

What's Black and White About the Grey Matter?

Rodney J. Douglas · Kevan A. C. Martin

© Springer Science+Business Media, LLC 2011

Abstract In 1873 Camillo Golgi discovered his eponymous stain, which he called *la reazione nera*. By adding to it the concepts of the Neuron Doctrine and the Law of Dynamic Polarisation, Santiago Ramon y Cajal was able to link the individual Golgi-stained neurons he saw down his microscope into circuits. This was revolutionary and we have all followed Cajal's winning strategy for over a century. We are now on the verge of a new revolution, which offers the prize of a far more comprehensive description of neural circuits and their operation. The hope is that we will exploit the power of computer vision algorithms and modern molecular biological techniques to acquire rapidly reconstructions of single neurons and synaptic circuits, and to control the function of selected types of neurons. Only one item is now conspicuous by its absence: the 21st century equivalent of the concepts of the Neuron Doctrine and the Law of Dynamic Polarisation. Without their equivalent we will inevitably struggle to make sense of our 21st century observations within the 19th and 20th century conceptual framework we have inherited.

Keywords Neuron Doctrine · Law of Dynamic Polarisation · Golgi stain · Canonical cortical circuits · High-throughput circuit reconstruction

Les Pensées

Anatomy tells you what could be, physiology tells you what is.

J Tony Movshon (personal T-shirt motto)

R. J. Douglas · K. A. C. Martin (✉)
Institute of Neuroinformatics, UZH/ETH,
Winterthurerstrasse 190,
8057 Zürich, Switzerland
e-mail: kevan@ini.phys.ethz.ch

Form ever follows function, and that is the rule. "By speaking generally, outward appearances resemble inner purposes"

Louis Sullivan (1896)

If you do not make headway understanding a complex system, then study its structure and knowledge of its function will follow automatically.

Francis HC Crick (1985)

We want to see how . . . the forms of living things, of the parts of living things, can be explained by physical considerations and to realize that in general no organic forms exist save such as are in conformity with physical and mathematical laws.

D'Arcy W Thompson (1971)

A microelectrode forces one to attend to what is at its tip, but there is no need for it to make one neglect what that structure does when one is not prodding it... We badly need all possible information on what one might call "principles and technology of neural engineering" and the only way to acquire it is to relate anatomical structures and cellular function to overall performance.

Horace Barlow (1977)

We always have to start with structure. All the books start with structure. You can't know how an organ works unless you know the structure of what is working. Yes, yes! But pages of anatomy are utterly indigestible unless one can appreciate what part the structure plays in the working of the organ. And to describe in detail what is there is so much easier that to discover the part it plays that the great chapters in minute anatomy—those deserts of detail without a living functional watercourse, only a mirage from unverified speculation—are nearly unreadable. That

chapter is like a dictionary, not to be perused from cover to cover, but to be consulted, a word at a time to throw light on some particular obscurity.

William AH Rushton (1977)

Introduction

'*La reazione nera*'—the black reaction discovered by Camillo Golgi in 1873 provided the typeface that allowed generations of anatomists to read the letters of the nervous system—the neurons and glia. Santiago Ramon y Cajal was well-primed to take advantage of the serendipity and soon discovered that *la reazione nera* worked even better in neonatal and foetal material. Both Golgi and Cajal stared down their microscopes at the same brain structures and drew them in detail. However, it was only Cajal who could take the single cells he saw and turn them into the circuits that we still use as references. The grammar he needed to link individual neurons into meaningful circuits was provided by two concepts: the Neuron Doctrine, which proposed that nerve cells were discrete, and the Law of Discrete Polarisation, which gave direction to the flow of activity. The Neuron Doctrine was first stated formally by Waldeyer-Hartz in 1891 on the basis of the observations of many anatomists, including household names like Golgi, Nissl, Forel, Koelliker and Cajal (Shepherd 1991). In his autobiography Cajal (1937) described in detail how he collaborated with the 'savant of Louvain' (Arthur von Gehuchten) to formulate the Law of Dynamic Polarisation. The Law declares that activity flowed from the dendrites to the soma and thence down the axon to the next neuron in the chain. Charles Sherrington provided the final missing links of this conceptual leap: 'synapse' and 'excitation' and 'inhibition'. These were the functional links between neurons that Ramon y Cajal could not see, but Sherrington could infer from physiology.

It was Sherrington more than anyone who provided the evidence that cemented the belief of every neuroscientist who followed him that the fundamental cause of behaviour is the generation and passage of nerve impulses from receptor to neuron and thence to muscle. We would all agree that the physiological processes by which this passage of nerve impulses happens are no longer mysterious. On the contrary, since the first electrophysiological measurements of 'action currents' in neurons by de Bois-Reymond in 1843, the generation and conduction of the nerve impulse has been a central area of investigation in neurophysiology. It is now beyond debate that the events that lead to the generation of nerve impulses, their conduction along axons, their translation as synaptic conductances, the integration of synaptic potentials in the

cables of dendrites and the regeneration of the axon potential at the soma and axon hillock are well-understood and have been found to exist in simple nervous systems as well as in the most complex. Indeed, the stages along this extraordinary intellectual journey of discovery are memorialised with milestones that mark the chapters of modern neurophysiology.

The theories that describe these physiological events are about as mathematically secure as anything gets in biology. The formalism originally developed by Alan Hodgkin and Andrew Huxley for the action potential, for example, continues to shape our thinking and has been of such wide-ranging power of explanation that it applies not only for the original case of Loligo giant axon, but for excitable cells in general, including muscle. Similarly the application by Wilfrid Rall of William Thomson's cable theory to neurons provided us with a second powerful conceptual tool for thinking about electric current flow in dendrites and axons. By coupling these formalisms to the power of modern digital computers it seems inevitable that it is only a matter of time before the heady ambition to simulate large portions of the brain of any animal of our choice will be realised. Already, sophisticated suites of off-the-shelf software are available so that from a standing start a student can, within a few months, run detailed simulations of models of any nerve cell. The portage of these tested programs to large parallel machines makes such relentless progress (e.g. Hines et al. 2008) that the exciting prospect offered by those involved is that within a decade we will be able to simulate the entire brain (<http://discovermagazine.com/2009/dec/05-discover-interview-the-man-who-builds-brains>). The only substantial issues that remain to be worked out is how to speed up the tracing of all the connections at synaptic resolution (see. e.g. Helmstaedter et al. 2008; Seung 2009, 2010) and to fill in the details of the molecular and cell biology of developing and adult brain. However, a brief reality check indicates that while superficially major problems of circuitry and function seem tantalisingly near solution, this is actually far from the case.

Revealing Circuits: A Short History

Our own major goal has been to explore the structural and functional implications of our hypothesis that there are canonical cortical circuits—basic circuits that can be found in all areas of the neocortex in all species. When we first developed this idea in the 1980's, the single unit was king, whether in theory, electrophysiology, or neuroanatomy. Against some resistance we predicted the future shift would be from analyses of the single neuron to the study of neuronal circuits. Since that review (Douglas and Martin 1991) there have been two major technique-driven shifts in

mainstream neuroscience. Firstly, the predictable, but nevertheless impressively fast expansion of molecular biological techniques and tools for physiology and circuit tracing (see reviews by Callaway 2008; Boyden et al. 2005) that has made molecular and genetic neuroscience a major growth area, so much so that the prediction is that we will now be able to ‘reverse engineer’ the brain of a rat or mouse (Kalisman et al. 2005; O’Connor et al. 2009). Secondly, the use of brain imaging, particularly fMRI, for exploring human perception and cognition, has led to an avalanche of correlative data. These two advances, however, pull in different directions. While molecular-genetic techniques focus attention down to the role of gene expression on neural development and behaviour, mainly in rodents and invertebrates, brain imaging magnifies the functional map of the human brain, and increasingly directs our attention to aspects of cognition and emotion. Paradoxically, neither of these two approaches address directly the issues of what all these brains are doing or how they do it, just that they do it. It is here that the analysis of neural circuits offers its greatest attraction: to be the bridge between the molecular and the cognitive.

