

2010 Special Issue

Synaptic rewiring for topographic mapping and receptive field development

Simeon A. Bamford^{a,*}, Alan F. Murray^b, David J. Willshaw^c^a Doctoral Training Centre for Neuroinformatics and Computational Neuroscience, University of Edinburgh, Edinburgh, United Kingdom^b Institute of Integrated Micro and Nano Systems, University of Edinburgh, Edinburgh, United Kingdom^c Institute of Adaptive and Neural Computation, University of Edinburgh, Edinburgh, United Kingdom

ARTICLE INFO

Article history:

Received 12 January 2009

Received in revised form 30 January 2010

Accepted 31 January 2010

Keywords:

Synapse formation

Synapse elimination

Synaptic rewiring

Synaptic plasticity

Spike-timing-dependent plasticity (STDP)

Activity dependent

Activity independent

Integrate-and-fire

Receptive field

Topographic map

Mapping

Map development

Ocular dominance

Topographic refinement

ABSTRACT

A model of topographic map refinement is presented which combines both weight plasticity and the formation and elimination of synapses, as well as both activity-dependent and activity-independent processes. The question of whether an activity-dependent process can refine a mapping created by an activity-independent process is addressed statistically. A new method of evaluating the quality of topographic projections is presented which allows independent consideration of the development of the centres and spatial variances of receptive fields for a projection. Synapse formation and elimination embed in the network topology changes in the weight distributions of synapses due to the activity-dependent learning rule used (spike-timing-dependent plasticity). In this model, the spatial variance of receptive fields can be reduced by an activity-dependent mechanism with or without spatially correlated inputs, but the accuracy of receptive field centres will not necessarily improve when synapses are formed based on distributions with on-average perfect topography.

© 2010 Elsevier Ltd. All rights reserved.

1. Introduction

The development of topographic mappings in the connections between brain areas is a subject that continues to occupy neuroscientists. There have been a number of theoretical investigations on the development of maps through networks with fixed connectivity and changes to synaptic weights (Goodhill, 1993; Miller, Keller, & Stryker, 1989; Song & Abbott, 2001; Willshaw, 2006; Willshaw & von der Malsburg, 1976). Other models have considered the formation and elimination of synapses with fixed weight (Elliott & Shadbolt, 1999). There have been few attempts to include both of these forms of plasticity in a model, i.e. both synaptic weight change and synaptic formation and elimination. Theories of topographic map formation can be divided by the extent to which activity-dependent processes, based on Hebbian reinforcement of the correlated activity of neighbouring cells, are deemed responsi-

ble for the formation of topography. Some assume that activity-independent processes, based on chemoaffinity (Sperry, 1963) provide an approximate mapping, which is then refined (Rut-hazer & Cline, 2004). Others (Willshaw, 2006) show how activity-independent processes may fully determine the basic topography. This paper presents a model of topographic map development, which combines both weight plasticity and the formation and elimination of synapses, as well as both activity-dependent and activity-independent processes. In Section 2, synaptic plasticity and models of topographic map development are briefly reviewed, in order to place the model in context. Section 3 then presents the model, developing it from a general to a more specific form. Section 4 describes the parameterisation of the model for the purpose of simulation, as well as describing a novel approach to analysing map quality. Simulation results are then presented in Section 5, and some interesting consequences of the model are explored. This work is part of a project to implement synaptic rewiring in neuromorphic VLSI (Bamford, Murray, & Willshaw, *in press*), however the results presented here are purely computational.

* Corresponding author.

E-mail address: simeon.bamford@iss.infn.it (S.A. Bamford).

2. Review

2.1. Synaptic plasticity

The term “synaptic plasticity” encompasses the formation and elimination of synapses and changes in their physiological strength. The growth of axons to form synapses between neurons of different brain areas is a prerequisite for the development of topographic maps. Synaptic connections can also be eliminated; a process which has been well studied, for example at the neuromuscular junction. In neonatal mammals, each muscle fibre is innervated by axons from several different motor neurons and then during development most of these synapses are eliminated so that in the adult, each muscle fibre is innervated by only one motor neuron. The process is known to be competitive and various mechanisms have been proposed to account for this (Buffelli, Busetto, Bidoia, Favero, & Cangiano, 2004), including Hebbian mechanisms. The formation and elimination of synapses (as well as the remodelling of axons and dendrites that underlies it) is collectively referred as synaptic rewiring (Chklovskii, Mel, & Svoboda, 2004).

Hebbian mechanisms are those in which changes in the strengths of synapses between neurons are related to the correlation of the neurons’ activity. In models of synaptic plasticity in which coincidence of pre- and post-synaptic activity causes potentiation, additional constraints are typically applied to prevent the run-away potentiation of synapses, such as global normalisation or decay of synaptic strength (Miller et al., 1989). However, Hebb’s (1949) original postulate implies causality; for a pre-synaptic spike to cause a post-synaptic neuron to fire it is necessary that the pre-synaptic spike precede the post-synaptic spike. Bi and Poo (1998) observed that in cultured hippocampal neurons, the potentiation or depression of a synapse was dependent on the temporal order of induced pre- and post-synaptic activity. In this study, pre-synaptic activity preceding post-synaptic activity caused potentiation (and *vice versa*) in accordance with the causality condition, though in other studies the opposite temporal dependence has been observed (Bell, Han, Sugawara, & Grant, 1997). Such Spike-Timing-Dependent Plasticity (STDP), as it has become known, has been investigated extensively in computational neuroscience. Song, Miller, and Abbott (2000) modelled STDP to show that in a neuron whose dendritic synapses implemented STDP, the synaptic weights would diverge into a strong group and weak group, with the effect that groups of synapses whose input spikes were more correlated, i.e. more likely to arrive within a narrow time window of each other, would be preferentially strengthened over synapses whose input spikes were less correlated. Thus, competition between inputs is implemented without the need for additional weight normalisation mechanisms.

There is ongoing debate about the nature of STDP, the molecular mechanisms that give rise to it and its relevance as a candidate mechanism for memory and learning. To give some example of the range of questions that exist: STDP-like behaviour can arise from a synaptic update rule dependent on post-synaptic membrane voltage rather than post-synaptic spikes (Brader, Senn, & Fusi, 2007); there are experiments which indicate that individual synapses may have binary strengths and experience all-or-nothing plasticity events (Petersen, Malenka, Nicoll, & Hopfield, 1998), which are apparently at odds with studies showing that synapses have unimodal distributions; and there are questions over how the contributions of different spike pairs should be combined (Butts, Kanold, & Shatz, 2007; Sjostrom, Turrigiano, & Nelson, 2001). Experiments demonstrating the nature of STDP have typically used *in vitro* preparations (Bi & Poo, 1998) or unrealistic levels of stimulation (Zhang, Tao, Holt, Harris, & Poo, 1998) leading to questions about their relevance to normal cellular processes. STDP models include weight update rules which are dependent on

the weight of a synapse (Gutig, Aharonov, Rotter, & Sompolinsky, 2003); arguably these better model experimental data (Morrison, Aertsen, & Diesmann, 2007).

Notwithstanding the above, weight-independent STDP rules similar to the formalism of Song et al. (2000) have been used to investigate: topographic map formation (Song & Abbott, 2001); the response to latency in inputs (Guyonneau, Van Rullen, & Thorpe, 2005); visual feature map learning (Masquelier & Thorpe, 2007); receptive field reorganisation (Young et al., 2007); learning cross-modal spatial transformations (Davison & Fregnac, 2006) etc. The study of Jun and Jin (2007) is notable as a study of the formation of synfire chains using a combination of STDP and a form of rewiring plasticity. In the present work, STDP is adopted as a form of competitive Hebbian plasticity, in line with the above body of work.

2.2. Topographic maps

A topographic map is an area of the brain where the response to input parameters varies continuously across the area. Where a sheet of neurons in one area (the “source” area) innervates a sheet of neurons in another (“target”) area, the mapping between the areas can be said to be topographic if neighbouring neurons in the target area are (maximally) responsive to the activity of neighbouring neurons in the source area (Udin & Fawcett, 1988). The receptive field of a neuron in a target area is a region of the source area in which stimulation causes activity in the neuron. The topographic maps present in the cortex receive input from both eyes. Although these projections are intermixed, tracing the connections from the eyes to V1 (via LGN) reveals that in many mammals there are alternating stripes in V1 in which cells predominantly receive input from one eye or the other (Hubel, Wiesel, & LeVay, 1977), known as patterns of ocular dominance.

The development of topographic maps between two brain areas requires that axons grow from the source area to the correct target area and then form synapses with neurons in the correct location. This paper excludes consideration of axon guidance to the correct brain area and additionally does not consider the questions of how to target the correct layer (in projections to layered tissue) or how to terminate on the correct part of a dendritic tree. Furthermore, it does not model the growth of neural areas and overall changes in topology during development. Rather it is primarily concerned with the development of receptive fields. The model assumes a mechanism for axons to find the correct topographic position.

