Generalization of learning by synchronous waves: from perceptual organization to invariant organization

David M. Alexander · Chris Trengove ·
Phillip E. Sheridan · Cees van Leeuwen

Received: 10 September 2010 / Revised: 9 November 2010 / Accepted: 9 November 2010 / Published online: 10 December 2010
© Springer Science+Business Media B.V. 2010

Abstract From a few presentations of an object, perceptual systems are able to extract invariant properties such that novel presentations are immediately recognized. This may be enabled by inferring the set of all representations equivalent under certain transformations. We implemented this principle in a neurodynamic model that stores activity patterns representing transformed versions of the same object in a distributed fashion within maps, such that translation across the map corresponds to the relevant transformation. When a pattern on the map is activated, this causes activity to spread out as a wave across the map, activating all the transformed versions represented. Computational studies illustrate the efficacy of the proposed mechanism. The model rapidly learns and successfully recognizes rotated and scaled versions of a visual representation from a few prior presentations. For topographical maps such as primary visual cortex, the mechanism simultaneously represents identity and variation of visual percepts whose features change through time.

Keywords Visual cortex · Learning · Topographic maps · Cortical dynamics

Introduction

The equivalent apparitions in all cases share a common figure and define a group of transformations that take the equivalents into one another but preserve the invariant figure. So, for example, the group of translations removes a square appearing at one place to other places; but the figure of a square it leaves invariant. These figures are the *geometric objects* of Cartan and Weyl, the *Gestalten* of Wertheimer and Köhler. We seek general methods for designing nervous nets which recognize figures in such a way as to produce the same output for every input belonging to the figure. We endeavour particularly to find those which fit the histology and physiology of the actual structure.

Pitts and McCulloch 1947—‘How do we know universals: the perception of auditory and visual forms’

Briefly, the characteristics of the nervous system are such that, when it is subject to any pattern of excitation, it may develop a pattern of activity, reduplicated throughout an entire functional area by spread of excitations, much as the surface of a liquid develops an interference pattern of spreading waves when it is disturbed at several points.

Lashley 1950—‘In search of the engram’

From a few presentations of an object, perceptual systems quickly build a representation that enables new presentations to be recognized. In the visual system, for instance, it

D. M. Alexander (✉) · C. van Leeuwen
Laboratory for Perceptual Dynamics, RIKEN Brain Science
Institute, Wako-shi, Saitama, Japan
e-mail: dalex@brain.riken.jp

C. Trengove
Brain and Neural Systems Team, RIKEN Computational Science
Research Program, Saitama, Japan

C. Trengove
Laboratory for Computational Neurophysics, RIKEN Brain
Science Institute, Wako-shi, Saitama, Japan

P. E. Sheridan
School of Information and Communication Technology,
Griffith University, Meadowbrook, QLD, Australia

is crucial that an object can be recognized independently of its orientation and extent. The problem is illustrated in Fig. 1, which shows a rectangle rotating and changing in size over time. Despite these changes, the visual system perceives the rectangle as a single object, not a collection of different objects. While this type of task may appear trivial, exactly how perceptual constancy of this sort is established with rapidity and maintained efficiently remains something of a mystery.

We introduce a model in which object invariance is realized through dynamic generalization. This mechanism in some respect is similar to formal approaches the problems of perceptual Goodness (Pitts and McCulloch 1947; Garner 1962; Hoffman and Dodwell 1985; van der Helm and Leeuwenberg 1996; Wagemans 1999; Olivers et al. 2004; van der Helm and Leeuwenberg 2004). The concept of Goodness has connotations of both “simple” and “easy to remember”, but is notoriously hard to define in general (for some debates, see e.g., Olivers et al. 2004; van der Helm and Leeuwenberg 1996, 2004; Wagemans 1999). Pitts and McCulloch (1947) envisaged, see the quote at the

Fig. 1 Sequential images (*top to bottom*) of a rectangle that transforms through time by rotating and changing size. Recognition of object constancy is achieved automatically and effortlessly by the visual system. The mechanism described in the present research is able to achieve this type of generalization from only a few examples. This process of generalization is indicated in the figure by the *gray points* forming the top-most rectangles, which become black by the fourth instance of the rectangle

start of this paper, a definition in terms of invariance under pattern transformation. Garner (1962, 1966) provided such a definition, based on the view that categorical representations are fundamental to perception. Categories are defined as equivalence sets; they include several specific items that can be regarded as versions of each other via a group operation. Garner’s fundamental idea is that good patterns have few alternatives within their equivalence sets. Garner and Clement (1963) tested this idea, using patterns exemplified in Fig. 2. They showed that their criterion of equivalence set size (ESS) could reliably predict observers’ Goodness ratings.

Garner (1966) considered invariance under transformation of rotation and reflection. Rotation will feature in our approach as a key example. Unlike Garner, we will not address reflection (symmetry), as this is understood more adequately as a pairwise relationship of pattern components (van der Helm and Leeuwenberg 1996). Instead of reflection, we will model scaling invariance. The main difference with Garner’s approach however is that, rather than set-theoretic, our description of the equivalence relation is dynamical, and embodied in cortical topography.

The proposed model integrates two critical findings in cortical neurophysiology. First, the human cortex contains over fifty well-defined cortical areas. As understanding of these areas increases, it has become evident that many of them display clear topographical structure related to their function (Catania 2002; Malach et al. 2002; Chklovskii and Koulakov 2004). The best studied area is the primary visual cortex, which represents the visual field as a retinotopic mapping (Schwartz 1980; Tootell et al. 1988c; Balasubramanian et al. 2002; Adams and Horton 2003).

Second, dynamically-oriented studies in neuroscience have revealed an astonishing complexity in spatiotemporal patterns of electrical activity of the cortex (Freeman and Barrie 2000; Wright et al. 2001). The dynamics of cortical activity encompass a multiple temporal and spatial scales: sufficient to include both fast but localized perceptual process (Nicolelis et al. 1995; Freeman and Barrie 2000;

ESS 1	ESS 4	ESS 8

Fig. 2 Garner patterns comprising the set of all 90 five-dot patterns that can be constructed on an imaginary 3×3 grid leaving neither rows nor columns empty. They fall into 17 disjunctive Equivalence Sets (ES) of patterns that can be transformed into each other by rotations in 90° steps and/or by reflections. Seven ES contain eight patterns, eight sets contain four patterns, and two sets consist of only one pattern (see Fig. 2). Garner and Clement (1963) proposed that the size of the ES determines Goodness, in the sense that the smaller the Equivalent Set Size of a pattern (ESS), the larger its Goodness

Eckhorn et al. 2001) and whole cortex event-related integration of cognition (Alexander et al. 2006a; Klimesch et al. 2007). At the local as well as the global scale, the spatio-temporal dynamics often have the appearance of transient traveling waves (Freeman and Barrie 2000; Alexander et al. 2006a; Ito et al. 2007).

Map-based invariance

Representational invariance has been observed in several visual and vision-related areas including the ventral intraparietal area, inferotemporal cortex, entorhinal cortex and hippocampal structures (Rolls 1992; Ito et al. 1995; Kovacs et al. 1995; Logothetis et al. 1995; Duhamel et al. 1997; Quiroga et al. 2005). These studies report invariant neural responses for partial occlusion, size, position, 3-D orientation, direction of eye-gaze as well as across multiple images of famous people and landmark objects. However, in general only a small fraction of neurons show invariant responses over the entire range of tested object variations. This suggests that localized representational schemes provide an incomplete account of invariance.

If individual neurons or small populations of neuron represent invariant features, the scientific problem can be posed as to how these features can be acquired by localized learning (Rolls 1995; Booth and Rolls 1998; Kording and König 2001; Ullman and Bart 2004). Solutions to this problem have been proposed, including the generation of explicit invariant representations through pre-processing (Bishop 1995; Bradski and Grossberg 1995; Tuytelaars and van Gool 2000; Mikolajczyk and Schmid 2002); integration of learning over neighbouring spatial locations (Fukushima 1988; Phillips et al. 1995; Stone and Bray 1995) or over the short-time scales in which the object varies (Földiák 1991; Rolls 1995; Stone and Bray 1995); and summation over units that each have invariant response over a limited range of feature variation (Poggio and Bizzi 2004). A shortcoming of all these models is that they are unable to generalize to a range of objects from only a few examples.

We propose, instead, that object invariance is achieved by creating a field of distributed representations, which in turn enables swift generalization of object identity. Our model generates and stores, in a non-local manner, object identity within a topographical map. Invariant representations are laid down in these maps from only a limited number of examples. The critical property of the topographic maps is that translated ‘copies’ of representations refer to different versions of the same pattern, depending on their location within the map. The creation of set-wise translations of the original representation enables the system to predict potential future transformations. The detection of object invariance thus amounts to the priming

of multiple, distributed representations that are consistent with future object variation.

The proposed neuro-dynamical mechanism operates on sparsely activated input patterns in the topographic maps, leading to synchronous activity between mutually connected foci of activity. Spatio-temporal waves emerge from these foci which then broadcast ‘copies’ of the representation throughout the map. Dynamic changes in synaptic gains, dependent on timing relationships between peaks in the wave-fronts, store these broadcast ‘copies’. In this manner, the system can represent familiarity with all the transformed (e.g. rotated) versions of an object subsequent to only a few instantiations of the representation of that object.

The mechanism shares features of localized approaches to the problem of representational invariance: the mapping of the input projections to the topographic map is a form of invariant pre-processing; broadcasting of ‘copies’ over a map performs a function analogous to integration of learning over object variation; and, finally, each ‘copy’ of the representation is able to contribute activity to a more abstract representation ‘upstream’ in the cortex. None of the previous approaches, however, used distributed representation of both object identity and variation; nor did their system dynamics enable the generalization of invariant object properties from only a few presentations.

Topographic-dynamic invariant object recognition draws together recent findings from a range of disciplines. From cognitive neuroscience it draws upon our understanding of cortical maps. Knowledge is stored within a systematically organized representational structure that both enables and limits the scope of cognitive ability. From neuroscience the theory makes use of recent findings in the study of neuro-dynamics at the level of mass neuronal populations, as well as the empirical study of topographic maps in the cortex. From mathematics it makes use of the long known—but little disseminated—properties of pseudo-invariant maps. The theory is pinned down by numerical simulations and a geometrical proof that together demonstrate the in principle efficacy of the proposed mechanism.

Generalization as pseudo-invariance

The first critical concept of the dynamic generalization theory is that of pseudo-invariance. By definition, invariant representations have the property that they encode all relevant variations of a percept as a single representation. Pseudo-invariant representations differ from invariant ones in that the variations of the representation are translations, or displaced copies, of each other (Sheridan et al. 2000). In the dynamic generalization theory, pseudo-invariance is hypothesized to be property of cortical maps. In other words, a class of nontrivial transformations of a percept can

be encoded by a set of representations within a cortical map related by mere translations of position.

