

Mapping the Matrix: The Ways of Neocortex

Rodney J. Douglas^{1,*} and Kevan A.C. Martin^{1,*}

¹Institute of Neuroinformatics, UZH/ETH, Winterthurerstrasse 190, 8057 Zurich, Switzerland

*Correspondence: rjd@ini.phys.ethz.ch (R.J.D.), kevan@ini.phys.ethz.ch (K.A.C.M.)

DOI 10.1016/j.neuron.2007.10.017

While we know that the neocortex occupies 85% of our brains and that its circuits allow an enormous flexibility and repertoire of behavior (not to mention unexplained phenomena like consciousness), a century after Cajal we have very little knowledge of the details of the cortical circuits or their mode of function. One simplifying hypothesis that has existed since Cajal is that the neocortex consists of repeated copies of the same fundamental circuit. However, finding that fundamental circuit has proved elusive, although partial drafts of a “canonical circuit” appear in many different guises of structure and function. Here, we review some critical stages in the history of this quest. In doing so, we consider the style of cortical computation in relation to the neuronal machinery that supports it. We conclude that the structure and function of cortex honors two major computational principles: “just-enough” and “just-in-time.”

Mappa Mundi

Maps are comforting. They reveal to us the fixed points of the known world and alert us to the regions that are “terra incognita.” However, maps themselves also map the changes in our perception of what is the “known world”—and these reveal our perceptions to be unstable. One famous example is the island of Buss (Figure 1), which was first “discovered” in 1578 somewhere between Ireland and Frisland and appeared on nautical charts from then on until it finally sank from consciousness after it last appeared on a chart in 1856. The island of Madya was the longest survivor of these phantoms. It first appeared on maps in about 1400, positioned in the north Atlantic to the southwest of Ireland. Over the centuries it moved more westward, so that by 1566 it was located near Newfoundland, and then took a turn south, and was last seen on a Rand McNally map of 1906 at the level of the West Indies. Claude Levi-Strauss (in contrast to William of Occam: “No more things should be presumed to exist than are absolutely necessary”) argued that every culture has a need for certain concepts and expressions to absorb any excess of existence that has not yet had a word coined for it. James Hamilton Paterson (1993) suggests that these phantom islands operated as Levi-Strauss’s “floating signifiers” to provide comfort and points of reference within the void of ignorance—the terra incognita.

The neocortex is one of the most elaborate maps we have. Not only does it contain many different areas, but these areas also contain within themselves multiple maps, which may reflect directly the sensory periphery or may appear as more abstract “cognitive maps.” The neocortex has its own floating signifiers, with words like “column,” “module,” “neural representation,” “cortical code,” and “consciousness,” which have been coined to absorb the enormous functional and structural excess of existence that is evident in every material record of the brain and, most particularly, in the cortical circuits about whose mode of organization and operation we are still greatly

ignorant. One fundamental question is whether the neocortex is a unitary structure with a grammar and a logic of construction and operation that can be understood in terms of the physical circuits and their physiology, or whether it is a collective of very many separate modules with their own specialist “trick” circuitry?

The Languages of Neocortex

To begin at the beginning: like the syntax of human languages, the structure of the neocortex appears equally complex in all land mammals. Just as there is no simple or prototype version of a human language in existence, a simple or primitive form of neocortex does not exist. Yet, it is as evident that, like different languages, the neocortex consists of different areas as defined by histological or with physiological methods. But just as with languages, we will claim here, the neocortical areas are also essentially the same and, like languages, can be translated, one into the other. Thus, in understanding one area, we can expect to understand another. It is this sameness that we have called a “canonical” property of cortical circuits (Douglas et al., 1989). The immediate challenge is the question, what defines “neocortex”? The usual answer is structural: that, unlike allocortex, which has fewer layers and is phylogenetically older, the neocortex possesses six layers. The number of layers would, of course, seem a rather fragile means of defining a structure that varies over five orders of magnitude in volume from shrew to whale, that supports the processing of input from an unlikely range of sensory systems allowing detection of electromagnetic radiation, vibration, temperature, sound, and chemicals, and that then provides output to an equally unlikely range of motor structures, allowing an animal to fly, swim, walk, jump, and run. In fact, the “six-layered” neocortex is something of a unicorn, for the number of layers that can be distinguished varies greatly between areas and the histological stains used to reveal the layers. Yet, somehow, neocortex is so instantly distinguishable from other laminated structures,