The purpose of topographic maps in the brain is a matter of debate. There are suggestions that they may serve to perform dimension reduction, that they may have arisen through reasons of wiring efficiency (Chklovskii & Koulakov, 2004), and that in some cases they serve as a basis for multimodal integration (Holmes & Spence, 2005). A distinction can be drawn between maps which perform some transformation between one layer and another and those which merely relay information without transforming it (Knudsen, du Lac, and Esterly (1987) defined the former as “Computational maps”). The topographic maps created by the model presented in this paper serve mainly to relay information; it may be that broad topographic mappings between two areas serve as a basis for computational transformations at a finer scale, but this will not be demonstrated in this paper.

2.2.1. Activity dependence vs. independence

Models of topographic map formation can be divided into those which require activity of the participant cells (that is to say, electrical or spiking activity) in order to form the map and those which do not. Sperry (1963) proposed that the target area be labelled by two orthogonal gradients of chemicals which incoming axons could use to be guided to the correct location

(this is known as the chemoaffinity hypothesis). Experimental evidence in favour of this hypothesis has followed recently, with the discovery of various candidate marker chemicals, e.g. the ephrin family of membrane-bound molecules and their associated receptors. Prestige and Willshaw (1975) distinguished between two ways the interactions may work. In Type I matching, source neurons have an affinity for a specific patch of the target area, whereas in Type II matching, all source neurons have maximum affinity for one end of the target area. By contrast, other models show how maps can form based on Hebbian reinforcement of the correlated activity of neighbouring cells. Willshaw and von der Malsburg (1976) presented a model in which two 2-dimensional sheets of neurons represent two brain areas between which a topographic mapping should form, with excitatory connections from each neuron of the source layer to each neuron of the target layer. Neighbouring source-layer neurons are simultaneously active and target-layer neurons excite neighbouring or nearby neurons whilst inhibiting distant neurons. With a Hebbian mechanism applied to the weights of the synapses, the mapping develops so that neighbouring source-layer neurons maintained strong connections with neighbouring target-layer neurons whilst other connections become weakened. In order for the maps to be oriented correctly, however, an initial bias in the correct direction is required. Therefore some element of the explanation of topography involves processes more fundamental than the activity-dependent forces invoked in this model. In an alternative activity-dependent mechanism (proposed by Elliott, Howarth, & Shadbolt, 1996; Elliott & Shadbolt, 1998, 1999), activity in axons induces the release of neurotrophic factor from tectal cells, which diffuses locally, causing formation of synapses from nearby axons.

The above models assume that activity in the source layer of a mapping is spatially correlated, that is, that neighbouring or nearby neurons in the source layer are more likely to be co-active. This would be a reasonable assumption for the visual system since positions close together in the visual field are more likely to receive similar stimuli. In mammals, retinotopic projections largely form prior to birth or to the opening of the eyes, when an organism would not receive any visual input, but in developing eyes which have not yet opened there are spontaneous waves of activity (Wong, 1999), which would provide spatial correlations in the absence of genuine visual input. The role of such correlations in forming mappings is a subject of continued debate. Butts et al. (2007) modelled a mechanism by which retinal wave activity could cause retinotopic refinement in LGN assuming a learning rule based on the timing of bursts (which they showed is not incompatible with various forms of STDP); this study is interesting in that the mechanism requires spatio-temporal correlations in the input structure, but this was used simply to reinforce existing preferences laid down by an activity-independent wiring mechanism. The model presented in this paper works on a similar principle.

The relative contribution of activity-dependent and activity-independent mechanisms continues to be a subject of debate. The model presented in this paper includes activity-dependent and activity-independent processes. Specifically, it looks at how an activity-dependent process may affect receptive field spread and ocular dominance pattern formation.

2.2.2. Weight vs. wiring plasticity in topographic map models

A number of models of map development use networks with fixed connectivity where synaptic weights are subject to change (Goodhill, 1993; Miller et al., 1989; Song & Abbott, 2001; Willshaw, 2006; Willshaw & von der Malsburg, 1976). In such models, a synaptic weight of zero is often interpreted as meaning that the synapse has been retracted. Other models have considered the formation and elimination of synapses with fixed weight (Elliott &

Shadbolt, 1999). A mathematical equivalence between such models has been demonstrated under certain conditions (Miller, 1998).

There have been few attempts to include both forms of plasticity in a model. Miiikkulainen, Bednar, Choe, and Sirosh (2005) predominantly used synaptic weight change but supplemented this with a synaptic elimination process, which occurs periodically for all weights below a certain threshold. Willshaw and von der Malsburg (1979) include both synaptic weight change and rewiring; a more recent version of the same model (Willshaw, 2006) used only synaptic weight change, noting that (p. 2708) “A synaptic strength can be interpreted as the probability of a given retinal axon contacting a given tectal cell.”

The model presented in this paper considers both of these processes. Notwithstanding any mathematical equivalence as aforementioned, there are compelling reasons to include both processes. Firstly, since both processes are known to exist and to operate alongside each other, any model which seeks to fully explain topographic map development must ultimately include the two together. Secondly, each synapse which exists has a cost to the organism in terms of the volume of brain it takes up and the energy required to maintain it and to carry it around. In these terms, all-to-all connectivity is prohibitively expensive, but rather only the most useful connections should exist. Chklovskii et al. (2004) argued that the ability of a brain to rewire its connections could substantially increase its capacity to store information (which in the context of topographic maps might be interpreted as the formation of patterned receptive fields which reflect the statistics of the input activity); they raised the question of how the brain might implement search over a space of possible network topologies. The model presented hereafter proposes a possible mechanism, within the domain of topographic map formation. As a related point, even setting aside the constraints of a physical neural network, not implementing all possible connections between neurons can allow greater computational efficiency in simulations.

As noted above, some models only suppose elimination of synapses without formation. Whilst for the development of retinotopy the dominant trend appears to be over-elaboration of axonal arbors followed by pruning, nevertheless new synapses are added. The consideration of how, at least in *Xenopus*, receptive field location may change during development due to the changing ways that retina and tectum areas develop (Gaze, Keating, & Chung, 1974) suggests that synapse formation is a vital element for explaining such phenomena.

Regarding the relationship between the two forms of plasticity, at the neuromuscular junction, a reduction in synaptic efficacy precedes synapse withdrawal (as judged by quantal release probability and post-synaptic receptor density, (Balice-Gordon & Lichtman, 1993; Colman, Nabekemura, & Lichtman, 1997)). One interpretation of this is that the weakness of a synapse is a causal factor in its elimination; this is the assumption on which the relationship between the two forms of plasticity in this model is based.

2.2.3. Lateral interactions in target layer

Linsker (1986a, 1986b, 1986c) used a correlational mechanism similar in style to Willshaw and von der Malsburg (1976) to demonstrate the formation of spatial opponent cells and orientation-specific cells, arranged in columns. Notably, this model assumed that axons from the source layer terminate in random distributions around a pre-defined and immutable location in the target layer; the model presented in this paper uses the same assumption. Linsker then used this framework as a basis for investigating the detailed formation of receptive fields. The model shows how such receptive fields can develop, at least in principle, without any spatial correlations in the input. A similar framework lies

behind LISSOM models (Miikkulainen et al., 2005). A major contribution of these models was to apply the same plasticity rules to synapses implementing lateral interactions between target-layer neurons as was applied to feed-forward synapses (from the source to the target layer). Song and Abbott (2001) applied STDP rules to excitatory though not inhibitory connections, in order to explore how the spike-timing dependence of their learning rule would affect the performance of the lateral connections. They found that excitatory connections could act as a guide towards the development of similar preferences between neighbours and then be weakened at a later point in development. They also found that, given an initial bias towards a desired topology, the STDP rule could act to refine the topology without the need for inhibitory connections.

The model which is presented in this paper assumes short-range excitatory lateral interactions but no longer-range inhibitory interactions. This follows from Song and Abbott's observation, as noted above. Intuitively, the purpose of long-range inhibitory interactions within an activity-dependent model is to ensure that different parts of the map develop different input preferences. However, if an activity-independent process is assumed to create and maintain a broad topography, different regions of the target layer are constrained to be innervated by different parts of the input space and therefore differing input preferences are enforced, rendering long-range inhibitory interactions redundant.