Historically, brain circuits have been studied post-mortem in slices of brain. Although electrophysiological and behavioural methods have made a very significant contribution, the classical neuroanatomical techniques of staining and tracing processes in the light microscope is still the gold standard for revealing the structure of brain circuits. The lure of the current program of detailed circuit reconstruction, both in the sense of high fidelity physical reconstruction of neurons and their connections (and to be able to do this on a large scale—the goal of the DIADEM Grand Challenge) and of the high fidelity simulation of the electrical structures of neurons and synapses, is that it seems to capture a wide range of relevant phenomena. These include all the usual suspects: coding, decoding, spike timing, linear and non-linear dendritic integration, stochastic synapses, synaptic plasticity, network learning, neural representations, and so forth. The product of such a program is at least encyclopaedic in cataloguing the known world of the circuit, but also makes the strong claim, *a la* Barlow (see below), that it is possible to bootstrap from single neurons to the high cognitive levels. However, a number of present attempts at a bottom-up only solution seem like programming a computer by setting the individual switches.¹

¹ While at first glance this looks like an interesting inversion of the flow diagram of discovery suggested by David Marr (theory → algorithm → implementation) in fact it goes much further, for to set the switches implies that one understands the goal of the computation. Here the (perhaps unconscious?) strategy is to avoid having to discover the goal of the computation altogether.

Satnav for the Brain?

We can pose the simple question: do we understand the paths nerve impulses take in their journey between input and output, even for simple nervous systems? Of course, for most people this question is absurd, for have anatomists not ‘solved’ many brain circuits, as in the examples of the entire nervous system of *C. elegans*, central pattern generators (CPGs) in crustacea, retinae of insect and vertebrates, and even the circuits of the central nervous system, like thalamic nuclei, cerebellum, olfactory bulb, and hippocampal formation? Ramon y Cajal was so sure about the polarity of the many circuits he saw that he annotated his circuit diagrams with arrows to show the direction of flow of the nerve impulses (Fig. 1). Even that formidable challenge, the neocortex, which ultimately defeated Cajal, now has partial solutions, albeit mainly from neurophysiologists.

Lorente de Nó, fresh from his studies of motor neuron physiology, provided the text for possibly the most cited book chapter in neuroscience (Chapter 15 of Fulton’s, ‘The Physiology of the Nervous System’). It contains his description of the cortical circuits of the mouse barrel cortex, which he misidentified as the acoustic cortex in his original 1922 paper (see Lorente de Nó, 1992). The 1922 paper is almost never cited, which is unfortunate, because it contains the most complete description of the neuronal contents of barrel cortex—a description that was unsurpassed in the 20th Century. Lorente’s medium was the Golgi technique, but unlike Cajal, his arrows had shrunk and provided only a minimalist indication of input or output (Fig. 2).

Lorente’s most insightful structural observations were that thalamic input entered at all levels in the cortical layers; that local cortical circuits were recurrent; and his embryonic notion of a cortical column. The notion of column was picked up by Mountcastle (1997), who made it a center point of his investigations of the somatosensory cortex. Looking down his microscope, Lorente thought he saw Sherrington’s reflex arcs, but curiously he, like Cajal, ignored Sherrington’s evidence that inhibition always acts in concert with excitation. Nevertheless, in deducing that the effect of impulses entering the cortex must depend on the activity currently circulating in the cortical circuits, Lorente arrived at a view very different from the concept of serial processing that was soon to emerge from physiological studies of the cat’s visual cortex.

It was the Two Circuiteers, David Hubel and Torsten Wiesel (1962), who showed the way with their text-book circuits for the visual cortex. Their logical ordering of interlaminar processing showed how raw retinal information is serially transformed to extract features of the visual world. Indeed, their cortical circuits, later elaborated by Charles Gilbert and Torsten Wiesel (1981; see Fig. 3), do

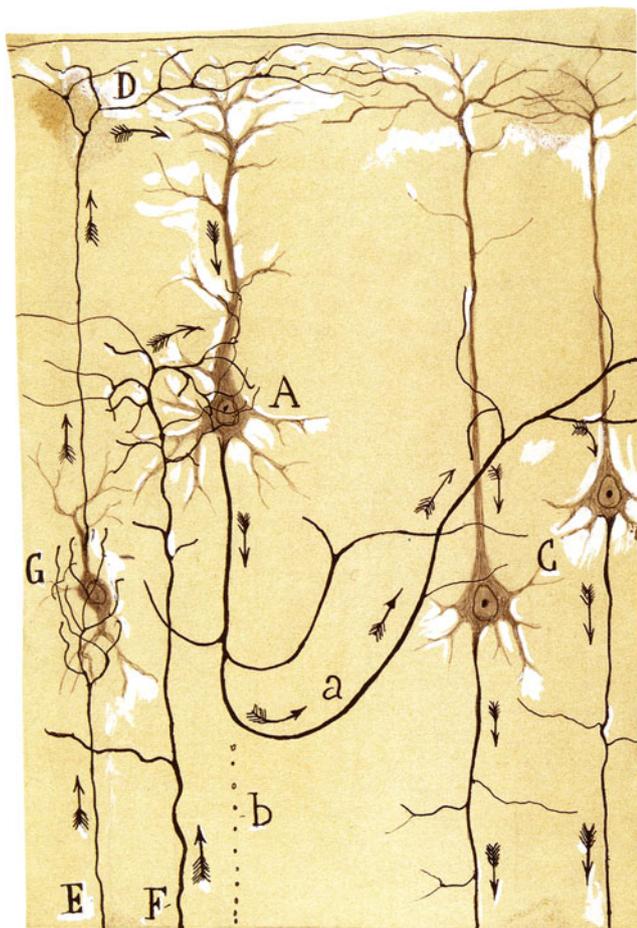


Fig. 1 Ramon y Cajal's observation and interpretation of part of the neocortical circuit. His legend reads: 'Schematic drawing showing the possible track of currents across the arciform pyramidal cells. A, archiform pyramidal cell; C, normal pyramidal cell; D, plexiform layer; E, F, afferent fibers; G, cell with an ascending axon, called Martinott's cell; a, hypertrophic collateral; part of the axon that has disappeared'. Reproduced, with permission, from Cajal, S. R. (1913-1914)

not look out of place alongside other brain circuits 'solved' by Cajal. Hubel and Wiesel's ability to conceptualise circuits whose emergent physiological properties were evident at a glance, set the scene for a much larger scale venture by David Van Essen and colleagues in which the entire cortical visual system was also ordered in a hierarchy of serial processing, crowned by a final convergence on the hippocampus (see Felleman and Van Essen 1991). Graph theoretical analysis has shown the macaque visual hierarchy to be indeterminate (Hilgetag et al. 1996). But this fact has not deterred the iconic status or utility of the Felleman-Van Essen schema as a visual mnemonic. All the instantiations of forward local circuits neglect the inescapable fact that every neuron is polyeuronally innervated. Yet for simplicity's sake we tend to assume that only one 'driving' pathway exists and all others are 'modulating'. Thus, the local interlaminar circuit of Gilbert and Wiesel see (Fig. 3

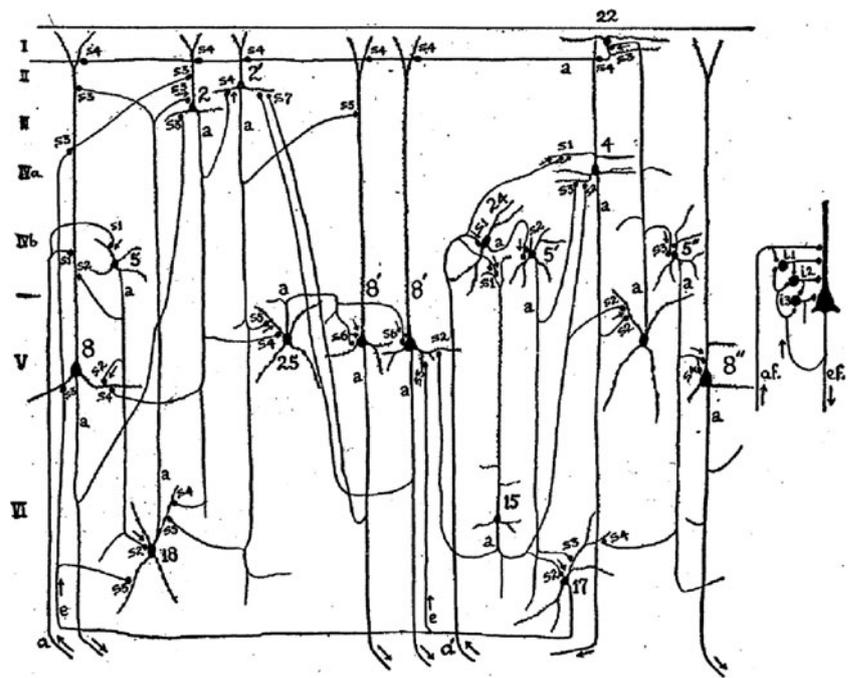
legend) accounts for only about 20% of all the excitatory synapses formed by local neurons (Binzegger et al. 2004). Similarly, the inconvenient truth is that any single area receives convergent input from many subcortical and cortical sources (Fig. 3) is usually conveniently put to one side.

