Long-range cortical connections give rise to a robust velocity map of V1

Author

Sheridan, Phillip

Published

2015

Journal Title

Neural Networks

Version

Accepted Manuscript (AM)

DOI

10.1016/j.neunet.2015.08.005

Long-range Cortical Connections Give Rise to a Robust Velocity Map of V1

Phillip Sheridan (corresponding author)

School of Information and Communication Technology, Griffith University, University Drive, Meadowbrook, Qld, Australia, p.sheridan@griffith.edu.au

Abstract

This paper proposes a two-dimensional velocity model (2DVM) of the primary visual cortex (V1). The model's novel aspect is that it specifies a particular pattern of long-range cortical temporal connections, via the Connection Algorithm, and shows how the addition of these connections to well-known spatial properties of V1 transforms V1 into a velocity map. The map implies a number of organisational properties of V1: 1) the singularity of each orientation pinwheel contributes to the detection of slow-moving spots across the visual field; 2) the speed component of neuronal velocity selectivity decreases monotonically across each joint orientation contour line for parallel motion and increases monotonically for orthogonal motion; 3) the cells that are direction selective to slow-moving objects are situated in the periphery of V1; and 4) neurons in distinct pinwheels tend to be connected to neurons with similar tuning preferences in other pinwheels. The model accounts for various types of known illusionary perceptions of human vision: perceptual filling-in, illusionary orientation and visual crowding. The three distinguishing features of 2DVM are: 1) it unifies the functional properties of the conventional energy model of V1; 2) it directly relates the functional properties to the known structure of the upper layers of V1; and 3) it implies that the spatial selectivity features of V1 are side effects of its more important role as a velocity map of the visual field.

Keywords: Gabor filter, contextual modulation, velocity model, primary visual cortex, orientation pinwheels, motion streaks

1 Introduction

The most casual observations of biological vision systems invariably lead to the conclusion that the ability to predict an object's velocity is fundamental to vision. In simple terms, in the absence of motion in the animal world, nobody eats or gets eaten. Ever since the Cambrian explosion, that period of earth history where the number of phylum increased from four to 38 and when the

genetic code for vision first appeared, velocity estimation from visual inputs has been driving an evolutionary arms race [1].

1.1 Visual Perception of Motion

The intensity of vision research over the past few decades has reflected the importance of understanding how biological vision systems reliably perceive motion. Although computational models of velocity have undergone substantial development, empirical experiments continue to reveal new properties of the vision system that are not accounted for by these models. Central to this and other papers is the concept of the receptive field (RF) which we interpret as: a position in the visual field that, if a stimulus is placed there, will excite a population of neurons in V1. Experiments on single neuron responses of cats to certain types of motion reported in [2] produced surprising response properties of V1 cells to two types of object motion: the first type had objects moving quickly parallel to the RF's orientation (parallel motion); the second type had the objects moving slowly orthogonally to the RF's orientation (orthogonal motion). Although they reported that cells had the same temporal frequency response to both types of motion, to the former type, cells tended to be selective to motion direction and band pass selective to spatial frequency, while to the latter, the same cell tended to be non-selective to motion direction and low pass selective to spatial frequency.

Two investigations of the orientation tuning curves of simple and complex cells in response to moving and flashing bars of different lengths produced results that were seemingly at odds with each other. The conclusion reported in [3] was that the orientation tuning curve was always unimodal (i.e. it exhibited just a single peak corresponding with the cells optimal orientation) whereas the shape of the orientation curve was markedly dependent on the length of the bar, the curve sharpening with increasing length. These results implied that, if the bar length was reduced to a moving spot, the optimal direction was unchanged and remained orthogonal to the longer axis of the RF. In contrast, results reported [4] that the optimal direction for a small spot was generally orthogonal to that of a long bar. Hence, a cell responded to two perpendicular axes of movement: the first for long bars, which was perpendicular to the preferred orientation, i.e. perpendicular to the long axis of the RF (orientation component); the second for short bars or spots, parallel to the preferred orientation (axial component). At an intermediate bar length, the two axes of movement might coexist. Two distinct models were employed to account for these apparent divergent results [5].

A third category of intriguing experiments investigated the perception of visual illusions. Results reported in [6] showed that long, drifting lines could activate

a V1 neuron even when the RF of that neuron was covered by a mask many times the size of the RF (perceptual filling in). Another set of experimental results reported in [7] revealed the perception of illusionary orientation. When a moving bar orientated obliquely to its parallel motion was presented to an observer, two related phenomenon were demonstrated: 1) with the bar length held constant, the perceived orientation approached the parallel direction of motion as the speed increased; 2) with the speed held constant, the perceived orientation approached the actual orientation of the bar with increasing bar length. Consequently, they concluded that V1 neurons did not act as feature detectors that isolate individual stimulus attributes. A third visual illusion is a variation on the visual crowding phenomenon reported in [8].

1.2 Overview of Model

It has been shown that models representing variations on the standard energy model [9] and [10] successfully demonstrated the inter-dependency between attributes of orientation, spatial frequency, temporal frequency, length, speed and direction. These spatial characteristics are used to predict purely functional aspects of V1 neurons. The model presented in this paper, named 2-D Velocity Model (2DVM), extends this energy paradigm by relating neuronal function to the structure of Layer 2/3 of V1 via a novel Long Range Cortical Connection Algorithm. Two important structural aspects of Layer 2/3 of V1 are central to 2DVM: 1) a local structure of patches of neurons known as pinwheels describe a pattern of neuronal orientation preferences and spatial frequency preferences; 2) a global structure is implied by the pattern of connectivity of long-range cortical connections between pinwheels. 2DVM thus represents a structural as well as a functional velocity model of the primary visual cortex.

The properties of 2DVM, to be discussed later, are summarized as follows:

- *Functional properties that 2DVM was designed to account for include*:
  - Velocity selectivity: different cells in a local map respond optimally to stimuli moving at unique velocities;
  - Suppression/facilitation: responses of cells lying along an orientation contour line are suppressed by cells in the same relative position in local maps orthogonal to the cells preferred orientation and tend to be excited by cells in the same relative position in local maps aligned with the cell's preferred contour line;
  - Orthogonal pair selectivity: cells are selective to orthogonal velocity pairs, fast moving in the parallel direction and slow moving in the orthogonal direction;
  - Orientation selectivity: different cells within a local map respond prefer-

entially to stimuli of different orientation. Cells responding to stimuli of the same orientation lie along iso-orientation lines radiating from a cell with no preferred orientation, known as the singularity;
  - Spatial frequency selectivity: different cells lying along an orientation contour line of a local map respond preferentially to different spatial frequency where the singularity has a preference for the lowest spatial frequencies;
  - Bar length selectivity: cells are selective to bar length.
- *Structural properties that 2DVM was designed to account for include*:
  - At the global scale, the location of each pinwheel is associated with a position in the visual field;
  - At the local scale, each point in a pinwheel is associated with a particular orientation, frequency and velocity.

In addition to the intended design effects, the model also produced the following:

- *Side effects of the model's design that coincide with known properties of V1 include*:
  - Neurons tend to be direction selective to parallel motion;
  - Neurons tend to be non-direction selective to orthogonal motion;
  - Perceptual filling in: drifting lines can activate a V1 neuron even when the neuron's associated RF is occluded by a mask many times its size;
  - Illusionary orientation: when a perceived orientation of a stimulus is attracted towards the direction of its motion.
- *Side effects of the model's design that predict structural properties of V1 that are as yet untested by empirical experiments*:
  - The speed component of neuronal velocity selectivity decreases monotonically across each joint orientation contour line for parallel motion and increases monotonically for orthogonal motion;
  - The neurons that are direction selective to slow-moving objects are situated in the periphery of V1.

Assumption: Theory predicts that when an image moves with sufficient speed across the visual field, the captured image (or neural image) becomes smeared due to the temporal integration properties of the retinal receptors. [11]. The perception of these effects, such as the patterns produced by waving a sparkler against the night sky, is known as motion streaks. 2DVM's design is predicated on the assumption that motion streaks are projected into V1.

Three key features of 2DVM are: 1) it unifies the functional properties of the conventional energy model of V1; 2) it directly relates the functional properties to the known structure of the upper layers of V1; and 3) it makes predictions about as yet untested hypotheses of V1.

1.3 Organization of Paper

This paper is structured as follows: in Section 2 we outline the background information that underpins the development of 2DVM, in Section 3 we provide the mathematical formulation of the model, and Section 4 focuses on what the model does by presenting simulation results for key response properties of 2DVM. Section 5 discusses implications of the model and the paper is summarized and concluded in Section 6.

2 Background

In this section we provide the background information that underpins the development of 2DVM. The spatial responses to oriented stimuli of neurons in the primate V1 have been known for some time [12]. Since then, it has been generally accepted that these tuning functions of RFs were largely context independent [13]. More recent research has demonstrated contextual influences from the region close to the RF [14], [15], [16]. The near surround region of an RF can modify RF responses through suppression, [17] by collinear facilitation [18] and by cross-orientation facilitation effects [19], [15], [20]. While the mechanisms underlying the modulating effects of suppression and facilitation immediately surrounding the region of an RF are pretty well understood, the modulating effects of long-range influences have not as yet received the same degree of attention. However, the evidence for long-range contextual modulation is increasing, e.g. [21], [22], [23], [24], and has important implications for the organization of V1. The Local-Global Map hypothesis [25] states that orientation pinwheels constitute fundamental building blocks in Layer 2/3, which together describe the global organization of cortical cell response properties. This concept was further developed in [26] where an abstraction of the orientation pinwheels, called local maps, that tile V1 at the scale of the pinwheel demonstrate how functional units of contextual modulation are related to orientation pinwheels. It was hypothesized that these local maps might also have significance for other response properties in V1. The velocity model presented in this paper extends the Local-Global Map hypothesis to show how the spatial properties predicted by the local map are fundamental building blocks to a velocity map of V1.

2DVM draws heavily on the concept of the Gabor function to implement this velocity map. The response properties of neurons in V1 have been shown to be well modelled by convolving the input image with a family of Gabor functions [27]. Further, research has shown that the upper layers of Area V1 are modelled well by a multi-dimensional Gabor space [28]. Most of the multi-dimensional Gabor space models of V1 described to date employed short-range convolu-

tions at the scale of the RF [29], [30], [31], [32]. By contrast, 2DVM extends the Gabor space paradigm by incorporating both long-range and short-range contextual modulation. Although there are many inputs via short-range connections from cells close to a particular local map, the inputs via long-range connections from distant ones are the means by which the high velocity selectivity embodied in this model is achieved.

3 Methods

This section presents a formal mathematical description of 2DVM. We will first describe the basic spatial model that employs one cycle of a plaid pattern to produce a Gabor response function. Then we will show how this can be extended to include temporal suppression and facilitation to describe a computational velocity model of V1.

3.1 Base (Spatial) Model

The classic Gabor response function is the product of a sinusoidal, known as the carrier (car), and a Gaussian, known as the envelope (env). 2DVM employs a variant of the classic Gabor function, named One-Cycle-Plaid-Gabor function (OCP-Gabor). The (car) of OCP-Gabor is a plaid pattern. This is the product of a pair of orthogonally orientated sinusoidal gratings. Fig. 1 displays an example of a plaid pattern, where the top two images are the orthogonal sinusoidal components of the plaid and the bottom left image is the plaid that results from the product of those top two images. A plaid pattern can be defined formally as follows:

$$S_{ew} = \cos(2\pi F_{ew}(x\cos(\omega_0) + y\sin(\omega_0)) + \phi_{ew})$$

$$S_{ns} = \cos(2\pi F_{ns}(x(-\sin(\omega_0)) + y\cos(\omega_0)) + \phi_{ns})$$

$$car(x, y, F_{ew}, F_{ns}, \omega_0, \phi_{ew}, \phi_{ns}) = S_{ew}S_{ns} \tag{1}$$

$env\,(x, y, \theta, \sigma_{ew}, \sigma_{ns},) = Kexp\,(-\pi(\sigma_{ew}^2(x - x_0)_\theta^2 + \sigma_{ns}^2(y - y_0)_\theta^2))$, where

$$(x - x_0)_\theta = (x - x_0)\cos(\theta) + (y - y_0)\sin(\theta)$$

$$(y - y_0)_\theta = -(x - x_0)\sin(\theta) + (y - y_0)\cos(\theta)$$

The response function of OCP-Gabor, $g\,()$, is then:

$$g\,(x, y, \theta, \sigma_{ew}, \sigma_{ns}, F_{ew}, F_{ns}, \omega_0, \phi_{ew}, \phi_{ns}) = env \times car \tag{2}$$

where:

- F_{ew}, F_{ns}: spatial frequencies of the east-west (ew) and north-south (ns) sinusoidal carrier
- ω_0: angle of orientation of the plaid pattern
- ϕ_{ew} and ϕ_{ns}: phase of the east-west (ew) and north-south (ns) components of the plaid carrier
- K: scale of the magnitude of the Gaussian envelope
- θ: rotation angle of Gaussian envelope
- (σ_{ew}, σ_{ns}): scale of the two axes of the Gaussian envelope
- (x_0, y_0): location of the peak of the Gaussian envelope

Fig. 1, displays an example of the OCP-Gabor where the bottom left image represents the carrier and the bottom right represents the product of the carrier with a Gaussian envelope. Note that this is subtly but importantly different from the classical Gabor function. While both the single sinusoidal carrier and one cycle of the plaid result in an OCP-Gabor function that has an orientation and spatial frequency preference, a one cycle plaid also induces a preference for bar length. This preference is a mathematical consequence of choosing the envelope so that it covers one cycle in the preferred direction of orientation. We note that our choice of the Gaussian parameters restricts the effect of the plaid to one cycle in the orthogonal direction. Consequently, the resulting Gabor function is only selective to one cycle of the plaid and not selective to the full plaid pattern. Therefore, OCP-Gabor does not contradict the well-known result that MT neurons are selective to the motion of complete plaid patterns and V1 neurons are not [33].