3. Model

3.1. Overview

In this section a model of map formation is presented. The model is intended to be general to the extent that it could apply equally to retinotectal, retinocollicular or retinogeniculate projections, and possibly others. In brief, this model proposes the following:

1. Activity-independent processes fully specify a topographic mapping between a source and target area and guide axons from the source area towards their "ideal" location in the target area, i.e. the location dictated by the topographic mapping. The mechanism that yields this mapping is unspecified; it could be thought of as a type I chemoaffinity mechanism with fixed affinities, though other mechanisms could be inserted. This mechanism acts both initially and continuously.
2. Axon branching leads to formation of synapses over an area surrounding the ideal topographic location (broadly in line with, for example, the innervation of the tectum (McLaughlin, Hindges, & O'Leary, 2003), though simplifying the directional overshoot of axons observed in the superior colliculus of chick and rodent).
3. Competitive Hebbian learning detects correlations in input patterns due to spatial proximity in the source area, such that synapses from more spatially clustered afferent neurons are strengthened at the expense of synapses from neurons which are more distant from other afferents. The effective spread of the receptive fields of target neurons in the source area is thereby reduced; this follows the model of Song and Abbott (2001). To the extent that receptive fields contain input-specific features, such as ocular dominance segregation, then these arise from this process.
4. Preferential elimination of weak synapses allows the reduction of spread to be embedded in the network topology, offering a possible cause for the reduction in axonal arbor spread seen by, for example, McLaughlin, Torborg, Feller, and O'Leary (2003).
5. To the extent that this process continues, with further creation and elimination of synapses, there is the potential for the spread to be reduced further.

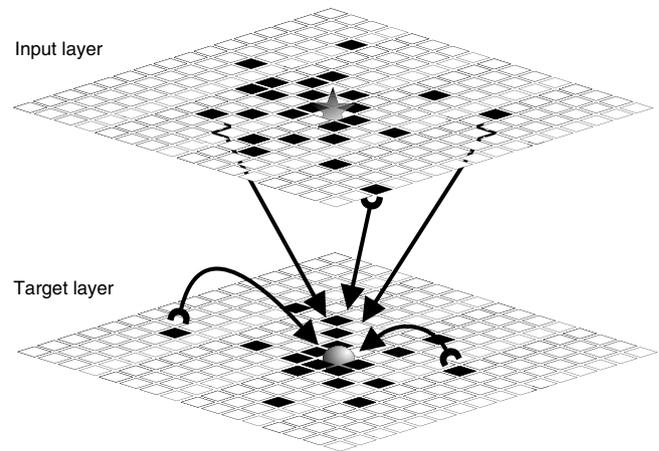


Fig. 1. The upper grid represents the input layer, whose neurons project to the neurons of the target layer (the lower grid). Incoming projections are indicated only for a selected neuron shown as a ball in the target layer. The neurons which make outgoing connections to that selected neuron are coloured black and arrows are indicative of these connections. Note that the selected neuron receives incoming connections from both layers. It may have more than one incoming synapse from the same neuron and it may also make a connection with itself. The coloured neurons in the target layer were sampled from a Gaussian distribution around the selected neuron, whereas the coloured neurons in the input layer were sampled from a Gaussian distribution around the location which corresponds to the position of the selected neuron in the target layer. This location is known as the ideal location and is marked with a star.

The model can be seen as a unique synthesis of existing ideas and elements of models. The primary purpose is to investigate the interplay between two types of plasticity, weight change and rewiring, as they relate to topographic mapping and receptive field development. The phenomena which are focused on are changes in the spread of receptive fields and the development of ocular dominance.

3.2. Details of the model

In this section, greater detail is given, and distinctions are drawn between those properties which are general and those specific ones which have been adopted in order to develop a system which is amenable to tractable simulation and analysis. Thus, a class of models is defined, though only a small subset have been simulated within this project.

There are two layers, the input layer and the target layer, see Fig. 1. Layers are 2D spaces on which neurons are located. Each location in one layer has a corresponding ideal location in the other, such that one layer maps smoothly and completely to the other. In general, the layers could be of any shape and the transformation that maps one layer to the other could be any that does not require discontinuities. For simulation, neural areas are square grids of neurons, the two layers are the same size, periodic boundaries are imposed to avoid edge artefacts.

Each cell in the target layer can receive a maximum number of afferent synapses. It can be said then that each cell has a certain synaptic capacity. For simulation, all target-layer neurons have the same synaptic capacity. A set of connections from one layer to another is referred to hereafter as a projection; this can also refer to a set of connections from one layer to itself. There are two excitatory projections, a feed-forward projection from the input layer to the target layer and a lateral projection from the target layer back to itself. Axons within these projections compete for the synaptic capacity of the target neurons. As noted in Section 2.2.3, for simplicity, inhibitory lateral interactions are not implemented in this model.

It is assumed that an unspecified activity-independent process is capable of guiding the formation of new synapses so that they are

distributed around their ideal locations. A Gaussian distribution is assumed, since a process which is initially directed towards a target site and then randomly branches on its way would yield a Gaussian distribution of terminations around the target site. To implement the Gaussian distributions, where a target neuron has fewer than its maximum number of dendritic synapses, the remaining slots are considered “potential synapses”. At a fixed rate, a synapse from the neurons of the target layer is randomly chosen. If it is a potential synapse, a possible pre-synaptic cell is randomly selected (for the simulations which follow, the last cell to have fired is used as a possible pre-synaptic partner) and synapse formation occurs when:

$$r < p_{form} e^{-\frac{\delta^2}{2\sigma_{form}^2}} \quad (1)$$

where r is a random number uniformly distributed in the range (0, 1), p_{form} is the peak formation probability, δ is the distance of the possible pre-synaptic cell from the ideal location of the post-synaptic cell and σ_{form}^2 is the variance of the receptive field. In other words, a synapse is formed when a uniform random number falls within the area defined by a Gaussian function of distance, scaled according to the peak probability of synapse formation, (which occurs at $\delta = 0$). This is a rejection sampling process.

Lateral connections are formed by the same means as feed-forward connections, though σ_{form} can be different for each projection. For simulation, p_{form} was set so as to allow the same overall probability of formation for each projection. This is because in the absence of a general rule for the relative numbers of feed-forward vs. lateral connections formed, starting with equal numbers of each is a good basis for observing the relative development of these projections.

If the synapse which has been selected from the neurons of the target area already exists (i.e. it is an actual synapse, rather than a potential one) then it is considered for elimination. In general it is proposed that the probability of elimination should be some monotonically decreasing function of weight. For simulation, due to the nature of the chosen learning rule (weight-independent spike-timing-dependent plasticity) which tends to deliver a bimodal weight distribution, the probability of elimination has been simplified to one of two values, with a higher value for synapses with weights below a certain threshold ($P_{elim-dep}$) and *vice versa* ($P_{elim-pot}$).

In general, synapses implement some competitive Hebbian learning rule, such that correlations in inputs to a given target neuron result in preferential strengthening of those synapses at the expense of the strength of other synapses. The neuron model and inputs used should be of types which support the chosen synaptic process. For simulation, the synapses, neurons and type of input are based on the model of Song and Abbott (2001), i.e. with integrate-and-fire neurons with synaptic modulation governed by STDP. The criticisms of STDP raised in Section 2.1 should be borne in mind; the model inevitably loses some generality due to this decision. Most models of map formation use more abstract learning rules. This may be partly motivated by computational constraints but also by a desire for simplicity. Nevertheless the model of Song and Abbott (2001) has been chosen because: (1) its applicability to the modelling of topographic map formation has already been demonstrated, such that existing results can be built on; (2) the use of STDP is attractive because it implements a competitive Hebbian learning rule without requiring additional processes for weight normalisation; (3) STDP is a form of weight change which is known to occur in biological neurons, and as such, has the potential to add biological realism to a model (although lack of knowledge of suitable parameters may undermine this advantage). The detail of the model used, including neuron and synapse dynamics, inputs and initial conditions, is given in algorithm 1.

Algorithm 1 Model summary

There are two layers of the same size, the *input* and *target* layers; each is a square grid of neurons with periodic boundaries, and the ideal location of each neuron in the input layer is the location with the same coordinates in the target area. Each target-layer neuron has the same number of *potential synapses*; these are dendritic locations in which actual synapses may form. Synapses can be with a pre-synaptic neuron from either the input or target layer, including the post-synaptic neuron itself.

Initial conditions: all potential synapses start formed, with conductance g_{max} .

Input: neurons are independent Poisson processes. A stimulus location s is randomly chosen and firing rates are set to $f_{base} + f_{peak} \exp(-d/2\sigma_{stim}^2)$, where d is the distance from s . With a period t_{stim} , s moves to a new random location and the process repeats.

Neuron dynamics (target layer): the membrane voltage V_{mem} is described by:

$$\tau_{mem} \frac{\delta V_{mem}}{\delta t} = V_{rest} - V_{mem} + g_{ex}(t) (E_{ex} - V_{mem})$$

E_{ex} = excitatory reversal potential; V_{rest} = resting potential; τ_{mem} = membrane time constant. Upon reaching a threshold V_{thr} , a spike occurs and V_{mem} is reset to V_{rest} . A pre-synaptic spike at time 0 causes a synaptic conductance at time $t \geq 0$ of $g_{ex}(t) = g_{ex} e^{-t/\tau_{ex}}$ (τ_{ex} = synaptic time constant); this is cumulative for all pre-synaptic spikes.

