Specification of Cortical Anatomy Within the Spiral Harmonic Mosaic Algebra

D. M. Alexander, Faculty of Information Technology, University of Technology Sydney, Australia.
P. Sheridan, School of Computing and Information Technology, Griffith University, Australia.
T. Hintz, Faculty of Information Technology, University of Technology Sydney, Australia.
J. J. Wright, The Liggins Institute, University of Auckland, New Zealand.

Abstract: We discuss the relationship between the local-global map hypothesis of V1 and Spiral Harmonic Mosaic algebra, illustrated with a simple model of V1. We describe long-range patchy connections by assuming they share a symmetry with the local connectivity. We also show that the elongated patterns of connectivity, implicated in the orientation response property, can be related to the patterns of connectivity seen with large injections of tracer. We predict a slight orientation anisotropy in macaque patchy connections.

1. Introduction

This paper discusses the relationships between the local-global map (LGM) hypothesis of the primary visual cortex [1,2] and Spiral Harmonic Mosaic (SHM) algebra [10,11]. The latter is described in detail in other papers in this volume. We begin with some basic definitions of the relevant physiology. A suitable companion paper, for those interested in the detailed anatomical discussions in this paper, is [6].

Neurons in the cerebral cortex are organized into columns. A column is approximately 210μm in diameter and contains around a thousand neurons. Columns are defined by so-called patchy connectivity that connects distant columns to each other within the same cortical area (e.g. within the primary visual cortex—area V1). Other column-sized features are revealed by staining techniques, such as cytochrome oxidase (CO) blobs. The cortical columns are assumed here to tile the cortical sheet as a noisy hexagonal lattice. A macrocolumn is a region of V1 approximately 500μm in diameter, or approximately 7 columns in area. A macrocolumn contains a complete representation of all possible orientation preferences, where orientation (of lines or textures) is the visual primitive to which V1 most strongly responds. The macrocolumn is a local representational map, and ~10,000 copies of this map tile the surface of V1 in the macaque. Various physiological features, such as inter-patch distance of patchy connections, and inter-blob distance of CO blobs, also tile V1 at a scale of ~500μm [6].

V1 is generally assumed to process only local features of the visual scene [9]. Such feature processing includes extracting orientation and spatial frequency components, in the absence of their more global meanings or contexts. The LGM hypothesis [1,2] of V1 diverges from this assumption. This hypothesis states that the global visual map in V1 is remapped to the macrocolumn scale in V1 in the form of a map of response properties e.g. orientation, spatial frequency preference and colour selectivity. These local maps tile the surface of V1 and each receive inputs from a large extent of the visual field. Because these local maps arise originally from global visual inputs, their organization reflects the properties of the visual field e.g. the CO blob in each local map responds preferentially to colour and low contrasts, reflecting the properties of the centre of the retina. So rather than the macrocolumn being simply a map of primitive visual features that apply to a point in visual space, the macrocolumn is a map of primitive visual features *as they arise in the organization of the the visual field and become relevant* to a point in visual space. Readers are

referred to [1,2] for a fuller discussion of this hypothesis.

Consider a simplified version of V1 in which the global image is represented twice: once as a global (identity) mapping of the visual field (in the input layer), and once as a local mapping of features in the visual field (in the feature layer). The local feature mapping need not be topologically identical to the input mapping. Such a scheme is described naturally within the framework of SHM. If the image is $10^n{}_7$ in size, we can project it to a $10^{2n}{}_7$ sized SHM, allowing each pixel in the image to expand to $10^n{}_7$ in size in the input layer. We can then perform some local feature extraction (such at edge detection) and multiply the resulting image by the inverse of $10^n{}_7$. The $10^n{}_7$ copies of the feature map are stored in the feature layer. This is illustrated in figure 1. Note that (a) every superpixel in the input layer has access to an entire feature map in the feature layer. But the local-global relationship is arbitrary. If we multiply both layers by the inverse of $10^n{}_7$ again, then the feature map is now the global map and the input image is now represented as a tiling of local maps. This property we call *local-global invertibility*, and is described in more detail in sections to follow. Since we multiply both layers by the same number, the relationships between corresponding, individual pixels in the two layers remain constant. It can therefore be seen that (b) every superpixel in the feature layer has access to the entire representation of the original image. This representational scheme enables us to have two simultaneous representations of the image, where each point in one representation has local access to all points in the other representation.