The present paper will focus on visual cortex. Despite the standard textbook account of early sensory cortex, it has become evident in recent years that widespread dynamic interactions play a critical role in the formation of perceptions (see Alexander and van Leeuwen 2010, for a review of long-range contextual modulation in the primary visual cortex). We will therefore consider the role of the primary visual cortex as an architecture for pseudo-invariance of two dimensional representations. In this case the stimulus field, S , is the two dimensional visual field described in retinotopic coordinates. The relevant topography of cortical maps is assumed to be well approximated by a two dimensional sheet. Denoting the cortical map as C , we can define the following mappings in the complex plane:

$$f : S \rightarrow S$$

where f is some transformation in S ,

$$M : S \rightarrow C$$

where M is the mapping of S into C , provided by the anatomy of the cortex, and

$$T : C \rightarrow C$$

where T is translation within C . Our interest lies in how some transformation f in S can be achieved within a cortical map by mere translation. This is given by the following mapping diagram:

$$\begin{array}{ccccc} & & M & & \\ & z_S & \rightarrow & z_C & \\ f \downarrow & & & & \downarrow T \\ & \omega_S & \rightarrow & \omega_C & \\ & & M & & \end{array}$$

This mapping diagram captures the property that the stimulus field and the cortical map co-exist as two different functional spaces. If M is a homeomorphic mapping, having the property that it can be reversed without loss of information, the property of pseudo-invariance can be expressed as

$$f(z) = M^{-1}(T(M(z))) \quad (1)$$

where M^{-1} is the inverse mapping of M . Equations 1, 2 describes in a simple way the property that some (useful) transformation of the stimulus field can be achieved by mere translations within the cortical map. Alternative formulations of the concept of pseudo-invariance have previously been explicated (Pitts and McCulloch 1947; Agu 1988; Sheridan et al. 2000).

We illustrate Eqs. 1, 2 in Fig. 3, for the case of the complex-logarithmic mapping, which is a homeomorphic

Fig. 3 Illustration of the complex-logarithmic mapping and its property of rotation invariance. *Left* of figure shows four *triangles* rotated about the origin. This rotation is denoted by the transformation $f(z)$. *Right* of figure shows the complex-logarithmic mapping of the z -plane, $M(z')$, in which the four *triangles* are also represented. While the shape of each *triangle* is distorted, the four versions are all translations, T , of each other. Translating the objects vertically within the complex-logarithmic map is equivalent to rotating about the origin in the untransformed space. If a translated version of the triangle in M is mapped back into the original space via M^{-1} , the result is a rotated version of the object. Likewise, translating horizontally in the complex-logarithmic map is equivalent to scaling (smaller/larger) about the origin in the untransformed space (see Schwartz 1980)

mapping providing rotational and scaling pseudo-invariance (Schwartz 1980; Balasubramanian et al. 2002). The critical geometric features of a pseudo-invariant mapping can be seen here: the original representation is distorted in M , but that these distorted versions can be merely translated around within M in order to achieve the transformation f . That is, if the distorted representation is mapped back to the original representational space, via M^{-1} , the representations regain their original undistorted shape but have been transformed according to f rather than merely translated.

Dynamic generalization through travelling wave activity

The second critical concept of the dynamic generalization theory is that of propagation of copies. This is a process by which a particular representation is copied to all positions within a cortical map within the map via dynamical mechanisms. ‘Copying’ means that all translated versions of the representation are primed. One way to achieve this end would simply be to allow the cortical map to be activated by all possible translations of a representation over many successive presentations. However, dynamic generalization requires that a small number of instances of the representation are sufficient to achieve the desired propagation of copies. This process therefore requires that all translations of a representation are primed after activation of the cortical map by only a limited number of translations of the representation. This priming could occur through temporary increases in synaptic gains within each of the

copies or by leaving a residual, ongoing trace of dynamic activity that creates internal links within each of the copies.

The process of propagation of copies is illustrated for a specific type of dynamical activation in Fig. 4. A representation is formed from discrete foci of activation in the cortical maps. Activated foci on the cortical map are assumed to become synchronously activated by their inputs and reciprocal interconnections (Wright et al. 2001). A wave of activity then propagates outwards from each point in the representation, like ripples on the surface of a pond. Short-term memories are stored between neurons in the cortex via coincident activations induced by the wave fronts. Consideration of Fig. 4 indicates that an arbitrary copy (translation) of the starting pattern of focal activations will also be stored as a short-term memory.

The computational potential of neurodynamics has long been understood (Agu 1988; Ermentrout and Kleinfeld 2001; Freeman 2003; Gong and van Leeuwen 2009). Recordings over large populations of neurons have shown several activity patches can occur simultaneously within or across cortical regions (Arieli et al. 1995; Kleinfeld and Delaney 1996; Dale et al. 2000; Freeman and Barrie 2000; Lam et al. 2000; Senseman and Robbins 2002; Derdikman et al. 2003; Freeman 2003; Kenet et al. 2003; Tucker and

Katz 2003; Fox et al. 2005; Ferezou et al. 2007; Vincent et al. 2007). They are not bound to specific locations but propagate, spread, drift, or move about across the space of the cortex (Arieli et al. 1995; Kleinfeld and Delaney 1996; Lam et al. 2000; Senseman and Robbins 2002; Derdikman et al. 2003; Kenet et al. 2003; Fox et al. 2005; Ferezou et al. 2007). The term “wave packets”, used in the context of olfactory, visual, auditory and somatosensory cortices of behaving rabbits (Freeman and Barrie 2000; Freeman 2003), nicely captures these aspects of the collective activity. Propagating activity patterns pervasively occur in multi-unit electrophysiological recording, EEG local field potential recording, MEG, optical imaging and fMRI imaging, both in spontaneous activity (Arieli et al. 1995; Fox et al. 2005; Vincent et al. 2007) and evoked responses (Ribary et al. 1991; Kleinfeld and Delaney 1996; Prechtl et al. 1997; Dale et al. 2000; Freeman and Barrie 2000; Lam et al. 2000; Senseman and Robbins 2002; Derdikman et al. 2003; Freeman 2003; Kenet et al. 2003; Tucker and Katz 2003; Roland et al. 2006; Rubino et al. 2006; Benucci et al. 2007; Ferezou et al. 2007; Xu et al. 2007). They are traditionally considered to be epiphenomena of network activity. However, their active role in entrainment of neural network activity has recently been shown (Frohlich and McCormick 2010). In other words, the propagating field activity builds a feedback loop with the activity of individual neurons, thus helping to establish a dynamic network structure. We interpret this phenomenon as functional for shaping a network structure able to perform generalization.

Methods

Definition of the system: We define a cortical topographic map, G , as an 8-tuple directed weighted graph:

$$L, p, R, t, k, v, W, w$$

where L denotes a lattice of points in the Cartesian plane, where each point $i \in L$ has an address associated with it. The symbol p denotes presentation number, $p \in \mathbb{N}$. A representation is defined as subset of points on the lattice, $R \subset L$, which change at each p . The symbol t denotes time within each p , and by definition $t_0 = 0$ at the onset of each presentation. Each point in the lattice is associated with a time-dependent variable, denoted $k_i(t)$, that can vary during the course of a given presentation. Each point on the lattice is also associated with a state variable that is updated at the beginning of each presentation, $v_j(p)$.

Every point j on the lattice receives inputs from other points by connections of $(|I| + 1)$ th order, W_{Ij} . Each such connection is defined by a set of input addresses $I \subset L$ and a single output address $j \in L$. The lattice is fully connected;

Fig. 4 The principle of reinforcement of map translations of a representation by travelling waves. *Upper* The original input pattern is shown as a set of red points. A translation of the representation is shown as a set of blue points. Waves emanating from each point the red pattern reinforce the pattern of all cells that are activated in phase (circles). The blue pattern is thereby reinforced. *Lower* The mutual connections (black lines) reinforced between the points in original input pattern, and mutual connections reinforced in the arbitrary translation of the representation

for each j , and given synapses of n th order, all possible $W_{I,j}$ exist. Each $W_{I,j}$ is associated with a weight value denoted by $w_{I,j}(p)$, which may be updated at each p .

Set-wise translation of a representation or ‘copy’: The translation of R by λ on L is

$$T_\lambda(R) = \{j + \lambda : j \in R, j + \lambda \in L\}; \lambda \in L \quad (2)$$

(see Sheridan (1996) for a formal account of the operation of translation as addition between two addresses on a discrete lattice). For convenience, we restrict the possible translations such that no points will be translated off the lattice. Henceforth, *copy* refers to this definition of a set-wise translation of a representation. Given the set R , the set of all possible translations of R is

$$D(R) = \{T_\lambda(R) : \lambda \in L\}$$

The goal connection set: The set of $(|L|+1)$ th order connections associated with a given $M \subset L$ is

$$B(M) = \{W_{I,j} : j \in M, I \subset M, j \notin I\}$$

Given some R , the set of connections associated with all possible translations of that R is

$$C(R) = \{B(T_\lambda(R)) : \lambda \in L\}$$

termed the *goal connection set*.

Synchronous activation and expanding wave-fronts: $j \in R$ are defined as being synchronously active from t_0 . From t_0 , wave-fronts expand from each $j \in R$ until they reach the edge of the lattice which has absorbing boundary conditions. The waves are modeled as a delta function; only the peak in the expanding wave-front is considered. We define wave-states on points of the lattice, $k_i(t)$, corresponding to peaks on circular expanding wave-fronts. The radius of the wave about point $j \in R$ after t is given by a function $\delta(t)$.

$$\begin{aligned} k_j(t) &= 1 \quad \text{if } j \in R, t = t_0 \\ k_i(t) &= 1 \quad \text{if } \exists j \in R : |j - i| = \delta(t); i \in L, t > t_0 \\ k_i(t) &= 0 \quad \text{otherwise} \end{aligned} \quad (3)$$

Reinforcement by synchronous activation: $S(t)$ is termed a *synchronous set* and comprises the set of points at t that are coincidentally engaged by the expanding wave-fronts.

$$S(t) = \{i : k_i(t) = 1\} \quad (4)$$

The learning mechanism is defined such that it increments the values of $w_{I,j}(p)$ if points in the lattice are coincidentally activated sometime during p . The set of connections that is reinforced is $B(S(t))$:

$$w_{I,j}(p) = w_{I,j}(p-1) + \Delta w \quad \text{if } \exists t, W_{I,j} \in B(S(t)) \quad (5)$$

where Δw is the amount of gain increase.

Resetting of previous reinforcement by asynchronous activation: Let $A(t) = \{i_A\} \cup S(t)$, where i_A is some point $k_i(t) = 0$. $A(t)$ is termed an *asynchronous set*. Connections

associated with the asynchronous set are reset according to following rule:

$$w_{I,j}(p) = 0 \quad \text{if } \exists t, W_{I,j} \in B(A(t)) \setminus B(S(t)) \quad (6)$$

$$\text{and } \forall t, W_{I,j} \notin B(S(t))$$

The proposed mechanism is sufficient to assure that all *translations of a representation are stored*. This is illustrated geometrically in Fig. 4. In words, for any translation $T_\lambda(R)$ of a pattern R there will be a time t when $\delta(t) = \lambda$. At that time the wave-front expanding from each $j \in R$ will activate the corresponding element $j + \lambda$ of $T_\lambda(R)$. Thus $T_\lambda(R) \subset S(t)$. Therefore all $W_{I,j} \in B(T_\lambda(R))$ get reinforced by Eq. 5 and will not be reset to zero by Eq. 6.