STDP: a pre-synaptic spike at time t_{pre} and post-synaptic spike at t_{post} modify the synaptic conductance by $g \rightarrow g + g_{max} F(\Delta t)$, where $\Delta t = t_{pre} - t_{post}$ and $F(\Delta t) = A_+ \exp(\Delta t/T_+)$ if $\Delta t < 0$, otherwise $F(\Delta t) = -A_- \exp(-\Delta t/T_-)$, where $A_{+/-}$ are magnitudes and $\tau_{+/-}$ are time constants for potentiation and depression respectively. This is cumulative for all pre- and post-synaptic spike pairs; g is bounded in (0, g_{max}).

Synaptic rewiring: At a fixed rate f_{rew} , a potential synapse is randomly chosen. If it is an actual synapse, the elimination rule is used, otherwise the formation rule is used.

Formation rule: A possible pre-synaptic neuron is randomly selected from either the input or target layer and synapse formation occurs if:

$$r < p_{form} e^{-\frac{\delta^2}{2\sigma_{form}^2}}$$

r = uniform random number in (0, 1); p_{form} = peak formation probability; δ = distance of possible pre-synaptic neuron from ideal location of post-synaptic neuron; σ_{form}^2 = variance of the connection field. p_{form} and σ_{form}^2 may differ based on which layer the possible pre-synaptic neuron is from.

Elimination rule: If the synapse's conductance is below $0.5g_{max}$ it is eliminated with probability $p_{elim-dep}$, otherwise probability $p_{elim-pot}$ is used.

4. Methods

4.1. Experimental parameters

In this section, the process by which the model was parameterised is explained. Parameters for the following simulations are given in Table 1.

Simulations were run with a C++ function, with initial conditions created and data analysis carried out with Matlab. Simulations used a time step of 0.1 ms and rewiring simulations typically settled within 5 min of simulated time. Full-scale simulations were computationally intensive, necessitating the use of relatively small numbers of neurons and synapses; this in turn necessitated a rigorous statistical approach to analysis of results. The size of the grids

representing neural layers was 16×16 (i.e. number of neurons in a layer, $N_{layer} = 256$), enough that discernible patterns of ocular dominance might be observed (c.f. 25×25 (Miller et al., 1989); 32×32 (Goodhill, 1993); 20×20 (Elliott & Shadbolt, 1999)). In determining the maximum fan-in, or number of potential afferent synapses per target neuron, it was found that as the fan-in reduced, the performance of STDP as a correlation detector degraded and the bimodal distributions generated were less extreme; this could be compensated to some extent by building stronger correlational cues into the inputs, as shown by Bofill-i Petit (2005) who achieved strong segregation between just 6 synapses with carefully constructed inputs. Therefore the choice of fan-in represents a compromise between amount of computation necessary and the desire to use more realistic input spike trains. Data is scarce on actual number of dendritic synapses in areas such as the tectum or superior colliculus, so it is difficult to say what a biologically realistic number might be – it may vary greatly between different organisms, brain areas and developmental stages. For most of the simulations in this paper the maximum fan-in (S_{max}) was 32, though it was increased to 64 for one set of experiments.

Data is scarce on appropriate values for the probabilities governing synapse formation and elimination. However, dendritic spines have been imaged extending and retracting over periods of hours compared with others stable over a month or more (Grutzendler, Kasthuri, & Gan, 2002; Trachtenberg et al., 2002). In the simulations, much higher rates were used so that synapses had several chances to rewire during the short periods for which it was tractable to run simulations, while maintaining a large difference between these probabilities (in practice a factor of 180 was used, representing the difference between 4 h and 1 month, i.e. $p_{elim-pot} = p_{elim-dep}/180$). The value of p_{form} works together with the rewiring rate ($f_{rew} = 10^4$ Hz, an arbitrary choice), the number of synapses ($16 \times 16 \times 32 = 8192$), σ_{form} , and the topology of the area to define the actual rate of formation. $\sigma_{form-ff}$ (i.e. for feed-forward synapses) was given a larger value than $\sigma_{form-lat}$ (i.e. for lateral synapses), in line with generic parameters given by Miikkulainen et al. (2005). Once a synapse has been eliminated there is no computational benefit from not being formed again as soon as possible, therefore $p_{form-lat} = 1$, so that if a possible pre-synaptic partner is presented whose ideal location matches the location of the post-synaptic neuron, then the match is accepted. Since $\sigma_{form-ff} > \sigma_{form-lat}$, $p_{form-ff}$ should be less than $p_{form-lat}$ in order to balance the overall probability of synapse formation with each afferent layer; in fact, to achieve this balance:

$$p_{form-ff} = p_{form-lat} \frac{\sigma_{form-lat}^2}{\sigma_{form-ff}^2}. \quad (2)$$

The mean formation rate can then be calculated. $p_{elim-dep}$ was set at half the mean formation rate so that weak synapses would be eliminated half as often as potential synapses became actual synapses, so that the majority of the potential synapses would be formed at any point. In practice, for the parameters given, depressed synapses were eliminated after an average of 33 s whereas strong synapses would only be eliminated with a probability of ≈ 0.05 within a 5 min simulation.

Regarding inputs, the stimulus location changed regularly every 0.02 s. This regularity is a move away from the model of Song and Abbott 2001 in which t_{stim} was chosen according to an exponential distribution; this was a necessary concession to provide stronger correlation cues (i.e. more effective symmetry breaking) given the smaller number of synapses per neuron. A further concession was the more extreme values of the base and peak firing frequencies, f_{base} and f_{peak} . The spread of the stimulus, σ_{stim} , was chosen to be between the values of $\sigma_{form-ff}$ and $\sigma_{form-lat}$ and f_{peak} was set so as to keep the overall mean firing rate at a value, f_{mean} , which was chosen to allow sufficient difference between f_{base} and f_{peak} .

Table 1
Simulation parameters.

Wiring	Inputs	Membrane	STDP
$N_{layer} = 16 \times 16$	$f_{mean} = 20$ Hz	$V_{rest} = -70$ mV	$A_+ = 0.1$
$S_{max} = 32$	$f_{base} = 5$ Hz	$E_{ext} = 0$ V	$B = 1.2$
$\sigma_{form-ff} = 2.5$	$f_{peak} = 152.8$ Hz	$V_{thr} = -54$ mV	$\tau_+ = 20$ ms
$\sigma_{form-lat} = 1$	$\sigma_{stim} = 2$	$g_{max} = 0.2$	$\tau_- = 64$ ms
$p_{form-lat} = 1$	$t_{stim} = 0.02$ s	$\tau_m = 20$ ms	
$p_{form-ff} = 0.16$		$\tau_{ex} = 5$ ms	
$p_{elim-dep} = 0.0245$			
$p_{elim-pot} = 1.36e^{-4}$			
$f_{rew} = 10^4$ Hz			

For the neuron and synapse dynamics, parameters were set starting from parameters given in Song and Abbott (2001). A_+ was increased 20-fold as a concession to limited computational resources for simulations (this should not qualitatively change the model since many plasticity events are still needed to potentiate a depressed synapse). Then key parameters were varied, in order to maintain key conditions, which were: The total weight should be approximately 50% of the maximum possible; the average target neuron firing rate should approximately match the average input firing rate; and the total weight of lateral synapses should roughly match the weight of feed-forward ones. The parameters which were varied are as follows. The peak synaptic conductivity, g_{max} , was varied, since this affects the amount of stimulus the neurons receive and thus their firing rates. The ratio of time constants for depression and potentiation, τ_-/τ_+ , was varied, since this affects the relative weights of feed-forward and lateral synapses, (since correlated feed-forward synapses benefit less from symmetry breaking when the ratio increases); in practice τ_+ was held constant whilst τ_- was varied, as in Song and Abbott (2001). The ratio of depression to potentiation, $B = A_- \tau_- / A_+ \tau_+$, was varied, since this affects the balance of weights; in practice, A_- was treated as the free parameter in order to vary B , however B is quoted, since its meaning is more intuitive. In the interests of simplicity, B was constrained to having the same value for different projections, feed-forward vs. lateral. In practice, it was difficult to find a good set of parameters, since they are interdependent. For example, varying the spike rate changes the balance of the weights, and *vice versa*. Moreover, a single set of parameters inevitably leads to different results depending on the nature of the inputs and depending on whether rewiring was implemented, so there are inevitably confounding factors when attempting to compare these different cases.

Initial placement of synapses was performed by iteratively generating a random pre-synaptic partner and carrying out the formation rule. Feed-forward and lateral connections were placed separately, up to their initial number of 16 synapses each.

4.2. Analysing topographic map quality

For calculating the centre of the receptive field for each target cell (hereafter referred to as the preferred location of a target cell), the use of the centre of mass measure as in Elliott and Shadbolt (1999) would be erroneous. The space is toroidal but, as the authors noted, the centre of mass is always calculated relative to perfect projections. Therefore the calculation of preferred location would be skewed by the choice of reference point from which synapses' coordinates are measured. This bias has been avoided by the novel method of searching for the location around which the afferent synapses have the lowest weighted variance (σ_{aff}^2), i.e.:

$$\sigma_{aff}^2 = \arg \min_{\bar{x}} \frac{\sum_i w_i |\bar{p}_{xi}|^2}{\sum_i w_i} \quad (3)$$

where i is a sum over synapses, \bar{x} is a candidate preferred location, $|\bar{p}_{xi}|$ is the minimum distance from that location of the afferent

Table 2

Summary of simulation results: case 1: rewiring and input correlations; case 2: input correlations and no rewiring; case 3: rewiring and no input correlations.