Figure 1. The Phantom Island of Buss
 First “discovered” in 1578, it disappeared from the Nautical charts after 1856. http://eaudrey.com/myth/Places/buss_island.htm.

such as the hippocampus or superior colliculus, that early anatomists referred to it as “isocortex.”

Although it is now clear that language comprehension and production involves much more of the brain than just the well-known regions first discovered by Broca and Wernicke, their 19th century idea of a compartmentalization of specific functions has reappeared in modern times, most prominently in evolutionary biology. The best known claim is captured by the “Swiss army knife” metaphor for the functional organization of cortex (Barkow et al., 1992). In this view, the brain has evolved a series of special-purpose modules, which, like the Swiss army knife, consist of individual components that have a specific function and are not designed to work together like the components of a machine. For humans, the language module is the most obvious of these special-purpose modules, but strong claims have been made that such specialized modules are the means by which the neocortex is organized and works (Fodor, 1983; Zeki, 1993). Implicit in this is the notion that the neural pathways in the brain subserving these areas are highly segregated and that there is no general-purpose architecture that carries out the neocortical part of the computations.

Written on the Brain (in Indian Ink)

The era of microcircuit analysis was launched by Camillo Golgi’s discovery of “la reazione nera,” which allowed individual neurons to be visualized, and by Santiago Ramon y Cajal’s law of dynamic polarization, which provided the

critical algorithm for identifying the input and output regions of individual neurons. Put together, these two advances made it possible for the first time to show the probable route of impulses from input to output for a given structure. As he recorded in his autobiography, the extraordinary claim that Cajal made was that even the highest center of the brain, the neocortex, was built of stereotyped circuits like those he had discovered in the retina, cerebellum, hippocampus, spinal cord, and other parts of the central nervous system (Cajal, 1937). Despite intense efforts on his part, however, he was unable to define the basic cortical circuit, but until the end of his life he nevertheless remained convinced that it existed.

When Cajal applied Golgi’s stain to neonatal brain, he was able to map, mostly correctly, significant circuits in the spinal cord, retina, and visual pathways, cerebellum, hippocampus, olfactory bulb, auditory nuclei, and others. From this he developed the notion of the “neural avalanche,” which was essentially the inverse of Sherrington’s “final common path.” It stated that the number of neurons involved in conducting impulses from a sensory receptor increases progressively from the periphery to the cortex (Cajal, 1937). De Kock et al. (2007) have calculated the neural avalanche in the rat barrel cortex, and they estimate that a single whisker deflection generates about 4000 impulses in the cortex. This avalanche grows further through the associated cortical areas, before it is funneled down the final common path to the motoneuron, but even with his great skills of preparation, observation, and imagination, Cajal was unable to trace the route from input to the cortex to its output. However, his efforts were not without reward, for he provided a comprehensive description of the different cell types that inhabit the neocortex of different animals and incorporated the earlier descriptions of cortical cell types of Retzius, Meynert, Betz, and others.

Lorente de Nó (1949) pursued Cajal’s dream, also with Golgi’s stain, and suggested that the functional unit of cortex consisted of a specific thalamocortical fiber and a cylindrical group of cells surrounding the fiber, some of which formed synapses with the thalamocortical fiber. With succeeding generations, however, this confidence in a basic circuit became less secure, and there was even a return, in the 1930s, to the idea that the neocortex was an equipotential network (Lashley, 1930), an idea demolished by Sperry (Sperry, 1947; Sperry et al., 1955), or that the connections between cortical neurons were not at all specific, but perhaps statistical (Sholl, 1956) or semi-random (Szentágothai, 1978) or even random (Braitenberg and Schüz, 1991). Thus, proposals that the local circuit of barrel cortex begins its life as a “tabula rasa” to be written on by experience (Jeanmonod et al., 1981; Le Be and Markram, 2006; Kalisman et al., 2005) are simply a continuance of a surprisingly long-lived hypothesis that the cortex really wants to be a randomly connected neural network.