Specific constraints on the system: We limit the size of R to n , such that $n_{\min} \leq n \leq n_{\max}$, where $n_{\min}$ and $n_{\max}$ are small, so typically $3 \leq |R| \leq 6$. $|I| = 1$ and $|I| = 2$ correspond to the cases of 2nd and 3rd order connections. R comprises the set of sites that provide driving inputs to the lattice. A number of points on the lattice may additionally be driven by noise inputs, N . Noise inputs, by definition, are randomly chosen sites in L at each p , $N \not\subset R$. Noise points evoke the same activation dynamics as points in R i.e. they also produce expanding wave fronts and thereby contribute to weight updates.

Calculating the output of the system: Let $d_j(p)$ denote the state at the beginning of the p th presentation due to the driving inputs alone:

$$\begin{aligned} d_j(p) &= 1 \quad \text{if } j \in R \cup N \\ d_j(p) &= 0, \quad \text{otherwise} \end{aligned}$$

$W_{I,j}$ act in a similar fashion to an AND gate, so that $w_{I,j}(p)$ only has effect if, for all $i \in I$, $d_i(p) = 1$. The *total activation* to site j on the p th representation is denoted by $v_j(p)$, and is calculated at t_0 :

$$v_j(p) = d_j(p) + \sum_I w_{I,j}(p) \prod_i d_i(p); i \in I \quad (7)$$

We assume that the effects of the driving inputs on activation are large compared to inputs via network weights, $w_{I,j}$. This means, for the purposes of assessing recognition we need only consider those points in the lattice $j \in R \cup N$. $v_j(p)$ is a dynamical state variable; it denotes the level of synchronous activation at j . The time-dependent nature of $v_j(p)$ is not explicitly modeled, but we assume that the amplitude of synchronous activity for a given $j \in R \cup N$ is a monotonically increasing function of $v_j(p)$.

Goal state of the system: Points j become synchronously activated by virtue of their mutual activation by external inputs (i.e. $j \in R \cup N$) and mutual interconnectivity, $B(R \cup N)$. Recognition is defined in terms of some threshold, θ_v , of the amplitude of synchronous activity at j . For the purposes of mathematical definitions to this point,

a representation was defined as $j \in R$. For the purposes of modeling terminology, it is useful to distinguish between the *representation*—and the process of recognition—which are a function of $v_j(p)$ values, and the driving *input pattern* which is specified solely by $d_j(p)$ values.

The desired system behavior at presentation p is that if each of the previously presented patterns were of the form $T_\lambda(R)$ plus noise, where all λ are unique, then upon presentation of

$$X \subset L, |X| = |R| + |N|$$

if $\exists R_p \subset X : J \in R_p \in D(R)$ then

$$v_j(p) \geq \theta_v \text{ otherwise } v_j(p) < \theta_v.$$

The goal state of the system is therefore that, after some previous number of presentations that are translations of the ‘same’ input pattern, the representation R_p is recognized as being the ‘same’ at the next presentation. Noise inputs that are not part of the representation are not to be recognized as such. ‘New’ input patterns that are not translations of the previous input pattern are also not recognized as being the ‘same’ representation.

Experimental tests

The purpose of the modeling is to demonstrate the basic geometrical properties of the proposed mechanism; that is, that synchronously expanding wave-fronts can in principle copy representations. This empirical confirmation is necessary because not all of the connections reinforced at each presentation *belong to a copy of the representation* (see Fig. 5, lower). That is, while the goal connection set $C(R)$ is always reinforced, connections outside this set, $B(L) \setminus C(R)$, may also be reinforced on a given presentation. This results in spurious activations in subsequent presentations. The pattern of interaction between correct and spurious activations is complex, and boundary conditions are highly influential. For example, the size of the lattice plays a role in the effectiveness of the mechanism. The results show that the model has graceful degradation in performance with noise and no negative effects of increasing lattice size. In order to address these issues, the modeling focused on the basic geometrical properties of the spatio-temporal waves, rather than more physiologically realistic modeling of the waves themselves. Details on modeling the underlying cortical dynamics can be found elsewhere (Liley et al. 1999; Wright et al. 2001; Chapman et al. 2002).

The experimental tests answer the central question: given the problem of spurious activations, how many prior presentations of an input pattern are required to successfully recognize it as being the same representation, regardless of where it may appear in the map? If only a limited number of prior presentations are required, then at each subsequent presentation the

Fig. 5 Modeling of spatio-temporal waves on a hexagonal lattice of columns. *Upper* When an input pattern is presented to the lattice (red) a set of synchronous spatio-temporal waves (black lines) emerge—one wave centred on each of the pattern’s points. The radius of the wave increases with increasing time during the presentation. The peak in spatio-temporal wave is the relevant active state for the learning rule ($k_i = 1$, blue) and points not at peak phase are inactive ($k_i = 0$, green). *Middle* The same wave-activated points as upper figure, with black interconnection lines showing that the synchronously activated points ($k_i = 1$, blue) form translations of the input pattern. All of the $n(n - 1)$ interconnections between n points in each copy of the input pattern are synchronously activated by the travelling wave. Not all $n(n - 1)$ connections are shown, for visualization purposes. *Lower* In addition, a large number of connections that are not within the set of $n(n - 1)$ interconnections between n points in each copy of the input pattern are also synchronous activated. For example, all connections between different copies of the input pattern are also synchronously activated. Not all connections between different copies are shown, for visualization purposes

translated input pattern can be *immediately* recognized as belonging to the same representation. If the map has the property of pseudo-invariance, the resultant property is immediate, ongoing recognition of objects whose features vary through time along those dimensions defined by the map (see Fig. 1).

The cortical map is modeled as a hexagonal lattice of columns. External inputs to n sites in the lattice create an n -point pattern of focal activations, R , at each presentation, p . The points in the input pattern are assumed to engage in synchronous activity. The spatio-temporal waves emerge as an expanding wave-front from each focal point of activation. Each wave is implemented on the lattice as a circle that expands at constant rate with time. Each circle is centered on one of the n lattice sites in the n -point input pattern. The spatio-temporal waves emerging from the set of points in the input pattern are synchronized in the sense that the wave-fronts emerge from their foci simultaneously, resulting in a set of expanding circles which all have the same radius at each point in time during the wave-cycle of the presentation. This is illustrated in Fig. 5. The radii of the circles expand with time until they pass the edge of the lattice. The short-range connectivity, >1 mm (Blasdel et al. 1985; Lund and Wu 1997), that underlies these dynamics in more neurophysiologically realistic modeling (Liley et al. 1999; Wright et al. 2001; Chapman et al. 2002) is not explicitly implemented in the present modeling.

Weight updates are calculated at the end of each presentation, according to Eqs. 5 and 6. Since the learning rule is formulated in terms of the output of the population of neurons at j rather than in terms the timing of individual synaptic inputs, the time offset in rules based on spike-timing dependent plasticity (STDP) is implicit (see section reviewing the neurophysiology).

The present modeling addresses three specific questions, which bear upon the issue of how many prior presentations are required to successfully recognize an object under transformation, that is, regardless of where its representation appears in the map.

1. What is the behavior of $B(L) \setminus C(R)$ over successive presentations?
2. What is the impact of $B(L) \setminus C(R)$ on recognition performance?
3. What effect do noise inputs, N , have on recognition performance?

To answer these questions, the modeling used

- a) four different input patterns,
- b) two different lattice sizes,
- c) two different presentation conditions, and
- d) four different levels of noise.¹

¹ Consideration of Eq. 7 indicates that the system will have more difficulty discriminating the *target representation*, $T_i(R)$, from the presence of additional noise compared to discriminating $T_i(R)$ from another representation which is a deformed version of $T_i(R)$. This is because $v_j(p)$ will tend to be larger for $j \in N$ in the case when $M \cup N$, $M \in D(R)$ is presented, compared to when $j \in O$, $O \in Q(R) \setminus D(R)$. More specifically, $v_j(p)$ will tend to be higher for noise points ($j \in N$ compared to $j \in O \cap M \cap O$) and higher for points in the

Fig. 6 Test input patterns used in the modeling. *Upper* The four input patterns used to test the generalization mechanism in the small lattice. This lattice has 343 hypercolumns (7^3). The three point, four point, five point and six point input patterns are shown *top-left*, *top-right*, *bottom-left* and *bottom-right*, respectively. *Lower* The four input patterns used to test the generalization mechanism in the large lattice. This lattice has 2,401 columns (7^4). The three point, four point, five point and six point input patterns are shown *top-left*, *top-right*, *bottom-left* and *bottom-right*, respectively. These four input patterns are identical to their small lattice counterparts, except that they are spread over seven times the area. Since the large lattice also has seven times the area, the effect is to preserve the sizes of the input pattern relative to the size of lattice, while the lattice resolution is seven times greater. The geometric-algebraic transformation that scales the input pattern sevenfold also rotates them (see Sheridan et al. (2000) for details)

The four different input pattern types and two lattice sizes are illustrated in Fig. 6. It shows the configuration of (a) the 3-, 4-, 5- and 6-point input patterns that were used, as presented on (b) the small (343 site) and large (2,401 site) lattices.

There were, moreover, two presentation conditions. Over successive presentations, input patterns were presented c) either at random locations in the map or in sweeping motions across the map. The effect of $B(L) \setminus C(R)$ on recognition performance was assessed on the large lattice by

Footnote 1 continued

representation ($j \in M$ compared to $j \in M \cap O$). This is true even when M and O differ by one point only, that is, one point is displaced within the target representation. In short, adding noise increases activation levels compared to displacing points within the target representation. For this reason we used additional noise rather than deformations of the target to assess recognition performance.

Fig. 7 An example of a presentation sequence for the measurement of spurious activations, over three successive presentations. The input pattern is swept across the lattice from 2 o'clock to 8 o'clock ($\pi/3$ to $4\pi/3$), with random jitter in position for each of the three placements. The extra noise element for each presentation is shown as a *gray dot*. Sweeps of the input pattern took place for each of 6 directions (0 to π , ..., $5\pi/3$ to $2\pi/3$)

introducing an additional noise-point to each presentation of the n -point input pattern. This additional noise-point was randomly placed in the vicinity of the input pattern, illustrated in Fig. 7. The activations of this additional random point provided an index of the effects of $B(L)\backslash C(R)$ on spurious activations in the network. In some experiments on the large lattice d) increasing amounts of random noise were added to each presentation.