Conceptual models of serial circuits have led a long and largely untrammelled life. From studies of reflexes, Sherrington was able to infer the circuits of monosynaptic spinal reflexes free of the complexity of the interneurons or the ascending and descending pathways (which he could remove by isolating the spinal cord from the brain). He also predicted the likely existence of hierarchies of processing in the neocortex, although the implication of a final pontifical or 'grandmother cell' was intellectually and aesthetically distasteful for him (Sherrington 1940). His pupil Jack Eccles, together with Masao Ito and Janos Szentagothai (1967), demonstrated the machine-like passage of impulses through the cerebellar circuit that Cajal had drafted. Indeed, the insistence that processing hierarchies in the visual system follow the arrows from the photoreceptors to the lateral geniculate nucleus and from thence up to the first stage of cortical processing and beyond to extrastriate cortex, turns out to solve effectively most problems of visual processing. The forward arrows continue right up to the inferotemporal cortex, in the case of object recognition (Kouh and Poggio 2008), and to far extrastriate areas like MT (Heeger et al. 1996), and the parietal cortex in the case of the coordinate transformations associated with spatial awareness (Cohen and Andersen 2002). While we perhaps do not yet know enough about the relevance of the ubiquitous 'feedback' arrows that point in the opposite direction, ignoring them has not led to any major difficulties of explanation. The major lack is that the models are not biologically as accurate as they could be (not 'need be').

Linking Single Neurons to Perception

In 1972, Barlow introduced his own 'Neuron Doctrine' to explain how high level perception can be mediated by a few thousand 'cardinal neurons', which respond selectively to high level features such as faces. This is not so far-fetched as it once seemed in the intervening era of 'connectionism', with its neural networks and distributed representations, for recordings from single neurons in human mediotemporal lobe (i.e. the top of the cortical hierarchy), indicate that a very high degree of selectivity and rapid learning of faces and objects exists at the level of single neurons (Quiroga et al. 2005; Mormann et al. 2008). Crucially, Barlow (1972) claimed that we really could understand how the brain works by studying the activity of individual cells. His central proposition was that perceptions are caused by a

Fig. 2 Diagram of intracortical chains of neurons. Lorente de Nó proposed that functionally the cortex is composed of vertical processing chains of neurons, rather than of layers. In his view, neurons in all layers except 1 and 2 receive input from the thalamus. The input is then propagated vertically to the most superficial layers. His arrows are barely visible. (In Fulton 1949, pg 307)



small number of active neurons embedded in an otherwise silent sea of neighbours. These active neurons reflect the specific clusters of features that represent the present object. By this simplification Barlow cut through the Gordian knot of details that all the combinatorial alternatives for representation of visual scenes produce. Implicit in this is the idea of efficient coding, where the goal is to represent sensory information with the fewest number of active neurons. Of course Barlow also considered the number of active neurons lower down the hierarchy—if there are K neurons active out of a population N , then N increases as one moves up the hierarchy (fan out), but how does K change? Quiroga et al. (2008) believe that K decreases, but not to the small numbers implied by the extreme ‘grandmother cell’ view. The result of Barlow’s simplification is impressive, because his Neuron Doctrine validates the idea that the circuits that create cardinal cells are functionally simple, even if they are composed of neurons that have complicated geometries and a plethora of receptors and ion channels. As Barlow pointed out, these details are interesting in and of themselves, but knowledge of them is not a prerequisite for developing concepts of how the brain works.

Fine Print—Again

This lack of pressure to engage with the details of circuits has a negative consequence however: that we are always at risk of falling back into the simple circuits that have a wide comfort factor and familiarity, as traditionally we have

‘understood’ Ramon y Cajal’s solutions to the circuits of retina and cerebellum. The Golgi technique, however, was best applied to immature material, because the myelinated parts of the axon were not impregnated. The brain was usually immersion fixed, usually in alcohol, a procedure that would certainly produce post-mortem artifacts (e.g. Tao-Cheng et al. 2007). It was only with the advent of modern perfusion fixation and *in vivo* labeling techniques that we first got an inkling of the true complexity and extent of individual axonal arbors. In the cat’s visual cortex for example, the reconstructions provided in the 1970’s by Ferster and LeVay (1978) and Gilbert and Wiesel (1979) were a revelation for those brought up on Cajal’s 1921 study of the cat’s visual cortex. We have also never properly got to grips with the fact that all neurons in the CNS are polyeuronally innervated and usually recurrently so. At any one moment the source of the input that drives a neuron to spike may be a mix of excitation and inhibition arising from many source neurons whose own firing has various degrees of correlation. The temptation in thinking about simple circuits is always to consider only one source to be the major excitatory drive, with inhibitory neurons providing a balance and other minor connections, if present, providing ‘modulation’. This train of thought is strongly reinforced by both the conceptual frameworks, e.g. hierarchies, parallel processing, and by the experimental methods, especially *in vitro* recordings.

In the case of nuclei in the thalamus, the sensory input from peripheral structures such as the retina, is the primary excitatory drive. Similarly for the neurons in primary sensory cortex, the thalamus is the primary excitatory

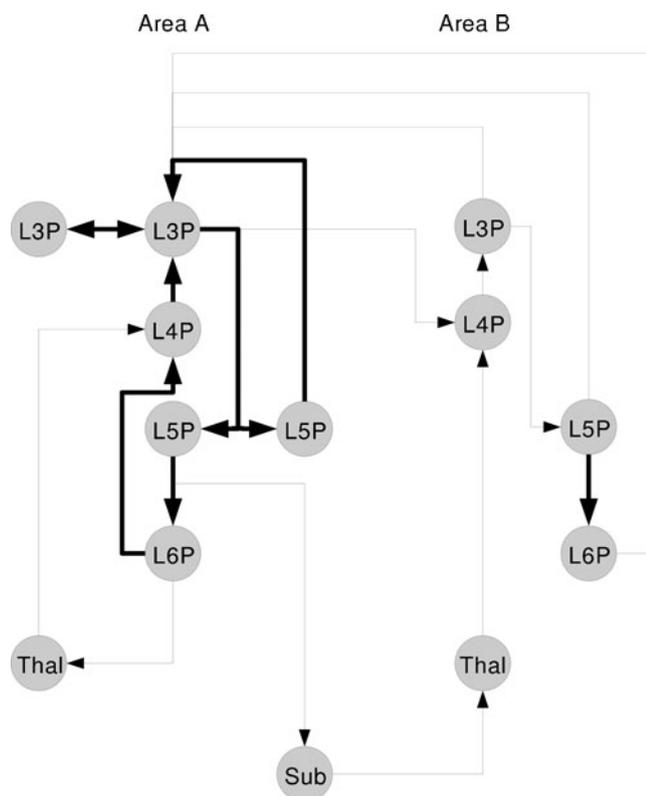


Fig. 3 Graph of connections of the excitatory cell types in neocortex and their cortical and sub-cortical relations. The nodes of the circuit graph are organised approximately spatially; vertical corresponds to the layers of cortex, and horizontal to its lateral extent. Directed edges (arrows) indicate the direction of excitatory action. Thick edges indicate the relations between excitatory neurons in a local patch of neocortex, which are essentially those described originally by Gilbert and Wiesel (1981) and Gilbert (1983) for visual cortex. Thin edges indicate excitatory connections to and from subcortical structures, and also inter-areal connections. Each node is labeled for its cell type. ‘Lx’ refers to the layer in which its soma is located. ‘P’ indicates an excitatory neuron (generally of pyramidal morphology). ‘Thal’ denotes the thalamus, and ‘Sub’ other sub-cortical structures, such as the basal ganglia