Case	1	2	3
Target neuron mean spike rate	24.7	17.4	10.5
Final mean number of feed-forward incoming synapses per target neuron	14.1	NA	12.5
Weight as proportion of max for the initial number of synapses	0.60	0.36	0.33
Mean $\sigma_{\text{aff-init}}$	2.36	2.36	2.36
Mean $\sigma_{\text{aff-fin-con-shuf}}$	2.32	NA	2.32
Mean $\sigma_{\text{aff-fin-con}}$	1.95	2.36	2.17
p (WSR $\sigma_{\text{aff-fin-con}}$ vs. $\sigma_{\text{aff-fin-con-shuf}}$)	2.4×10^{-25}	NA	5.0×10^{-6}
Mean $\sigma_{\text{aff-fin-weight-shuf}}$	1.88	2.10	1.99
Mean $\sigma_{\text{aff-fin-weight}}$	1.70	1.98	1.95
p (WSR $\sigma_{\text{aff-fin-weight}}$ vs. $\sigma_{\text{aff-fin-weight-shuf}}$)	2.7×10^{-27}	8.7×10^{-6}	0.028
Mean AD_{init}	0.78	0.78	0.78
Mean $AD_{\text{fin-con-shuf}}$	0.89	NA	0.90
Mean $AD_{\text{fin-con}}$	0.83	0.78	0.93
p (WSR $AD_{\text{fin-con}}$ vs. $AD_{\text{fin-con-shuf}}$)	0.31	NA	–
Mean $AD_{\text{fin-weight-shuf}}$	0.92	1.36	1.21
Mean $AD_{\text{fin-weight}}$	0.95	1.58	1.34
p (WSR $AD_{\text{fin-weight}}$ vs. $AD_{\text{fin-weight-shuf}}$)	0.48	0.0012	–

for synapse i and w_i is the weight of the synapse (if connectivity is evaluated without reference to weights, synapses have unitary weight). This has been implemented with an iterative search over each whole number location in each dimension and then a further iteration to locate the preferred location to 1/10th of a unit of distance (the unit is the distance between two adjacent neurons). Note that in the non-toroidal case this location is equivalent to the centre of mass, as used in Goodhill (1993).

Having calculated the preferred location for each neuron in the target layer, the mean of the distance of the preferred location from the ideal location was taken to give a mean Absolute Deviation (AD) for the projection. By reporting both mean AD and mean σ_{aff} for a projection there is a basis for separating the spread of the receptive fields from the deviation of their preferred locations from their ideal locations. However AD and σ_{aff} are both dependent on the numbers and strengths of synapses and these can change during development. Therefore to observe the effect of the activity-dependent development mechanism irrespective of changes in synapse number and strength, comparison was made in two ways. Firstly, for evaluating change in mapping quality based only on changes in connectivity without considering the weights of synapses, a new map was created by taking the final number of synapses for each target neuron and randomly placing them in the same way as the initial synapses were placed. σ_{aff} and AD were then calculated for each neuron in each of the maps and the means of these (i.e. mean σ_{aff} and mean AD) were compared, applying significance tests between the values of two populations of neurons, i.e. all the neurons on the final map vs. all those on the reconstructed map. Having established what effect there was on connectivity, the additional contribution of weight changes was considered, by creating a new map with the same topology, taking the final weights of synapses for each target neuron and randomly reassigning these weights amongst the existing synapses for that neuron. The two maps were then compared as described above.

5. Results and discussion

Three main experiments were carried out: case 1 had both rewiring and input correlations, as described in Section 3; case 2 had input correlations but no rewiring; case 3 had rewiring but no input correlations (i.e. all input neurons had rate f_{mean}). The results are given in Table 2.

For comparisons, mean σ_{aff} and mean AD were each calculated for the feed-forward connections of the following networks: (a) the initial state (with all weights initially maximised) – these results are suffixed “*init*”, i.e. mean AD_{init} ; (b) the final (“*fin*”) network with weights not considered but only connectivity (“*con*”) with all synapses weighted equally, i.e. mean $AD_{\text{fin-con}}$; (c) for comparison

with mean $AD_{\text{fin-con}}$, the final number of synapses for each target neuron, randomly placed (“*shuf*”) in the same way as the initial synapses (not applicable for simulations with no rewiring), i.e. mean $AD_{\text{fin-con-shuf}}$; (d) the final network including weights, i.e. mean $AD_{\text{fin-weight}}$; (e) for comparison with mean $AD_{\text{fin-weight}}$, the final connectivity for each target neuron with the actual weights of the final synapses for each target neuron randomly reassigned amongst the existing synapses, i.e. mean $AD_{\text{fin-weight-shuf}}$. Results were compared using Wilcoxon Signed-Rank (WSR) tests on AD and σ_{aff} for incoming connections for each target neuron over the whole target layer for a single simulation of each of the two conditions under consideration.

5.1. Receptive field spread and the effect of rewiring

The effect of rewiring can be seen by comparing case 1 (with rewiring) and case 2 (without rewiring). Considering topology change, in case 1 mean $\sigma_{\text{aff-fin-con}}$ drops to 1.95, c.f. 2.32 for mean $\sigma_{\text{aff-fin-con-shuf}}$; this drop is significant. In case 2 mean $\sigma_{\text{aff-fin-con}}$ is constrained to remain at mean $\sigma_{\text{aff-init}} = 2.36$. Considering weight change, in case 1, mean $\sigma_{\text{aff-fin-weight}}$ drops to 1.70, c.f. 1.88 for mean $\sigma_{\text{aff-fin-weight-shuf}}$. In case 2, mean $\sigma_{\text{aff-fin-weight}}$ drops to 1.98, c.f. 2.10 for mean $\sigma_{\text{aff-fin-weight-shuf}}$. Both drops are significant.

Mean $\sigma_{\text{aff-fin-weight}}$ appears to be lower in case 1 than case 2. It is not possible to say for sure that this superior reduction of variance is due to the effect of the rewiring mechanism because the different numbers and weights of final synapses in each case make a comparison impossible. However, there is a good reason to believe that this is so: the drop in mean $\sigma_{\text{aff-fin-con}}$. This drop on its own indicates that the rewiring mechanism has helped to reduce variance and would also lay the groundwork for different final measures of σ_{aff} when weights are considered.

It can be seen then that (a) the weight-changing learning rule causes some reduction in the variance of the receptive fields, and (b) when the rewiring mechanism is applied, the network topology develops such that a variance reduction can be observed in the placement of the synapses, irrespective of their weight. Since the rewiring mechanism on its own can only generate synapse distributions according to the variance used by the formation rule it has no means to reduce this variance except the influence from the effect of the weight-change mechanism, whereby outlying synapses are weakened and become subject to preferential elimination. Thus, the variance reduction is caused by the weight-change mechanism and becomes embedded in the network topology as a result of the rewiring mechanism.

It can also be seen qualitatively that the effect of rewiring is to embed in the connectivity of the network input preferences which arise through the weight changes mediated by the learning

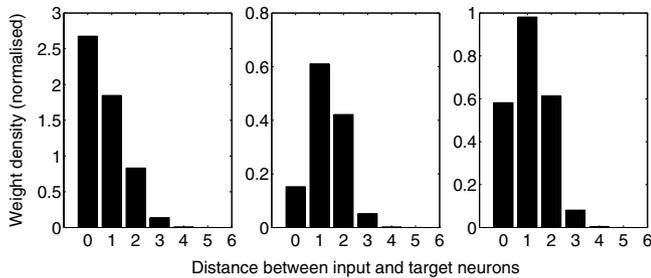


Fig. 2. Normalised weight density of incoming lateral synapses (weight/unit area; y-axis) radially sampled and interpolated at given distances of pre-synaptic neuron from post-synaptic neuron (x -axis), averaged across population. Left: Initial connectivity (weights maximised); middle: final connectivity, weighted by final synaptic weights; right: final connectivity, not considering weights, i.e. each synapse is considered to have unity weight. Autapses (those synapses with a neuron at zero distance) are disproportionately depressed and this weakness leads to preferential elimination.

rule. STDP favours causal inputs with the lowest latency and local excitatory lateral connections tend to lose the competition with excitatory feed-forward connections as they have a higher latency (Song & Abbott, 2001). The extreme of this effect can be seen in synapses from a target neuron back to itself (“autapses”). The placement rule allows these synapses to form, but they only ever receive a pre-synaptic spike immediately following a post-synaptic spike and therefore they have a strong tendency to be depressed by the learning rule. Fig. 2-left shows the initial density of incoming lateral synapses from pre-synaptic partners at given distances out from the post-synaptic neuron. It can be seen that the average neuron receives more synapses from itself (those with zero distance of pre-synaptic neuron from post-synaptic neuron) than from any of its closest neighbours. Fig. 2-middle shows the final distribution where synapses are weighted. The autapses have been depressed much more than their neighbours. Fig. 2-right shows the final distribution only considering numbers of synapses and not their weights. The proportion of autapses to lateral synapses with neighbours has reduced from the initial state, due to the preferential elimination of the weak autapses.