Simple Hebbian-type synapses (2nd order connections) were modeled for the sake of computational convenience. This meant the network had only $|L|(|L| - 1)$ connections. However, higher order connections could have been used or the complete list of all n -tuplets of synchronous activations stored as a list, without changing the substantive conclusions.²

² The 2nd order synapses, as implemented in the present network, bring about two pseudo-problems which, however, do not reflect upon the properties of the generalization mechanism presently of interest. First, 2nd order synapses are not able to distinguish between representations that are rotated by π in relation to each other. Every correctly reinforced 2nd order synapse will also contribute to the 'recognition' of a pi-rotated version of the representation. If 3rd order synapses were used in the network, these spurious generalizations would not occur. The second pseudo-problem concerns the implementation of additional noise points in the representation. The generalization mechanism treats pairs of points with the same vector length and direction as identical, since they are translations of each other. That is, any representation that contains points a and b , gives rise a set of connections (with updated gains) of the form $B(T_\lambda\{a,b\})$, each element of which corresponds to the connection vector, $\underline{ab}$. For this reason, the noise points, c , added to each representation, while

In many neural network applications, the sum of the synaptic gains is normalized in some fashion and/or non-linear activation functions applied. Here, the network activity is calculated without any scaling of total weights or nonlinear activation functions. This makes it easier to interpret our results in answering the question of how many *prior presentations* are required for the retrieval process to correctly distinguish whether the set of points in a new representation is a set-wise translation of the previously presented input patterns. A simple threshold rule enables the network to recognize whether the set of points in the new representation is a set-wise translation of the previous input patterns and to ignore any sub-threshold noise. For ease of exposition, hereafter we express the threshold in units of p rather than $v_j(p)$:

$$\theta_p = \frac{\theta_v - d_j}{\Delta w_{lj}(n - 1)}; d_j = 1; |I| = 1; \theta_p \in \mathbb{N} \quad (8)$$

where n is the number of points in R (not including noise inputs) and θ_v is the activation threshold as defined earlier in the methods section. The threshold is applied at t_0 and therefore takes into account learning up to, but not including, the immediate presentation. For a given set of experimental conditions, the threshold was calculated as the number of prior presentations required to correctly detect the target representation while failing to detect any additional noise points.

Results

Figure 5 shows the four-point input pattern (red) being presented to the small lattice, and subsequent activation of waves states $k_i = 1$ (blue) at one time during the wave cycle. The figure also shows that the hypercolumns with $k_i = 1$ form copies of the original four point pattern. Due to the learning rule, therefore, each set of connections storing copies of R are reinforced by the travelling waves. The lower part of the figure shows some of the connections within $B(L)\backslash C(R)$ that were also reinforced.

Figure 8 shows the results of a single experiment in which a four point pattern was presented at 50 random locations. The figure shows a histogram of the number of updated connections, categorized by the weight of the connection. The probability that a connection within $B(L)\backslash C(R)$ will be consistently reinforced over successive

Footnote 2 continued

otherwise random, were chosen such that they did not create connection vectors that were translations of the set of the existing set of connection vectors within the pattern. That is, $\underline{ab} \neq \underline{ca}$ and $\underline{ab} \neq \underline{ac}$ for each a and b in the n -point representation and for each noise point, c . Use of higher order synapses, while computationally more intensive, greatly reduces this effect, allowing for arbitrary noise points.

Fig. 8 Histogram of weights for all connections in the network after 50 random presentations. The weights are expressed in units of Δw . The 4 point input pattern was used on the small lattice (343 sites). The weight profile of connections outside the goal connection set, $B(L) \setminus C(R)$, shows an approximately exponential decline in numbers of connections as a function of weight. That is, at the 50th presentation, the numbers of connections in $B(L) \setminus C(R)$ with weight $4 \Delta w$ was negligible compared to the number of connections in $B(L) \setminus C(R)$ with weight 0 or Δw . The connections of the goal connection set, $C(R)$, all had weights of $50 \Delta w$

presentations declines (approximately) exponentially with each successive presentation. For example, the number of connections in $B(L) \setminus C(R)$ that have been updated four times in a row is very small compared to the number of connections in $B(L) \setminus C(R)$ updated twice in a row. The weights of $C(R)$, which store each and every copy of the four point input pattern, are *always* increased. This can be seen in the last bar of the histogram.

Figures 9 and 10 show the averaged results of all the experiments in which the test patterns were presented at random locations. In a range of experiments, the size of the lattice and the number of points in the input pattern were varied. As in the previous figure an exponential weight profile of $B(L) \setminus C(R)$ results, shown here on a logarithmic scale for ease of interpretation. The slope of the exponential weight profile of $B(L) \setminus C(R)$ becomes flatter as a function of the number of points in the input pattern. However, the slope becomes steeper as a function of lattice size. Similar results were found using input patterns that were swept systematically across the map.

The averaged profiles of weights in $B(L) \setminus C(R)$ are a ‘snapshot’ of the typical profile of weights in the network. However, this histogram also reveals the rate at which connections in a *particular subset* of $B(L) \setminus C(R)$, reinforced

Fig. 9 Effects of input pattern size on $B(L) \setminus C(R)$ with non-zero weights, as a function of connection weight (expressed in units of Δw). Each line in the graph shows the averaged results of 100 experiments in which input patterns were presented randomly to 50 different locations on the small (343 site) lattice. The number of connections in $B(L) \setminus C(R)$ of a given weight is represented as a proportion of the number of connections in $C(R)$ to provide an index of signal to noise. The graph shows an average snapshot of the proportion of $B(L) \setminus C(R)$ of a given weight, as a function of input pattern size. The graph can also be interpreted in terms of the rate in which a particular subset of $B(L) \setminus C(R)$ with non-zero weights, once introduced into the network on a *specific* presentation, are then removed from the network upon *successive* presentations. The rate of removal of $B(L) \setminus C(R)$ with non-zero weights is slower for larger input patterns

Fig. 10 Effects of input pattern size on $B(L) \setminus C(R)$ with non-zero weights, as a function of connection weight (expressed in units of Δw). Each line in the graph shows the averaged results of 50 experiments in which input patterns were presented randomly to 50 different locations on the large (2,401 site) lattice. The number of connections in $B(L) \setminus C(R)$ of a given weight is represented as a proportion of the number of connections in $C(R)$. The same overall pattern is found as for the smaller lattice, but the rate of decay of non-zero weights in $B(L) \setminus C(R)$ is much steeper (note change in scale of x-axis). The maximum number of successively reinforced $B(L) \setminus C(R)$ is also smaller in a large lattice, particularly for larger input patterns

to have weights of Δw on a *specific presentation*, decline in number over *successive presentations*. This interpretation of the $B(L) \setminus C(R)$ weight profile is referred to as the *decay rate* in $B(L) \setminus C(R)$ of non-zero weights. Since $B(L) \setminus C(R)$ of

Fig. 11 Activation due to non-zero weights in $B(L)C(R)$, shown by activation measured at a random point added to each presentation. The gray line near $y = 0$ shows the activation at each successive presentation for the additional random point. The black diagonal line shows the activation level given in units of θ_p , as defined in Eq. 8. It can be seen that for most presentations, the additional random point was activated at a level below the activation threshold of $\theta_p = 1$. This data was generated using the large lattice, using the three point input pattern

non-zero weight behave as background ‘noise’ to the foreground ‘signal’ of $C(R)$, this decay rate can be taken as a measure of performance. Performance is defined in terms of the number of prior presentations required to correctly recognize the translated representation, and this in turn is a direct consequence of the decay rate of non-zero weights in $B(L)C(R)$. This means performance declines with pattern size, but improves with lattice size.

The effects of $B(L)C(R)$ on performance were verified directly via calculation of the recognition threshold, θ_p , under a variety of experimental conditions. Figure 11 shows the results of one such experiment, on the large lattice when one random point of noise was added to the 3-point pattern, which was swept systematically across the map. The extra point is randomly positioned in the vicinity of the input pattern at each presentation. The grey line shows the activation of the randomly placed point at each presentation, calculated according to Eq. 7. Setting θ_p to 1 means that for the majority of presentations the extra random point will not be mistakenly identified as belonging to the representation. Setting the threshold to $\theta_p = 3$ means that the extra random point is never mistakenly identified as an element of $D(R)$ over these 50 presentations.

Figure 12 shows the combined results of a series of experiments using the large lattice, showing the effects of pattern size on the setting of θ_p . The graph shows the number of prior presentations required for correct recognition, given one extra random point at each presentation. For $3 \geq |R| \geq 6$, $\theta_p = 3$ was sufficient to discriminate the representation from the extra random input. In other words, no matter where the pattern was presented to the lattice, only three prior presentations were required to prevent an

Fig. 12 Number of prior presentations required to detect the representation without error. The data here are equivalent to Fig. 10, but shown in summary form for the 4 input pattern types. The histograms show the number of prior presentations that would be required to successfully detect the representations in the presence of the additional noise point. For example, setting the threshold $\theta_p = 1$ results in correct detection of the representation 80–90% of the time, depending on input pattern size. Setting the threshold to $\theta_p = 3$ results in correct detection of the representation 100% of the time. There is a slight degradation of performance with larger input pattern sizes. These data are from the large lattice, aggregated over 100 presentations per input pattern type

Fig. 13 Number of prior presentations required to detect the representation without error in the presence of increasing noise. The conventions are the same as Fig. 11. The figure shows results for the 5 point input pattern with four levels of noise. This histograms show the number of prior presentations that were required to successfully detect the representation in the presence of each additional noise point. For example, setting the threshold $\theta_p = 1$ results in correct detection of the representation 35–80% of the time, depending on the amount of additional noise. There is a degradation of performance with increasing noise levels. Setting the threshold to $\theta_p = 5$ would result in correct detection of the representation 100% of the time at the maximum noise levels tested. These data are from the large lattice, aggregated over 100 presentations per input pattern type

extra randomly activated point being mistakenly identified as belonging to the representation.

Figure 13 shows the results of adding increasing amounts of noise during presentations of the 5-point

pattern in the large lattice. The required threshold for successful discrimination increased gently with $|N|$. For the maximal levels of noise tested (4 randomly positioned points), $\theta_p = 5$ was sufficient to successfully discriminate the pattern from the noise. In other words, even when $|N|$ approached $|R|$, only five prior presentations were required to discriminate pattern and noise effectively. Similar results to those shown in Figs. 11, 12 and 13 were obtained for presentation of input patterns at random locations in the map.

The results show that for patterns with reasonably small R and for reasonably large L , the translated versions of a representation can be recognized *after only a few previous presentations*. Taking the visual system as exemplar, a presentation can be assumed to last of the order 100 ms (Gawne and Martin 2002; Angelucci and Bullier 2003). Taking into account eye, head and body movements, as well as motion in the visual field, ‘a few presentations’ involve durations of the order 500 ms. These short time scales make the proposed mechanism ideally suitable for fast detection of perceptual invariance, which may be required for visual object tracking.

Neurophysiological mechanisms

We will now discuss the neurophysiological underpinnings of the proposed cortical generalization mechanism, taking the primary visual cortex as exemplar.

Cortical representations occur within maps

It is widely known that topographical maps exist at a number of spatial scales in the cortex. These range from the global scale of the cortex, through the scale of the Brodmann areas, to the sub-millimetre scale. At the global scale of the cortex, the Brodmann areas are organized in a systematic fashion, for example with perceptual regions located posteriorly and laterally, and executive control regions located anteriorly and medially. At the scale of the individual Brodmann area, the most well-known maps are the retinotopic organization of the visual areas, tonotopic organization of auditory areas and somatotopic organization of the somato-sensory areas (see Chklovskii and Koulakov (2004) and Catania (2002) for reviews). Topographical organization at the sub-millimetre scale is apparent from the study of visual areas—for example the upper layers of the primary visual cortex (V1) display patterns of response properties, such as orientation, contrast and spatial frequency preference and ocular dominance that vary at a spatial scale of $\sim 350 \mu\text{m}$ (Tootell et al. 1988a, b, c; Bartfeld and Grinvald 1992; Hubener et al. 1997; Vanduffel et al. 2002; Alexander

et al. 2004a), in register with the patch to patch distance of long-range intrinsic connections (Livingstone and Hubel 1984; Yoshioka et al. 1996; Bosking et al. 1997; Lund et al. 2003).