drive. However, in both cases the primary sensory afferents form only 5–10% of the synapses on their target neurons. So, although we are surprised, perhaps we should not be, when we find that optical imaging of cortex by voltage sensitive dyes reveals rapidly changing ‘spontaneous’ patterns of activity that cannot be explained simply by patterns of thalamic activation (Tsodyks et al. 1999). Lorente de Nó’s (1949) notion of the state dependence of the effects of the afferent impulses reemerged with the realisation from the work with voltage sensitive dyes that if much of brain activity is self-generated then the role of the sensory periphery is not so much as driving the target structure, but in dynamically biasing this activity towards a certain learned/imprinted pattern sets. The computational question is: how many different sets?

Reconstruction = Nirvana?

The contemporary programs to reconstruct the cortical column in all its glorious detail offer a comprehensive cure for the temptation to sit with simple circuits (Markram 2006, 2010; Helmstaedter et al., 2007). However, there is a major conceptual difficulty: what kind of thing would the result of such a program be? Inevitably it will be incomplete, given that no-one can actually agree on what constitutes a cortical column or even whether the column is still useful in the 21st century, even as a concept (da Costa and Martin 2010; Rockland 2010; Horton and Adams 2005; Markram 2008). Even if the volume of cortex to be reconstructed can be defined, the major hurdle is the practical difficulty of getting structural and physiological data that are at all comprehensive. Over the past decade or more researchers in well-funded institutes (MPI Heidelberg, EPF Lausanne, HHI Janelia Farm) have devoted their considerable efforts to solve the ‘cortical column’ of the rat barrel cortex through an attempt to convert the cottage industry of circuit analysis into an ‘industrial-scale’ process of gathering data. Yet the industrial scale has not been realised for any aspect of the process: the enterprise has remained garage-scale and many bottle-necks remain to be solved before we have a production line of even Henry Ford quality. Not least there is the problem of gaining high quality, high resolution structural data from *in vivo* preparations. The present databases are mostly derived from *in vitro* recordings from brain slices of immature/juvenile cortex and a smaller number of more complete neurons stained *in vivo* and reconstructed at light microscope resolution. Thus the number of reconstructed neurons remains in the thousands, rather than the tens of thousand, and the connections are assumed on the basis of paired recordings and light microscope evidence of contacts between axons and dendrites. This is a hazardous assumption since our own correlated light and electron microscope estimate was that only 20% of the contacts seen a light microscope level turn out to form synapses (da Costa and Martin 2009). Despite the intensive work, a consensus picture of the quantitative local circuit has yet to be unveiled for the barrel cortex and since most of these neurons are obtained from different animals these data provide an average for each type of neuron. Even when this cataloging is done there remain major gaps in the structural analysis: e.g., the essential quantitative estimates of numbers of thalamic afferents, synaptic densities, identified projection neurons etc. These hard won data are there for the cat, but are only now emerging for barrel cortex from one of these groups (Meyer et al. 2010a, b; Wimmer et al. 2010).

For the same reasons that motivated the DIADEM Grand Challenge, there are intensive efforts to develop ‘high-throughput’ EM methods for synaptic resolution reconstruc-

tions. These efforts are to be strongly encouraged—they will be a boon for many applications—but even if wholly successful according to their present dreams, they will not provide for reconstructions of more than a few cubic mm of tissue (Helmstaedter et al. 2007; 2008; Hayworth et al. 2006; Kasthuri and Lichtman 2010; Lichtman et al. 2008). Despite the fact that many seem to be tripping over themselves in the gold rush for the ‘connectome’ (Kasthuri and Lichtman 2010; Seung 2009), or the ‘projectome’ (Kasthuri and Lichtman 2007), or even worse, making hubristic promises that we will soon be able to describe a ‘diseasome’ for the brain, any data will be partial at best. A circuit at synaptic resolution for the whole neocortex will inevitably be very significantly incomplete, whether in mouse, cat, or monkey, let alone the particularly large brain of *Homo sapiens*.

In the present flood of enthusiasm for high throughput EM reconstruction, one should also not lose sight of the fact that the major advances in circuit reconstruction for the foreseeable future will remain at the light microscope level, as they have always been. It is here that we have a major bottleneck, for although we have many ways of labeling neurons *in vitro* and *in vivo*, including intracellular labeling, electroporation, virus vectors, transgenic expression of Green Fluorescent Protein in restricted neuronal types, etc., we still have no means of automated reconstruction of individual neurons. Even with its computer assistance, the method we currently use for much of the wide field reconstruction is essentially no different from that used by Cajal and Golgi. There are no grand prizes for guessing why DIADEM is termed a Grand Challenge: We have not yet found simpler route than theirs through to the high-throughput light microscopic reconstructions that form such an essential part of our explorations of the nervous system. Nevertheless, aided by Moore’s Law and a growing armory of very clever algorithms, such as those on display at DIADEM (Bas and Erdogmus 2011; Chothani et al. 2011; Narayanaswamy et al. 2011; Turetken et al. 2011; Wang et al. 2011; Zhao et al. 2011), we can expect a rapid improvement in the methods and speed of tracing axons and dendrites (see for example, the many novel tracing methods described elsewhere in this collection). But, what will we do with our knowledge of all this wire?

Possible, Art of,

For those of us working on the large neocortex of higher mammals, such as cat and monkey, our aim is not to attempt a *C. elegans* circuit diagram, where the origin and destination of every wire is known. Instead our aim is to achieve a solution that is rule-based and probabilistic. What this means is that we understand the rules of connections between different types of neurons, that we have a census

of neuron and synapse numbers, and that we understand the 3-D pattern of dendrites and axons of the different cell types. Presently a computer-assisted 3-D reconstruction of a single cortical neuron takes our very skilled operators a minimum of 100 h (Fig. 4), so if the aims of DIADEM were fully realised this would greatly speed our investigations. By coupling a statistical approach with analyses of the rules that organize the circuits, we also have a means of predicting what circuits could/should exist, given accurate knowledge about the cells types in play.

The statistical approach pioneered by Donald Sholl (1956) and by Valentino Braitenberg (Braitenberg and Lauria 1960) is half a century old. In its modern incarnation it is better known in its eponymous form of ‘Peters’ Rule’ (the term coined by Valentino Braitenberg and Almut Schüz in their classic 1991 duograph). Peters’ Rule is the principle that pre- and post-synaptic elements connect in the same proportion in which they appear in the neuropil. Alan Peters and colleagues applied Peters’ Rule (before it was named) to the connections between the lateral geniculate nucleus and cortical neurons in the rat visual cortex (Peters and Feldman 1976). As yet there has been only one attempt to apply Peters’ Rule to estimate the total synaptic connectivity for all neurons in a local circuit of neocortex and this was by Binzegger et al. (2004), who made this estimate for cat area 17, because it is the only cortical area for which sufficient quantitative information had accumulated over the decades. These data include the cell types, number of neurons and synapses in each layer, thalamic inputs, projection neurons, and detailed reconstructions from *in vivo* material to provide an estimate of the contribution of area 17 neurons to the total synaptic complement of area 17. Of course, known exceptions to Peters’ Rule were allowed for: for example chandelier cells only form synapses with pyramidal cells.