We define a Gabor field, $GF()$, as the convolution of an input signal, $i(x, y)$, with a OCP-Gabor response function as follows:

$$GF_{i,\theta,\sigma_{ew},\sigma_{ns},F_{ew},F_{ns},\omega_0,\phi_{ew},\phi_{ns}}(x, y) = i(x, y) * g(x, y, \theta, \sigma_{ew}, \sigma_{ns}, F_{ew}, F_{ns}, \omega_0, \phi_{ew}, \phi_{ns})$$

With this OCP-Gabor field defined, we now assemble a family of Gabor fields employing the following notation:

- $roi(p, q)$: represent a disc of a fixed radius centred on p and q
- I: a set of input signals (visual field)
- Ω: finite and discrete set of coordinate pairs (p, q)
- Θ: a set of orientation angles represented as real values
- B: a set of Gaussian scales represented as real values.
- Φ: $((0, \pi/2), (0, -\pi/2), (\pi/2, 0), (-\pi/2, 0))$, a set of four ordered pairs of real values that represent the phase of the east-west and north-south components of the plaid carrier

From this we define a mapping:

$$2DVM : (I, \Omega, \Theta, B, \Phi) \Rightarrow R^4$$

as

$$2DVM(i, p, q, \theta, \sigma_{ew}, \sigma_{ns}, F_{ew}, F_{ns}, \phi_{ew}, \phi_{ns}) = \sum_{x,y \in roi(p,q)} GF_{i,\theta,\sigma_{ew},\sigma_{ns},F_{ew},F_{ns},\omega_0,\phi_{ew},\phi_{ns}}(x, y) \quad (3)$$

Throughout the remainder of this paper, we will adopt the convention that ω_0 has the same value as θ with $x_0 = y_0 = 0$.

Although the 2DVM output space possesses four dimensions (two prescribing the internal spatial layout of the local maps, and two prescribing the spatial layout of the collection of local maps), it can be represented very naturally in two dimensions. This is similar to Layer 2/3 of Area V1, which is conceptually a multi-dimensional space but it is organized physically as a two-dimensional sheet of neurons [34], [35], [36]. Layer 2/3 appears as a patchwork of structures known as orientation pinwheels or local maps, a structure that is integral to this model. Each orientation pinwheel has a singularity from which iso-orientation contour lines project. The term singularity identifies the confluence of the pinwheel's iso-orientation contours and thus possesses no preferred orientation. Each orientation pinwheel is associated with an RF [37]. When an orientated stimulus is presented to the visual field, orientation pinwheels associated with the RF that the stimulus input signal is projected onto will a evoke strong response along the preferred iso-orientation contour [38]. Layer 2/3 of Area V1 is composed of approximately 10,000 orientation pinwheels [39]. Fig. 2 depicts this organizational structure where it can be observed that the orientation contours span 180 degrees. We adopt the convention in this paper that when we refer to the physical collection of neurons that are organized as a pinwheel-like structure we employ the term 'pinwheel' and when we refer to the properties attributed to these neurons, we use the term 'local map'. It was suggested in [25] that these pinwheel contour lines be modelled by a uniform spread of angles over 360 degrees. Taking a local map in isolation, the contour lines of the pinwheels possess 180 degree symmetry (Fig. 2). However, in the context of the global map and for the purposes of the velocity model of this paper, we require an additional property that is implied in a figure reported in [40] and displayed in Fig. 2, one of 90 degree symmetry globally. To this end, we divide a circle into four quadrants. Label the four sectors as West, North, East and South to reflect their relative positions to each other. In the West sector, we place the contour lines depicted in the left half of Fig. 2 so that the angle of each contour line as determined by the coordinate system coincides with its preferred orientation. In the East sector, (right half of Fig. 2) the angle of each contour line as determined by the coordinate system is orthogonal to its preferred orientation. Note that the North and South sectors are devoid

of contour lines and that the West and East segments of every straight line passing through the singularity represent orthogonal orientations. This is an important property of our model that we will return to shortly. A discrete form of these relations is depicted in Fig. 3 with the following interpretation:

- The singularity is depicted by the darkened cell in the centre
- Each of the cells containing a pair of numbers represents a point along a contour line where:
  - the first number represents its position out from the singularity
  - the second number represents the angle (degrees) of the contour line it lies on
- The cells that contain a dash (–) are inactive.

In 2DVM, we associate each pair of numbers in the cells of Fig. 3 with a particular OCP-Gabor function in the following manner: the specified angle of each cell determines the angle of orientation of the carrier and envelope, ω_0 of Eq. 2. The other number of the pair, which specifies the cell's position out from the singularity, determines the values of the four parameters, two governing the plaid's spatial frequency and two, the Gaussian's spread, as follows: let $1/k$ denote a spatial frequency for some positive integer k. Let F^i_{ew} and F^i_{ns} represent the F_{ew} and F_{ns} and let σ^i_{ew} and σ^i_{ns} represent the σ_{ew} and σ_{ns} respectively of Eq. 2, of the *ith* cell out from the singularity. Let the Few and F_{ns} values in the West sector be given by:

$$F^i_{ew} = 1/(k2^i),\ F^i_{ns} = 1/(k2^{-i})$$

$$\sigma^i_{ew} = k2^{i+1},\ \sigma^i_{ns} = k2^{-(i+1)})$$

and their counterparts in the East sector be given by:

$$F^i_{ew} = 1/(k2^{-i}),\ F^i_{ns} = 1/(k2^i)$$

$$\sigma^i_{ew} = k2^{-(i+1)},\ \sigma^i_{ns} = k2^{(i+1)})$$

This specification of parameters has the effect of associating a unique OCP-Gabor response function with each cell of a local map. A few examples of the OCP-Gabor function associated with cells of a local map are displayed in Fig. 4 where each of the six images represents an intensity map of an OCP-Gabor function. In this figure, the six sub-images all display distinct preferences: *Top*: left for the narrow light/dark edge of a long motion streak; right for its associated wide light/dark edge; *Middle*: left has preference for the narrow light/dark edge of a shorter motion streak; right is its associated low pass wide light/dark edge; *Bottom*: left for the light/dark edge of a horizontal pair of spots; right is its associated light/dark vertical edge; *top two pairs*: for horizontal orientation; *bottom pair*: no orientation preference; *left-hand filters of top two pairs*: band-pass; *right-hand filters of those*: low-pass; *bottom of three*

pairs: will give the strongest relative response to the lowest spatial frequency. These spatial preferences are consistent with the empirical results derived in [41]; [42].

In summary of the main features of the spatial aspects of 2DVM, we observe that: 1) the novel aspect of the spatial component is that it employs a plaid function as the carrier of its OCP-Gabor functions; 2) the global organization of the model is a tiling of pinwheels that mimic the topological organization of the visual field.

3.2 Velocity (Temporal) Extension

Our base model is a purely spatial model of Layer 2/3. The cells of the local maps have been designed to be selective to orientation, spatial frequency and bar length. In addition, the parallel component of each cell is a band-pass filter for frequency and the orthogonal component is a low-pass filter. 2DVM transforms this spatial model into a velocity model using a novel Connection Algorithm. To this end, consider a pattern of inter-local map connectivity described as follows: each cell along each contour of each local map possesses connections to cells in the same relative position in four other local maps. The four connected local maps are determined by the orientation line and its distance from its singularity, i.e. the two quantities specified in the cells of Fig 3. The end points of each set of four connections are determined by the following Connection Algorithm.

Connection Algorithm Given a cell (current cell) lying on a particular orientation contour line of a particular local map (current local map):

(1) Enumerate the cells along the contour line out from the local map's singularity.
(2) Let n represent the number of cells in the orientation contour line and p the position of the current cell out from the singularity.
(3) Enumerate the local maps on either side of the current local map that lie along a line passing through the singularity and orientated as the cell's orientation contour line.
(4) Count (n+p) local maps out from the current local map in both directions and connect the current cell to a cell in a similar position in each of these two local maps.
(5) Enumerate the local maps on either side of the current local map that lie along a line passing through the singularity and orientated orthogonally to the cell's preferred angle of orientation.
(6) Count (n-p+1) local maps out from the current local map in both directions and connect the current cell to a cell in a similar position in each

of these two local maps.

An example of the connectivity pattern generated by the Connection Algorithm is depicted in Fig 5. The figure displays the simplest case of the connections associated with a single contour line that consists of five cells. An important feature of the connectivity pattern established by this algorithm is that it relates the global distribution of local maps to the local structure of each local map in a disciplined manner. Specifically, if we consider the sequence of cells that span a given joint contour line of a local map, it can be observed that the cell located at the left-hand end is connected to the local map furthest away in the parallel direction. (In Fig. 5, note the yellow line connecting Cell 5A to 1A.) The second cell from the end is connected to the second furthest local map in the parallel direction. This pattern is continued for each subsequent cell along the line across the singularity to the opposite end of this joint contour line. At the extreme right-hand end of this sequence, the last cell is connected to the closest local map in the parallel direction. (In Fig. 5, see red line connecting 5E to 4E.) As the mirror image of the connectivity pattern also applies, each cell possesses four-way connectivity, thus the local map's connectivity possesses 90 degree symmetry.

The model also specifies these connections as being excitatory over one unit of time. That is, each cell at the current time receives inputs, through these connections, from four cells at the preceding time slice, two from the parallel direction and the other two from the orthogonal direction. The parallel connections transmit the response to a moving bar's narrow edges and the orthogonal connections transmit the responses to a moving bar's long edges. The non-singularity edge detectors have aspect ratios greater than one, while the singularity has an aspect ratio value of one and a preference for the lowest spatial frequencies in its local map. In our model, each singularity is associated with four OCP-Gabor filters, each associated with a phase pair, (ϕ_{ew}, ϕ_{ns}), as specified in Φ. Although the net effect is a non-preferential spatial orientation response, the effects on the velocity maps are significant. Individually, each filter will respond differently to the edge bordering a light/dark or dark/light pair of spots aligned horizontally or vertically. The role of this set of four Gabor filters in our model is to detect slow-moving spots that would otherwise be undetectable by the filters with aspect ratios greater than one. The need to include this property results from the anatomical evidence that shows that intrinsic connections to and from singularities are predominantly of short-range, whereas long-range intrinsic connections arise between iso-orientation domains [43]; [44].

Another issue the Connection Algorithm addresses is keeping track of the individual stimuli when multiple stimuli move across the visual field at the same velocity, such as each bird in a flock or each plane flying in a formation, the correspondence problem. Our model deals with this by employing the well-

known phenomenon of suppression/facilitation. Surround suppression is the name given to the inhibitory modulating effects of similarly-orientated stimuli in the near surround of an RF. Collinear facilitation is the facilitatory effect that co-orientated stimuli situated in the near surround of an RF can have on the neuron associated with that RF. We specify an additional set of inhibitory/excitatory connections with the Connection Algorithm to address this aspect of the correspondence problem. Specifically, as well as propagating the signal to a local map in the subsequent time slice, the additional connections also propagate the signal to the local map at the same time slice. The cells lying along parallel contour lines are propagated in an inhibitory manner while cells along aligned contour lines are propagated in an excitatory manner. The parallel components induce surround suppression and the aligned components induce collinear facilitation. In order to ensure that only internal members of a group of stimuli are suppressed/facilitated, the weights of these connections are determined dynamically depending on the state of the surround area. There are many ways in which these weights can be formulated [29]. In this paper, we employ a simple mathematical tool to produce a maximum weight (inhibitory/excitatory) when the two incoming signals are perfectly balanced achieving a value of 1/4. The weight's magnitude decreases with increased imbalance. This connection weight is expressed mathematically as:

Let r_1 represent the response received from one of the orthogonal connections and r_2 the response from the opposite connection. The weight is given by $((r_1/(r_1 + r_2))(r_2/(r_1 + r_2)))$. It can be observed that this product is a maximum, with a value of 0.25, when $r_1 = r_2$ and is a minimum when r_1 or $r_2 = 0$.

The final local step for each cell receiving inputs from two distinct time slices, before summing these inputs, is to apply the function pLog() to the inputs, which returns the logarithm of the input if the input is greater than one and zero otherwise. The effect of pLog() ensures that unusually large values do not exert too large an influence on the velocity computation. The final global step is to perform a half-wave rectification on the signal. This process has the effect of ignoring any small values and is a component of many models of V1 [2].

3.3 Summary of Velocity (Temporal) Connections

The Connection Algorithm was employed to specify two sets of inter-local map connections within V1. One set connected local maps across time slices in a purely excitatory manner. The other set connected local maps within the same time slice and was a combination of inhibitory orthogonal connections and

excitatory parallel connections. The combined effects of the spatial selectivity properties with the time-synchronized excitatory connections transmitting the outputs of bar/edge detectors from the base spatial model are summarized as follows. The presentation of a stimulus to an RF will evoke a response by the cells in the RF's associated local map. However, the strongest response will be by the cell whose preferred orientation and preferred spatial frequency is closest to that of the presented stimulus. The temporal connections will then propagate these responses to the population of local maps according to the connectivity pattern previously described. If the stimulus then moves to another RF in the visual field in the next unit of time, a similar response at this new local map will be made. At this point in time, the local map's cells receive inputs via their connections and this value is summed with its current value. All non-zero values represent a measure of the stimulus change in location (distance and direction) over one unit of time. Of these, the cell with the peak response establishes the stimulus velocity at the current time. Moreover, as a result of the 90 degree symmetry of connectivity, the cell with the strongest response will represent an estimate of four possible stimulus velocities.

The key features of the temporal aspect of 2DVM are: 1) the Connection Algorithm, which specifies a pattern of connectivity between cells in distinct local maps; 2) this pattern also specifies the global organization of V1.

3.4 Implementation

In this section, we briefly discuss an important feature of the particular implementation of 2DVM employed to generate the simulation results reported in the next section. Although the organization of V1 can be represented in polar coordinates, a more convenient and efficient data structure is Spiral Honeycomb Image Algebra (SHIA) [45], [46]. SHIA is finite and discrete, allowing it to more closely model the properties of the layered structure of V1. These two properties allow SHIA to eloquently capture and manipulate the multi-dimensional aspects of the layers of V1 that are awkwardly achieved in polar coordinates. Another attractive feature of SHIA pertinent to 2DVM is that it facilitates the interactions inherent in the global-local structure. That is, both the relationships between cells as part of the global map and the cell as part of a local map are captured with an algebraic operation of vector multiplication. This feature is particularly useful when manipulating the pattern of inter-local map connectivity as described above. Note, the shape displayed in the boundary of the local and global maps of 2DVM is a side-effect of the hexagonal data structure and is not directly pertinent to the modelling presented in this paper. Consequently, the reader is cautioned not to be side-tracked by an unintended meaning of the patterns produced by these borders.