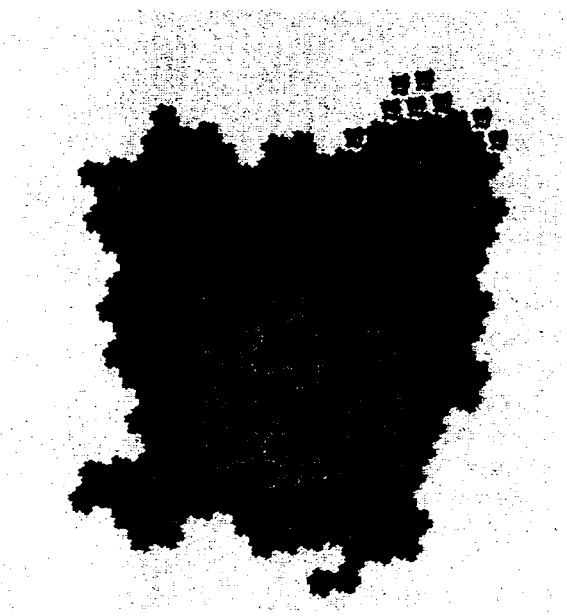

Figure 1: *Double mapping of the visual field to the input and feature layers of model V1. The global image of the bear is tiled with local feature maps (shown only top right).*

In our model V1, the P^{th} superpixel in the input layer has an adjacent feature map in the feature layer. There is some pixel, p_0, in this second map that corresponds to P. This p_0 pixel acts as a reference point in the local feature map, so that the features are defined as a set of relative relationships. So each superpixel in the input layer has access to its own feature role, relative to the other features in the visual scene e.g. P's activity is part (p_0) of a line (other p's) that is oriented globally at 60 degrees. Similarly, by argument of local-global invertibility, each superpixel in the feature layer has access to its role relative to the entire original image.

2. Specifying patchy connectivity

The mappings into each layer and other operations such as feature detection are carried out by sets of connections between neurons in the cortical tissue. These connections coordinate the neurodynamics responsible for signal transmission and transformation. Anatomical connections within V1 are revealed by injecting tracer into the grey matter. This method enables the patterns of monosynaptic connections to or from a particular set of neurons to be mapped. The tracer may be injected into a small focal region of the cortex (<200μm) to reveal fine-scale connectivity relationships, or into large regions of cortex (>500μm) to reveal large-scale patterns of connectivity. Two classes of connections can be distinguished within cortical grey matter when focal injections are used: 1. short-range intrinsic connections that are very dense around the injection site out to ~300μm radius (approximated by a 2-D Gaussian); and, 2. long-range patchy connections that arise for several mm around

the injection site, with patches ~200μm in diameter and an inter-patch distance of ~500μm. The 2-D Gaussian exerts a strong local excitation, and a ring of inhibition around this [13]. The basic effect of the 2-D Gaussian is to enhance and sharpen local maxima of excitation, while suppressing neighboring activity. Distortions of the Gaussian connectivity provide the properties required for feature detection.

In this paper we describe a novel interpretation of the patchiness of the long-range connections seen in V1. We assume the property of local-global invertibility, that applies to the input and feature maps, explains the connectivity patterns. Spatially separated (but otherwise identical) feature pixels will have average correlations in activity that decrease monotonically with distance. The connectivity structure reflects this, with the consequence that, at the global scale, long-range connections are patchy, while local connections have a Gaussian form. The LGM structure can be understood as a projection, onto the 2-D cortical sheet, of a 4-D representational space. Each local feature map forms a plane that intersects this 4-D space and each version of the input image forms a plane in the orthogonal dimensions. We assume that the long-range patchy connections are simply the projection of 2-D Gaussian connectivity, present in the planes within the 4-D space, onto the 2-D cortical surface.