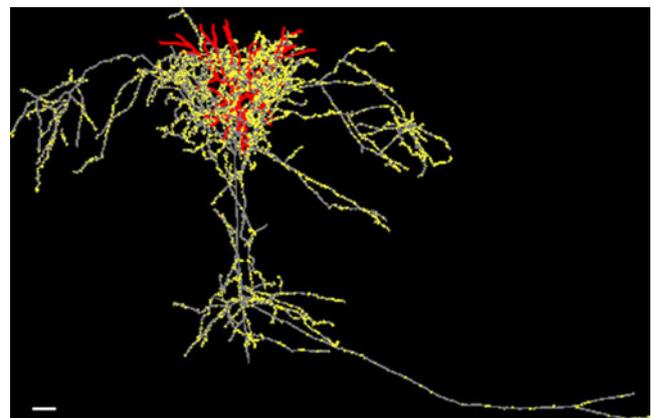


Fig. 4 Example pyramidal cell of layer 3 of cat visual cortex, recorded *in vivo* and injected intracellularly with horseradish peroxidase. Red dendrites, grey axon, yellow synaptic boutons. Scale bar: 100 microns (Anderson, Douglas & Martin, unpublished image)

The exercise told us three things: firstly, it provided a quantitative standard for assessing whether any connection was random or specific. Violations of Peters' Rule would indicate that the connection was specific, as in the case of the chandelier cells. Secondly, it provided an estimate of the relative weight of connections between the different neuronal types in the different layers. Two very interesting numbers are that the thalamic input provides less than 1% of the excitatory synapses of area 17, while the layer 3 pyramidal cells provide almost 25%. These two numbers point strongly to the necessity of computational modeling to understand these extraordinary differences in synaptic weighting. Thirdly, the exercise provided us with an estimate of what is still missing. Here the results are also very interesting. The synaptic complement of layers 3 and 5 can be largely explained by inputs from the axon collaterals of other neurons in area 17. However, in layer 4 only about 70% of the synapses can be accounted for by known local sources, and in layers 1 and 6, only a minority of inputs can be accounted for by the local circuits of area 17 neurons. The problem of the 'dark matter' of cortex is most severe in layer 1, where synapses of unknown parenthood formed over 90% of the excitatory synapses. Of course, one side of the equation is clear, for the targets of these inputs to layer 1 are mainly the apical dendrites of the local pyramidal neurons. This unexpected finding points up the fact that even in the best studied cortical area we cannot yet assign large populations of synapses to their source.

The fourth important lesson from this exercise concerns the overall architecture of the cortical circuit. Our cortical circuit (Fig. 4) does not resemble one of the standard artificial neural network models having either feed forward or recurrent connectivity. Instead these two characteristics are intimately mixed, and so require a new interpretation of the operation of cortical circuits. This is a stark reminder that if we are genuinely to achieve a comprehensive functional understanding, then we have to go beyond the local circuits to a quantitative analysis of all the long distance inputs to the local circuit and make a comprehensive identification of the different targets to which the circuit projects. This is beyond the realms of even the most optimistic promises of high-throughput EM methods.

Theory Guided Structure

Horace Barlow (1980) offered this hypothesis about what the neocortex 'does': 'The hypothesis postulates that the whole cortex is engaged in one type of task, one that is requires in complex judgments as well as simple perception'. His idea ties in well with our notion of canonical circuits for neocortex, but of course such ideas are not confined to neocortex for all we need are repeated

structures, such as those found in the cerebellar cortex, where the repeated patterns are extreme, but where interestingly, many fMRI studies show the presence of localised spots of relatively increased activity. Does this localized activity just reflect the mapping of climbing fibers? In neocortex we are discovering that there are rules of mapping and connection that seem to scale between cortical areas in single species (e.g. In the monkey cortex, the dendritic trees of pyramidal cells in occipital cortex are small and their basal dendrites are simply branched, while in prefrontal cortex the pyramidal cells are amongst the largest found in primate (see Elston et al. 2001; Elston 2002).

Quantitative studies of the interareal connections in the primate by Henry Kennedy and colleagues (Barone et al. 2000; Markov et al. 2010) show that the convergence in prefrontal areas such as area 46 is far greater than that for occipital areas, yet the rules of feedforward and feedback connectivity described by Kathy Rockland and Deepak Pandya still holds (Rockland and Pandya 1979; Rockland 1997). These observations seem to indicate that the neocortex is a single sheet of neurons connected by rules that express graded changes across the entire cortical surface, rather than of piecemeal rules across one modality or one cytoarchitectonic area. One could take this argument further and say that only by developing theory on this higher level can one begin to understand the relevance of the different components of the circuits. For example, we generally assume that all neurons in a given circuit are involved in a given 'computation', and that there are none just doing housekeeping, whether it be maintaining the core processors in balance or in range, or providing a supporting role in span calibration, or just waiting in silence. This assumption may be wrong.

Canonical Circuits—What Do They Predict?

Predictive structural analysis means that we have a theory about what circuits we should find in a given region of the brain. In the case of neocortex, the notion of canonical circuits involving both local and long distance pathways has provided us with a clear target for experimental investigations. In the case of the local circuits of cat area 17, our Mark I canonical circuit of 1989 (Fig. 5, top left) was not based simply on anatomy. Instead we adopted an engineering-like approach in which we applied pulse stimuli to the thalamocortical afferents and measured the cortical response over a few hundred milliseconds by intracellular recordings from neurons sampled throughout the cortical layers. In addition, we dissected subcircuits involved in the response using ionophoretically applied neurotransmitter agonists and antagonists. By this means,