The retinotopic mapping of the visual field to the monkey primary visual cortex may serve as the prototype of a cortical map for the present research. The input layers of V1 are organized as a map of the visual field (Tootell et al. 1988c; Adams and Horton 2003). The mapping of the visual field to monkey V1 is distorted in a manner that can be approximately described by the complex-logarithmic function illustrated in Fig. 3 (Schwartz 1980; Balasubramanian et al. 2002).

Representations within maps are spatially sparse

Firing patterns of neurons are highly selective for particular stimulus features. Because of this, neuronal coding in the cortex may be considered sparse (Vinje and Gallant 2000; Guyonneau et al. 2004). Studies on the effects of naturalistic stimulation have shown that activity in the visual field from well beyond the classically defined surround further increases the sparseness of neuronal coding in V1, by further refining the stimulus selectivity of neurons (Vinje and Gallant 2000; Terashima and Hosoya 2009; Haider et al. 2010). The spatial pattern of activity in V1 neurons in the upper layers has an excitatory-centre/inhibitory-surround structure (Blakemore and Tobin 1972; Nelson and Frost 1978; Allman et al. 1985). This spatial pattern of activation means that neighboring neurons tend to code for similar response properties and to be co-activated by those same stimulus properties, while suppressing sub-optimally activated neurons in the surrounding region (Swindale 1996). When viewed in spatial terms of the total activity in a cortical map, the sparseness of the neuronal code means that sites of high activity in the map in response to a particular stimulus are relatively small in total area compared to the size of the map; activated representations within cortical maps are *spatially* sparse.

The present modeling therefore assumes that representations within a given map contain relatively few points of high activation. Representational complexity arises from combinations of activity across multiple maps. Each point in a representation can be considered as an activated feature, where the feature can in turn refer to some other set of features represented in another map.

Cortical maps can be modeled as a discrete lattice of points

Lattice points in the model may be considered to represent approximately one hypercolumn (Hubel and Wiesel 1974) or column (Lund et al. 2003) in size, depending on the

lattice size (see Fig. 5). Columnar structures are ubiquitous in the cortex (for a recent review on the primary visual cortex see Lund et al. 2003).

Representations are recognized on successive presentations by virtue of short-term increases in gains between their mutual connections

The mechanism described has potential to explain the formation of cortical maps and has consequences for permanent storage of memories (see discussion). The present research only considers short-term memory; that is, temporary increases in synaptic gain due to previous activation by an input pattern. The proposed mechanism assumes rapid increases in synaptic efficacy. A consequence of this is a selective increase in response evoked by repeated exposure to preferred stimuli, and this effect have been recently observed in a study of cat visual cortex (Yao et al. 2007; Hua et al. 2010).

Broadcasting the representation all positions in a pseudo-invariant map is equivalent to generalizing over all transformations allowed by the map

The mapping of the visual field into V1 can be described in an idealized form as the complex-logarithmic mapping, which has the pseudo-invariant properties described in Fig. 3 (Schwartz 1980; Sheridan et al. 2000; Balasubramanian et al. 2002; Hinds et al. 2008). Another example of a pseudo-invariant map occurs in the frequency domain following a Fourier transformation. Translations of objects represented within a mapping of the frequency domain result in a frequency shifts for the representation, while preserving the relative frequency relationships within the representation. A pseudo-invariant mapping of the Fourier domain would have the appearance of a tonotopic map, similar to those found in the primary auditory cortex (Pantev and Lutkenhoner 2000; Schreiner et al. 2000).

A set of activated foci are connected through synchronization

Synchronous patterns of firing have been suggested as a mechanism by which the cortex binds disparate elements of mental objects together (von der Malsburg and Schneider 1986). Synchrony has been demonstrated in V1 between sites separated by up to 6 mm millimetres, when the recording sites are stimulated by moving bars of the same orientation (Eckhorn et al. 1993; Singer and Gray 1995; Livingstone 1996). Modeling of synchrony using realistic assumptions about structure and function of the cortex has revealed that synchrony can emerge rapidly between distant cortical sites, provided those sites are simultaneously

activated and mutually connected (Wright et al. 2001). In V1, long-range horizontal fibres link sites in the upper layers at distances of up to 3 mm (Stettler et al. 2002) in the lower layer at distances of up to 8 mm (Rockland and Knutson 2001).

Synchronous spatio-temporal waves emerge from each point in the synchronous set of activated points

Synchrony phenomena can take the form of spatio-temporal waves in the cortex, as experiments using wide-field stimuli have demonstrated (Freeman and Barrie 2000; Eckhorn et al. 2001). These waves propagate across a distance of up to 8 mm (the maximum extent of the recording array) at modal speeds of 0.4 m/s in monkey V1 (Eckhorn et al. 2001) and similar results have been found in cat (Benucci et al. 2007; Nauhaus et al. 2009) and rat (Xu et al. 2007) V1 using more focal visual stimuli. Synchrony has been measured at a more local scale in V1 between pairs of neurons when one neuron is activated by a stimulus of preferred orientation and spatial frequency and a second neuron activated sub-optimally by the same stimulus (König et al. 1995; Maldonado et al. 2000). The former neuron leads the latter in phase-coupled synchrony. These experiments results can be interpreted as spatio-temporal waves arising in V1 at the $\sim 350 \mu\text{m}$ scale, the distance over which local response properties vary. Patterns of phase delays have also been recorded in V1 across different depths within the same cortical column (Livingstone 1996). Taken together, these results suggest wave activity can arise at different scales within the primary visual cortex.

The period of the spatio-temporal waves is less than or equal to the size of the map

In a survey of travelling wave phenomena across multiple species, Ermentrout and Kleinfeld (2001) note that the phase gradient across the measurable region does not generally exceed π (i.e. substantially less than one full cycle of the wave). This enables phase to uniquely encode network states (Ermentrout and Kleinfeld 2001). Long-wavelength oscillatory travelling waves are functionally equivalent—with regard to the proposed generalization mechanism—to the single, transient wave fronts assumed in the present modeling.

The synchronous components of the spatio-temporal waves do not dissipate or deform as they pass through each other

Modeling of spatio-temporal wave phenomena within physiologically realistic simulations of cortex shows

damped travelling waves emerging from sites of focal activation (Liley et al. 1999; Wright et al. 2001). The modeling shows that dynamical activity of the cortex can be described by linear wave equations obeying the principle of superposition (Chapman et al. 2002; Robinson 2006). Consistent with this principle, measurement of phase cones in the neocortex reveals multiple, overlapping waves exist at any given time (Freeman and Barrie 2000). Accordingly, in our model, wave-fronts take the form of concentric circles around the set of activated points, allowing the waves to pass through each other.

Onset of synchrony and interaction between waves is fast

Modeling of synchrony phenomena in the cortex has shown the onset of synchrony between distant sites to be fast; almost as fast as axonal conduction times (Chapman et al. 2002). This is because the onset of synchrony in these models takes place as a phase transition (Freeman and Barrie 2000), rather than depending on the slow build-up of neural interactions. The latter mechanism requires multiple iterations of neuronal activation processes governed by the relatively long time constants of dendritic membranes. The similarity of time-constants between axonal conduction velocities and synchrony onset allows that STDP may likewise be rapid over long connection distances, as assumed in the present modeling.

Changes in synaptic gain are timing dependent

Synaptic potentiation (both long term and fast-acting) is assumed in the present research context to be associative (i.e. Hebbian-like) and to involve STDP. The mechanism of STDP follows from the observation that pre-synaptic triggering of EPSP that immediately precedes firing or cell depolarization can induce long term potentiation, whereas reversing the order of these events causes long term depression of synaptic gains (Abbott and Nelson 2000; Feldman 2000). STDP involves an interaction of NMDA receptor activation and back-propagation of action potentials up the dendritic tree of the post-synaptic neuron (Magee and Johnston 1997; Markram et al. 1997). STDP-based learning rules have been used to model theta precession in the hippocampus (Wu and Yamaguchi 2004); related learning rules that more explicitly depend on the phase of activity have also been used to model theta precession (Scarpetta and Marinaro 2005). The present modeling assumes a population analogy to STDP effects, based on sensitivity to coincidence, similar to the learning rule of Scarpetta and Marinaro (2005).

Fast acting synaptic changes are associative

The focus of the present research is on fast-acting changes in synaptic gain (0.1–1 s), long thought to underlie perceptual processes and short-term memory (Konen and von der Malsburg 1993; von der Malsburg 1994; Sandberg et al. 2003). While evidence for fast-acting associative learning remains limited, one potential class of mechanisms involves retrograde messengers that are released from dendrites to act on presynaptic terminals, regulating the future release of neurotransmitter (Abbott and Nelson 2000). For example, the release of endocannabinoids leads to an inhibition of neurotransmitter release for tens of seconds (Ohno-Shosaku et al. 2001; Wilson and Nicoll 2001). Another potential mechanism for fast-acting potentiation involves voltage gated calcium channels (Zucker and Regehr 2002). In this case, increases in synaptic efficacy are due to membrane potential depolarization over a temporal window 100–200 ms prior to EPSP initiation, and are independent of an actual change in synaptic transmission (Deisz et al. 1991; Fregnac et al. 1996). Since voltage gated calcium channels are not associative, per se, this mechanism would require interactions between populations of neurons to achieve associative-like effects. Learning models that include both tonic activation and individual spiking within the formulation of the STDP rule show hetero-associative properties (Bush et al. 2010).

Discussion

Summary and implications

Consistent with the geometrical proof provided in Fig. 4, modeling showed that cortical representations can be broadcast to all points in a cortical map by means of spatio-temporal waves of activity. The experimental results show that unambiguous detection of a pattern occurs at any position in the lattice after a limited number of previous presentations, depending on the size of the lattice, the complexity of the input pattern and the level of additional noise. *This means that a pattern will be recognized as being the same anywhere on the lattice, after only a few presentations elsewhere on the lattice.* The relevance of these results to representational invariance derives from the proposed pseudo-invariant nature of the mappings. Pseudo-invariant maps transform computationally intensive symmetries, such as the multiplicative symmetries of rotation and scaling, into the additive symmetry of translation (see Fig. 3).

Each aspect of the mechanism: synchronous activation of input patterns; travelling waves; timing-dependent learning; and pseudo-invariant mapping, are supported by

known neurophysiological observations. Input patterns are primed through a mechanism of short-term increases in gain in the connections of the model, a mechanism equivalent to short-term memory (Konen and von der Malsburg 1993; von der Malsburg 1994; Sandberg et al. 2003). The generalization that results is likewise short-term. We sought to show the *in principle* efficacy of this mechanism, albeit in an idealized form. The representations are stored in a distributed, non-local fashion and the representations are broadcast by waves. The mechanism can therefore be conceptualized as a form of field computation (Agu 1988; MacLennan 1999).