As a further demonstration of the effect of rewiring, a simulation was carried out with the input neurons divided into two groups, mimicking the effect of binocular inputs. The groups were interspersed in a chequered pattern, i.e. each input neuron was in the opposite group to its 4 orthogonally adjacent neurons; the stimulus location switched between the two groups every time it changed. To keep the overall input rate the same, the peak firing rate was doubled. These input patterns are referred to as binocular inputs, as opposed to the monocular inputs used previously. In Fig. 3, each raster is an ocularity preference map, wherein each cell represents a target neuron, and is shaded on a scale from white to black according to the (weighted) proportion of its afferent synapses which are from one of two separately intra-correlated input spaces interspersed in the input space. Given the small models used, the visual differences are not very striking, therefore a quantitative measure of ocularity preference was calculated. The ‘ocularity’ of a target neuron, considering the weights of its incoming synapses, is defined as:

$$\text{Ocularity} = \left| \frac{\sum_{i=1}^{n_1} w_{i1} - \sum_{i=1}^{n_2} w_{i2}}{n} \right| \quad (4)$$

where n_x is number of its incoming synapses from the x th input space (where the input spaces are arbitrarily numbered from 1 to 2), n is the total number of incoming synapses from both input spaces, and w_{ix} is the weight of the i th synapse from the x th

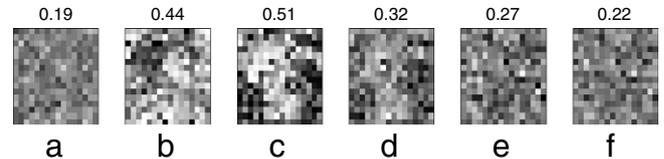


Fig. 3. Ocular preference maps for a set of three experiments. Above each raster is the mean ocularity measure, as defined in Eq. (4). (a): initial connectivity for all experiments (weights maximised); (b): result with binocular inputs but no rewiring, showing the unchanged connectivity weighted by final synaptic weights; (c–d) results with binocular inputs and with rewiring; (c) final connectivity weighted by final synaptic weights; (d): final connectivity, not considering weights, i.e. each synapse is considered to have unity weight; (e–f) results with rewiring but with monocular inputs, showing as (c–d) respectively. The pattern which developed in the weights in (c) has become embedded to some extent in the connectivity (d).

input space. Where ocularity is defined based only on connectivity without considering the weights of synapses, the same expression is used with all weights set to unity. Ocularity therefore gives a real value in the range $[0, 1]$, where low values indicate a balanced input from both input spaces and high values indicate a strong preference for one input space or the other.

Fig. 3(a) shows the initial preferences of each target neuron for input neurons in the two groups. (c) shows the final ocularity preference map where synapses are weighted. Although the space used was too small and the result of the learning rule with a small number of synapses too random for familiar striped ocularity dominance patterns to emerge, ocularity preference zones can be seen. This pattern is reflected in the final map of connectivity in (d), where synaptic weights are not considered. For comparison, the results in (b) are from an experiment which differed only in that rewiring was not performed; in this case, although a pattern of ocularity preference arose in the weights of the connections to a similar extent, this pattern could not be transferred to the network topology (thus the final connectivity is as in (a)). This, therefore, is another example of weight patterns caused by input activity becoming embedded in connectivity patterns by the rewiring mechanism. (e) and (f) are the final ocularity preference maps for weights and connectivity respectively for a control experiment with monocular inputs (in fact, this is case 1 from Table 2). To demonstrate that the apparent increase in ocularity in the final connectivity map in the rewiring case with binocular inputs (d) is not just an artefact of the simulation, a shuffled version of the map was created, in which all the pre-synaptic neurons from all the synapses were randomly reassigned amongst all the existing synapses, whilst the same number of dendritic synapses was maintained for each target neuron. The ocularity was then recalculated for each of the target neurons and this set was compared to the ocularities of the target neurons in the unshuffled map, using a WSR significance test. The change in ocularity in the final connectivity map in the rewiring case with binocular inputs (d) is significant ($p = 7.5 \times 10^{-5}$), whereas for comparison, the change in ocularity in the final connectivity map with rewiring but with monocular inputs is not ($p = 0.98$). Although the mean ocularity in (b) is higher than in (c), conclusions should not be drawn from this, as the numbers and overall weights of incoming synapses may independently affect the ocularity measure, thus a direct comparison is not possible.

5.2. Receptive field centres

Considering the effect of the algorithm on mean AD, in case 2 mean $AD_{fin-weight}$ is significantly increased c.f. mean $AD_{fin-weight-shuf}$. In case 1 the corresponding change is not significant. In case 1 the drop in mean $AD_{fin-con}$ c.f. mean $AD_{fin-con-shuf}$ is not significant.

The basic action of weight-independent STDP on a set of incoming synapses for a single neuron is to deliver a bimodal weight

distribution (Song & Abbott, 2001). Where there are input correlations these cause the more correlated inputs to be maximised and the less or uncorrelated inputs to be minimised. The effect of both the input correlations and the local excitatory lateral synapses on each individual incoming connection field then should be to cause a patch of neighbouring synapses to become potentiated and for outliers from this patch to be depressed. This could be more simply thought of as choosing a subset of the synapses. Ideally the subset which is chosen will be the subset which is most tightly clustered, in other words, the subset with the lowest variance. This is also true if the sample variance measure is used instead of the population variance. It can be proven that for samples drawn from normally distributed data, the sample variance is independent from the sample mean. That is to say, the centre of mass of the tightest cluster is no more likely to be located towards the ideal location than if the same number of synapses were placed randomly according to the initial distribution. This is only true for samples drawn from a normal distribution; however, in this case the population of synapses from which the sample is drawn has a finite number of data points and therefore only approximates to a normal distribution. The effect of this is to introduce an additional error in AD , such that the final value of AD is likely to be slightly higher than if the corresponding number of synapses is distributed according to the initial distribution. This increase though is only slight, such that it does not necessarily pass the significance tests (as the number of afferent synapses for a neuron increases, this increase should tend to zero). Rewiring cannot be expected to do anything to eliminate this error since it can only enhance existing trends.

The result of the learning rule is not therefore to drive the preferred location towards the ideal. This suggests that, given the assumptions in this model, any improvement in the quality of topography, as judged by reduction in the distance of preferred locations from ideal locations (as opposed to reduction in the spread of receptive fields) observed in the development of maps in biology, is likely to be due to such activity-independent mechanisms as exist.

The final results in this respect are not inconsistent with biological topographic maps. In V1 receptive field centres of cells within a single cortical column have been found to be distributed randomly with a standard deviation which is comparable to the areas of the receptive fields (thus, in the terms used here, comparable to the spreads of the receptive fields, as judged by the standard deviation) (Creutzfeldt, Innocenti, & Brooks, 1974; Hubel & Wiesel, 1968).

5.3. The role of input correlations

Considering the role of input correlations, in case 3 (rewiring but no input correlations) mean $\sigma_{\text{aff-fin-con}} = 2.17$, vs. 2.31 for mean $\sigma_{\text{aff-fin-con-shuf}}$; this is significant. Mean $\sigma_{\text{aff-fin-weight}} = 1.95$ vs. 1.99 for mean $\sigma_{\text{aff-fin-weight-shuf}}$; this is significant.

The slight drop in mean $\sigma_{\text{aff-fin-weight}}$ is a sufficient cue to drive the narrowing of the incoming connection fields, as evidenced by the drop in mean $\sigma_{\text{aff-fin-con}}$. It was shown (Linsker, 1986b; Miikkulainen et al., 2005) that functional architecture can form in the absence of any input except uncorrelated random noise. Here a complementary result can be seen in which receptive field spread reduces without input correlations. However, the mechanisms are different, since Linsker's result was described in terms which require lateral inhibition, which is not present in this model. Rather there are two possible mechanisms for this. Firstly, a target neuron may receive two or more synapses from a single input neuron. These synapses will always be correlated and are likely to reinforce each other, becoming more likely to be amongst the potentiated neurons and thus narrowing the spatial variance. Secondly, a spike from a single input neuron will excite a given target neuron and any other of its neighbours which have a synapse from that input. Thus the neuron will also tend to receive some excitation from

lateral connections because of that spike. The smaller range of $\sigma_{\text{form-lat}}$ should selectively enhance the input from a smaller range of locations, leading to a reduction in variance.