The implementation of 2DVM that produced the results reported in this paper along with Java source code is available in a software package named VisioLab from the following website: http://www.ict.griffith.edu.au/sheridan/.

4 Simulation results

In this section we present simulation results to demonstrate various aspects of 2DVM. In the first subsection we show that 2DVM achieves its key design goal of modelling V1 as a velocity map. The second subsection presents simulation results that demonstrate how the model addresses three types of visual illusion important to human vision.

4.1 Velocity Selectivity

The results below demonstrate three important features of the model: 1) orthogonal pair selectivity: cells are selective to orthogonal velocity pairs, fast moving in the parallel direction and slow moving orthogonally, 2) velocity tuning varies monotonically across joint iso-contour lines, and 3) direction selectivity: where cells that respond to parallel motion tend to be selective to direction while cells responding to orthogonal motion do not.

Two figures presented in this section, 6 and 7, display distinct velocity responses but have a common format. As this common format is somewhat complex, the reader is advised to carefully consider the remainder of this paragraph to understand the intended features of the simulation results displayed. Each row of three sub-images of each figure has the following interpretation: the image on the left, the fractal-like border of the dark region represents the boundary of the simulated visual field, and the circles identify the location of 49 simulated RFs that we denote with rf in the model. The image displays two oriented bars which are half a cycle of a sinusoid of frequency 1/32 cycles/pixel wide. This frequency is set to the singularity's preferred spatial frequency. The bar, labelled A, is positioned at one of the rfs and represents the initial position of the simulated stimulus. The bar, labelled B, represents its new location one unit of time later. Therefore, we interpret the difference of the two locations as velocity. For example, the upper left image of Fig. 6 represents a change of stimulus position from the rf located at the 10 o'clock position to 4 o'clock. Note that this distance represents the length of the straight line segment connecting the two rfs and intersects three other rfs. Consequently, we say that in this case the stimulus has moved a distance spanning five rfs. The challenge of simulating a motion streak is dealt with using an elongated bar. We define fast moving as more than two rfs. Consequently, when the motion is more

than two rfs, we choose to represent this as an orientated bar parallel to the direction of motion.

The parameterization of 2DVM employed to generate the results discussed in this section is: number of local maps = 49; number of cells in a local map = 49; singularity spatial frequency = 1/32 cycles/pixel. For this parameterization, the second and third images of each row display the model output from the associated input. In these images, the fractal-like borders identify the boundaries of the simulated global velocity map and simulated local map. The centre image displays the activation levels of the cells of a particular local map. For example, in the middle image Row 1 of Fig. 6 it can be observed that the cell with the highest activation level within the local map has value 11.1 and is located in the 9 o'clock position two cells out from the centre cell. All other cells of that local map possess values less than 6.0. The right image depicts the global output signal (49 local maps) and is an intensity map where the brightness of a pixel represents the activation level: light pixels indicate high activation and dark pixels indicate low activation. All 49 local maps are displayed with each map's singularity depicted by a mark. For example, in the right image of Fig. 6 Row 1, observe that the most activated local map is centred on the mark in the 4 o'clock position of the global map. The activation values of this local map are displayed in the centre image where, it can be seen, the brightest pixel in the global map corresponds to the cell with activation value 11.1. The pixels displaying low brightness correspond to the activation levels in the range (2.3, 5.8). All other cells of the global map are dark, which mean their activation levels are close to zero.

Although 2DVM can be applied to an arbitrary number of rfs, the examples have been parameterized to 49 rfs and therefore 49 local maps. It can also be observed that this parameterization implies that each local map is also composed of 49 cells, 24 of which are identified with orientation contour lines. In the simulated velocity map, the singularities correspond on a one-to-one basis with the rf of the visual field and possess the same spacing as the rf. This spacing property is purely intended to facilitate the reader's ability to relate each local map with the appropriate rf and is not a property of the model it depicts. The local maps are organized such that the speed component of velocity tuning varies monotonically across a joint orientation contour line whereas direction varies according to a cells preferred orientation. This local-global structure is revealed in our investigation of the following four pairs of velocity scenarios.

The first set to be considered is that of a stimulus moving horizontally at the highest detectable speed as displayed in Fig. 6 Row 1, and of a bar moving orthogonally at the slowest detectable speed as displayed in Fig 6 Row 2. The image in Row 1, left, of Fig 6 indicates that the stimulus bar is in an rf at the left edge of the field. After one unit of time, the image indicates

that the stimulus has arrived at an rf at the right-hand edge of the field. The activation levels on the local map associated with this rf (middle image) indicate that the cell with the highest value (11.1, shaded for clarity) is located at the left end of the horizontal contour line. The intensity values of this local map and the other 48 local maps are displayed in the global velocity map, right. In this parameterization of the model, there are four cells, not counting the singularity, lying along the joint horizontal contour line. The fact that the most activated cell (value 11.1) lies at the end of the contour line means that, according to our model, the stimulus has either moved four rfs in the parallel direction or one rf orthogonally. To illustrate this point, consider the velocity scenario depicted in Row 2 of Fig. 6. In the left sub-image, it can be observed that the stimulus has moved one rf orthogonally to the cell's preferred orientation. The middle and right sub-images indicate that the most activated cell is the same one as displayed in Row 1. Therefore, this particular cell has responded to both a fast-moving spot in the parallel direction and a slow-moving bar in the orthogonal direction and thus constitutes evidence in support of one of the intended design features of the model. The most activated local maps in the middle images of Rows 1 and 2 also provide examples of cells that are selective to motion direction. Recall that each cell in the model receives inputs from up to four distinct local maps. However, because this particular local map does not possess neighbouring local maps horizontally to its right or orthogonally above, all the cells lying along this contour line will be selective to motion direction.

Our second set of velocity scenarios to be discussed is displayed in Rows 3 and 4 of Fig. 6. These two scenarios explore the case where the stimulus is moving slightly slower horizontally and slightly faster vertically than was discussed in the first set. More specifically, comparing these two motion scenarios, the simulated horizontal motion depicted in Row 3 left spans four rfs, while the vertical motion depicted in Row 4 left spans three rfs. The most activated cell in the local map in both rows lies on the horizontal contour line one cell left of the singularity (shaded). Accordingly, the most activated cell in each of the two global maps is the one corresponding to the biggest value of the local maps. In other words, the cell on the horizontal contour line one cell left of the singularity has responded to both of these distinct velocities. Importantly, no other cell in either global map has responded to these two velocities. Once again, the intended design feature of the model–that a cell responds to fast parallel motion and slow orthogonal motion–has been illustrated. As was the case with our first pair of velocity scenarios, the local map on the third row possesses only one neighbouring local map to its right thus the active cell is also selective to parallel motion. However, this cell is not selective to orthogonal motion because there are local maps on both its orthogonal sides.

The third pair of velocity scenarios to be considered is displayed in Rows 1 and 2 of Fig. 7. These scenarios explore the case where, once again, the

horizontal motion is slightly slower and the vertical motion is slightly faster than the case discussed in the previous scenario pair. The first row depicts the stimulus initially located one rf to the left of the centre while the second row depicts the stimulus initially located in an rf at the bottom of the field. In both rows the stimulus is positioned one unit of time later in the rf to the right of the centre. Thus Row 1 depicts the simulation of slow horizontal motion over a distance spanning three rfs, while Row 2 depicts the simulation of fast vertical motion spanning four rfs. Observe that, in contrast to the first two pairs of scenarios, here the faster speed is vertical and the slower is horizontal. Therefore, according to the design of our model and the pinwheel configuration that it mimics as displayed in Fig. 2, we should expect that the contour line displaying the highest activation should be the one projecting horizontally to the right of the singularity. These expectations are confirmed as can be observed in the local map of each row. Moreover, the cell with the highest response to these two velocities is on the horizontal contour line one cell out from the singularity. It can also be observed that this scenario pair provides an example of a cell that is non-selective to both parallel and orthogonal motion, as there exist local maps at all required distances.

The fourth pair of velocity scenarios to be discussed is displayed in Rows 3 and 4 of Fig. 7. This is the reverse of the first scenario set in that the simulation of the highest detectable velocity is vertical and the slowest detectable velocity is horizontal. Once again, according to the design of our model and the pinwheel configuration that it mimics (Fig. 2), we expect that the cell exhibiting the highest activation should be at the opposite end of the horizontal contour line to that of the active cell in the first scenario. The local maps displayed in Rows 3 and 4 of this figure confirm this expectation.

To summarize these simulation results: the speed component of velocity tuning decreases monotonically across the joint contour line from left to right for horizontal motions, and increases monotonically for vertical. Given the parameterization of the model employed (49 rfs and 49 local maps) we assert that the four velocity scenarios presented constitute all possible detectable speeds in a horizontal and vertical direction. Although we have only presented velocity variation for stimulus moving horizontally and vertically, the response pattern observed in the velocity map will be identical for all other detectable directions of motion differing only in the orientation contour line that is activated [47].

An aside regarding spatial properties may be useful here. As it is well known that Gabor space models are selective to orientation and frequency, we make this claim for 2DVM without explicitly presenting simulation results. However, our claim that 2DVM is selective to stimulus length requires justification. We assert that this property is a consequence of the OCP-Gabor. In particular, the length of a cycle of the plaid in the parallel direction is proportional to

the length of the stimulus that it is selective to. Simulation experiments were conducted where stimuli of various lengths were presented to two different plaid functions. A representative sample of the data from these experiments are presented in the graphs of Fig 8, where it can be observed that the empirical results are consistent with the theoretical expectations of the OCP-Gabor function.

4.2 Illusionary Effects Induced by the Model

In this section we present results that demonstrate how 2DVM models three types of visual illusions experienced by the human vision system.

The first type, visual crowding, is the inability to recognize objects in a clutter and represents a fundamental limit on conscious visual perception and object recognition [48]. One manifestation of this phenomenon is an observer's inability to distinguish individual items in a moving clutter. 2DVM mimics this effect in two ways: 1) a sequence of slow-moving stimuli aligned in orthogonal motion; and 2) fast-moving stimuli aligned in parallel motion. Two simple examples that simulate this effect are presented in Fig. 9, which displays three stimuli aligned in orthogonal motion, and Fig. 10, which displays three stimuli aligned in parallel motion. In both cases, the response to the middle stimulus has been suppressed and consequently did not initiate a velocity response. The result is that the velocity of the remaining two stimuli were correctly identified, albeit at the expense of not 'seeing' the middle stimulus.

The second type is known as perceptual filling-in. An experiment reported in [6] involved measuring the response of cortical cells of V1 to RFs associated with the blind spot created by the optic disc in the retina. Orientated moving bars were presented to the visual field and the cortical cells of V1 associated with the optic disc accurately interpolated the position of the bar at the time the bar passed over the blind spot. They also artificially produced a blind spot by obstructing an RF and found that in all cases the cells associated with the occluded RF responded in the same way as when the RF was not occluded. We have simulated this experiment in 2DVM by taking two disjointed bars that were aligned and passing them over the model's simulated visual field. The results produced were consistent with the data reported in [6]. These results are displayed in Fig. 11, where it can be observed that the local map associated with the gap between the two bars responded with a similar velocity to the local maps on either of its sides. This illusionary velocity resulted from the effects of collinear facilitation at the rf as the two co-aligned bars passed over the region.

The third type is that of illusionary orientation induced by fast moving stimuli,

motion streak. (Discussed in Section 1.) A simulation of this effect is displayed in Fig 12, where it can be observed that the longer the length of the motion streak the greater is the illusionary orientation. In particular, the top left sub-image depicts a stationary bar oriented at 30 degrees counter clockwise from the vertical. Its associated local map (right) reveals that the maximum energy occurs, as expected, on the contour line of 30 degrees (no motion streak). The middle row represents simulated motion of the bar depicted in the first row. The simulated motion streak takes the form of a parallelogram, which would result from the smearing caused by the bar's motion. The associated local map (right) indicates that the maximum energy occurs at both the contour line of 30 degrees and 60 degrees. This means that the perceived orientation is somewhere between 30 degrees and 60 degrees. The bottom row of the figure simulates increased speed as revealed by the elongated parallelogram. The local map of this scenario reveals that the maximum energy unambiguously occurs on the contour line of 60 degrees. Together, the three pairs of images demonstrate that illusionary orientation is induced by increased speed.

5 Discussion

In this section we interpret the model and simulation results, then discuss how these might impact future investigations of V1.

5.1 Role of Visual Illusion

Surround suppression, as a spatial characteristic, has been recognized since the 1970s [49]. Although it is still a debate as to where in the vision system the surround suppression actually takes place, experiments reported in [50] demonstrated that it resulted from inhibition in the visual cortex. Collinear facilitation, as a spatial characteristic, has been recognized since the 1980s [18]. The suppression/ facilitation effects of surround suppression and collinear facilitation in classical models of Area V1 tend to account for purely spatial properties [29], [51]. They also hypothesized that the spatial function of these effects was to produce a pop-out effect and as such, direct attention to important features in the visual field. Although 2DVM possesses the pop-out effect, we hypothesize that the critical role of suppression/ facilitation is not merely pop-out but to enhance robustness in the velocity map. In particular, surround suppression and collinear facilitation address the correspondence problem associated with determining velocities of individuals moving in the same direction as a group. This is important because a vision system's inability to adequately address the correspondence problem in the context of velocity estimation would tend to introduce errors. This also has implications for the

neural mechanism that underpins visual crowding.

5.2 Predictions of 2DVM

We observe that the OCP-Gabor space model presented in this paper has implications for the manner in which contextual modulation in general, and velocity in particular, manifest in V1. The parameterization of the OCP-Gabor functions accounted for the empirically-derived data that described the velocity preference bandwidth response on points within local maps. However, the same parameterization implies four predictions not yet tested empirically. The foregoing computational analysis of long-range contextual modulation and velocity in V1 leads to an interesting prediction about the structure of the local maps in V1 and, in particular, the storage of contextual modulation parameters within these local maps. The model parameters σ_{ew}, σ_{ns}, F_{ew} and F_{ns} of $2DVM(i, p, q, \theta, \sigma_{ew}, \sigma_{ns}, F_{ew}, F_{ns})$, in Equation 3, imply a spatial structure for the storage of long-range contextual information, i.e. a spatial structure that is measurable at the scale of a local map.