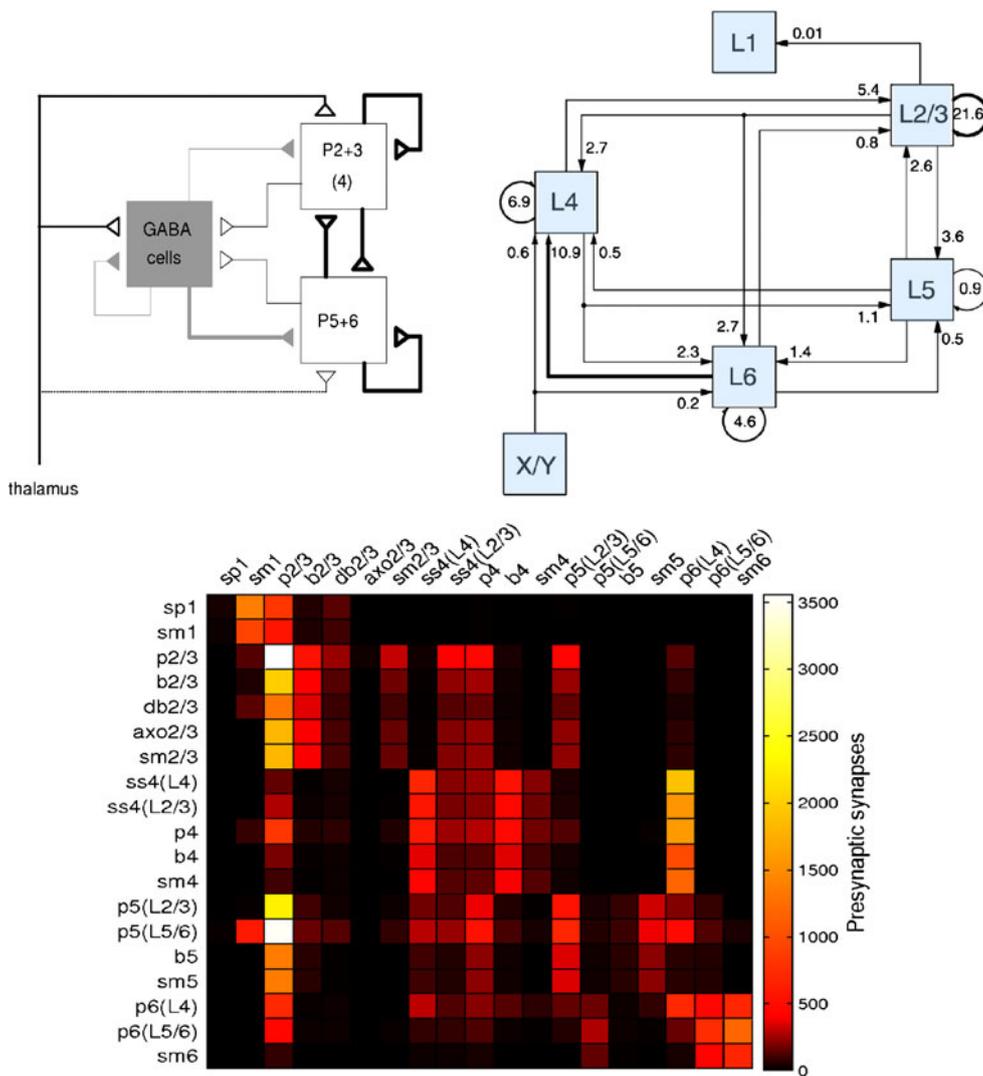


Fig. 5 Top Left. Canonical microcircuit for the neocortex proposed in 1989 to explain the intracellular responses observed in cortical neurons following electrical pulse activation of thalamic afferents. The circuit components underlying these responses were dissected by ionophoretic application of neurotransmitter agonists and antagonists. The circuit is composed of three populations of neurons, which interact with one another: one population is inhibitory (smooth cells, filled synapses), and two are excitatory (open synapses), representing superficial (P2+3) and deep (P5+6) layer pyramidal neurones. The layer 4 spiny stellate cells (4) are incorporated with the superficial group of pyramidal cells. Some neurons within each population receive excitatory input from the thalamus. Continuous versus dashed lines indicate that thalamic drive to the superficial group is stronger. The inhibitory inputs activate both GABAA and GABAB receptors on pyramidal cells. The thick continuous line connecting smooth cells to P5+6 indicate that the inhibitory input to the deep pyramidal population is relatively greater than that to the superficial population. However, the increased inhibition is due to enhanced GABAA drive only. The GABAB inputs to P5+6 are similar to those applied to P2+

3. Top Right. Quantitative map of anatomical connections between the major excitatory and inhibitory neuron types in area 17 of the cat, including the X-type and Y-type afferents to area 17 from the dorsal LGN. Each arrow is labeled with a number indicating the proportion of all the synapses that are formed between excitatory neurons (B), from excitatory onto inhibitory neurons (C), from inhibitory onto excitatory neurons (D), and between inhibitory neurons (D). For details see Binzegger et al. (2004). Bottom: Connection matrix for cat area 17. Colour codes for the number of synapses that all presynaptic neurons of a cell type (abscissae) form with an individual postsynaptic neuron of a given cell type (ordinate). Color bar is indicated to the right. Cell type abbreviations are as follows: ‘b2/3’, ‘b4’, ‘b5’ basket cells in layer 2/3, 4 and 5; ‘db2/3’ double bouquet cell in layer 2/3; ‘p2/3’, ‘p4’, ‘p5’, ‘p6’ pyramidal cells in layer 2/3, 4, 5 and 6. ‘ss4’ spiny stellate cells in layer 4. Spiny stellate cells and pyramidal cells in layer 5 and 6 were further distinguished by the preferred layer of the axonal innervation (‘ss4(L4)’ (not shown), ‘ss4(L2/3)’, ‘p5(L2/3)’, ‘p5(L5/6)’, ‘p6(L4)’ and ‘p6(L5/6)’)

we were able to estimate the spatial (in cortical depth) and temporal dynamics of spike events and transmembrane potentials over the important time window during which cortex resolves perception. Of course, electrical pulse

stimuli are not the same as sensory stimuli, but they do offer a probe of the fundamental dynamics of the cortical circuitry. Using this systems identification approach we were able to identify a minimal circuit capable of

generating those observed dynamics, and verify its behaviour by simulation.

This ‘canonical microcircuit for neocortex’ (Douglas et al. 1989) has been remarkably successful as a stimulus and guide to our experimental investigations. In particular, it led to our exploration of the detailed anatomical circuit of neocortex, now expressed in the ‘Binzegger Circuit’, which was the first attempt at a complete quantitative description of the average connectivity of any cortical area (Binzegger et al. 2004; Fig. 5, top right). These two circuits, the dynamic canonical circuit, and the static Binzegger circuit, offer constraints on the architecture and organisation of computation in cortex. However these are not separate paths—the two routes of exploration merge in our analyses of the constraints on dynamical stability of the excitation and inhibition in the circuit and how the topology of the circuit can be configured dynamically through average inhibition (Binzegger et al. 2009; Fig. 5, bottom). The detailed synaptic mapping has also been essential for constraining our detailed biophysical models of subcircuits of cat cortex (Banitt et al. 2007). These realisations of circuits highlight the importance of accurate and quantitative structural data to constrain network models. The Binzegger circuit is also proving significant as a generic computational circuit for theoreticians who are interested in large-scale simulations (Izhikevich and Edelman 2008; Ananthanarayanan et al. 2009; Haeusler et al. 2009). These large-scale simulations show clearly the complex dynamics arising from these circuits, but they provide few insights into the reasons for the circuits we have discovered or into the functional properties that have been richly documented in cat visual cortex.

The Binzegger circuit of cat visual cortex (Binzegger et al. 2004; Fig. 5) now has progeny that extend even to prefrontal cortical areas of primates. For example, we have shown that they can be virtually transplanted to the frontal eye field area of the monkey where, implemented as a spiking neuron model with only slight modifications in the circuit, it captures the dynamics of neuronal discharges recorded in awake behaving monkeys who are performing saccadic eye movements (Heinzle et al. 2007). From this it was a natural step to show that the model could also be successfully applied to the very precise control of saccadic eye movements deployed during human reading (Heinzle et al. 2010). The potential of this detailed model for structure is great, for it tells us where to look and what kind and number of connections should exist and what their value (excitatory, inhibitory) should be. One observation, that the neurons of the superficial layers are strongly recurrently interconnected, provides a more general mechanism for generating and controlling the gain required for generic processing mechanisms such as signal selection and restoration (Douglas and Martin 2007), and the construction of state-

machines (Rutishauser and Douglas 2009). These examples provide important steps in bringing the observed canonical circuit closer to computationally relevant processes. .

Whither?

The introduction of the black reaction by Golgi was the most revolutionary technique introduced into structural neuroscience before the advent of the electron microscope. Although its practitioners have ever found it capricious, its quality of randomness was a huge advantage for the early anatomists, who had barely any inkling of what lay within the nervous systems of the various species they examined. With the Golgi stain and a ‘beginner’s mind’ they described whatever they saw down their microscopes and discovered the variety of forms of neurons and glia and intracellular structures that are so familiar to us today. The new technology demanded new concepts—and they arrived. In particular by applying the Law of Dynamic Polarisation, Cajal was able to link the independent neurons he saw into circuits and his strategy was general and could be applied in any region of the nervous system in any species. Now that we have our own new revolution of molecular technologies to apply to the neural circuits, claims are frequently heard that the precise targeting offered by molecular genetics will transform our understanding in a way that scattergun techniques like the Golgi stain never could. Our gentle rejoinder is that of course new technology always brings with it new views, but we should not feel at all superior to our masters of the past. While the Golgi technique is now seen as a relic, arguably our field would have been held back if it had been possible to target the Golgi stain more precisely. One simply has to leaf through the two volumes of Cajal’s ‘Histology of the Nervous System of Man and Vertebrates’ to discover how narrow our modern vision has become in what we now deem as relevant areas of research in the nervous system—and also what a mountain we have to climb if we aspire to match the number of new discoveries that Cajal made with his application of Golgi’s capricious stain. Indeed, even with all the panoply of modern techniques the promise of ‘precise’ targeting has still to be realised. None of these new techniques have yet achieved precision beyond broad divisions of neurons—and finding markers for those few has been a trial-and-error process, because we do not know much about what actually differentiates one neuron from another. It is clear however, that there are many exciting technical developments in physiological tools that, coupled to solutions for structure such as that offered by the Diadem Grand Challenge, will surely change the game.

In our view, the real barrier to progress lies in our lack of any equivalent of a Neuron Doctrine or a Law of Dynamic

Polarisation. While we are well on the road to building automated techniques to reconstruct neural circuits and to having optogenetic tools that will allow us to switch on or off selected subtypes of neurons, we clearly lack a conceptual framework to interpret the results. Conceptual progress, history shows, does not come from the semi-industrial, high throughput methods that are currently in vogue, but through the traditional cottage industry of deep thought and clear insight. But to discover how a brain built of neurons actually makes these conceptual leaps—that is the ultimate Grand Challenge.