5.4. The effect of rewiring on pattern stability

It is natural to enquire about the utility of embedding a learnt pattern in the connectivity of the network. Although a comprehensive answer will not be provided, a further set of experiments with ocular preference maps provides a suggestion. Billings and van Rossum (2009) investigated the persistence of learned patterns stored in weights, in feed-forward networks with plasticity governed by STDP. One of their theoretical findings was that, in the case of weight-independent STDP as used in this paper, the amount of time for which a learnt pattern will remain stable (as measured by the proportion of synapses that stay in the potentiated or depressed state) is exponentially dependent on the inverse of the magnitude of conductance change. In the experiments presented here, the magnitude of conductance change ($A_{+/-}$) was raised in order to reduce simulation time, as stated in Section 4.1, thus the weights change at unrealistically high speeds. Rewiring speeds are also unrealistically fast but, importantly, synapses rewire at a lower rate than they undergo weight plasticity, as appears to be the case in biological neural networks. A pair of experiments were carried out in which binocular inputs were used, as in the ocular preference experiments in Section 5.1. One experiment was with rewiring and the other without. In these experiments, double the number of synapses were used (i.e. 64 in total), and some parameters were modified to optimise for the production of ocular dominance segregation ($g_{\text{max}} = 0.1$, $p_{\text{elim-dep}} = 0.01$, $p_{\text{elim-pot}} = 5.56 \times 10^{-5}$, $\sigma_{\text{form-ff}} = 1.99$, $\sigma_{\text{form-lat}} = 2.49$, Initial number of *ff* synapses per target neuron = 25, Initial number of *lat* synapses = 39). Both experiments had exactly the same initial conditions and exactly the same input spike trains. The results are shown in Fig. 4. Without rewiring, ocular preference patterns form but they are quickly replaced with other patterns as the inputs change. With rewiring, patterns in the weights of connections are initially changeable but as preferences in the weights start to be transferred to the pattern of connections, the changeability reduces. From around 90s onwards, the broad pattern in the weights remains the same with only gradual drift of the pattern; meanwhile the strength of the pattern embedded in the connectivity continues to increase, until it visually reflects the pattern in the weights. In this demonstration, the patterns that form are arbitrary, since there is nothing in the inputs to bias the target neurons to form a preference towards one input space rather than the other. Nevertheless this serves as a demonstration of the potential utility of synaptic rewiring as a means of prolonging the retention of learnt patterns.

5.5. Limitations of the model and future directions

Although a reduction in the spread of receptive fields is seen, it is on a small scale. It is possible that the small numbers of synapses limit the effect. Though figures are not available for the results of Song and Abbott (2001, figure 6), the variance reduction achieved appears to be larger (with 200 feed-forward connections cf. ≈ 16 in these simulations, albeit in a 1D mapping scenario). They also noted that in their simulations, altering the range of the spatial correlations could affect the tightness of the final projection. Quantitative information on the variance reduction seen in biology is sparse. However, whilst axonal arbor spread reduces quite substantially in rodents (McLaughlin, Torborg et al., 2003) and chicks, in fish and amphibians it does not necessarily reduce in absolute size, only relatively to the area of the tectum, which expands during development (McLaughlin, Hindges et al., 2003). Thus, it may be counter-productive to judge the observed variance

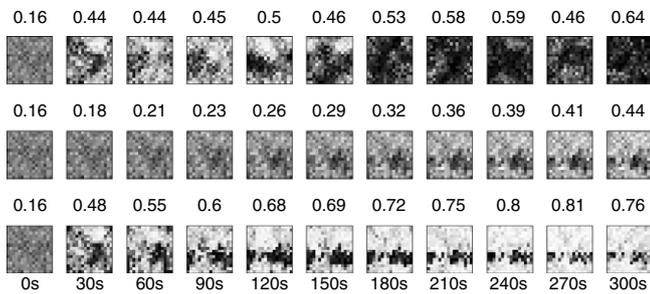


Fig. 4. Ocularity time series for two experiments. Each raster is an ocularity preference map, superscripted with the mean ocularity measure, as in Fig. 3. Left to right: samples at increasing times through the simulation, from 0 s, incrementing each 30s until 300s (as subscripted). Top row: no rewiring, ocularity including weights; Middle row: with rewiring, ocularity for connectivity only; Bottom row: with rewiring, ocularity including weights. Note that a connectivity-only set is not given for the no-rewiring case, since this does not change but rather remains exactly as in the first (0s) image in row 1 (no rewiring, ocularity including weights), at which point the weights are all maximised and the preferences are therefore due to the connectivity. While the pattern in the weights without rewiring changes as the inputs change, in the rewiring case the pattern in the weights becomes embedded in the connectivity and this helps to stabilise the pattern in the weights.

reduction quantitatively in a model with intended generality, but rather sufficient simply to observe that such an effect is possible. Nevertheless further work would be needed to fully characterise the performance of the variance reduction phenomenon with different parameters.

A possible way this model could be extended would be to allow axon branching to be guided by the existence of axons, such that an input neuron is more likely to form synapses with target cells which are close to target cells which it is already innervating. This might be expected to model axonal arbor development more accurately. In addition, there is no mechanism in this model for the preferential sprouting of synapses due to potentiation (Toni, Buchs, Nikonenko, Bron, & Muller, 1999)—a complementary way in which the two types of plasticity could interact.

As is common in this field, point neurons have been modelled, that is to say, there has been no consideration of the spatial and temporal summation and filtering performed by transmission of post-synaptic potentials through dendritic trees. Such considerations are likely to be important for a complete understanding of map formation, especially since it is known that different temporal learning windows can exist for synapses at different locations on the dendritic tree (Kampa, Letzkus, & Stuart, 2007).

The development of a mapping from the start has not been simulated, nor has the growth of areas been addressed. Rather, this model assumes a starting time at some point during development, in order to assess the effect of the proposed learning rules. This is reasonable, as it mimics the progress from an initially diffuse mapping towards a final mapping, which is thought to be at least partly dependent on activity (Simon, Prusky, O’Leary, & Constantine-Paton, 1992). Nevertheless, this focus limits the scope of the model. Additionally, an unjustified assumption, used by Song and Abbott (2001) and adopted here, is that new synapses start strong and then get weakened; the opposite case seems more likely when the process of synapse formation is considered. This assumption has been used for simplicity because it avoids the need for any additional homeostatic mechanisms to kick-start the activity of the network. If after initial maximisation, newly formed synapses are added with zero strength, the simulations function similarly, although mean weights are slightly lower.

As noted above, this work is part of a project to implement synaptic rewiring in neuromorphic VLSI. An implementation of this model has been fabricated; the interested reader is directed to Bamford et al. (in press) for details. The computational efficiency which may be gained by not implementing synapses between all possible pairs of neurons was noted above, as was the related

saving of space and energy in the brain. The hardware implementation of this model, in which each synapse has a physical instantiation, has space and energy costs analogous to those of the brain. Meanwhile, STDP, whilst more computationally intensive to simulate than certain more abstract learning rules, has a neat implementation in pulse-based neuromorphic hardware. These considerations suggest that computational modelling using the elements in this model may have benefits in the development of new potentially useful computational hardware.

6. Conclusions

A model of topographic development has been presented which includes both weight and wiring plasticity. There are three key assumptions: (a) synapses preferentially form in locations to which their axons are guided, (b) weights of dendritic synapses of a neuron are modified according to a competitive Hebbian learning rule, and (c) weaker synapses are more likely to be eliminated. In order to instantiate the model, more assumptions have been made, the main one being that the weight-change mechanism is a form of STDP.

It has been found that whilst spatially correlated inputs help to create patterns of synaptic weights which favour narrower projections, spatial correlations are not necessary for some reduction of variance to occur. A weight-change mechanism and a rewiring mechanism can work together such that the rewiring mechanism acts to embed patterns of synaptic strengths in the network topology; this is as one would expect, though it has not been demonstrated quantitatively before. There has also been a qualitative demonstration of the possibility of synaptic rewiring increasing the stability of learnt patterns. The accuracy of preferred locations for target neurons will not necessarily improve when synapses are initially distributed around ideal locations. The division of mapping quality into the quantities of mean σ_{aff} and mean AD is a useful means for investigating these effects, and a method of applying statistical significance tests has been demonstrated which avoids possible biases in order to extract highly significant effects from small-scale simulations.

Acknowledgements

This work was funded by EPSRC. We are grateful to Guy Billings for providing the basis of the simulator code, and to Adrian Haith and Chris Williams, as well as others at the Institute of Adaptive and Neural Computation, for mathematical insights.