We approximate the structure of a macrocolumn as a superhexagon made up of 7

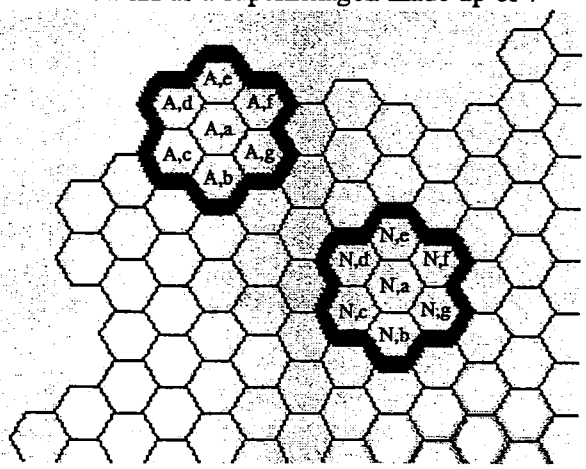

Figure 2: *Local and global coordinate system shown in relation to hexagonal lattice*

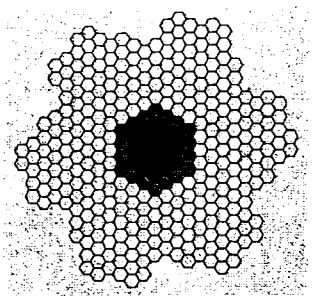

Figure 3: *Connectivity diagram into the centre point of SHM.*

columns. Each column is itself made up of approximately 49 minicolumns, but taking a column as the base unit is a useful approximation for the present purposes. Each column within SHM is assigned a global and local indice (G,l). This is achieved by breaking the SHM address in two, e.g. for hexagon address 234, G=23 and l=4 i.e the 4th column in the 23rd local map. All the global indices with the same value e.g. '*A*', are adjacent to each other. All the local indices of the same value e.g. '*a*', are separated from each other. This is illustrated in figure 2. We can make a 1-D cortex analog of this type local-global mapping, where the global coordinates are given by letters and the local coordinates by numbers:

```
A1 A2 A3 A4 A5 B1 B2 B3 B4 B5 C1 C2 C3
...C4 C5 D1 D2 D3 D4 D5 E1 E2 E3 E4 E5
```

The local Gaussian connectivity is projected onto the SHM, shown in figure 3. The grey area represents the local 2-D Gaussian connectivity extending out from the central site. All connections are reciprocal, so it is also a map of connectivity from all sites to the central site. Therefore, in each SHM connectivity map to follow, the grey values represent a set of weights w_{ij} and w_{ji}, for address i=0 and for all addresses j. Net excitatory effects are shown in dark grey, and net inhibitory effects are shown in light grey. Similarly we can introduce local Gaussian connectivity to the 1-D case (excitatory and inhibitory values are ignored here):

```
A1 A2 A3 A4 A5 B1 B2 B3 B4 B5 C1 C2 C3
...C4 C5 D1 D2 D3 D4 D5 E1 E2 E3 E4 E5
```

We now swap local indices with global indices (G,l) -> (l,G), invoking the assumed property of local-global invertibility. In the SHM connectivity scheme described here, this