Information Sharing Statement

Specifications of the ‘Binzegger circuit’ (see Fig. 4) are available upon application to the authors.

Acknowledgements We thank our colleagues in the INI for unrelenting discussions and unsparing debates. This review formed the basis of a lecture given by KACM at the Diadem Grand Challenge final at HHMI Janelia Farm. Supported by EU SECO grant 216593 to both authors and Human Frontiers Science Program grant RG 0123/2000-B to KACM.

References

- Ananthanarayanan, R., Esser, S. K., Simon, H. D., & Modha, D. S. (2009). The cat is out of the bag: Cortical Simulations with 10^9 neurons and 10^{13} synapses. *Supercomputing 09: Proceedings of the ACM/IEEE SC2009 Conference on High Performance Networking and Computing*, Nov 14–20, 2009, Portland, OR.
- Banitt, Y., Martin, K. A., & Segev, I. (2007). A biologically realistic model of contrast invariant orientation tuning by thalamocortical synaptic depression. *The Journal of Neuroscience*, *27*, 10230–10239.
- Barlow, H. B. (1972). A neuron doctrine for perceptual psychology. *Perception*, *1*, 371–394.
- Barlow, H. B. (1977). Performance, perception, dark-light, and gain boxes. In E. Pöppel, R. Held, & J. E. Dowling (Eds.), *Neuronal mechanisms in visual perception* (pp. 394–397). Cambridge: Neurosciences Research Program Bulletin 15, MIT Press.
- Barlow, H. B. (1980). Cortical function: A tentative theory and preliminary test. In D. McFadden (Ed.), *Neural mechanisms in behaviour* (pp. 143–171). New York: Springer.
- Barone, P., Batardiere, A., Knoblauch, K., & Kennedy, H. (2000). Laminar distribution of neurons in extrastriate areas projecting to visual areas V1 and V4 correlates with the hierarchical rank and indicates the operation of a distance rule. *The Journal of Neuroscience*, *20*, 3263–3281.
- Bas, E., & Erdogmus, D. (2011). Principal curves as skeletons of tubular objects: locally characterizing the structures of axons. *Neuroinformatics*, THIS ISSUE.
- Binzegger, T., Douglas, R. J., & Martin, K. A. C. (2004). A quantitative map of the circuit of cat primary visual cortex. *The Journal of Neuroscience*, *24*, 8441–8453.
- Binzegger, T., Douglas, R. J., & Martin, K. A. C. (2009). Topology and dynamics of the canonical circuit of cat V1. *Neural Networks*, *22*, 1071–1078.
- Boyden, E. S., Zhang, F., Bamberg, E., Nagel, G., & Deisseroth, K. (2005). Millisecond-timescale, genetically targeted optical control of neural activity. *Nature Neuroscience*, *8*, 1263–1268.
- Braitenberg, V., & Lauria, F. (1960). Toward a mathematical description of the grey substance of nervous systems. *Nuovo Cimento*, *18*, 1135–1151.
- Braitenberg, V., & Schüz, A. (1991). Peters’ rule and White’s exceptions. In: *Anatomy of the cortex* (pp. 109–112). Berlin: Springer.
- Cajal, S. R. (1913–1914). Estudios Sobre la Degeneración y Regeneración del Sistema Nervioso, Moya. ‘Cajal’s Degeneration and Regeneration of the Nervous System’, OUP New York. Reprinted and edited with additional translations by J DeFelipe, J. and EG Jones, (1991).
- Cajal, S. R. (1921). Histology of the visual cortex of the cat. *Archivos de Neurobiologia*, *2*, 338–362. In J. DeFelipe & E. G. Jones (Eds.), *Cajal of the cerebral cortex* (pp. 495–523). New York: OUP.
- Cajal, S. R. (1937). Recollections of my life. Translated by EH Craigie, J Cano. 1989 Philadelphia PA Am Philos. Soc.
- Callaway, E. M. (2008). Transneuronal circuit tracing with neurotropic viruses. *Current Opinion in Neurobiology*, *18*, 617–623.
- Chothani, P., Mehta, V., & Stepanyants, A. (2011). Automated tracing of neurites from light microscopy stacks of images. *Neuroinformatics*, THIS ISSUE.
- Cohen, Y. E., & Andersen, R. A. (2002). A common reference frame for movement plans in the posterior parietal cortex. *Nature Reviews. Neuroscience*, *3*, 553–562.
- Crick, F. H. C. In V. S. Ramachandran (1985). The neurobiology of perception. *Perception* *14*, 97–103.
- da Costa, N. M., & Martin, K. A. C. (2009). Selective targeting of the dendrites of corticothalamic cells by thalamic afferents in area 17 of the cat. *The Journal of Neuroscience*, *29*, 13919–13928.
- da Costa, N. M., & Martin, K. A. C. (2010). Whose cortical column would that be? *Frontiers Neuroanat*, *4*, 16.
- Douglas, R. J., & Martin, K. A. C. (1991). Opening the grey box. *TINS*, *14*, 286–293.
- Douglas, R. J., & Martin, K. A. C. (2007). Recurrent neuronal circuits in the neocortex. *Current Biology*, *17*, R496–R500.
- Douglas, R. J., Martin, K. A. C., & Whitteridge, D. (1989). A canonical microcircuit for neocortex. *Neural Computation*, *1*, 480–488.
- Eccles, J. C., Ito, M., & Szentágothai, J. (1967). *The Cerebellum as a Neuronal Machine*. New York: Springer. 343 pp.
- Elston, G. N. (2002). Cortical heterogeneity: implications for visual processing and polysensory integration. *Journal of Neurocytology*, *31*, 317–335.
- Elston, G. N., Benavides-Piccione, R., & DeFelipe, J. (2001). The pyramidal cell in cognition: a comparative study in human and monkey. *Journal of Neuroscience*, *21*, RC163, 1–5.
- Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, *1*, 1–47.
- Ferster, D., & LeVay, S. (1978). The axonal arborizations of lateral geniculate neurons in the striate cortex of the cat. *The Journal of Comparative Neurology*, *182*, 923–944.
- Fulton, J. F. (1949). *Physiology of the nervous system* (pp. 288–312). New York: OUP.
- Gilbert, C.D. (1983). Microcircuitry of the visual cortex. *Annu Rev Neurosci*. *6*, 217–247.
- Gilbert, C. D., & Wiesel, T. N. (1979). Morphology and intracortical projections of functionally characterized neurons in the cat visual cortex. *Nature*, *280*, 120–125.
- Gilbert, C. D., & Wiesel, T. N. (1981). Laminar specialization and intracortical connections in cat primary visual cortex. In F. O. Schmitt, F. G. Worden, G. Adelman, & S. G. Dennis (Eds.), *The organization of the cerebral cortex* (pp. 163–191). Cambridge: MIT.