In the face of such enormous numbers of possible circuits that could potentially arise from such random neural networks, to pursue the concept that all of neocortex has

a uniform basic structure and performs some basic uniform operations would seem to be setting oneself up for yet another instance of Thomas Huxley's great tragedy of Science—"the slaying of a beautiful hypothesis by an ugly fact." Yet, the decade that began in a trough with Karl Lashley's despairing account of his search for the engram (Lashley, 1950) ended on a much more positive peak for the neocortex, with seminal studies of the physiology of sensory cortex (Mountcastle, 1957; Hubel and Wiesel, 1959).

Architects of the New Cortex

The long hoped for evidence that neocortex had a specific architecture did not come from anatomists, but from physiologists working *in vivo*, who provided the major new insights into cortical organization. Mountcastle (1957) recorded from the somatosensory cortex of cats and monkeys and found that neurons with common functional properties lay in a radial column of cells, extending from white matter to cortical surface. With this discovery, the anatomists were once again brought into play. Powell, working at Mountcastle's side, suggested that the vertical palisades of cells he saw in stained sections of sensory cortex were the elementary units that formed of the functional columns (Powell and Mountcastle, 1959). This concept of a functional column was not lost on their neighbors at John's Hopkins, who were just then plotting their first receptive fields of visual cortical neurons on bed sheets hung from the ceiling.

Hubel and Wiesel's ice-cube diagram (Hubel and Wiesel, 1972, 1977) summarized their basic findings about the functional architecture of area 17. What had first seemed to them irregularly shaped "columns," sometimes ordered, sometimes not, sometimes continuous, sometimes not (Hubel and Wiesel, 1963, 1968), were now schematized as parallel slabs, with the thin orientation slabs cutting the coarser ocular dominance columns at right angles. Since orientation selectivity is a property of single cells, it is inevitably discretized, but trivially so, for in all other respects the map of orientation space appears to be continuous. In the case of monkey visual cortex, Hubel and Wiesel (1974b) found that movements of the microelectrode as small as 25–50 μm could produce just-detectable shifts in orientation preference (10°) in monkey V1. From this they reckoned that the orientation slabs were not discrete, or if they were, they had a width of less than 25 μm , i.e., the diameter of a large cell soma. To Hubel and Wiesel this indicated that there was no correlation between the fine-grain of the physiological slabs and the coarser grain offered by the orders of magnitudes larger dimensions of the soma, dendrite, and axon. The recent two-photon imaging of the neurons that form the orientation map in the cat (Ohki et al., 2005, 2006) confirms in 2D the precision of the progression in orientation preference seen with the microelectrode and confronts us again with the puzzle of how this comes about, given what we now know about the underlying anatomy.

The anatomical basis of these functional domains was baffling to Hubel and Wiesel, since the radial fascicles seen in the light microscope seemed too ubiquitous and orderly to account for the picture they saw through their microelectrode (Hubel and Wiesel, 1968). Yet the causal relation between the two is undeniable: the neurons create the functional architecture, the functional architecture has regularities, and so how do the underlying circuits do it? The iso-orientation slabs were about 30 μm wide, but the ocular dominance slabs were 500 μm wide. For other cortical areas, the picture was even less clear. With respect to the whisker representation in rodents, Hubel and Wiesel stated, "Whether they [barrels] should be considered columns seems a matter of taste and semantics" (Hubel and Wiesel, 1974b). The whisker map in rodents is similar to retinotopic or tonotopic representations in other sensory cortices, but this map is not the same as higher-order maps of properties like orientation selectivity, whose topography is not predictable at the periphery. Nevertheless, the concept that neocortex consists of vertical arrangements of cells that are interconnected and have functional properties in common is almost universal. With the enormous success of the cortical physiologists in defining a basic architecture of cortex, the ball was back in the anatomists' court to explain the underlying structure.