Besides representing invariance through translation in a pseudo-invariant map, $M(z')$, the *amount* by which a representation is translated, T , parameterizes the transformation $f(z)$ (see Fig. 3). Thus, pseudo-invariant maps simultaneously represent object identity and object variation. In a recent scene rotation study using fMRI, the angle of rotation varied parametrically between a test and comparison scene (Nakatani et al. 2005). In posterior (retinotopic) visual areas the size of the activated regions increased with rotation angle while activation strength remained the same. This effect suggests that, the larger the extent of scene rotation, the wider the area of distributed activity evoked in these areas. Assuming the evoked perfusion reflects the range of dynamical interactions, this result can be interpreted as revealing the degree to which the scene representation was translated within retinotopically organized visual cortex.

The performance of the spatio-temporal wave mechanism improves with network size. This contrasts with various other learning algorithms, in which performance (e.g. time to convergence or probability of convergence) declines with network size (Bizzarri 1991; Erwin et al. 1992; Fine and Mukherjee 1999; Jordanov and Brown 2000). The trade-off here is the resolution of the phase variable implicit in both the wave fronts and the learning rule. The maximum realizable resolution of these phase variables sets the upper limit of performance improvement with lattice size.

All the assumptions of our model were shown in the previous section to be realizable through strictly local mechanisms in the brain. Synchronous spatio-temporal waves in the cortex are a function of short-range connectivity; they arise from the interplay of local excitatory and inhibitory connectivity as well as intrinsic oscillatory properties of individual neurons (D'Antuono et al. 2001; Wright et al. 2001; Chapman et al. 2002). This means that the mechanism of cortical generalization proposed here requires no global supervision, neither a coordinating clock, nor global error calculation. The functionality of the global system behavior is thereby entirely a product of self-organization.

Extensions of the model

In brain activity recorded with large-scale EEG arrays, we observe that travelling waves sometimes take the form of spiral waves (Ito et al. 2007) and spiral waves have also been observed in the visual cortex of turtles (Prechtl et al. 1997). Spiral waves have generalization properties similar to the 'bullseye' waves used in the present simulations. The property arises if the spiral waves are of long wavelength i.e. the phase of the wave uniquely encodes the spatial relationships (Ermentrout and Kleinfeld 2001).

The ability to detect object constancy during *temporary occlusion of features* can be easily introduced into the model, by allowing the gains to decay with a longer time constant than the typical duration of occlusion. In this case the broadcast copies of the input pattern will not be immediately removed from the system when parts or the whole of the input pattern are missing for a few presentations, and the map will remain primed for the whole object during the temporary occlusion. There is evidence for visual areas specialized for dealing with temporary occlusions. Plomp and van Leeuwen (2006) have shown the perception of occluded objects can be reinforced through priming the unoccluded object. Liu et al. (2006) used tomographic and statistical parametric mapping of the areas involved in the priming process and found the priming effect consistently localized in the right fusiform cortex between 120 and 200 ms post-stimulus. The area is one for which topographical mapping has been proposed (Malach et al. 2002); the time scale of these activities corresponds, roughly, to that of our model.

The mechanism requires a pseudo-invariant mapping in order for generalization of learning to occur. Several candidate maps are apparent from the literature, such as the primary visual cortex and primary auditory cortex. We may consider extending our approach to cases in which pseudo-invariant mapping has no mathematically tractable form. Such maps could be used for building classification schemes or taxonomies for sets of real world objects (Sutcliffe 1986), an organization that has been demonstrated for the occipito-temporal object-related visual areas (Malach et al. 2002). The construction of such pseudo-invariant mappings is possible through long term potentiation. Even when the input mapping is initially random, the short term increases in synaptic gain, described in this paper, enforce a pseudo-invariant representational scheme onto the map. Successive representations within the map will be translations of previous representations, *ceteris paribus*, due to the preferential reinforcement of $C(R)$. The long term potentiation applies to the mapping of inputs into the cortical map. So transformations in the initial domain, $f(z)$, tend to be mapped to translations, T , within the map M . Over time, long term

potentiation will build a patterning of inputs into M , such that it has the property of pseudo-invariance. Long-term potentiation thus provides the model with a mechanism to create a pseudo-invariant map from a blank slate state. Although slower than short-term generalization within an existing map of the sort described in this paper, this process of map formation is still much faster than requiring multiple presentations of every possible object under all possible transformations defined by $f(z)$.

A possible further development of the model would introduce higher-order connections to the storage network. Theoretically, we may expect system behaviour to improve when higher order connections are introduced. In short, this is because instances of non-zero weights in $B(L) \setminus C(R)$ are more common for 2nd order connections compared to non-zero weights in $B(L) \setminus C(R)$ for the case when 3rd order connections are used. 3rd order connections can be formed into equivalent sets of 2nd order connections, but the obverse is not always true. For example, a representation with four activated foci can be stored as four 3rd order or twelve 2nd order connections. But if only pairs of neurons $i, j \notin D(R)$ are coincidentally activated by wave fronts, no 3rd order connections in $B(L) \setminus C(R)$ can be formed. This effect will tend to predominate over multiple presentations, since coincidental activation of a triplet of neurons $i, j, k \in D(R)$ over p successive presentations will be a vanishingly rare occurrence compared to coincidental activation of a pair of neurons $i, j \notin D(R)$ over p successive presentations.

A further generalization of the model will involve parallel processing of multiple pseudo-invariant transformations. The concept of a pseudo-invariant map can be applied to higher dimensional spaces, in which multiple transformations are simultaneously generated via translations in multiple dimensions. The primary visual cortex, for example, can be conceived as a four dimensional mapping of visual properties. Alexander et al. (2004b) have described the global retinotopic mapping and the local mapping of response properties at the scale of $\sim 350 \mu\text{m}$ in terms of a four dimensional representational space. The different classes of intrinsic cortical connectivity in V1, short-range and long-range (patchy) connections, can be formulated as equivalent connection systems within this framework of a four dimensional representational space (Alexander et al. 2004b). Correspondingly, as reviewed in the previous section, spatio-temporal wave phenomena may arise at both the scale of the Brodmann area (Eckhorn et al. 2001) and the local scale of $\sim 350 \mu\text{m}$ response property maps in V1 (König et al. 1995). Rotation and scaling of visual objects comprise the feature transformations performed at the global scale of V1, while the orientation and spatio-temporal frequency of texture elements comprise the feature set at the local scale of $\sim 350 \mu\text{m}$. Translations of

representations within the combined higher dimensional feature space enable simultaneous generalization along each of these transformational dimensions.

The proposed mechanism of generalization of learning can also be applied to the larger scale of the whole visual cortex. Grill-Spector and Malach (2004) have argued that the various visual cortical areas form a continuum rather than discrete, specialized processing modules. Responsiveness to retinotopy, color, depth and primitive features vs. complex objects gradually change over streams of related visual areas. The entire visual cortex can be hypothesized as a global pseudo-invariant map which generalizes visual experience over the more specific invariances extracted in each individual visual area. The requisite dynamical mechanisms for this process have been demonstrated. Visually evoked P1 and N1 event-related potentials have been shown to coincide with occipital to parietal spatio-temporal waves (Klimesch et al. 2007) and more global spatio-temporal waves arise at the time of the visually evoked P2 and N2 (Alexander et al. 2007) that are predominantly posterior to anterior in direction (Alexander et al. 2006a, b).

Conclusions

In the present modeling the waves expand without distortion or noise on an idealized geometry. The proposed mechanism and its empirical confirmation are therefore limited by these idealized assumptions. However, the mechanism does not require a precise physical organization to the waves, but rather is a topological property of the relationship between the dynamics and the connectivity, bound together by the coincidence-sensitive learning rule. The positions of each column could be randomly perturbed, with corresponding changes to the connectivity structure and wave motion. The resulting simulation would appear noisier, but would essentially behave identically from the point of view of learning. The critical aspect of the mechanism is therefore the systematic relationship between coincidence-sensitive learning, synchronously generated spatio-temporal waves and map topography. When synchronous spatio-temporal waves and long-range connectivity are present in combination with a synchrony-sensitive learning rule, learning of pseudo-invariant maps and generalization of cortical representations within those maps may be a *generic* property of the cortical system.

The presence of spatio-temporal dynamics at multiple scales of cortex (global, Brodmann area, local networks) suggests that the proposed mechanism is widely applicable throughout the cortex at these multiple scales. The representations generated within cortical maps vary meaningfully

in topographical location (i.e. according to the transformation f) and are available as inputs to localized invariant representations elsewhere in cortex. The approach is therefore offered as complementary to localized learning of perceptual invariances. The critical novelty in the present research is that the generation and learning of perceptual invariances is achieved in a non-local fashion. Flowing from this property is the ability to generalize from only a limited number of examples.

The mechanism has a number of immediate implications for understanding perception and cognition, in addition to object invariance and generalization of learning. As an object varies in features over time, its pseudo-invariant representation is continuously broadcast throughout the map. This primes the map to future variation, but, paradoxically, with invariant representations. The mechanism therefore provides a powerful basis for detection of *object constancy* from moment to moment.

If the transformations are smooth in time, the mechanism enables objects to be *tracked* as their invariant representations translate smoothly along the relevant dimensions defined the pseudo-invariant map. The mechanism can also be seen as a method of forming *abstractions*. The representation is abstracted over the parameter space encoded by the pseudo-invariant map. The mechanism of abstraction is applicable to other visual areas, in addition to abstraction over the simpler features that strongly activate V1. When applied to the visual cortex as a whole, high level abstractions may be achieved via generalization across multiple visual areas.