- Haeusler, S., Schuch, K., & Maass, W. (2009). Motif distribution, dynamical properties, and computational performance of two data-based cortical microcircuit templates. *Journal of Physiology - Paris*, *103*, 73–87.
- Hayworth, K. J., Kasthuri, N., Schalek, R., & Lichtman, J. W. (2006). Automating the collection of ultrathin serial sections for large volume TEM reconstructions. *Microscopy and Microanalysis*, *12*, 86–87. CUP.
- Heeger, D. J., Simoncelli, E. P., & Movshon, J. A. (1996). Computational models of cortical visual processing. *Proceedings of the National Academy of Sciences of the United States of America*, *93*, 623–627.
- Heinzle, J., Hepp, K., & Martin, K. A. C. (2007). A microcircuit model of the frontal eye fields. *The Journal of Neuroscience*, *27*, 9341–9353.
- Heinzle, J., Hepp, K., & Martin, K. A. C. (2010). A biologically realistic cortical model of eye movement control in reading. *Psychological Review*, *117*, 808–830.
- Helmstaedter, M., de Kock, C. P., Feldmeyer, D., Bruno, R. M., & Sakmann, B. (2007). Reconstruction of an average cortical column in silico. *Brain Research Reviews*, *55*, 193–203.
- Helmstaedter, M., Briggman, K. L., & Denk, W. (2008). 3D structural imaging of the brain with photons and electrons. *Current Opinion in Neurobiology*, *18*, 633–641.
- Hilgetag, C. C., O'Neill, M. A., & Young, M. P. (1996). Indeterminate organization of the visual system. *Science*, *271*, 776–777.
- Hines, M. L., Markram, H., & Schürmann, F. (2008). Fully implicit parallel stimulation of single neurons. *Journal of Computational Neuroscience*, *25*, 439–448.
- Horton, J. C., & Adams, D. L. (2005). The cortical column: a structure without a function. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *360*, 837–862.
- Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.*, *160*, 106–154.
- Izhikevich, E. M., & Edelman, G. M. (2008). Large-scale model of mammalian thalamocortical systems. *PNAS*, *105*, 3593–3598.
- Kalisman, N., Silberberg, G., & Markram, H. (2005). The neocortical microcircuit as a tabula rasa. *Proceedings of the National Academy of Sciences of the United States of America*, *102*, 880–885.
- Kasthuri, N., & Lichtman, J. W. (2007). The rise of the 'projectome'. *Nature Methods*, *4*, 307–308.
- Kasthuri, N., & Lichtman, J. W. (2010). Neurocartography. *Neuropsychopharmacology*, *35*, 342–343.
- Kouh, M., & Poggio, T. (2008). A canonical neural circuit for cortical nonlinear operations. *Neural Computation*, *20*, 1427–1451.
- Lichtman, J. W., Livet, J., & Sanes, J. R. (2008). A technicolour approach to the connectome. *Nature Reviews. Neuroscience*, *9*, 417–422.
- Lorente de Nó, R. (1949). "Cerebral cortex: architecture, intracortical connections, motor projections," in *Physiology of the Nervous System*, ed. J. F. Fulton (New York: Oxford University Press), 288–312.
- Lorente de Nó, R. (1992). The cerebral cortex of the mouse, a first contribution—the acoustic cortex. Somatosensory and Motor Research. *9*, 3–36, translated by A. Fairen, J. Regidor, L. Kruger, from the original, La corteza cerebral raton, primera contribucion- la corteza acustica, Trabajos del Laboratorio de Investigaciones Biologicas de la Universidad de Madrid, *20* (1922) 41–78.
- Markov, N. T., Misery, P., Falchier, A., Lamy, C., Vezoli, J., Quilodran, R., et al. (2010). Weight consistency specifies regularities of Macaque Cortical Networks. *Cerebral Cortex*, (in press).
- Markram, H. (2006). The blue brain project. *Nature Reviews. Neuroscience*, *7*, 153–160.
- Markram, H. (2008). Fixing the location and dimensions of functional neocortical columns. *HFSP Journal*, *2*, 132–135.
- Markram, H. (2010). http://www.ted.com/talks/henry_markram_super_computing_the_brain_s_secrets.html.
- Meyer, H. S., Wimmer, V. C., Hemberger, M., Bruno, R. M., de Kock, C. P., Frick, A., et al. (2010). Cell type-specific thalamic innervation in a column of rat vibrissal cortex. *Cerebral Cortex*, *20*, 2287–2303.
- Meyer, H. S., Wimmer, V. C., Oberlaender, M., de Kock, C. P., Sakmann, B., & Helmstaedter, M. (2010). Number and laminar distribution of neurons in a thalamocortical projection column of rat vibrissal cortex. *Cerebral Cortex*, *20*, 2277–2286.
- Mormann, F., Kornblith, S., Quiroga, R. Q., Kraskov, A., Cerf, M., Fried, I., et al. (2008). Latency and selectivity of single neurons indicate hierarchical processing in the human medial temporal lobe. *The Journal of Neuroscience*, *28*, 8865–8872.
- Mountcastle, V. B. (1997). The columnar organization of the neocortex. *Brain*, *120*(Pt 4), 701–722.
- Narayananwamy, A., Wang, Y., & Roysam, B. (2011). A preprocessing pipeline to enhance 3-D images of neuronal arbors. *Neuroinformatics*, THIS ISSUE.
- O'Connor, D. H., Huber, D., & Svoboda, K. (2009). Reverse engineering the mouse brain. *Nature*, *461*, 923–929.
- Peters, A., & Feldman, M. L. (1976). The projection of the lateral geniculate nucleus to area 17 of the rat cerebral cortex. I. General description. *Journal of Neurocytology*, *5*, 63–84.
- Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C., & Fried, I. (2005). Invariant visual representation by single neurons in the human brain. *Nature*, *435*, 1102–1107.
- Quiroga, R. Q., Kreiman, G., Koch, C., & Fried, I. (2008). Sparse but not 'grandmother-cell' coding in the medial temporal lobe. *Trends in Cognitive Sciences*, *12*, 87–91.
- Rockland, K. S. (1997). Elements of cortical architecture. Hierarchy revisited. In K. S. Rockland, J. H. Kaas, & A. Peters (Eds.), *Cerebral cortex, vol. 12* (pp. 243–293). New York: Plenum.
- Rockland, K. S. (2010). Five points on columns. *Frontiers in Neuroanatomy*, *4*, 1–10.
- Rockland, K. S., & Pandya, D. N. (1979). Laminar origins and terminations of cortical connections to the occipital lobe in the rhesus monkey. *Brain Research*, *179*, 3–20.
- Rushton, W. A. H. (1977). Some memories of visual research in the past 50 years. In *The pursuit of nature* (pp. 85–104). Cambridge: CUP.
- Rutishauser, U., & Douglas, R. J. (2009). State dependent computation using coupled recurrent networks. *Neural Computation*, *21*, 478–509.
- Seung, H. S. (2009). Reading the book of memory: sparse sampling versus dense mapping of connectomes. *Neuron*, *62*, 17–29.
- Seung, H. S. (2010). http://www.ted.com/talks/sebastian_seung.html.
- Shepherd, G. M. (1991). *Foundations of the neuron doctrine*. New York: Oxford University Press.
- Sherrington, C. (1940). *Man on his Nature*. Cambridge. CUP. p. 277.
- Sholl, D. (1956). *The organisation of the cerebral cortex*. London: Methuen.
- Sullivan, L. (1896). The tall office building artistically considered. *Lippincotts Magazine* (March). In I. Athey (Ed.), *Kindergarten chats and other writings*. New York: Dover Publications. 1979.
- Tao-Cheng, J. H., Gallant, P. E., Brightman, M. W., Dosemeci, A., & Reese, T. S. (2007). Structural changes at synapses after delayed perfusion fixation in different regions of the mouse brain. *The Journal of Comparative Neurology*, *501*, 731–740.

- Thompson, D. W. (1971). *On growth and form. Abridged Edition*. Cambridge: CUP.
- Tsodyks, M., Kenet, T., Grinvald, A., & Arieli, A. (1999). Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science*, *286*, 1943–1946.
- Turetken, E., Gonzalez, G., Blum, C., Fua, P. (2011). Automated reconstruction of dendritic and axonal trees by global optimization with geometric priors. *Neuroinformatics*, THIS ISSUE.
- Wang, Y., Narayanaswamy, A., Tsai, C., Roysam, B. (2011). Generally Applicable 3D >Neuron Tracing Approach and System based on Open-Curve Snake. *Neuroinformatics*, THIS ISSUE.
- Wiesel, T. N., & Gilbert, C. D. (1983). The Sharpey-Schafer lecture. Morphological basis of visual cortical function. *Quarterly Journal of Experimental Physiology*, *68*, 525–543.
- Wimmer, V. C., Bruno, R. M., de Kock, C. P., Kuner, T., & Sakmann, B. (2010). Dimensions of a projection column and architecture of VPM and POm axons in rat vibrissal cortex. *Cerebral Cortex*, *20*, 2265–2276.
- Zhao, T., Xie, J., Amat, F., Clack, N., Ahammad, P., Peng, H., Long, F., Myers, E. (2011). Automated reconstruction of neuronal morphology based on local geometrical and global structural models. *Neuroinformatics*, THIS ISSUE.