Anatomists in the Garden

Although Golgi's stain had provided generations of anatomists with a powerful tool, it also had a major limitation, which Cajal had recognized and used to his advantage: it worked best in immature material. Because multiple cells were impregnated, it was virtually impossible to trace the same axon from one section to the next. Thus, the cell structures illustrated were obtained from reconstructions of single sections of perhaps 100 μm thick. A hint that this picture of the cortical axons was very incomplete came from degeneration studies in which small lesions had been made within the cortical gray matter. When the resulting degeneration was traced, it was clear that intracortical axons of pyramidal cells could extend over several mm (Fisken et al., 1975; Gatter and Powell, 1978; Creutzfeldt et al., 1977). What is most surprising was how similar the pattern of fiber degeneration was across different species and different areas. Figure 2 shows the close similarities between the striate visual cortex and the primary motor cortex of monkey cortex, which, in cytoarchitectonic and functional respects, differ the most. It was evidence like this that encouraged Powell to pursue the concept of cortical uniformity. Later, bulk injections of tracers like horseradish peroxidase into the cortex confirmed this pattern of spread (Blasdel et al., 1985) and also revealed the existence of widespread lateral connections of pyramidal cells that formed dense patches of boutons (Rockland and Lund, 1982).

The lateral fibers detected in the degeneration studies and the first complete picture of mature cortical neurons came from studies where single neurons had been injected

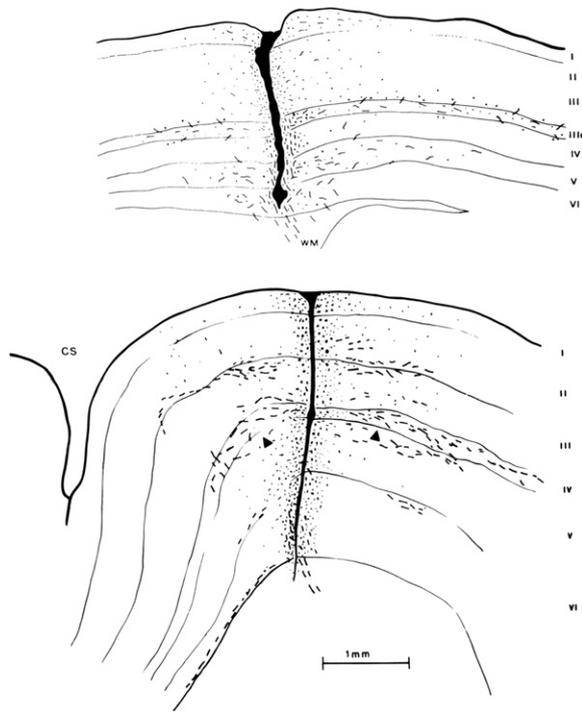


Figure 2. Similarity in the Patterns of Local Degeneration after a Narrow Cut Through the Layers of Cortex in Area 17 (Top) and Area 4 (Bottom) in the Monkey
From Fiske et al. (1975) and Gatter and Powell (1978).

intracellularly with horseradish peroxidase (HRP; Gilbert and Wiesel, 1979, 1983; Martin and Whitteridge, 1984). These pictures were a revelation (Figure 3).

It was as if the Golgi stain had been given growth hormone, for the black spindly trees had now extended their branches and sprouted bushy terminal thickets. It was evident that even the “projection neurons” were substantial players in the local cortical circuits. However, what constituted a local circuit was now even less clear. What was the elementary unit of structure that formed the vertical functional column? The palisades of cell bodies still formed their neat columns, but not only did the dendrites originating from these cell bodies spread well beyond these elementary columns, so did their axonal arbors. It became impossible on anatomical grounds to define the columnar structure of a given area, let alone explain how the visible structure gave rise to the functional phenomenon of columns. The difficulty of defining the column has also generated the opinion that the concept has failed as a unifying principle for cerebral cortex (Purves et al., 1992; Swindale, 1998; Horton and Adams, 2005). However, the idea of elementary modules that could, by repetition, generate an entire cortex, could not so easily be quashed by these brutal facts.