Prediction 1: The middle range of velocity selectivity should take place in the neighbourhood of local map singularities, whereas the extreme velocity selectivity (fastest and slowest) should be stored and processed by neurons located furthest from singularities. This represents only one plausible potential parameterization of 2DVM, as noted in Section 3. The core prediction is that different velocity ranges will be stored in different regions within each local map. One way this prediction could be tested empirically is with methods similar to standard optical imaging experiments by expanding the size of the orientation field from one RF to the entire visual field.

Prediction 2: It is a well-known fact that many V1 cells are selective to motion direction [2]. However, there is an open question as to why this is so. Indeed, numerous experiments aimed at explaining why certain cells have a direction selectivity have been reported [37]. In an investigation reported in [52] intracellular recordings from V1 simple cells favoured models of direction selectivity based on excitation from the lateral geniculate nucleus (LGN) rather than cortical inhibition. They concluded that it seemed unlikely that inhibition from the null direction was the primary mechanism for creating direction selectivity.

The cells modelled in 2DVM possess a mix of direction and non-direction motion selectivity. However, the structural modelling of 2DVM implies a clear prediction of which neurons are direction selective and why this occurs. 2DVM predicts that direction selectivity for slow-moving objects should only occur in neurons in local maps near the periphery of the global map. Direction

selectivity requires that local maps must have at least one null direction. The only way this can occur for slow-moving objects is if the local map is near the periphery of the global map. The converse should also apply. That is, direction selectivity to slow-moving objects should be a rare occurrence for neurons situated in local maps near the global map's centre. The model's explanation for direction selectivity for fast-moving objects is similar. But its prediction as to the location of such neurons is less precise. 2DVM indicates that the greater the velocity, the greater the separation of local maps. If the required separation means that the local maps are on different sides of the global map, then selectivity must ensue. Consequently, 2DVM predicts that the higher a neuron's peak velocity response, the more likely it is to be selective to motion direction.

Prediction 3: The Connection Algorithm presented in this paper prescribed a pattern of connectivity between cells located in different local maps with similar tuning properties. Although this pattern of connectivity underpins the design features of the model, it also raises two other inter-related issues. On the one hand, it implies a hypothesis about the V1 connectivity that is empirically untested. The specific testable prediction is that individual cells within orientation pinwheels are connected to specific cells in neighbouring pinwheels that have similar tuning preferences. On the other hand, if the prescribed connections exist in V1, then one would also expect that small variations to the pattern should only result in small effects in the accuracy of the velocity map. This expectation could be tested computationally with the use of genetic algorithms to investigate whether the pattern represents a computationally convergent property.

Prediction 4: Although the Velocity Model presented in this paper predicts two-dimensional (2-D) motion, it says nothing about 3-D motion. It seems very likely that the property of smooth velocity variation across a joint contour line has implications for the perception of 3-D motion and how the two eyes work together to produce this perception.

5.3 Limitations

The mathematical formulation of 2DVM and its implementation as presented in this paper are scale independent. However, the simulation results that were presented represented only a minimal non-trivial parameterization of the RF space. That is, we reported on the case where the number of RFs is 49 and the size of a local map is 49 cells. Although these results demonstrated the features of 2DVM on a small scale and on synthetically generated images, no evidence was provided that the model applies at the actual scale of V1 or on natural images. This would entail a parameterization of many more local

maps, possibly up to 10,000, which is the estimated number of orientation pinwheels in V1 [39]. Such a parameterization would involve a more elaborate consideration of the simulated visual field and a realistic distribution of RFs. While this was outside the scope of this paper, it would constitute a worthy investigation that might reveal other properties of the functional architecture of V1. For example, an interesting hypothesis is a connection between motion streaks and the perception of Glass patterns [53]. Although it is expected that 2DVM would produce a velocity response to Glass patterns, the current implementation possesses insufficient local maps to explore this hypothesis.

Another reason for implementing 2DVM with a realistic number of local maps is to be able to interpret the velocities of 2DVM in terms of the more commonly used dimensions, degrees per second, as employed in the literature. The current choice of dimensions, rfs per unit of time, is appropriate in the current context but it lacks a connection to the way velocity perception is measured in psychophysical experiments with humans.

This paper focused on the design issues of the model. However, it did not provide any insight as to how such a model might actually be implemented in cortical architecture. For a discussion of this issue, the reader is referred to [54].

5.4 Points of Controversy

The first point of controversy pertains to the motion streak hypothesis, which states that motion streak patterns are not only projected into V1, as assumed in this paper, but they are also used by the vision system to predict direction of motion. Psychophysical experiments aimed at investigating the motion streak hypothesis have produced a mixed bag of results. Evidence that rejected the motion streak hypothesis was reported in [55]; [56]. However, other experiments on single neuron responses of cats to certain types of motion supported this hypothesis [2]; [9]; [10]; [57]; [58]; [59]; [60]; [61]. Given that motion streaks are projected into V1, the results of this paper support the conclusion that they are also used to determine direction of motion. In addition, the results of this paper also demonstrate how motion streaks contribute to the V1 neurons estimate of stimulus speed. Consequently, 2DVM proposes a stronger version of the motion streak hypothesis, i.e. motion streaks are used by the vision system to predict velocity.

A second point of controversy relates to the question of whether or not the fundamental role of V1 neurons represents individual feature attribute detectors. Most of the earlier work in this area cited in this paper supports this proposition. However, the results reported in [7], concluded that V1 neurons did not

represent individual feature attribute detectors. Although 2DVM agrees with that conclusion up to the formulation of velocity, it markedly diverges from their conclusion at that point by asserting that V1 neurons are velocity detectors and these, taken together, form a velocity map of the visual field. The reason for the attention given to the idea that individual neurons in V1 are primarily feature detectors may have its roots in historical matters. That is, the technology of the day with its limitations may have constrained the types of hypotheses that were feasible to investigate at the time. Many significant discoveries were made over this 50 year period. In the interests of further discussion, we draw on an analogy to put these discoveries in perspective. Just as the stone blocks that make up the great pyramids of Egypt, with their precise engineering and placement, are individually significant in their own right, so also are the individual spatial and temporal selective properties of V1 neurons. However, as the importance of the individual stone blocks pales when compared with the pyramid structure as a whole, it may also turn out that the individual spatial and temporal properties pale into insignificance when placed in the context of V1 considered as a whole, i.e. a velocity map. Such a paradigm shift in the role of V1 would require considerable empirical evidence before being embraced by the vision community. Empirical proof is currently lacking but maybe now is the time to start looking for the required evidence.

6 Summary and Conclusion

At the outset of this paper, we discussed the fundamental importance of velocity estimation in biological vision. The novelty of 2DVM is in indicating the importance of V1s role in this process. Our model demonstrates how the tessellated pattern of orientation pinwheels, together with a pattern of long-range connections between these as specified by the Connection Algorithm, implies a global-local velocity map of the upper layers of V1. The statement is elaborated in the following three points:

(1) The observable architecture of the orientation pinwheels arranges spatial selectivity properties in a manner that is consistent with the spatial aspects of velocity computation.
(2) The inter-orientation pinwheel connectivity, as specified in the Connection Algorithm, over the time domain transforms the spatial properties into a velocity map of V1.
(3) The input signal to 2DVM is a representation of the entire visual field. This implies that the velocity response outputs are reliable estimates up to the constraints imposed by the resolution of the representation.

2DVM was formulated in two stages. First, the base model emerged directly from known spatial properties of V1 cells, such as orientation tuning and

frequency tuning. These properties were captured in an OCP-Gabor space. Second, in the extension to the base model, a pattern of temporal connections was specified via the Connection Algorithm. The primary purpose of the temporal connections was to account for the empirically observable response properties of V1 cells to two types of object motion: the first type is to objects moving quickly parallel to the RF's orientation (parallel motion); the second type is to objects moving slowly orthogonally to the RFs orientation (orthogonal motion). The simulation results presented in this paper demonstrated that 2DVM successfully achieved its design goal of accounting for V1 cell responses to these two types of 2-D motion. In addition, simulation results also demonstrated how the model accounted for various types of visual illusion and thereby suggested mechanisms that may account for this phenomenon in the human vision system. In particular, results indicated how the well-known contextual modulation effects of spatial suppression/facilitation contributed to the vision systems ability to achieve robustness in the velocity map. It was also argued that the occurrence of illusionary orientation associated with stimuli moving quickly obliquely to its direction of motion is a consequence of motion streaks. The addition of the temporal connections to the model implied a significant role for singularities in the velocity map. Recall that each singularity (confluence of all iso-orientation contour lines in a local map) is associated with four OCP-Gabor filters in the model. Although the net effect of these filters on the spatial orientation is a non-preferential orientation response, it was shown that singularities contributed to the detection of slow-moving spots that would otherwise be undetectable by the non-singularity cells of a local map. The net effect of the model's formulation also produced two modelling outcomes consistent with empirically-derived data: 1) to parallel motion, the model indicated that cells tended to be selective to motion direction and band pass selective to spatial frequency; 2) to orthogonal motion, the model indicated the same cell tended to be non-selective to motion direction and low pass selective to spatial frequency.

The intended design features of the model for known empirical results aside, the model also produced three hypotheses that would warrant the focus of future research. First, the speed component of neuronal velocity selectivity decreases monotonically across each joint orientation contour line for parallel motion and increases monotonically for orthogonal motion. Second, the cells that are direction selective to slow-moving objects are situated in the periphery of V1. Third, neurons in distinct pinwheels tend to be connected to neurons with similar tuning preferences in other pinwheels.

2DVM has extended all other cited models of V1 by demonstrating how many of the functional properties of V1 emerge from known structural properties. This paper has also contributed to the discussion on the role motion streaks played in the vision process. In particular, we have extended the motion streak hypothesis by arguing that motion streaks are integral to the vision system's

estimates of velocity and not merely, as previously conjectured, contributors to direction detection. We also demonstrated how motion streaks contributed to the phenomenon of illusionary orientation in the human vision system. Further, this paper demonstrated that the long-range cortical connections within Layer 2/3 of V1, when combined with the neuronal spatial tuning properties, result in a robust velocity map of the visual field. Consequently, it is the conclusion of this paper that the spatial selectivity features of V1 are side-effects of its more important role as a velocity map of the visual field and that empirical experiments aimed at investigating this are warranted.

7 Acknowledgements

We thank Valerie Christian McDougall for her work on improving the presentation of this manuscript.

The source code that implements 2DVM is available in a software package named VisioLab from the following website: http://www.ict.griffith.edu.au/sheridan/

References

[1] Parker, A. (2004). In the Blink of an Eye. New York: Basic Books.

[2] Geisler, W.S., Albrecht, D., Crane, A.M. and Stern, L. (2001). Motion direction signals in the primary visual cortex of cat and monkey. Visual Neurosci., 18, 501-16.

[3] Henry, G.H., Dreher, B. and Bishop, P.A. (1974). Orientation specificity of cells in cat striate cortex. J. Neurophysiol., 37, 1394.

[4] Worgotter, F., Muche, T. and Eysel, U.T. (1991). Correlations between directional and orientational tuning of cells in cat striate cortex. Exp. Brain Res., 83, 665.

[5] La Cara, G. and Ursino, M. (2007). Direction selectivity of simple cells in the primary visual cortex: comparison of two alternative mathematical models. II: Velocity tuning and response to moving bars. Computers in Biology and Medicine. 37, 5, 598-610.

[6] Fioroni, M. Jr., Rosa, M.G.P., Gattass, R. and Rocha-Miranda, C.E. (1992). Dynamic surrounds of receptive fields in primate striate cortex: a physiological basis for perceptual completion? Proc. Natl. Acad. Sci. USA 89, 8547-8551.

[7] Basole, A., White, L. E. and Fitzpatrick, D. (2003). Mapping multiple features in the population response of visual cortex. Nature, 423(6943), 986-990.

[8] Bex P.J., Dakin S.C. and Simmers A.J. (2003).The shape and size of crowding for moving targets. Vision Research. 43, 28952904.

[9] Mante, V., and Carandini, M. (2005). Mapping of stimulus energy in primary visual cortex. J. Neurophysiol., 94(1), 788-798.

[10] Baker, T.I., and Issa, N.P. (2005). Cortical maps of separable tuning properties predict population responses to complex visual stimuli. Journal of Neurophysiology, 94(1), 775-787.

[11] Burr, D. C. (1980). Motion smear. Nature, 284, 164-165

[12] Schiller, P.H., Finlay, B.L., and Volman, S.F. (1976). Quantitative studies of single-cell properties in monkey striate cortex, III, Spatial frequency. J. Neurophysiol. 39, 6, 1334-1351.

[13] De Valois, K.K., de Valois, R.L. and Yund, E.W. (1979). Responses of striate cortex cells to grating and checkerboard patterns. J. Physiol. (Lond) 291, 483-505.

[14] Sceniak, M.P., Hawken, M.P., and Shapley, R. (2001). Visual spatial characterization of macaque V1 neurons. J. Neurophysiol., 85, 5, 1873-1887.

[15] Cavanaugh, J.R., Bair, W., and Movshon, J.A. (2002). Selectivity and spatial distribution of signals from the receptive field surround in macaque V1 neurons. J. Neurophysiol., 88, 5, 2547-2556.

[16] Bair, W., and Movshon, J.A. (2004). Adaptive temporal integration of motion in direction-selective neurons in macaque visual cortex. J. Neurosci., 24, 33, 7305-7323.

[17] Blakemore, C. and Tobin, E.A. (1972). Lateral inhibition between orientation detectors in cats visual-cortex. Exp. Brain Res., 15, 4, 439-441.

[18] Nelson, J.I. and Frost, B.J. (1985). Intracortical facilitation among co-oriented, co-axially aligned simple cells in cat striate cortex. Exp. Brain Res., 61, 11, 54-61.

[19] Sillito, A.M. and Jones, H.E. (1996). Context-dependent interactions and visual processing in V1. J. Physiol. Paris 90, 3-4, 205-209.

[20] Kimura,R. and Ohzawa, I. (2009). Time course of cross-orientation suppression in the early visual cortex. J. Neurophysiol., 101, 3, 1463-1479.

[21] Lamme, V.A.F. (1995). The neurophysiology of figure-ground segmentation in primary visual cortex. J. Neurosci., 15, 2, 1605-1615.