References

- Ballice-Gordon, R., & Lichtman, J. (1993). In-vivo observations of pre- and postsynaptic changes during the transition from multiple to single innervation at developing neuromuscular junctions. *Journal of Neuroscience*, 13(2), 834–855.
- Bamford, S., Murray, A., & Willshaw, D. (2010). Large developing receptive fields using a distributed and locally reprogrammable address-event receiver. *Neural Networks, IEEE Transactions on* (in press).
- Bell, C., Han, V., Sugawara, Y., & Grant, K. (1997). Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature*, 387, 278–281.
- Bi, G., & Poo, M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience*, 18, 10464–10472.
- Billings, G., & van Rossum, M. (2009). Memory retention and spike timing dependent plasticity. *Journal of Neurophysiology*, 101, 2775–2788.
- Bofill-i Petit, A. (2005). An analogue VLSI study of temporally-asymmetric Hebbian learning. *Ph.D. thesis*. University of Edinburgh.
- Brader, J., Senn, W., & Fusi, S. (2007). Learning real world stimuli in a neural network with spike-driven synaptic dynamics. *Neural Computation*, 19(11), 2881–2912.
- Buffelli, M., Busetto, G., Bidoia, C., Favero, M., & Cangiano, A. (2004). Activity-dependent synaptic competition at mammalian neuromuscular junctions. *News in Physiological Sciences*, 19, 85–91.
- Butts, D., Kanold, P., & Shatz, C. (2007). A burst-based “Hebbian” learning rule at retinogeniculate synapses links retinal waves to activity-dependent refinement. *PLoS Biology*, 5(3), e61.
- Chklovskii, D., & Koulakov, A. (2004). Maps in the brain: What can we learn from them? *Annual Review of Neuroscience*, 27, 369–392.
- Chklovskii, D., Mel, B., & Svoboda, K. (2004). Cortical rewiring and information storage. *Nature*, 431, 782–788.
- Colman, H., Nabekura, J., & Lichtman, J. (1997). Alterations in synaptic strength preceding axon withdrawal. *Science*, 275, 356–361.

- Creutzfeldt, O., Innocenti, G., & Brooks, D. (1974). Vertical organization in the visual cortex (area 17) of the cat. *Experimental Brain Research*, 21, 313–336.
- Davison, A., & Fregnac, Y. (2006). Learning cross-modal spatial transformations through spike-timing-dependent plasticity. *Journal of Neuroscience*, 26, 5604–5615.
- Elliott, T., Howarth, C., & Shadbolt, N. (1996). Axonal processes and neural plasticity I: Ocular dominance columns. *Cerebral Cortex*, 6, 781–788.
- Elliott, T., & Shadbolt, N. (1998). Competition for neurotrophic factors: Ocular dominance columns. *Journal of Neuroscience*, 18, 5850–5858.
- Elliott, T., & Shadbolt, N. (1999). A neurotrophic model of the development of the retinogeniculocortical pathway induced by spontaneous retinal waves. *Journal of Neuroscience*, 19, 7951–7970.
- Gaze, R., Keating, M., & Chung, S. (1974). The evolution of the retinotectal map during development in xenopus. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 185, 301–330.
- Goodhill, G. (1993). Topography and ocular dominance: A model exploring positive correlations. *Biological Cybernetics*, 69, 109–118.
- Grutzendler, J., Kasthuri, N., & Gan, W. (2002). Long-term dendritic spine stability in the adult cortex. *Nature*, 420, 812–816.
- Gutig, R., Aharonov, R., Rotter, S., & Sompolinsky, H. (2003). Learning input correlations through nonlinear temporally asymmetric Hebbian plasticity. *Journal of Neuroscience*, 23(9), 3697–3714.
- Guyonneau, R., Van Rullen, R., & Thorpe, S. (2005). Neurons tune to the earliest spikes through STDP. *Neural Computation*, 17, 859–879.
- Hebb, D. (1949). *The organization of behavior: A neuropsychological theory*. New York: Wiley.
- Holmes, N., & Spence, C. (2005). Multisensory integration: Space, time and superadditivity. *Current Biology*, 15, 762–764.
- Hubel, D., & Wiesel, T. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, 195, 215–243.
- Hubel, D., Wiesel, T., & LeVay, S. (1977). Plasticity of ocular dominance columns in monkey striate cortex. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 278(961), 377–409.
- Jun, J., & Jin, D. (2007). Development of neural circuitry for precise temporal sequences through spontaneous activity, axon remodeling, and synaptic plasticity. *PLoS One*, 8.
- Kampa, B., Letzkus, J., & Stuart, G. (2007). Dendritic mechanisms controlling spike-timing-dependent synaptic plasticity. *Trends in Neurosciences*, 30(9), 456–463.
- Knudsen, E., du Lac, S., & Esterly, S. (1987). Computational maps in the brain. *Annual Review of Neuroscience*, 10, 41–65.
- Linsker, R. (1986a). From basic network principles to neural architecture: Emergence of orientation columns. *Proceedings of the National Academy of Sciences of the United States of America*, 83, 8779–8783.
- Linsker, R. (1986b). From basic network principles to neural architecture: Emergence of orientation selective cells. *Proceedings of the National Academy of Sciences of the United States of America*, 83, 8390–8394.
- Linsker, R. (1986c). From basic network principles to neural architecture: Emergence of spatial-opponent cells. *Proceedings of the National Academy of Sciences of the United States of America*, 83, 7508–7512.
- Masquelier, T., & Thorpe, S. (2007). Unsupervised learning of visual features through spike timing dependent plasticity. *PLoS Computational Biology*, 3, e31.
- McLaughlin, T., Hindges, R., & O'Leary, D. (2003). Regulation of axial patterning of the retina and its topographic mapping in the brain. *Current Opinion in Neurobiology*, 13(1), 57–69.
- McLaughlin, T., Torborg, C., Feller, M., & O'Leary, D. (2003). Retinotopic map refinement requires spontaneous retinal waves during a brief critical period of development. *Neuron*, 40, 1147–1160.
- Miikkulainen, R., Bednar, J., Choe, Y., & Sirosh, J. (2005). *Computational maps in the visual cortex*. New York: Springer.
- Miller, K. (1998). Equivalence of a sprouting-and-retraction model and correlation-based plasticity models of neural development. *Neural Computation*, 10, 529–547.
- Miller, K., Keller, J., & Stryker, M. (1989). Ocular dominance column development: Analysis and simulation. *Science*, 245, 605–615.
- Morrison, A., Aertsen, A., & Diesmann, M. (2007). Spike-timing-dependent plasticity in balanced random networks. *Neural Computation*, 19, 1437–1467.
- Petersen, C., Malenka, R., Nicoll, R., & Hopfield, J. (1998). All-or-none potentiation at CA3-CA1 synapses. *Proceedings of the National Academy of Sciences of the United States of America*, 95, 4732–4737.
- Prestige, M., & Willshaw, D. (1975). On a role for competition in the formation of patterned neural connexions. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 190, 77–98.
- Ruthazer, E., & Cline, H. (2004). Insights into activity-dependent map formation from the retinotectal system: A middle-of-the-brain perspective. *Journal of Neurobiology*, 59, 134–146.
- Simon, D., Prusky, G., O'Leary, D., & Constantine-Paton, M. (1992). N-Methyl-D-aspartate receptor antagonists disrupt the formation of a mammalian neural map. *Proceedings of the National Academy of Sciences of the United States of America*, 89, 10593–10597.
- Sjostrom, P., Turrigiano, G., & Nelson, S. (2001). Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, 32, 1149–1164.
- Song, S., & Abbott, L. (2001). Cortical development and remapping through spike timing-dependent plasticity. *Neuron*, 32, 339–350.
- Song, S., Miller, K., & Abbott, L. (2000). Competitive hebbian learning through spike-timing-dependent synaptic plasticity. *Nature*, 3, 919–926.
- Sperry, R. (1963). Chemoaffinity in the orderly growth of nerve fiber patterns and connections. *Proceedings of the National Academy of Sciences of the United States of America*, 50, 703–709.
- Toni, N., Buchs, P., Nikonenko, I., Bron, C., & Muller, D. (1999). LTP promotes formation of multiple spine synapses between a single axon terminal and a dendrite. *Nature*, 402, 421–425.
- Trachtenberg, J., Chen, B., Knott, G., Feng, G., Sanes, J., Welker, E., et al. (2002). Long-term in vivo imaging of experience-dependent synaptic plasticity in adult cortex. *Nature*, 420, 788–794.
- Udin, S., & Fawcett, J. (1988). Formation of topographic maps. *Annual Review of Neuroscience*, 11, 289–297.
- Willshaw, D. (2006). Analysis of mouse EphA knockins and knockouts suggests that retinal axons reprogram target cells to form ordered retinotopic maps. *Development*, 133, 2705–2717.
- Willshaw, D., & von der Malsburg, C. (1976). How patterned neural connections can be set up by self-organisation. *Proceedings of the Royal Society of London. Series B*, 194, 431–445.
- Willshaw, D., & von der Malsburg, C. (1979). A marker induction mechanism for the establishment of ordered neural mappings: Its application to the retinotectal problem. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 287, 203–243.
- Wong, R. (1999). Retinal waves and visual system development. *Annual Review of Neuroscience*, 22, 29–47.
- Young, J., Waleszczyk, W., Wang, C., Calford, M., Dreher, B., & Obermayer, K. (2007). Cortical reorganisation consistent with spike timing- but not correlation-dependent plasticity. *Nature Neuroscience*, 10, 887–895.
- Zhang, L., Tao, H., Holt, C., Harris, W., & Poo, M. (1998). A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, 395, 37–44.