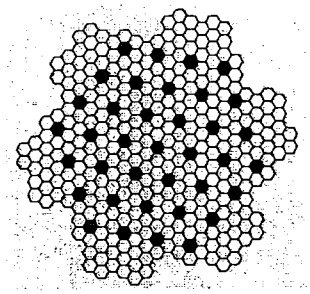

Figure 4: *Connectivity diagram after multiplication by 10^{-1}_{7}*

is equivalent to spiral multiplication by the inverse of 10_7. This results in a mapping of the local Gaussian connectivity onto the global map (i.e. enlargement to the global scale), while also breaking the connectivity up into discrete columns. Now the same-valued local addresses are all adjacent to each other, and the global addresses are separated. This transformation is therefore equivalent to inverting the local and global scales in our simple model V1 outlined earlier; the difference being that base units on this connectivity SHM are whole columns. To see how this works, consider the simpler 1-D case, after swapping the indices:

A1 B1 **C1** D1 E1 A2 B2 **C2** D2 E2 A3 B3 **C3**
...D3 E3 A4 B4 **C4** D4 E4 A5 B5 **C5** D5 E5

The addresses are still ordered, but now the local map scale is represented globally, and vice versa. We can represent the 1-D map in two dimensions, showing the connections as x's:

	GLOBAL				
	A	B	C	D	E
1	0	0	x	0	0
local 2	0	0	x	0	0
3	0	0	x	0	0
4	0	0	x	0	0
5	0	0	x	0	0

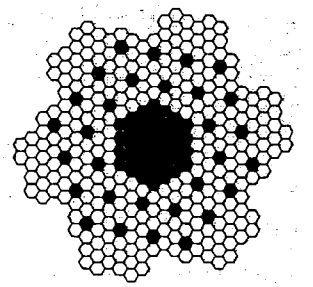

Figure 5: *Two connectivity patterns, one local and one global*

In accord with our assumption of an equivalence between connectivities in the local and global maps, we introduce a second set of local Gaussian connectivity to the connectivity SHM, while still in the (l,G) space. This is shown in figure 5. The same step in the 1-D cortex, but shown in the 2-D representation:

	GLOBAL				
	A	B	C	D	E
1	0	0	x	0	0
local 2	0	0	x	0	0
3	0	0	x	0	0
4	x	x	x	x	x
5	0	0	x	0	0

This is a connectivity diagram for site C4, i.e. it shows the connections to and from this column. The symmetry assumed for the two sets of connections is evident, appearing as a vertical and a horizontal row of x's in this 2-D representation. The sets of connections take up the entire width and height of this diagram for the sake of representational convenience, but *in vivo* may spread over only a part of the map.

We now swap local and global indices, thus returning to the (G,l) space. Our original local Gaussian connections are reconstituted, and the second set of Gaussian connectivity now becomes the long-range, patchy connections. This results in an identical mapping to figure 5 since we have ignored other asymmetries of the local and global mappings. This point can be made clearer by considering the 1-D case. We could break up the 2-D representation of this map in two ways, either representing the 'local' or 'global' dimension globally. In either case you would end up with a local Gaussian pattern and a global patchy pattern. Thus the local short-range connections and the long-range patchy connections, seen *in vivo*, are isomorphic in this scheme. They display local-global invertibility.

This symmetry is disturbed in real V1 in at least two ways. First consider local maps that are tiled in a mirror symmetric fashion. We break up the 2-D representation like weft threads in a rug:

```
                GLOBAL
              A   B   C   D   E

        1     0     0 x     0 0
local   2     0   0   x   0   0
        3     0   0   x   0   0
        4     x   x   x   x   x
        5     0 0     x 0     0
```

Our local maps are now mirrored about their borders. Now our original local connections are reconstituted (C1-C5), and the second set of connections appear as patchy, but not evenly distributed because of the mirroring of local maps:

```
A1 A2 A3 A4 A5 B5 B4 B3 B2 B1 C1 C2 C3
...C4 C5 D5 D4 D3 D2 D1 E1 E2 E3 E4 E5
```

Another way to disturb the symmetry between local and global connections is if the connections are not isotropic. For example, elongated patterns of connectivity are the presumed source of the orientation response prevalent in V1 [12]. The logarithmic distortion of the overall shape of patchy connections in the layers 2/3 of V1 [3] is the presumed source of rotation and scaling invariance hypothesized for the global V1 map [8,11].