[22] Zipser, K., Lamme, V.A.F., and Schiller, P.H. (1996). Contextual modulation in primary visual cortex. J. Neurosci., 16, 22, 736-738.

[23] Lamme, V.A.F., Zipser, K. and Spekreijse, H. (1998). Figure-ground activity in primary visual cortex is suppressed by anesthesia. Proc. Natl. Acad. Sci. USA, 95, 6, 3263-3268.

[24] Lee, T.S., Mumford, D., Romero, R., and Lamme, V.A.F. (1998). The role of the primary visual cortex in higher level vision. Vision Res., 38, 15-16, 2429-2454.

[25] Alexander, D.M., Bourke, P.D., Sheridan, P., Konstandatos, O., and Wright, J. (2004). Intrinsic connections in tree shrew V1 imply a global to local mapping. Vision Res. 44, 9, 857-876.

[26] Alexander, D.M. and Van Leeuwen, C. (2010). Mapping of contextual modulation in the population response of primary visual cortex. Cogn. Neurodyn. 4, 1, 1-24.

[27] T.D. Sanger, T.D. (1988). Stereo disparity computation using Gabor filters. Biol. Cybern. 59, 405-418.

[28] Grigorescu, C., Petkov, N. and Westenberg, M.A. (2003). Contour detection based on nonclassical receptive field inhibition. IEEE Trans. Image Process., 12, 7, 729-739.

[29] Huang, W., Jia, L., and Jia, J. (2008). Modeling contextual modulation in the primary visual cortex. Neural Networks 21, 8, 1182-1196.

[30] Lee, H.C. and Choe, Y. (2003). Detecting salient contours using orientation energy distribution. Proc. Internat. Joint Conf. on Neural Networks, 1, 206-211.

[31] Ursine, M. and La Cara, G.E. (2004). A model of contextual interactions and contour detection in primary visual cortex. Neural Networks 17, 5-6, 719-735.

[32] Tang, Q., Sang, N. and Zheng, T. (2007). Extraction of salient contours from cluttered scenes. Pattern Recognition 40, 11, 3100-3109.

[33] Welch, L. (1989). The perception of moving plaids reveals two motion-processing stages. Nature. Feb 23, 337(6209), 734-6.

[34] Hubel, D.H. and Wiesel, T.N. (1977). Functional architecture of macaque monkey visual-cortex. Proc. R. Soc. Lond. B Biol. Sci. 198, 1130, 1-59.

[35] Bonhoeffer, T. and Grinvald, A. (1991). Iso-orientation domains in cat visual-cortex are arranged in pinwheel-like patterns. Nature 353, 6343, 429-431.

[36] Blasdel, G.G. (1992). Orientation selectivity, preference, and continuity in monkey striate cortex. J. Neurosci. 12, 8, 3139-3161.

[37] Yao, X., Jin L., and Hu, H. (2007). Pinwheel patterns give rise to the direction selectivity of complex cells in the primary visual cortex, Brain Res. 1170, 140-146.

[38] K. Obermayer, K. and G.G. Blasdel, G.G. (1993). Geometry of orientation and ocular dominance columns in monkey striate cortex. J. Neurosci. 13, 4114-4129.

[39] Murphy, K.M., Jones, D.G. et al. (1998). Spacing of cytochrome oxidase blobs in visual cortex of normal and strabismic monkeys. Cerebral Cortex 8, 3, 237-244.

[40] Bruce, V., Green, P.R., and Georgeson, M.A. (2003). Visual Perception: Physiology, Psychology and Ecology. (4th ed.). London: Psychology Press.

[41] Schummers, J., Marino, J., and Sur, M. (2002). Synaptic integration by V1 neurons depends on location within the orientation map. Neuron 36, 5, 969-978.

[42] Marino, J., Schummers, J., Lyon, D.C., Schwabe, L., Beck, O., Wiesing, P., Obermayer, K., and Sur, M. (2005). Invariant computations in local cortical networks with balanced excitation and inhibition. Nat Neurosci. 8, 2, 194-201.

[43] Yousef, T., Toth, E., Rausch, M., Eysel, U.T., and Kisvarday, Z.F. (2001). Topography of orientation centre connections in the primary visual cortex of the cat. Neuroreport. 12, 8, 1693-1699.

[44] Lund, J.S., Angelucci, A., and Bressloff, P.C. (2003). Anatomical substrates for functional columns in macaque monkey primary visual cortex, Cereb Cortex. 13, 1, 15-24.

[45] Sheridan, P., Hintz, T., and Alexander, D. (2000). Pseudo invariant transformations on a hexagonal lattice. Journal of Image and Vision Computing, 18, 907-917.

[46] Sheridan, P.E. (2007). A method to perform a fast Fourier transform with primitive image transformations. IEEE Transactions on Image Processing, 16, 5, 1355-1369.

[47] Adelson, E. H. and Movshon, J. A. (1982). Phenomenal coherence of moving visual patterns. Nature 300, 523-525.

[48] Whitney, D., Levi, D.M. and Wills, H. (2011). Visual Crowding: a fundamental limit on conscious perception and object recognition. Trends Cogn Sci. Apr; 15(4), 160-160.

[49] Moffei, L. and Fiorentinni, A. (1976). The unresponsive regions of visual cortical receptive fields. Vision Research 16, 1131-1139.

[50] Durand, S., Freeman, T.C.B., and Carandini, M. (2007). Temporal properties of surround suppression in cat primary visual cortex. Visual Neurosci. 24, 5, 679-690.

[51] Sheridan, P.E. (2012). Contextual modulation in the primary visual cortex via low-level processing. Image and Vision Computing, 30, 4-5, 367-377.

[52] Priebe, N.J., Lampl, I., and Ferster, D. (2001). Mechanisms of direction selectivity in cat primary visual cortex as revealed by visual adaptation. J Neurophysiol., ISSN 0022-3077, 104, 5, 2615-2623.

[53] Ross, J., Badcock, D. R., and Hayes, A. (2000). Coherent global motion in the absence of coherent velocity signals. Current Biology, 10(11), 679-682.

[54] P.E. Sheridan, P.E. (2012). Cortical Specification of a Fast Fourier Transform Supports a Convolution Model of Visual Perception, in S.M. Salih (Ed.), Fourier Transform Applications, (pp. 181-204). Croatia: InTech.

[55] Ball, K. and Sekuler R. (1987). Direction-specific improvement in motion discrimination. Vision Research, Vol 27, Issue 6, 953-965

[56] Matthews N. and Allen J. (2005). The role of speed lines in subtle direction judgments. Vision Res. Jun, 45(12), 1629-40.

[57] Apthorp, D., and Alais, D. (2009). Tilt aftereffects and tilt illusions induced by fast translational motion: Evidence for motion streaks. J Vis, 9(1), 1-11.

[58] Apthorp, D.,Wenderoth, P., and Alais, D. (2009). Motion streaks in fast motion rivalry cause orientation-selective suppression. J Vis, 9(5). doi: 10.1167/9.5.10.

[59] Apthorp, D., Cass, J., and Alais, D. (2010). Orientation tuning of contrast masking caused by motion streaks. J Vis, 10(10), 1-13. doi: 10.1167/10.10.11.

[60] Apthorp, D., Cass, J., and Alais, D. (2011). The spatial tuning of 'motion streak' mechanisms revealed by masking and adaptation. J Vis, 11(7), 1-16. doi: 10.1167/11.7.17

[61] Apthorp, D., Samuel Schwarzkopf, D., Kaul, C., Bahrami, B., Alais, D. and Rees, G. (2013). Direct evidence for encoding of motion streaks in human visual cortex. Proc. R. Soc. B, 280 (1752).

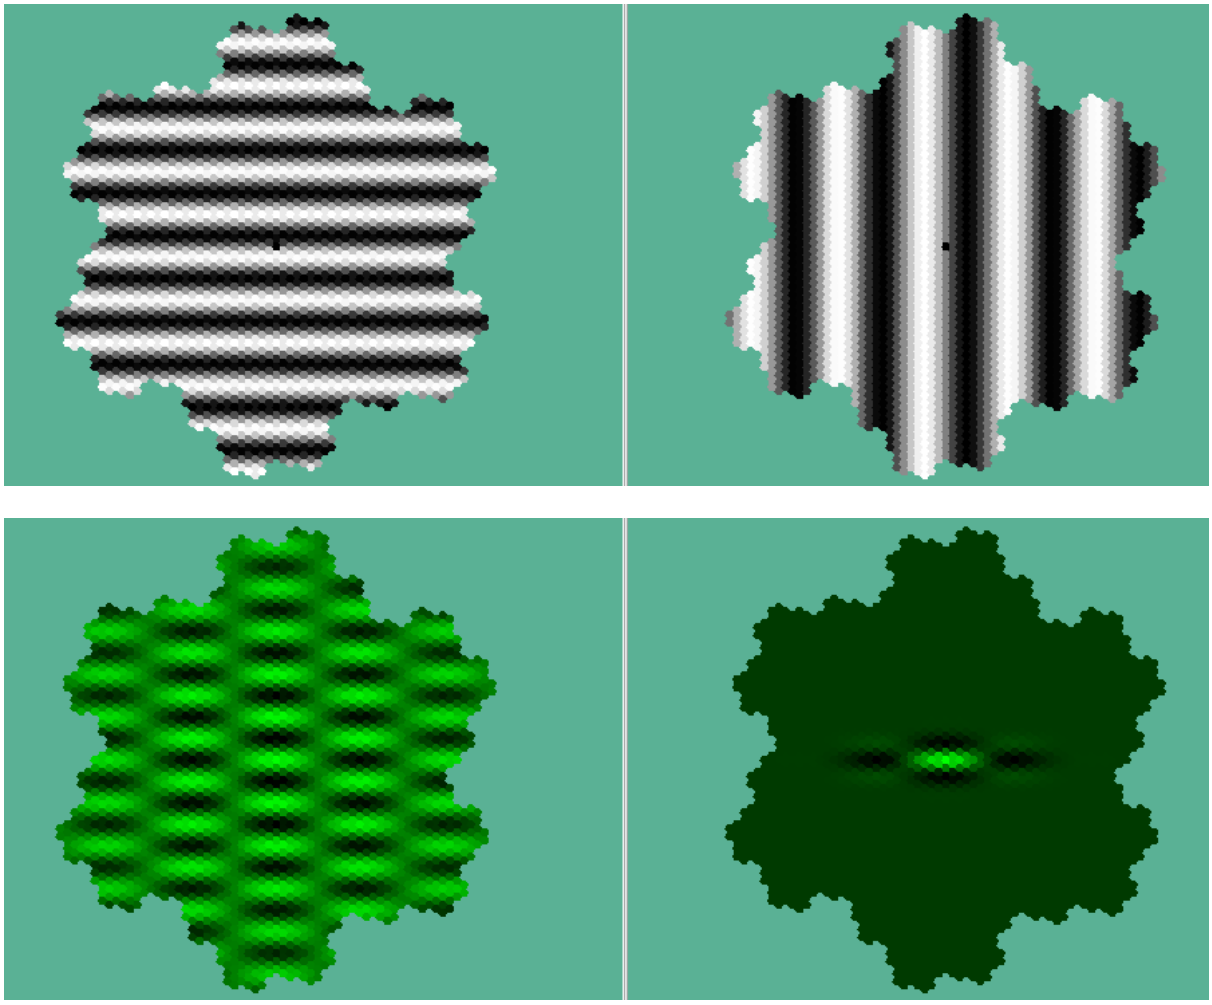

Fig. 1. Displays the components of an OCP-Gabor filter. Upper left is a horizontally-orientated sinusoid while upper right is a vertically-orientated sinusoid. Lower left is the product of these two sinusoids. Lower right is the product of the plaid (lower left) with a Gaussian and illustrates an OCP-Gabor filter.

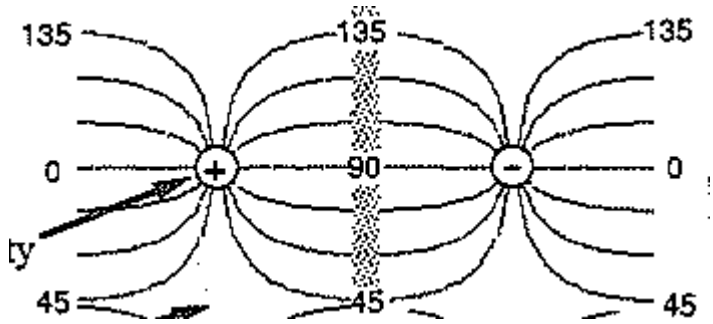

Fig. 2. Displays a schematic diagram of the pinwheel-like structures of Area V1, extracted from Figure 10 page 43 of [40].

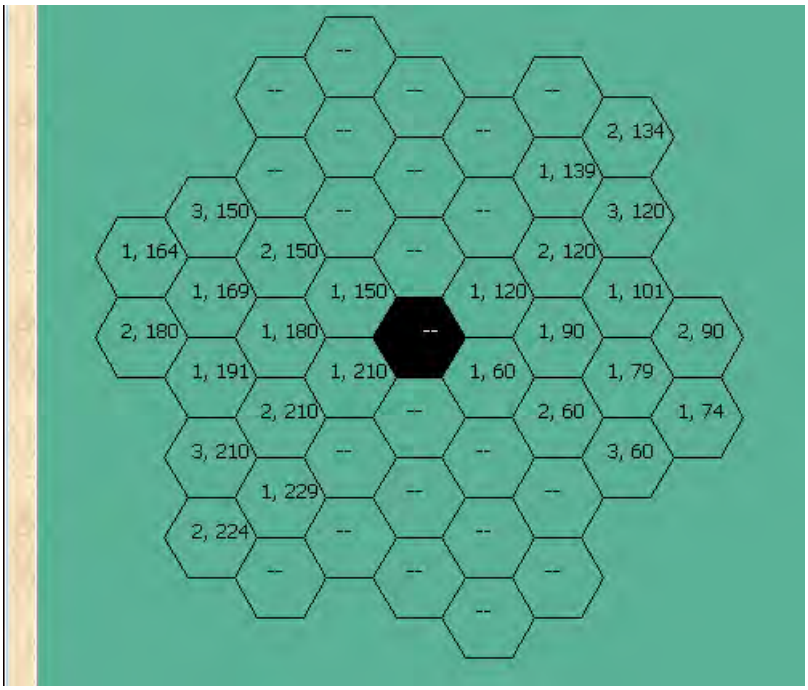

Fig. 3. Displays a discretised version of a pinwheel for a 49-cell local map of 2DVM. The darkened cell represents the singularity. Where cells display two integers: 1) the first represents a count of the cell's position out from the singularity; 2) the second represents the angle of orientation of the orientation contour passing through it. (Note that the orientation angle and the angle (+/-)180 represent the same contour line.) The cells displaying a dash (-) represent cells that do not lie on an orientation contour line.