Modules and Maps

Critics of the column concept have to deal with the reality that the very long history of microelectrode recordings

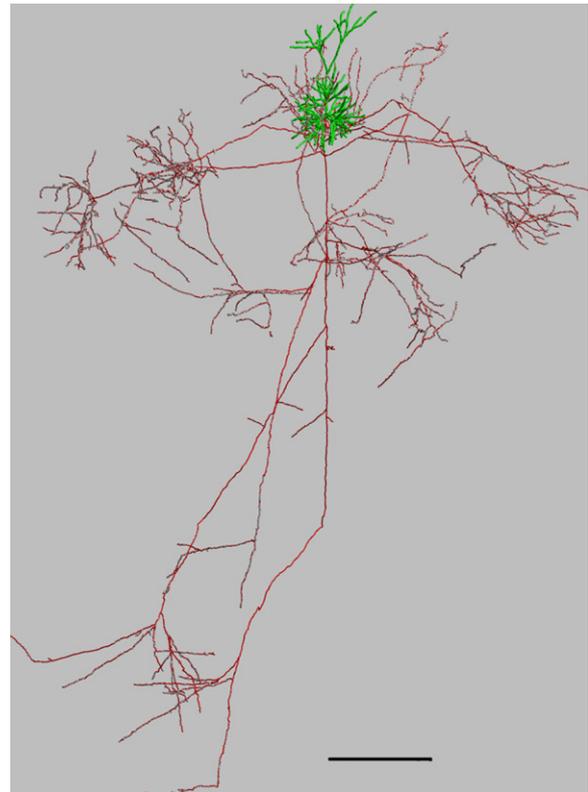


Figure 3. Pyramidal Cell of Layer 3 of Cat Visual Cortex Showing Dendrite (Green) and Axon (Red) Forming Multiple Clusters of Boutons (Black) in Layer 3 and 5

Scale bar, 500 μm . The neuron was intracellularly injected with horseradish peroxidase and reconstructed in 3D (Anderson, Binzegger, Douglas, and Martin, unpublished data).

from the primary visual cortex seemed to give the same results, regardless of where exactly the electrode was placed. If it were not so, progress would have been achingly slow, and the visual cortex would never have become the model system it has. This reliable repetitiveness could only be a result of an underlying uniformity. The case is most convincingly made for the many areas that contain topographic maps of the sensory periphery. Here, the cortical area represents in 2D a map of audible frequencies, or a whisker, or one patch of visual field, and the actual shape of the cortical surface is determined by this primary map (Daniel and Whitteridge, 1961). Within this map is the machinery studied so intensively by recording from single neurons and analyzing the receptive field properties. This realization that within the topographic map there was another dimension, the vertical distance between pial surface and white matter, prompted the deep thought that “the machinery may be roughly uniform over the whole striate cortex, the difference being in the inputs. A given region of cortex simply digests what is brought to it, and the process is the same everywhere.... It may be that there is a great developmental advantage in designing such machinery once only, and repeating it over and over

monotonously, like a crystal” (Hubel and Wiesel, 1974a, in the paper that they considered their most important after their 1962 paper).

Contrariwise, Hübener et al. (1997), on the basis of 2D optical imaging of the intrinsic signal evoked by various stimuli, suggested instead that the cortex “could not be considered a crystalline structure built from identical modules, but rather it is composed of ‘mosaics’ of functional domains for the different properties.” This latter view of functional mosaics, however, does not capture the fact that all recording methods show that there is continuity in many of the maps and that each neuron expresses not just one property, but a number of properties. These multidimensional receptive fields of single neurons means that different properties that map onto a given cortical surface can never be simply segregated into separate “modules.” However, the same multidimensionality is represented in orderly 3D mappings, forcing us back to the central conundrum of how this multidimensionality is generated in the physical circuits.