3. Specifying orientation anisotropy

Focal injections into layers 2/3 of V1 (analogous to our feature layer) have revealed different patterns of long-range patchy connectivity across different species. In the tree shrew (a primitive primate) the connections are very elongated, and the axis of elongation matches the orientation preference of the injection site [4]. Similar results have been found for the squirrel monkey, but the connections show less elongation [8]. In the macaque, a lack of anisotropy associated with orientation preference has been reported for the intrinsic patchy connections in layers 2/3 [3]. Large injections of tracer reveal somewhat different patterns of tracer uptake, nevertheless related to the patterns seen with focal injections. In the macaque, the solid area of tracer injection gives way to regions with holes, or lacunae, followed by peninsulas of tracer, then a few islands at the furthest extent of uptake [6, figure 4A]. In the squirrel monkey, the peninsulas and islands are more prominent [6, figure 4B]. This trend is greatest in the tree shrew, where long arms of tracer uptake and patchy regions extend to almost the entire extent of V1 [7, figure 1B].

We now describe the effects of orientation anisotropy with SHM. Our strategy is to mathematically describe the shape of focal injections and then to show their relationship to the pattern of connectivity seen when large injections of tracer are made. The focal injection set $\{F\}$ consists of an adjacent group of addresses centred on SHM address 0. The shape of $\{F\}$ can be circular or oval, and is specified by a width and length, (x,y), in columnar units. Here we ignore the excitatory and inhibitory values of the 2-D Gaussian connectivity and just consider the overall shape of the connectivity.

We further introduce the assumption of [6], that "there exists a population of pyramidal neurons ... that does not make long laterally spreading patchy connections, but confines their axon arbors largely to within 600μm of their cell body" (see also [2]). This assumption is applied by specifying a smaller central column connection set $\{F_0\}$, consisting of an adjacent group of addresses, with a circular shape and of radius r. For the purposes of the current discussion, we ignore inhomogeneities in the connectivity pattern introduced by the tiling of local maps that are reflected about their common borders. For a version of tree shrew connectivity that includes reflected local maps, see [2].

Each macrocolumn m consists of 7 columns, p_k, with local indices {0,1,2...6}. The

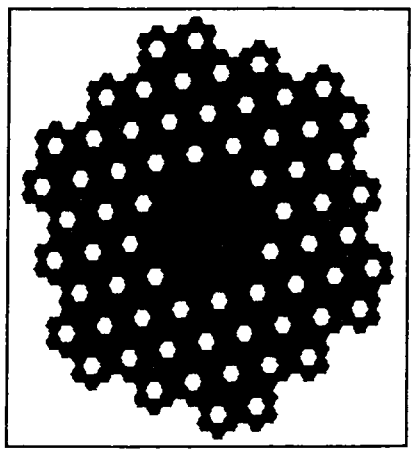

Figure 6: *Macaque connectivity diagram for large amount of tracer*

connectivity matrix for a given column in m is

$\{C_{p_k}\} = \{F\} \otimes p_k^{1/2} \otimes 10^{-1}{}_7 \oplus p_k$; for $0<k\leq 6$

$\{C_{p_0}\} = \{F_0\} \otimes 10^{-1}{}_7$

where $\otimes$ is spiral multiplication, $\oplus$ is spiral addition and $p_k^{1/2}$ is the address that performs pure rotation at half the angle value of p_k. The total set of connections for the macrocolumn is:

$\{C_m\} = \{C_{p_0}\} \cup \{C_{p_1}\} \cup \{C_{p_2}\} \cup ... \cup \{C_{p_6}\}$

The approximate anatomical parameters for different animals are as follows [3,4,9]:

Animal	x	y	r
Tree shrew	1	4	1
Squirrel monkey	2	4	1
Macaque	4	4	1

The model results for a large injection of tracer in the macaque are shown in figure 6. The site of injection is shown in black, sites of tracer up-take are shown in grey. A pattern of lacunae results. However, other features of the patterns seen in [6] are not present: the peninsulas and islands. If we introduce a further assumption of slight orientation anisotropy i.e. x=3, we get the pattern shown in figure 7, which is a better match to the macaque data. Note that anisotropy due solely to ocular dominance bands or the global logarithmic mapping of V1, and not orientation preference, will not produce arms. Instead it will produce an elongated version of diagram 6. If we introduce slightly more anisotropic connections (x=2), a pattern of tracer uptake similar to the squirrel monkey results. This is shown in figure 8. Larger 'injections' of tracer result in qualitatively similar patterns for each animal as the 7 column injections.

The above geometrical argument suggests elongation of patchy connectivity due to orientation preference exists in the macaque but is much slighter than in the squirrel monkey. The two competing sources of anisotropy, noted in [3], would make a slight orientation anisotropy difficult to detect. One way to demonstrate orientation related anisotropy is to inject 2 or 3 different types of tracer, one column apart, in layers 2/3 i.e. try to access different parts of the lattice that are added together in the preceding geometrical argument. This is shown in figure 9. Three injection sites are shown in dark grey and the tracer uptake is shown in lighter tones. Anisotropy due to other sources is not shown but does not change underlying argument. As can be seen, the slight orientation anisotropy assumed to fit the macaque data means the connection lattices fail to overlap by only one patch. It is possible that the slight anistropy predicted here for orientation preference is due to other sources, however, such as

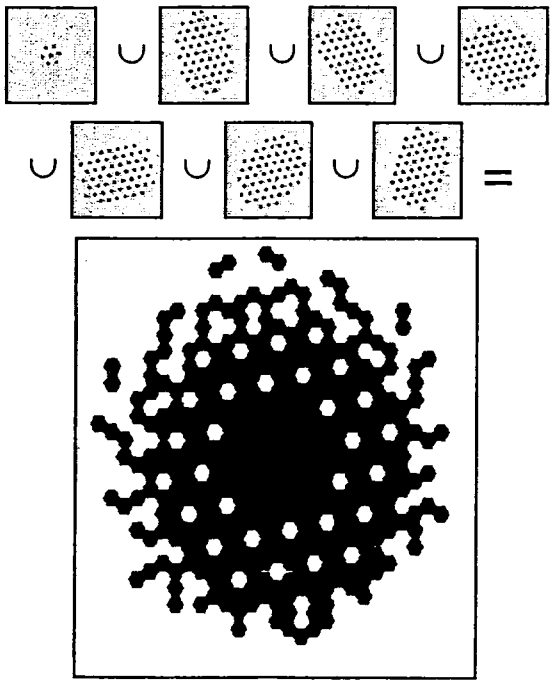

Figure 7: *Macaque connectivity diagram, assuming slight anisotropy (x=3), showing full set of focal connections as well*