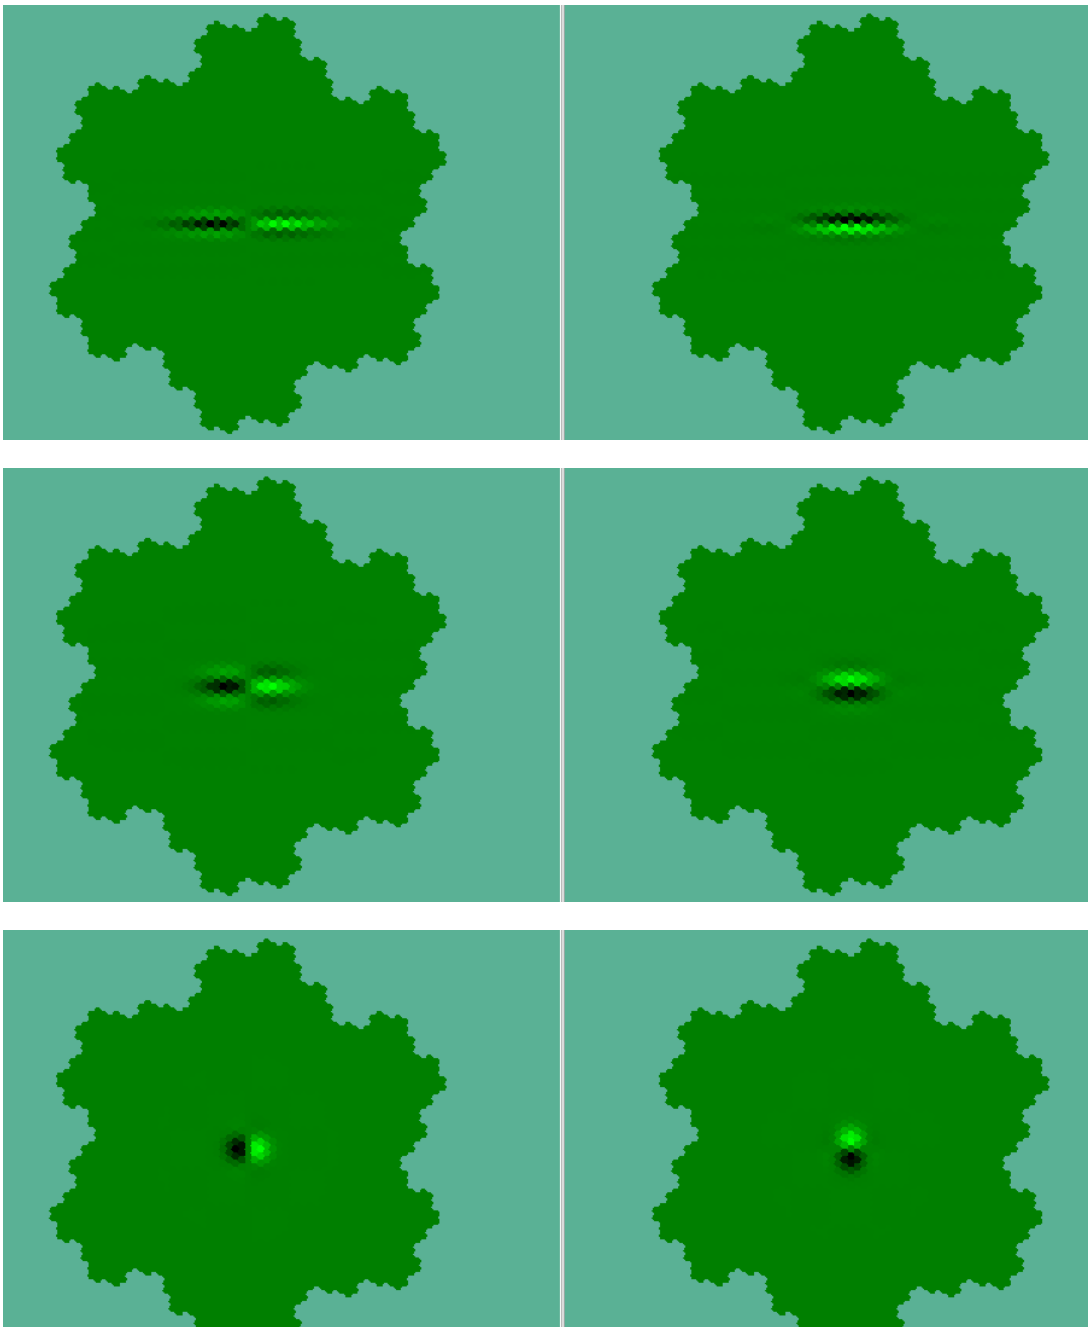

Fig. 4. Displays examples of OCP-Gabor response functions as intensity maps. In Rows 1 & 2, the left images are horizontally phase-shifted for edge selectivity while the right are vertically phase shifted. Row 1 shows the highest vertical and lowest horizontal frequencies while Row 2 shows slightly lower vertical and slightly higher horizontal frequencies. The bottom pair (Row 3) has two sinusoidal components with equal frequency: left provides preference for the light/dark edge of a horizontal pair of spots; right is its associated light/dark vertical edge.

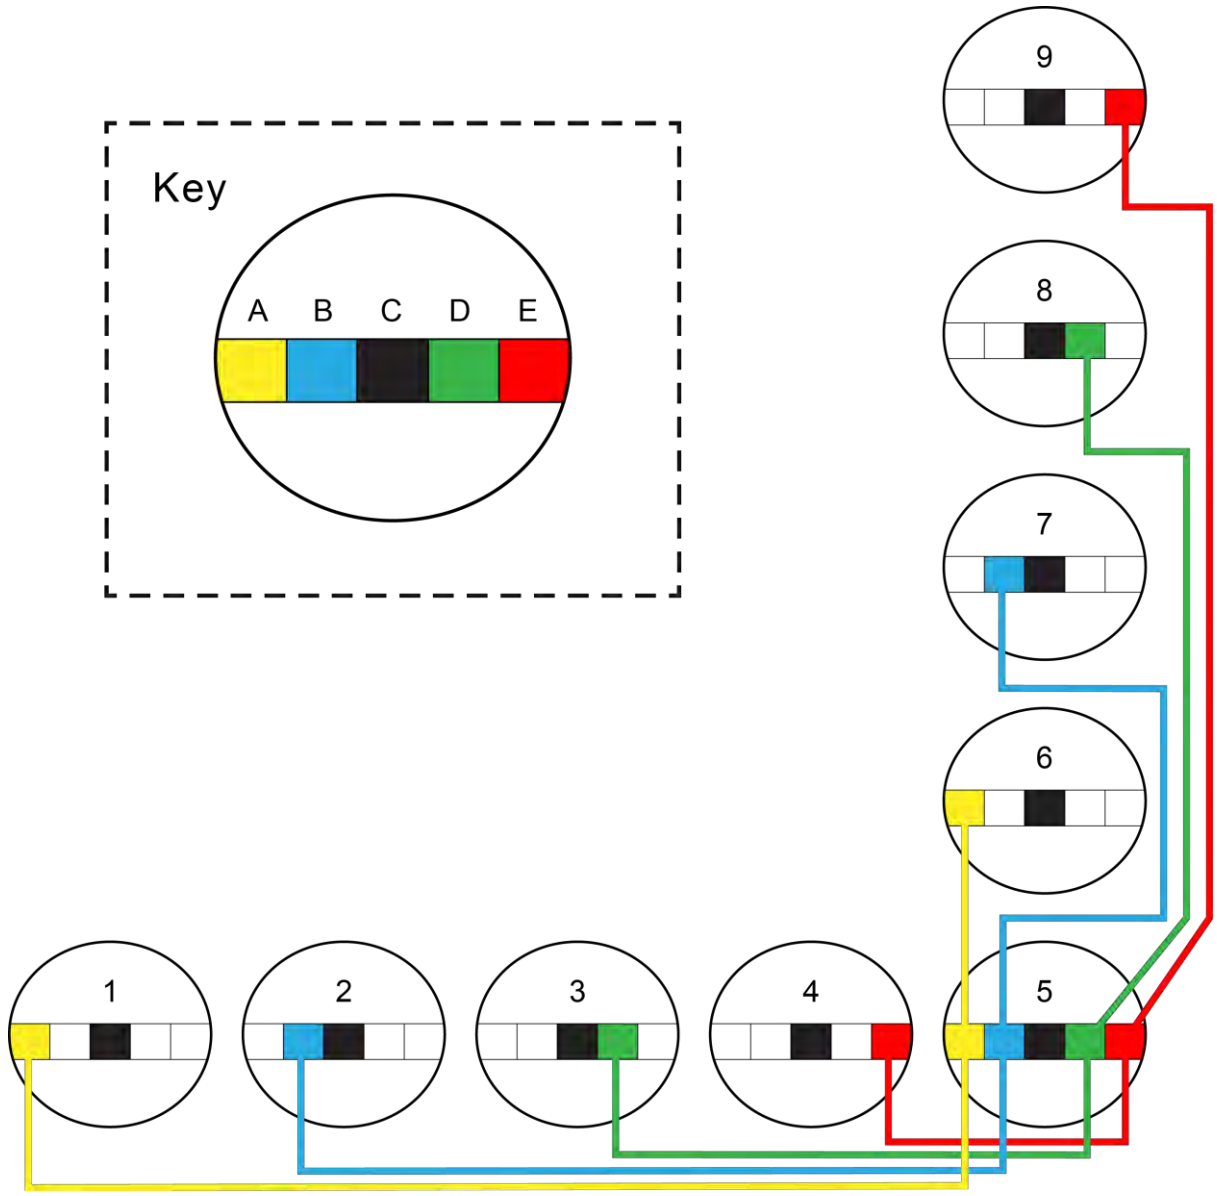

Fig. 5. This is an example of the connections generated by the Connection Algorithm. The circles (numbered) represent local maps; the five cells within each local map (we name A to E) represent the horizontal contour line of the map, with the darkened centre cell (C) indicating the singularity. The lines connecting individual cells in one map to another represent the temporal cortical connections as generated by the algorithm. NOTE: There is only one map with connections to all the others–the 'Destination Map' (Map 5). Further, note the symmetry of how the lines connect with specific cells in each direction. For instance, Cell 5A connects horizontally to cell 1A and vertically to Cell 6A; Cell 5B connects horizontally to Cell 2B and vertically to Cell 7B. (Diagram numbered and coloured for ease of explanation only.)

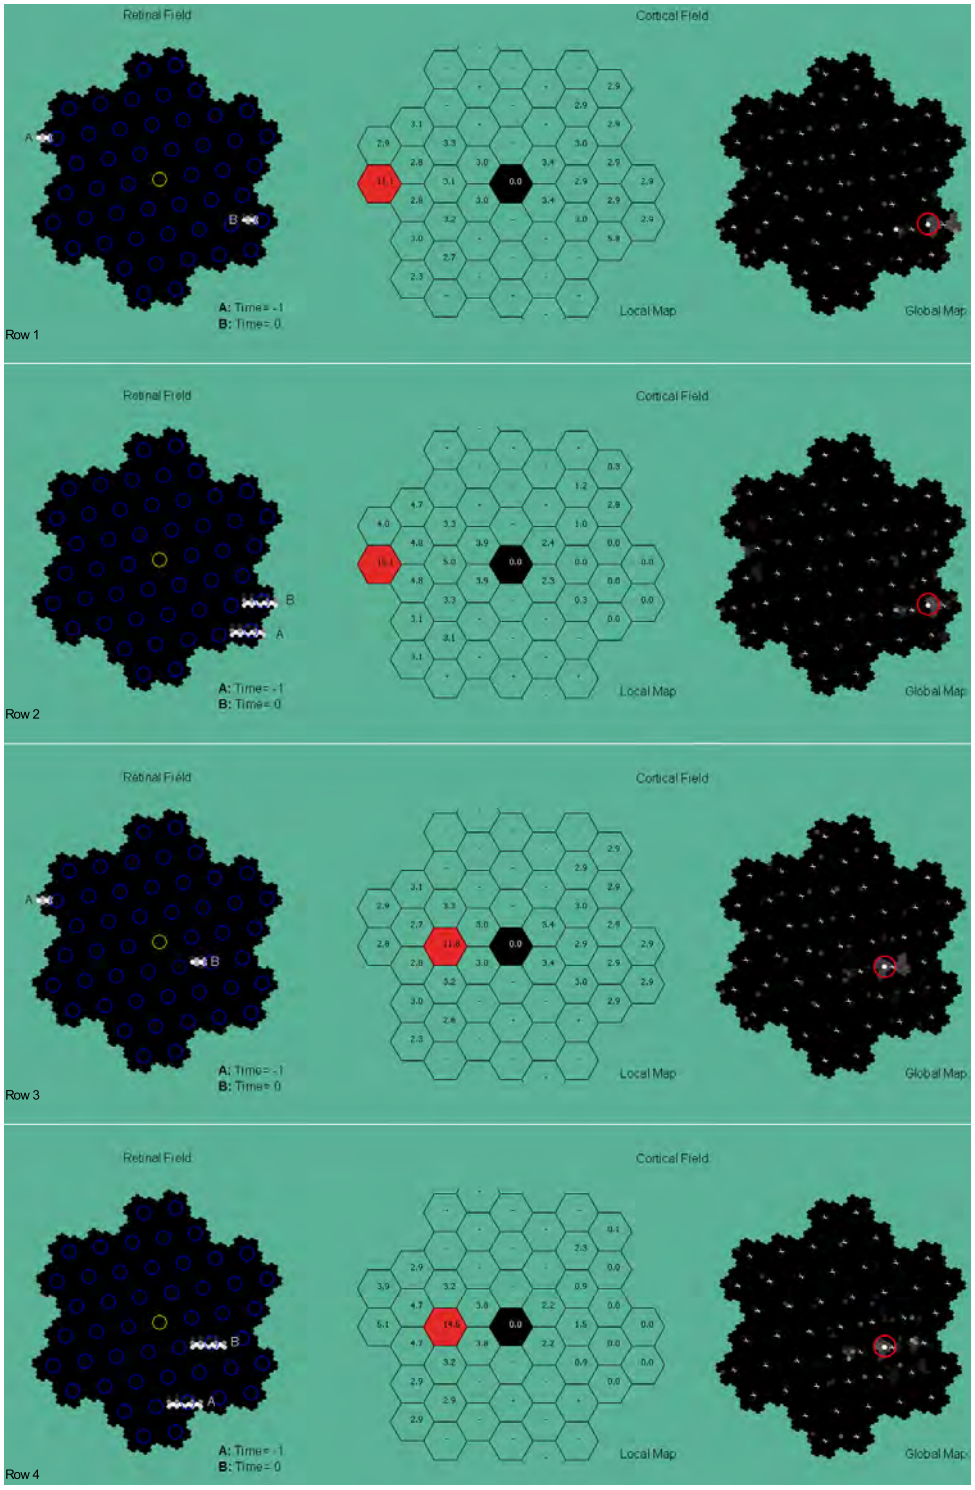

Fig. 6. Shows 4 velocity scenarios. Left images depict a stimulus at two points in time. The circles identify the rfs. The right images depict 49 local maps as intensity maps, with each singularity marked. Bright pixels represent high activating neurons and dark pixels, low. The brightest, circled, are shown as local maps (centre) with the numbers representing activation levels. Row 1 depicts the fastest horizontal motion, Row 2, the slowest vertical motion, Row 3, the second fastest horizontal motion, and Row 4, the second slowest vertical motion. The highest activating cells are shaded.

