. Rather than hard wiring a neural circuit, developmental processes adapt to the local environment and create networks that are robust to environmental differences. Moreover, once the network is compromised, different features of network wiring (e.g. excitatory and inhibitory connections) allow the network to self-adjust, restoring function. These principles allow the nervous system to develop and function under a wide range of conditions. We suggest that such principles may also be of benefit to artificial systems, such as wireless networks.

Acknowledgements. Dina Kronhaus is supported by the Heller Research Fellowship, St Catharine's College, Cambridge. Thanks to Dr Lucia Galli-Resta, Prof. Heinz Wässle and Prof. Randy McIntosh for providing experimental data.

References

- Hodgkin, A.L.: The Croonian lecture: Ionic movements and electrical activity in giant nerve fibres. *Proc. R. Soc. Lond. B* 148, 1–37 (1958)
- Sejnowski, T.J., Koch, C., Churchland, P.S.: Computational neuroscience. *Science* 241, 1299–1306 (1988)
- Nathan, R., Muller-Landau, H.C.: Spatial patterns of seed dispersal, their determinants and consequences for recruitment. *Trends Ecol. Evol.* 15, 278–285 (2000)
- Sejnowski, T.J., Rosenberg, C.R.: Parallel networks that learn to pronounce English text. *Complex Systems* 1, 145–168 (1987)
- Cook, J.E., Chalupa, L.M.: Retinal mosaics: new insights into an old concept. *Trends Neurosci.* 23, 26–34 (2000)
- Wässle, H., Riemann, H.J.: The mosaic of nerve cells in the mammalian retina. *Proc. R. Soc. Lond. B* 200, 441–461 (1978)
- Wässle, H., Boycott, B.B., Illing, R.B.: Morphology and mosaic of on-beta and off-beta cells in the cat retina and some functional considerations. *Proc. R. Soc. Lond. B* 212, 177–195 (1981)
- Galli-Resta, L., Resta, G., Tan, S.-S., Reese, B.E.: Mosaics of Islet-1-expressing amacrine cells assembled by short-range cellular interactions. *J. Neurosci.* 17, 7831–7838 (1997)
- Eglén, S.J., van Ooyen, A., Willshaw, D.J.: Lateral cell movement driven by dendritic interactions is sufficient to form retinal mosaics. *Network: Comput. Neural Syst.* 11, 103–118 (2000)
- Eglén, S.J., Willshaw, D.J.: Influence of cell fate mechanisms upon retinal mosaic formation: a modelling study. *Development* 129, 5399–5408 (2002)
- Eglén, S.J., Diggle, P.J., Troy, J.B.: Homotypic constraints dominate positioning of on- and off-centre beta retinal ganglion cells. *Vis. Neurosci.* 22, 859–871 (2005)
- Kronhaus, D.M., Willshaw, D.J.: The cingulate as a catalyst region for global dysfunction: a dynamical modelling paradigm. *Cereb. Cortex* 16, 1212–1224 (2006)
- Haxby, J.V., Ungerleider, L.G., Horwitz, B., Rapoport, S.I., Grady, C.L.: Hemispheric differences in neural systems for face working memory: a PET rCBF study. *Hum. Brain Mapp.* 3, 68–82 (1995)
- McIntosh, A.R., Grady, C.L., Haxby, J.V., Ungerleider, L.G., Horwitz, B.: Changes in limbic and prefrontal functional interactions in a working memory task for faces. *Cereb. Cortex* 6, 571–584 (1996)
- Shajahan, P.M., Glabus, M.F., Gooding, P.A., Shah, P.J., Ebmeier, K.P.: Reduced cortical excitability in depression. Impaired post-exercise motor facilitation with transcranial magnetic stimulation. *Br. J. Psychiatry* 174, 449–454 (1999)
- Kronhaus, D.M.: Neuroinformatics approaches to understanding affective disorders. PhD thesis, University of Edinburgh (2004)
- Drevets, W.C., Price, J.L., Simpson Jr., J.R., Todd, R.D., Reich, T., Vannier, M., Raichle, M.E.: Subgenual prefrontal cortex abnormalities in mood disorders. *Nature* 386, 824–827 (1997)