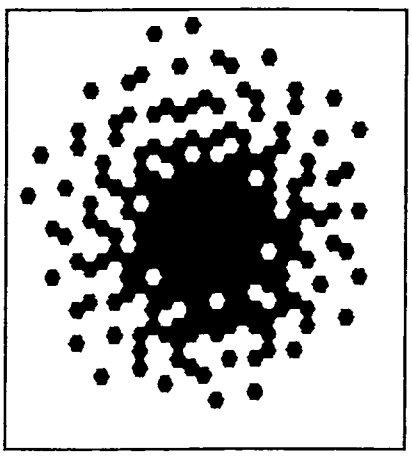

Figure 8: *Squirrel monkey connectivity diagram for large amount of tracer*

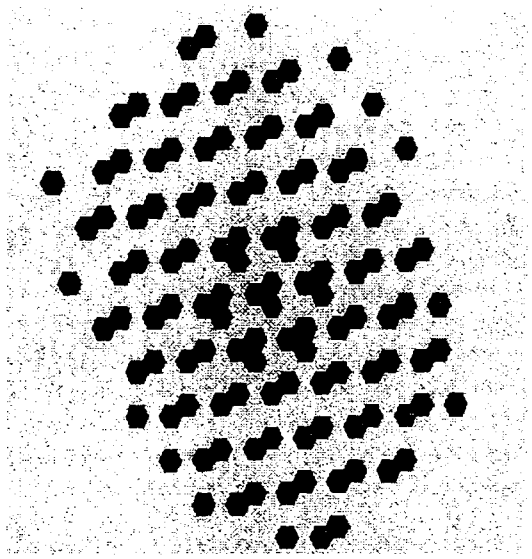

Figure 9: *Injection of multiple tracer types to test macaque orientation anistropy*

inhomogeneities in tracer uptake.

4. Conclusions

We have introduced a framework for describing long-range patchy connectivity seen in V1. The SHM provides a useful tool for characterizing these connections and enabling predictions to be made regarding their form. We have assumed that the property of local-global invertibility applies to V1 and that this implies an equivalent symmetry between the local and patchy connectivity. This assumption can be justified on the grounds of simplicity and elegance: we assume only one kind of functional connectivity and reproduce both types of connectivity. Future work will focus on an SHM version of the mapping from the input layer to layers 2/3 [2] as well as extending our work on the global logarithmic mapping [11] to develop an adjacency preserving version of this map.

5. References

[1] Alexander DM. Bourke PD. Sheridan P. Konstandatos O. & Wright JJ. (1998) Emergent symmetry of local and global maps in the primary visual cortex: Self-organization of orientation preference. Complexity International, 6, http://www.complexity.org.au/ci/vol06/alexander/

[2] Alexander DM. Bourke PD. Sheridan P. Konstandatos O. & Wright JJ. (2004) Intrinsic connections in tree shrew V1 imply a global to local mapping. Vision Research, 44: 857–876.

[3] Angelucci A. Levitt JB. Walton EJ. Hupe JM. Bullier J. & Lund JS. (2002) Circuits for local and global signal integration in primary visual cortex. Journal of Neuroscience, 22(19), 8633–8646.

[4] Bosking WH. Zhang Y. Schofield B. & Fitzpatrick D. (1997) Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. Journal of Neuroscience, 17(6), 2112–2127.

[5] De Valois KK. De Valois RL. & Yund EW. Responses of striate cortex cells to grating and checkerboard patterns. Journal of Physiology. 291:483-505.

[6] Lund JS. Angelucci A. & Bressloff PC. (2003) Anatomical substrates for functional columns in macaque monkey primary visual cortex. Cerebral Cortex, 13(1), 15–24.

[7] Rockland KS. & Lund JS. (1982). Widespread periodic intrinsic connections in the tree shrew visual cortex. Science, 215(4539), 1532–1534.

[8] Schwartz E. (1980) Computational anatomy and functional architecture of striate cortex: a spatial mapping approach to perceptual coding, Vision Research 20: 645–669.

[9] Sincich LC. & Blasdel GG. (2001) Oriented axon projections in primary visual cortex of the monkey. Journal of Neuroscience. 21(12):4416-26.

[10] Sheridan P. (1996) Spiral architecture for machine vision, PhD thesis, University of Technology Sydney.

[11] Sheridan P. Hintz T. & Alexander DM. (2000) Pseudo-invariant image transformations on a hexagonal lattice. Image and Vision Computing 18: 907–917.

[12] Somers DC. Nelson S. & Sur M. (1995) An emergent model of orientation selectivity in cat visual cortical simple cells. J Neurosci 15:5448–5465.

[13] Wright JJ. Robinson PA. Rennie CJ. Gordon E. Bourke PD. Chapman CL. Hawthorn N. Lees G.J. & Alexander DM. (2001). Toward an integrated continuum model of cerebral dynamics: The cerebral rhythms, synchronous oscillation and cortical stability. Biosystems, 63, 71–88.