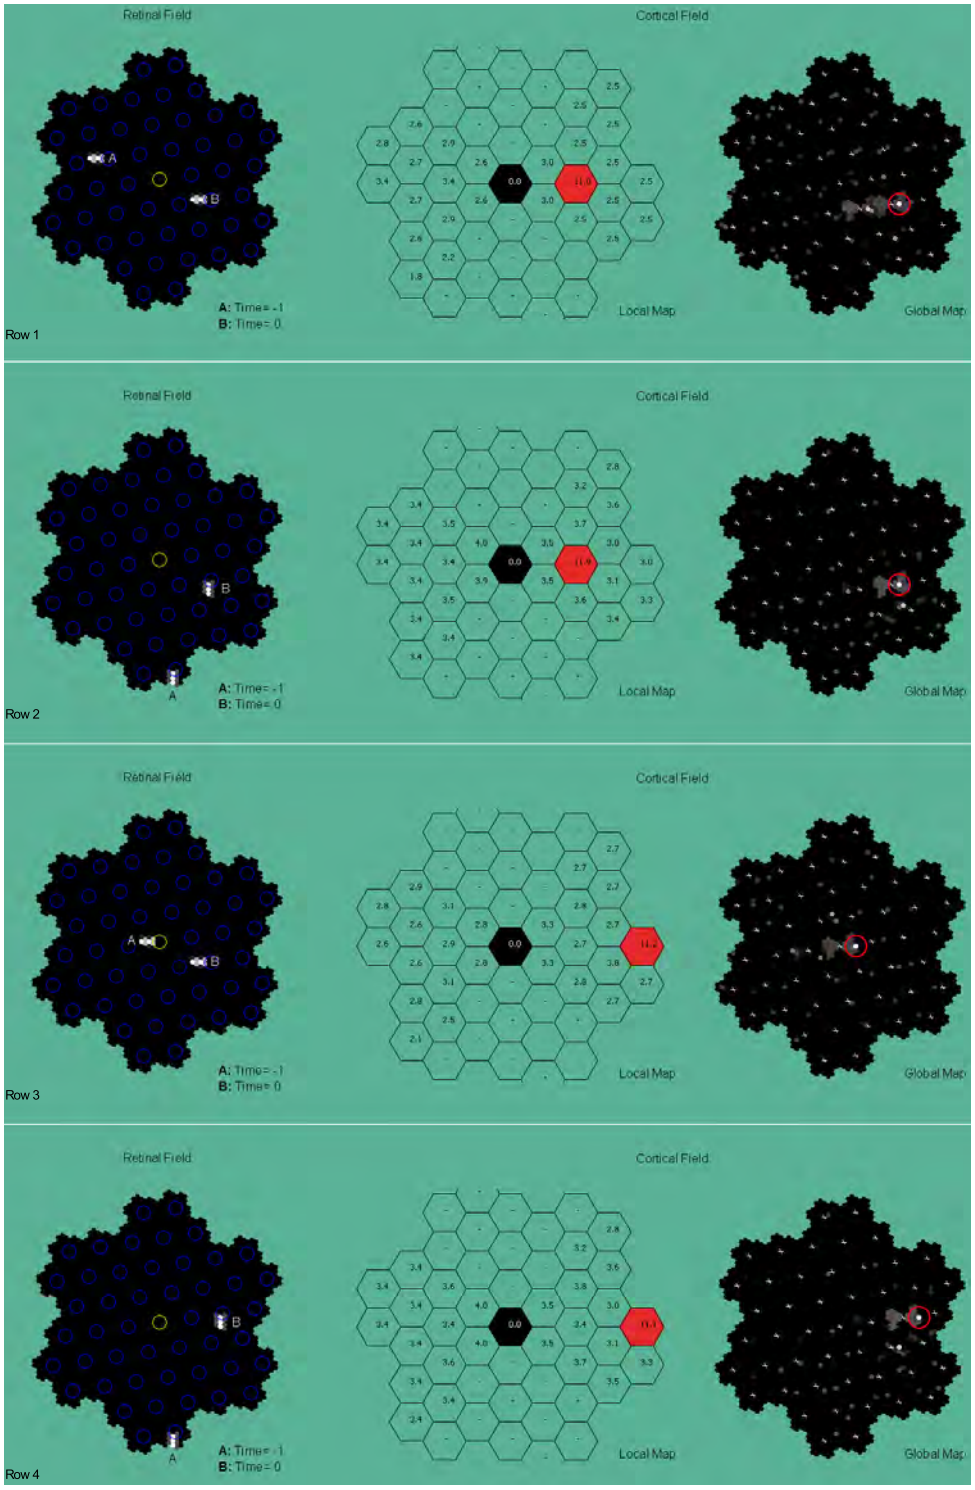

Fig. 7. Shows 4 velocity scenarios. The left-images depict a stimulus at two points in time. The circles identify the rfs. The right images depict 49 local maps as intensity maps, with each singularity marked. Bright pixels represent high activating neurons and dark pixels, low. The brightest, circled, are shown as local maps (centre) with the numbers displaying activation levels. Row 1 depicts the second slowest horizontal motion, Row 2, the second fastest vertical motion; Row 3, the slowest horizontal motion and Row 4, the fastest vertical motion. The highest activating cells are shaded.

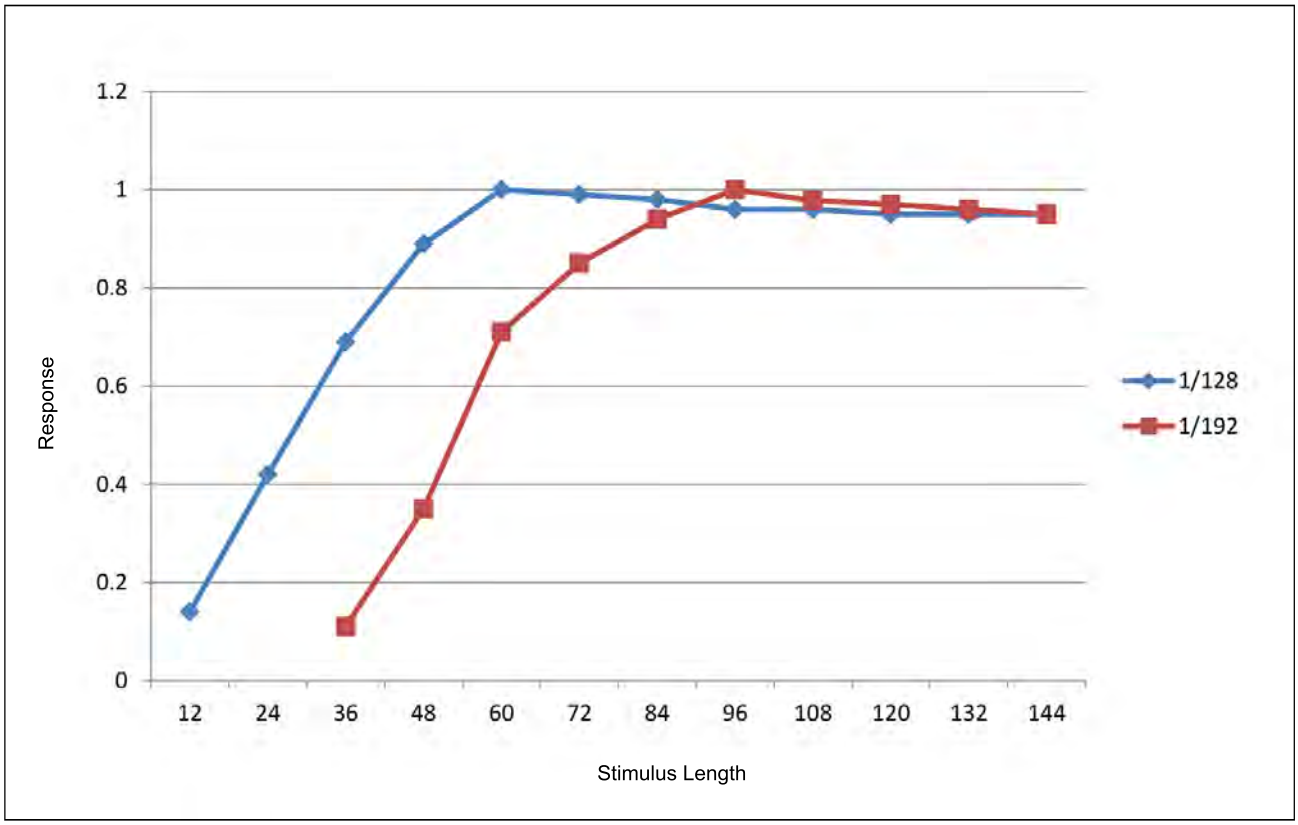

Fig. 8. Displays the response profiles of two different OCP-Gabor functions to stimuli of differing lengths. It can be observed that the OCP-Gabor function with parallel frequency of 1/128 reaches its optimum tuning at a stimulus of approximately 60 units and the 1/192 parallel frequency OCP-Gabor function is tuned to 96 units.

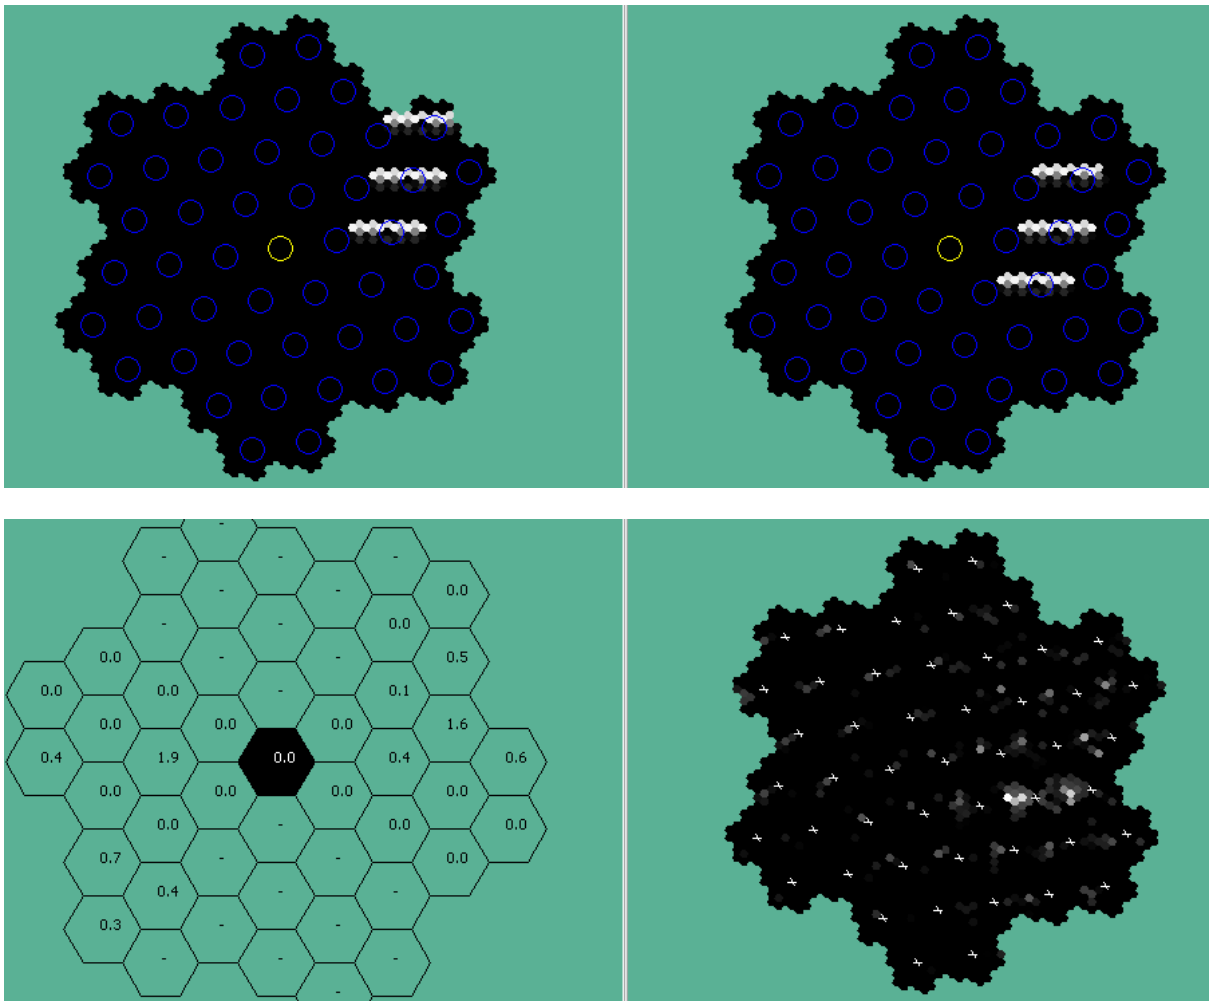

Fig. 9. Example of simulation results of the suppression effects on a group of slow-moving stimuli. The upper two images display the groups position at two distinct times and jointly they depict velocity. The circles indicate rfs. The three orientated bars represent orthogonal motion. Upper left indicates that these three stimuli are at the right-hand edge of the visual field. Upper right depicts their position one unit of time later. The 2DVM response to this velocity is displayed in the lower two images. The local map (left) displays the activation levels associated with the middle stimulus. The global map (right) reveals the suppression effect on the other two stimuli.