References

- Abbott LF, Nelson SB (2000) Synaptic plasticity: taming the beast. *Nat Neurosci* 3 Suppl:1178–1183
- Adams DL, Horton JC (2003) The representation of retinal blood vessels in primate striate cortex. *J Neurosci* 23(14):5984–5997
- Agu M (1988) Field theory of pattern identification. *Physics Review A* 37(11):4415–4418
- Alexander DM, Van Leeuwen C (2010) Mapping of contextual modulation in the population response of primary visual cortex. *Cogn Neurodyn* 4(1):1–24
- Alexander DM, Bourke PD, Sheridan P, Konstandatos O, Wright JJ (2004a) Intrinsic connections in tree shrew V1 imply a global to local mapping. *Vis Res* 44(9):857–876
- Alexander DM, Sheridan P, Hintz T, Wright JJ (2004b) Specification of cortical anatomy within the spiral harmonic mosaic algebra, *Proceedings of the 2004 International Conference on Imaging Science, Systems and Technology (CISST'04)*, pp 50–56
- Alexander DM, Arns MW, Paul RH, Rowe DL, Cooper N, Esser AH, Fallahpour K, Stephan BC, Heesen E, Breteler R, Williams LM, Gordon E (2006a) EEG markers for cognitive decline in elderly subjects with subjective memory complaints. *Journal of Integrative Neuroscience* 5(1):49–74
- Alexander DM, Trengove C, Wright JJ, Boord PR, Gordon E (2006b) Measurement of phase gradients in the EEG. *J Neurosci Methods* 156(1–2):111–128
- Alexander DM, Williams LM, Gatt JM, Dobson-Stone C, Kuan SA, Todd EG, Schofield PR, Cooper NJ, Gordon E (2007) The contribution of apolipoprotein E alleles on cognitive performance and dynamic neural activity over six decades. *Biol Psychol* 75(3):229–238
- Allman J, Miezin F, McGuinness E (1985) Stimulus specific responses from beyond the classical receptive field: neurophysiological mechanisms for local–global comparisons in visual neurons. *Annu Rev Neurosci* 8:407–430
- Angelucci A, Bullier J (2003) Reaching beyond the classical receptive field of V1 neurons: horizontal or feedback axons? *Journal of Physiology, Paris* 97(2–3):141–154
- Arieli A, Shoham D, Hildesheim R, Grinvald A (1995) Coherent spatiotemporal patterns of ongoing activity revealed by real-time optical imaging coupled with single-unit recording in the cat visual-cortex. *J Neurophysiol* 73(5):2072–2093
- Balasubramanian M, Polimeni J, Schwartz EL (2002) The V1–V2–V3 complex: quasiconformal dipole maps in primate striate and extra-striate cortex. *Neural Network* 15(10):1157–1163
- Bartfeld E, Grinvald A (1992) Relationships between orientation-preference pinwheels, cytochrome oxidase blobs, and ocular-dominance columns in primate striate cortex. *Proc Natl Acad Sci USA* 89(24):11905–11909
- Benucci A, Frazor RA, Carandini M (2007) Standing waves and traveling waves distinguish two circuits in visual cortex. *Neuron* 55(1):103–117
- Bishop CM (1995) *Neural networks for pattern recognition*. Oxford University Press, Oxford
- Bizzarri AR (1991) Convergence properties of a modified Hopfield–Tank model. *Biol Cybern* 64(4):293–300
- Blakemore C, Tobin EA (1972) Lateral inhibition between orientation detectors in the cat's visual cortex. *Exp Brain Res* 15(4):439–440
- Blasdel GG, Lund JS, Fitzpatrick D (1985) Intrinsic connections of macaque striate cortex: axonal projections of cells outside lamina 4C. *J Neurosci* 5(12):3350–3369
- Booth MC, Rolls ET (1998) View-invariant representations of familiar objects by neurons in the inferior temporal visual cortex. *Cereb Cortex* 8(6):510–523
- Bosking WH, Zhang Y, Schofield B, Fitzpatrick D (1997) Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *J Neurosci* 17(6):2112–2127
- Bradski G, Grossberg S (1995) Fast learning VIEWNET architectures for recognizing 3-D objects from multiple 2-D views. *Neural Networks* 8:1053–1080
- Bush D, Philippides A, Husbands P, O'Shea M (2010) Dual coding with STDP in a spiking recurrent neural network model of the hippocampus. *Plos Comput Biol* 6(7):e1000839. doi:10.1371/journal.pcbi.1000839
- Catania KC (2002) Barrels, stripes, and fingerprints in the brain—implications for theories of cortical organization. *J Neurocytol* 31(3–5):347–358
- Chapman CL, Wright JJ, Bourke PD (2002) Spatial eigenmodes and synchronous oscillation: co-incidence detection in simulated cerebral cortex. *J Math Biol* 45(1):57–78
- Chklovskii DB, Koulakov AA (2004) Maps in the brain: what can we learn from them? *Annu Rev Neurosci* 27:369–392
- D'Antuono M, Biagini G, Tancredi V, Avoli M (2001) Electrophysiology of regular firing cells in the rat perirhinal cortex. *Hippocampus* 11(6):662–672
- Dale AM, Liu AK, Fischl BR, Buckner RL, Belliveau JW, Lewine JD, Halgren E (2000) Dynamic statistical parametric mapping: combining fMRI and MEG for high-resolution imaging of cortical activity. *Neuron* 26(1):55–67
- Deisz RA, Fortin G, Zieglgansberger W (1991) Voltage dependence of excitatory postsynaptic potentials of rat neocortical neurons. *J Neurophysiol* 65(2):371–382

- Derdikman D, Hildesheim R, Ahissar E, Arieli A, Grinvald A (2003) Imaging spatiotemporal dynamics of surround inhibition in the barrels somatosensory cortex. *J Neurosci* 23(8):3100–3105
- Duhamel JR, Bremmer F, BenHamed S, Graf W (1997) Spatial invariance of visual receptive fields in parietal cortex neurons. *Nature* 389(6653):845–848
- Eckhorn R, Frien A, Bauer R, Woelbern T, Kehr H (1993) High frequency (60–90 Hz) oscillations in primary visual cortex of awake monkey. *Neuroreport* 4(3):243–246
- Eckhorn R, Bruns A, Saam M, Gail A, Gabriel A, Brinksmeier HJ (2001) Flexible cortical gamma-band correlations suggest neural principles of visual processing. *Visual Cognition* 8(3/4/5):519–530
- Ermentrout GB, Kleinfeld D (2001) Traveling electrical waves in cortex: insights from phase dynamics and speculation on a computational role. *Neuron* 29(1):33–44
- Erwin E, Obermayer K, Schulten K (1992) Self-organizing maps: stationary states, metastability and convergence rate. *Biol Cybern* 67(1):35–45
- Feldman DE (2000) Timing-based LTP and LTD at vertical inputs to layer II/III pyramidal cells in rat barrel cortex. *Neuron* 27(1):45–56
- Ferezou I, Haiss F, Gentet LJ, Aronoff R, Weber B, Petersen CCH (2007) Spatiotemporal dynamics of cortical sensorimotor integration in behaving mice. *Neuron* 56(5):907–923
- Fine TL, Mukherjee S (1999) Parameter convergence and learning curves for neural networks. *Neural Comput* 11(3):747–770
- Földiák P (1991) Learning invariance from transformation sequences. *Neural Comput* 3:194–200
- Fox MD, Snyder AZ, Vincent JL, Corbetta M, Van Essen DC, Raichle ME (2005) The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc Natl Acad Sci USA* 102(27):9673–9678
- Freeman WJ (2003) A neurobiological theory of meaning in perception. Part II: spatial patterns of phase in gamma EEGs from primary sensory cortices reveal the dynamics of mesoscopic wave packets. *International Journal of Bifurcation and Chaos* 13(9):2513–2535
- Freeman WJ, Barrie JM (2000) Analysis of spatial patterns of phase in neocortical gamma EEGs in rabbit. *J Neurophysiol* 84(3):1266–1278
- Fregnac Y, Bringuier V, Chavane F (1996) Synaptic integration fields and associative plasticity of visual cortical cells in vivo. *Journal of Physiology, Paris* 90(5–6):367–372
- Frohlich F, McCormick DA (2010) Endogenous electric fields may guide neocortical network activity. *Neuron* 67(1):129–143
- Fukushima K (1988) Neocognitron: a hierarchical neural network capable of visual pattern recognition. *Neural Networks* 1(119–130)
- Garner WR (1962) Uncertainty and structure as psychological concepts. Wiley, New York
- Garner WR (1966) To perceive is to know. *Am Psychol* 21(1):11
- Garner WR, Clement DE (1963) Goodness of pattern and pattern uncertainty. *Journal of Verbal Learning and Verbal Behavior* 2:446–452
- Gawne TJ, Martin JM (2002) Responses of primate visual cortical neurons to stimuli presented by flash, saccade, blink, and external darkening. *J Neurophysiol* 88(5):2178–2186
- Gong PL, van Leeuwen C (2009) Distributed dynamical computation in neural circuits with propagating coherent activity patterns. *Plos Comput Biol* 5(12)
- Grill-Spector K, Malach R (2004) The human visual cortex. *Annu Rev Neurosci* 27:649–677
- Guyonneau R, Vanrullen R, Thorpe SJ (2004) Temporal codes and sparse representations: a key to understanding rapid processing in the visual system. *Journal of Physiology, Paris* 98(4–6):487–497
- Haider B, Krause MR, Duque A, Yu YG, Touryan J, Mazer JA, McCormick DA (2010) Synaptic and network mechanisms of sparse and reliable visual cortical activity during nonclassical receptive field stimulation. *Neuron* 65(1):107–121
- Hinds O, Polimeni JR, Rajendran N, Balasubramanian M, Wald LL, Augustinack JC, Wiggins G, Rosas HD, Fischl B, Schwartz EL (2008) The intrinsic shape of human and macaque primary visual cortex. *Cereb Cortex* 18(11):2586–2595
- Hoffman WC, Dodwell PC (1985) Geometric psychology generates the visual gestalt. *Canadian Journal of Psychology* 39:491–528
- Hua TM, Bao PL, Huang CB, Wang ZH, Xu JW, Zhou YF, Lu ZL (2010) Perceptual learning improves contrast sensitivity of V1 neurons in cats. *Curr Biol* 20(10):887–894
- Hubel DH, Wiesel TN (1974) Uniformity of monkey striate cortex: a parallel relationship between field size, scatter, and magnification factor. *J Comp Neurol* 158(3):295–305
- Hubener M, Shoham D, Grinvald A, Bonhoeffer T (1997) Spatial relationships among three columnar systems in cat area 17. *J Neurosci* 17(23):9270–9284
- Ito M, Tamura H, Fujita I, Tanaka K (1995) Size and position invariance of neuronal responses in monkey inferotemporal cortex. *J Neurophysiol* 73(1):218–226
- Ito J, Nikolaev AR, van Leeuwen C (2007) Dynamics of spontaneous transitions between global brain states. *Hum Brain Mapp* 28(9):904–913
- Jordanov I, Brown R (2000) Optimal design of connectivity in neural network training. *Biomed Sci Instrum* 36:27–32
- Kenet T, Bibitchkov D, Tsodyks M, Grinvald A, Arieli A (2003) Spontaneously emerging cortical representations of visual attributes. *Nature* 425(6961):954–956
- Kleinfeld D, Delaney KR (1996) Distributed representation of vibrissa movement in the upper layers of somatosensory cortex revealed with voltage-sensitive dyes. *J Comp Neurol* 375(1):89–108
- Klimesch W, Hanslmayr S, Sauseng P, Gruber WR, Doppelmayr M (2007) P1 and traveling alpha waves: evidence for evoked oscillations. *J Neurophysiol* 97(2):1311–1318
- Konen W, von der Malsburg C (1993) Learning to generalize from single examples in the dynamic link architecture. *Neural Comput* 5(5):719–735
- König P, Engel AK, Roelfsema PR, Singer W (1995) How precise is neuronal synchronization? *Neural Comput* 7:469–485
- Kording KP, König P (2001) Neurons with two sites of synaptic integration learn invariant representations. *Neural Comput* 13:2823–2849
- Kovacs G, Vogels R, Orban GA (1995) Selectivity of macaque inferior temporal neurons for partially occluded shapes. *J Neurosci* 15(3 Pt 1):1984–1997
- Lam YW, Cohen LB, Wachowiak M, Zochowski MR (2000) Odors elicit three different oscillations in the turtle olfactory bulb. *J Neurosci* 20(2):749–762
- Lashley KS (1950) In search of the engram. Society of experimental biology symposium, no. 4: Psychological mechanisms in animal behavior. Cambridge University Press, Cambridge, pp 454–480
- Liley DT, Alexander DM, Wright JJ, Aldous MD (1999) Alpha rhythm emerges from large-scale networks of realistically coupled multicompartmental model cortical neurons. *Network* 10(1):79–92
- Liu LC, Plomp G, van Leeuwen C, Ioannides AA (2006) Neural correlates of priming on occluded figure interpretation in human fusiform cortex. *Neuroscience* 141(3):1585–1597
- Livingstone MS (1996) Oscillatory firing and interneuronal correlations in squirrel monkey striate cortex. *J Neurophysiol* 75(6):2467–2485
- Livingstone MS, Hubel DH (1984) Specificity of intrinsic connections in primate primary visual cortex. *J Neurosci* 4(11):2830–2835