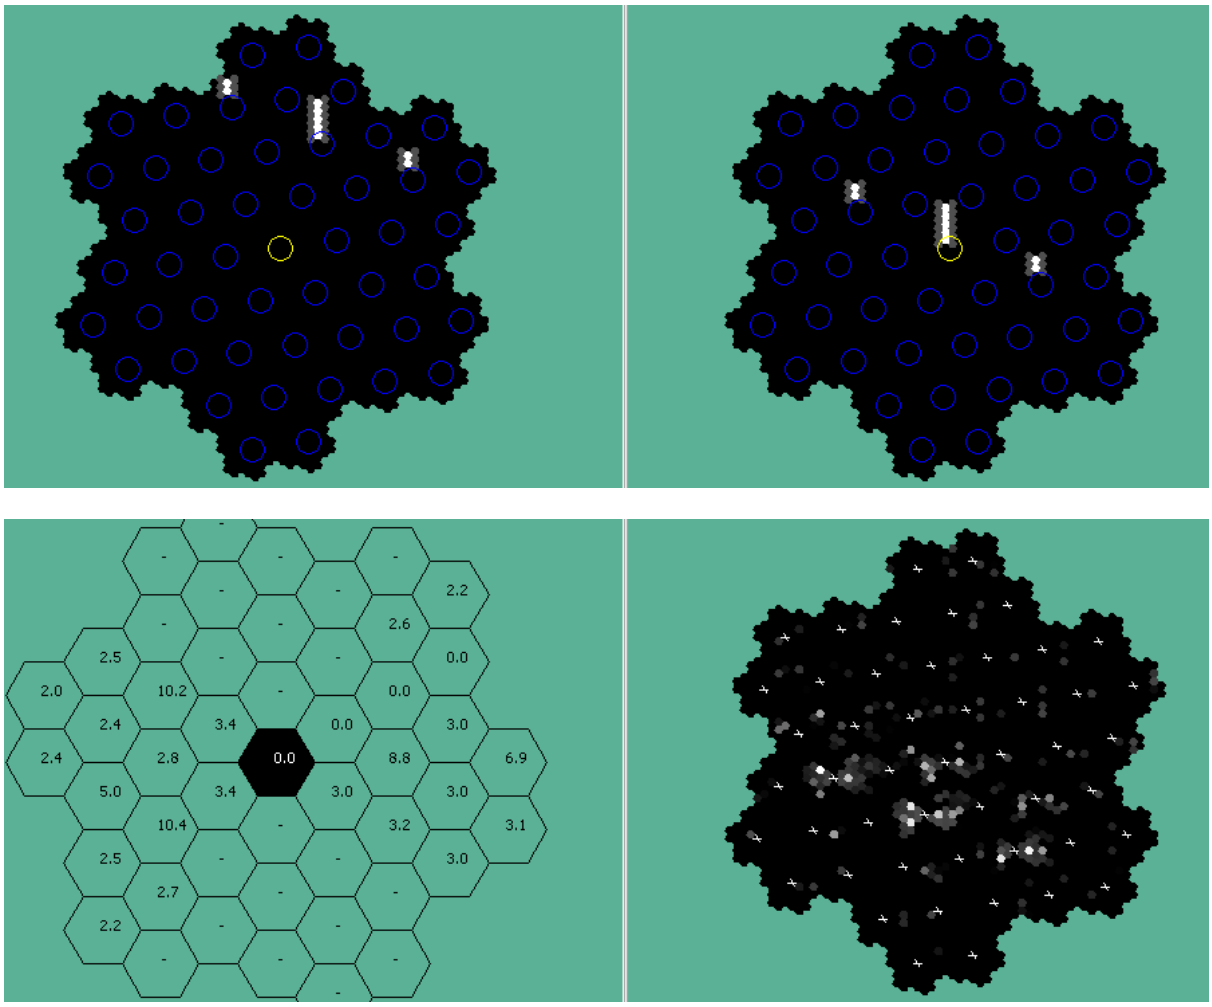

Fig. 10. Example of simulation results of the suppression effects on a group of fast-moving stimuli. The upper two images display the groups position at two distinct times and jointly they depict velocity. The circles indicate rfs. The three orientated bars represent parallel motion. Upper left indicates that these three stimuli are near the upper edge of the visual field. Upper right depicts their position one unit of time later. The 2DVM response to this velocity is displayed in the lower two images. The local map (left) displays the activation levels associated with the middle stimulus. The global map (right) reveals the suppression effect on the other two stimuli.

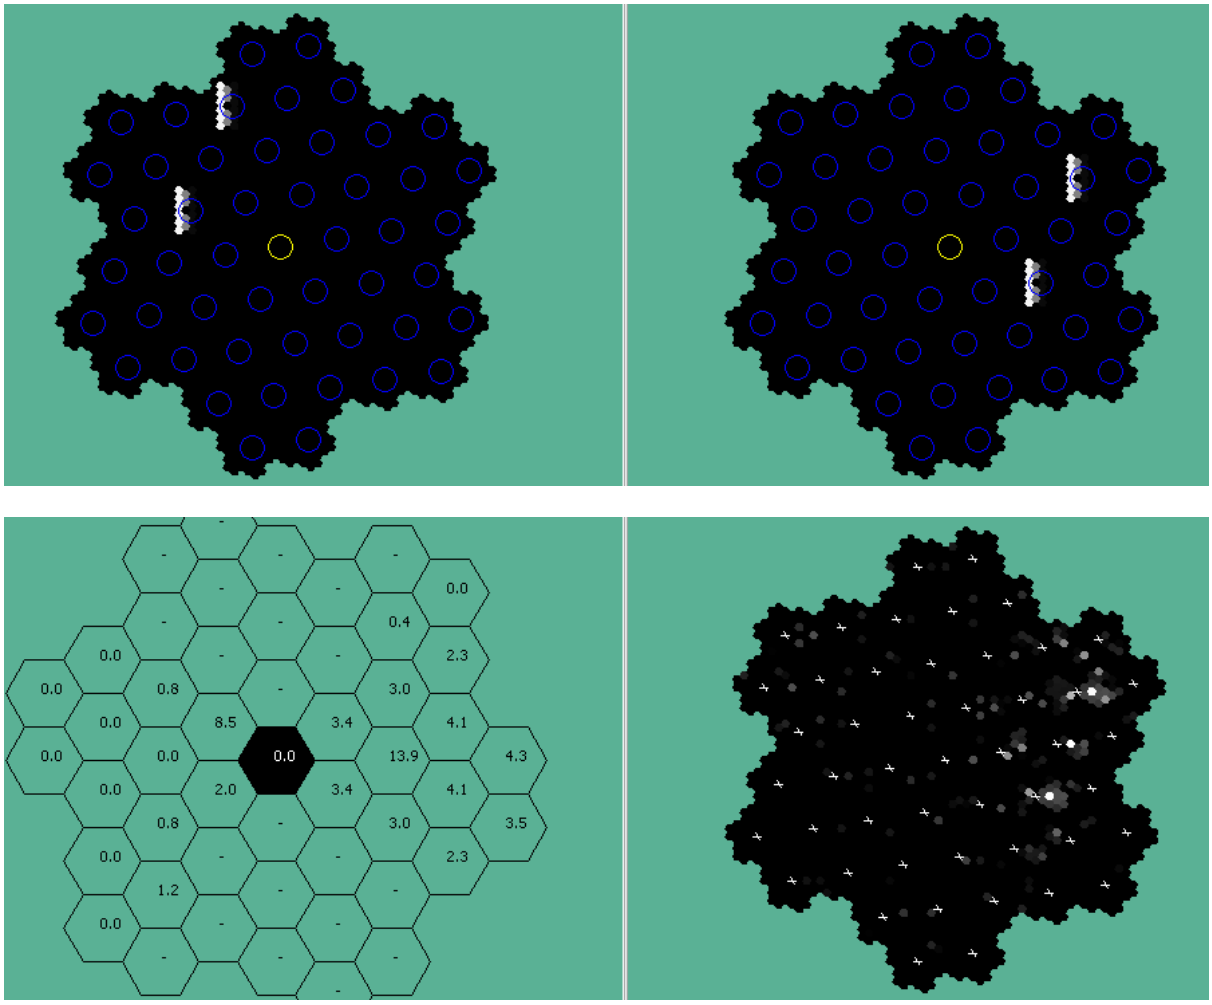

Fig. 11. Example of simulation results of the facilitation effects of two stimuli separated by a gap and moving at the same velocity. The upper two images display the groups position at two distinct times and jointly depict velocity. The circles indicate rfs. The two oriented bars separated by a gap represent orthogonal motion. The 2DVM response to this velocity is displayed in the lower two images. The local map (left) displays the activation levels associated with the gap. The global map (right) displays the facilitation effect on that local map.

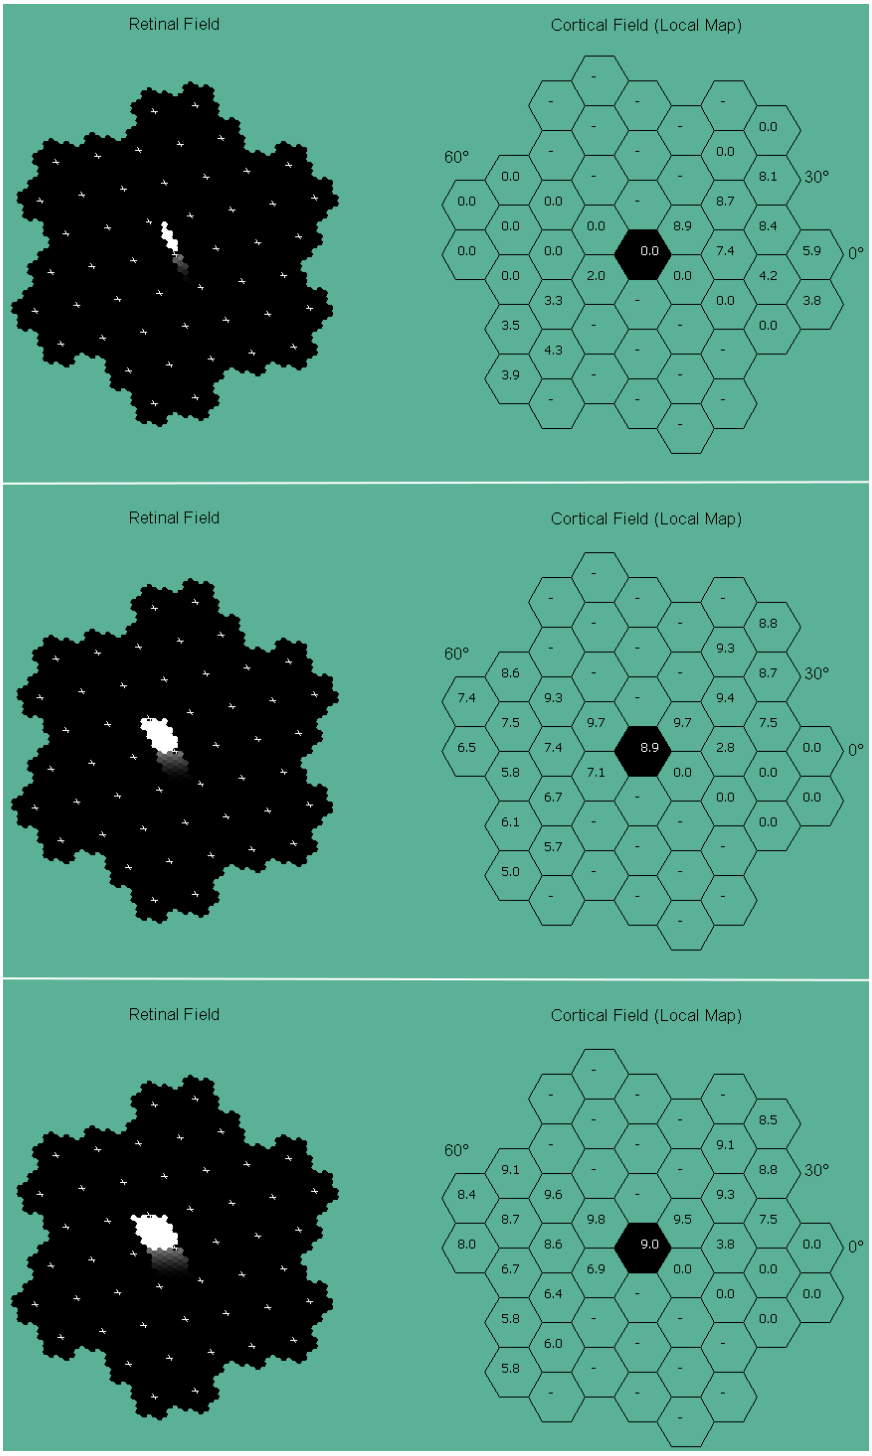

Fig. 12. Depicts illusionary orientation associated with increasing velocity induced by motion streaks. Each pair of sub-images represents a bar orientated at 30 degrees. The left images represent the projection of the bar onto the rf. The right images represent the local maps associated with the activated rf. Top pair: (left) zero velocity; (right) orientation contour line (OCL) at 30 degrees most activated; Middle: (left) small motion streak indicates a moderate speed; (right) OCL at 30 degrees and 60 degrees equally activated; Bottom: (left) longer motion streak indicating a higher velocity; (right) OCL at 60 degrees most activated.