- Logothetis NK, Pauls J, Poggio T (1995) Shape representation in the inferior temporal cortex of monkeys. *Curr Biol* 5(5):552–563
- Lund JS, Wu CQ (1997) Local circuit neurons of macaque monkey striate cortex: IV. Neurons of laminae 1–3A. *J Comp Neurol* 384(1):109–126
- Lund JS, Angelucci A, Bressloff PC (2003) Anatomical substrates for functional columns in macaque monkey primary visual cortex. *Cereb Cortex* 13(1):15–24
- MacLennan BJ (1999) Field computation in natural and artificial intelligence. *Inf Sci* 119(1–2):73–89
- Magee JC, Johnston D (1997) A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. *Science* 275(5297):209–213
- Malach R, Levy I, Hasson U (2002) The topography of high-order human object areas. *Trends Cognition Science* 6(4):176–184
- Maldonado PE, Friedman-Hill S, Gray CM (2000) Dynamics of striate cortical activity in the alert macaque: II. Fast time scale synchronization. *Cereb Cortex* 10(11):1117–1131
- Markram H, Lubke J, Frotscher M, Sakmann B (1997) Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275(5297):213–215
- Mikolajczyk K, Schmid C (2002) An affine invariant interest point detector. In: *Proceedings, European conference on computer vision*, pp 128–142
- Nakatani C, Ueno K, van Leeuwen C, Tanaka K, Cheng K (2005) Successive recruitment of brain areas in a parameterized scene rotation task: an fMRI study. In: *11th Annual organization for human brain mapping meeting*
- Nauhaus I, Busse L, Carandini M, Ringach DL (2009) Stimulus contrast modulates functional connectivity in visual cortex. *Nat Neurosci* 12(1):70–76
- Nelson JJ, Frost BJ (1978) Orientation-selective inhibition from beyond the classic visual receptive field. *Brain Res* 139(2):359–365
- Nicolelis MAL, Baccala LA, Lin RCS, Chapin JK (1995) Sensorimotor encoding by synchronous neural ensemble activity at multiple levels of the somatosensory system. *Science* 268(5215):1353–1358
- Ohno-Shosaku T, Maejima T, Kano M (2001) Endogenous cannabinoids mediate retrograde signals from depolarized postsynaptic neurons to presynaptic terminals. *Neuron* 29(3):729–738
- Olivers CNL, Chater N, Watson DG (2004) Holography does not account for goodness: a critique of van der Helm and Leeuwenberg (1996). *Psychol Rev* 111(1):242–260
- Pantev C, Lutkenhoner B (2000) Magnetoencephalographic studies of functional organization and plasticity of the human auditory cortex. *J Clin Neurophysiol* 17(2):130–142
- Phillips WA, Kay J, Smyth D (1995) The discovery of structure by multistream networks of local processors with contextual guidance. *Network: Computing Neural System* 6:225–246
- Pitts W, McCulloch WS (1947) How we know universals the perception of auditory and visual forms. *Bull Math Biol* 9:127–147
- Plomp G, van Leeuwen C (2006) Asymmetric priming effects in visual processing of occlusion patterns. *Percept Psychophysics* 68(6):946–958
- Poggio T, Bizzi E (2004) Generalization in vision and motor control. *Nature* 431(7010):768–774
- Prechtl JC, Cohen LB, Pesaran B, Mitra PP, Kleinfeld D (1997) Visual stimuli induce waves of electrical activity in turtle cortex. *Proc Natl Acad Sci USA* 94(14):7621–7626
- Quiroga RQ, Reddy L, Kreiman G, Koch C, Fried I (2005) Invariant visual representation by single neurons in the human brain. *Nature* 435(7045):1102–1107
- Ribary U, Ioannides AA, Singh KD, Hasson R, Bolton JPR, Lado F, Mogilner A, Llinas R (1991) Magnetic-field tomography of coherent thalamocortical 40-Hz oscillations in humans. *Proc Natl Acad Sci USA* 88(24):11037–11041
- Robinson PA (2006) Gamma oscillations and visual binding. In: *15th Australian society for psychophysiology conference, University of Wollongong, Clinical EEG and Neuroscience*
- Rockland KS, Knutson T (2001) Axon collaterals of Meynert cells diverge over large portions of area V1 in the macaque monkey. *J Comp Neurol* 441(2):134–147
- Roland PE, Hanazawa A, Undeman C, Eriksson D, Tompa T, Nakamura H, Valentiniene S, Ahmed B (2006) Cortical feedback depolarization waves: a mechanism of top-down influence on early visual areas. *Proc Natl Acad Sci USA* 103(33):12586–12591
- Rolls ET (1992) Neurophysiological mechanisms underlying face processing within and beyond the temporal cortical visual areas. *Philos Trans R Soc Lond B Biol Sci* 335(1273):11–20 (discussion 20–1)
- Rolls ET (1995) Learning mechanisms in the temporal lobe visual cortex. *Behav Brain Res* 66(1–2):177–185
- Rubino D, Robbins KA, Hatsopoulos NG (2006) Propagating waves mediate information transfer in the motor cortex. *Nat Neurosci* 9(12):1549–1557
- Sandberg A, Tegner J, Lansner A (2003) A working memory model based on fast Hebbian learning. *Network* 14(4):789–802
- Scarpetta S, Marinaro M (2005) A learning rule for place fields in a cortical model: theta phase precession as a network effect. *Hippocampus* 15(7):979–989
- Schreiner CE, Read HL, Sutter ML (2000) Modular organization of frequency integration in primary auditory cortex. *Annu Rev Neurosci* 23:501–529
- Schwartz EL (1980) Computational anatomy and functional architecture of striate cortex: a spatial mapping approach to perceptual coding. *Vis Res* 20(8):645–669
- Senseman DM, Robbins KA (2002) High-speed VSD Imaging of visually evoked cortical waves: decomposition into intra- and intercortical wave motions. *J Neurophysiol* 87(3):1499–1514
- Sheridan P (1996) Spiral architecture for machine vision. University of Technology, Sydney. Ph.D
- Sheridan P, Hintz T, Alexander DM (2000) Pseudo-invariant image transformations on a hexagonal lattice. *Image Vis Comput* 18:907–918
- Singer W, Gray CM (1995) Visual feature integration and the temporal correlation hypothesis. *Annu Rev Neurosci* 18:555–586
- Stettler DD, Das A, Bennett J, Gilbert CD (2002) Lateral connectivity and contextual interactions in macaque primary visual cortex. *Neuron* 36(4):739–750
- Stone JV, Bray AJ (1995) A learning rule for extracting spatio-temporal invariances. *Network: Computing Neural Syst* 6:429–436
- Sutcliffe JP (1986) Differential ordering of objects and attributes. *Psychometrika* 51(2):209–240
- Swindale NV (1996) The development of topography in the visual cortex: a review of models. *Network* 7(2):161–247
- Terashima H, Hosoya H (2009) Sparse codes of harmonic natural sounds and their modulatory interactions. *Network Computation in Neural Systems* 20(4):253–267
- Tootell RB, Hamilton SL, Switkes E (1988a) Functional anatomy of macaque striate cortex. IV. Contrast and magno-parvo streams. *J Neurosci* 8(5):1594–1609
- Tootell RB, Silverman MS, Hamilton SL, Switkes E, De Valois RL (1988b) Functional anatomy of macaque striate cortex. V. Spatial frequency. *J Neurosci* 8(5):1610–1624
- Tootell RB, Switkes E, Silverman MS, Hamilton SL (1988c) Functional anatomy of macaque striate cortex. II. Retinotopic organization. *J Neurosci* 8(5):1531–1568

- Tucker TR, Katz LC (2003) Spatiotemporal patterns of excitation and inhibition evoked by the horizontal network in layer 2/3 of ferret visual cortex. *J Neurophysiol* 89(1):488–500
- Tuytelaars T, van Gool LV (2000) Wide baseline stereo matching based on local, affinely invariant regions. British machine vision conference
- Ullman S, Bart E (2004) Recognition invariance obtained by extended and invariant features. *Neural Network* 17(5–6):833–848
- van der Helm PA, Leeuwenberg ELJ (1996) Goodness of visual regularities: a nontransformational approach. *Psychol Rev* 103(3):429–456
- van der Helm PA, Leeuwenberg ELJ (2004) Holographic goodness is not that bad: reply to Olivers, Chater, and Watson (2004). *Psychol Rev* 111:261–273
- Vanduffel W, Tootell RB, Schoups AA, Orban GA (2002) The organization of orientation selectivity throughout macaque visual cortex. *Cereb Cortex* 12(6):647–662
- Vincent JL, Patel GH, Fox MD, Snyder AZ, Baker JT, Van Essen DC, Zempel JM, Snyder LH, Corbetta M, Raichle ME (2007) Intrinsic functional architecture in the anaesthetized monkey brain. *Nature* 447(7140):83–84
- Vinje WE, Gallant JL (2000) Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287(5456):1273–1276
- von der Malsburg C (1994) The correlation theory of brain function. In: Domany E, van Hemmen JL, Schulten K (eds) *Models of neural networks*, vol 2. Springer, Berlin, pp 95–119
- von der Malsburg C, Schneider W (1986) A neural cocktail-party processor. *Biol Cybern* 54(1):29–40
- Wagemans J (1999) Toward a better approach to goodness: Comments on Van der Helm and Leeuwenberg (1996). *Psychol Rev* 106(3):610–621
- Wilson RI, Nicoll RA (2001) Endogenous cannabinoids mediate retrograde signalling at hippocampal synapses. *Nature* 410(6828):588–592
- Wright JJ, Robinson PA, Rennie CJ, Gordon E, Bourke PD, Chapman CL, Hawthorn N, Lees GJ, Alexander D (2001) Toward an integrated continuum model of cerebral dynamics: the cerebral rhythms, synchronous oscillation and cortical stability. *Biosystems* 63(1–3):71–88
- Wu Z, Yamaguchi Y (2004) Input-dependent learning rule for the memory of spatiotemporal sequences in hippocampal network with theta phase precession. *Biol Cybern* 90(2):113–124
- Xu W, Huang X, Takagaki K, Wu JY (2007) Compression and reflection of visually evoked cortical waves. *Neuron* 55(1):119–129
- Yao H, Shi L, Han F, Gao H, Dan Y (2007) Rapid learning in cortical coding of visual scenes. *Nat Neurosci* 10(6):772–778
- Yoshioka T, Blasdel GG, Levitt JB, Lund JS (1996) Relation between patterns of intrinsic lateral connectivity, ocular dominance, and cytochrome oxidase-reactive regions in macaque monkey striate cortex. *Cereb Cortex* 6(2):297–310
- Zucker RS, Regehr WG (2002) Short-term synaptic plasticity. *Annu Rev Physiol* 64:355–405