The Mismatch of the Minicolumns

Although the concept of an identifiable “minicolumn” has a wide currency (Mountcastle, 1997; Peters and Yilmaz, 1993; Rockland and Ichinohe, 2004), to suppose that there is some acreage of gray matter that sensible scientists will agree contains an essence of their “cortical column” seems doubtful. There is simply a mismatch between the anatomy and the functional maps. The notion of the minicolumn or module does not properly capture the granularity or the vertical and lateral interdigitation of component neurons that seems to be the essence of the cortical circuit. This means that even in a highly specialized “column” such as those evident in the rodent barrel cortex, one cannot simply clip out a cylinder of tissue that contains the whole local circuit, for later reconstruction “in silico” (Markram, 2006; Helmstaedter et al., 2007). This is not to say that detailed “bottom-up” models should not be pursued, or that there is no local circuit with a definable connection matrix. If you don’t know what is there, then you cannot hope to simplify intelligently for your models. We are convinced that there is a local circuit, one that we have referred to as a “canonical” circuit for neocortex (Figure 4) (Douglas et al., 1989; Douglas and Martin, 2004). This elemental circuit has been elaborated quantitatively for cat visual cortex (Binzegger et al., 2004, see below). It has also been demonstrated through detailed simulations that the quantitative circuit derived from cat visual cortex can, with few modifications, perform the operations of the prefrontal cortical area called frontal eye fields (FEF) in the monkey brain (Heinzle et al., 2007), thus importantly demonstrating the property of generality.

The matrix of cortical neurons scales in interesting ways across the cortical mantle: in the old world monkey, for example, Elston and colleagues have shown a rostro-caudal gradient in the size of pyramidal cells and their dendritic elaborations (Elston, 2002). An intriguing scaling across species is also apparent for the lateral collaterals of the

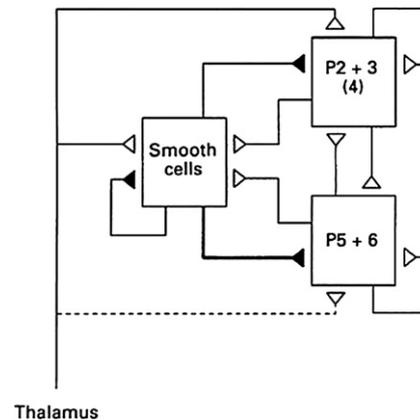


Figure 4. Canonical Cortical Circuit Based on Electrophysiological and Modeling Studies in the Cat Visual Cortex

From Douglas and Martin (1991).

layer 3 pyramidal cells, which collectively form a series of small clusters, which we have called a “daisy” (Douglas and Martin, 2004). Daisies seem to be present in all cortical areas in all species with the exception of rodents. From their work in the old world primate, Lund and colleagues (Lund et al., 1993; Pucak et al., 1996) pointed out that there is a simple correlation between the spacing between the clusters and the cluster size in the macaque monkey cortex. It turns out that this relationship may be universal across areas and species (Douglas and Martin, 2004). This is a remarkable and unexpected example of some underlying organizing principle, but supports the idea that there are common rules whereby the cortical circuits organize themselves in 3D.

Cajal and the Quantitative Circuits

There are many different levels at which we can begin to understand the cortical circuit. One of the simplest steps is to catalog the elements of the circuits: which types of neurons exist and in which layers and their relations. However, this in itself is not definitive; for the neocortex, it has allowed the full spectrum of interpretations of the local circuit, from tabula rasa to specified circuits. A second step, and one that has proved exceedingly difficult, is determining who talks to whom and how much. Thus, there are at least two levels of connection diagrams that can be distinguished. One connection diagram, let's call it the “Cajal circuit,” shows the average pattern of connections between the different neuronal types that make up the circuit. This circuit does not take into account any quantitative aspects and shows only the most essential functional connections of the circuit. So for the retina, for example, the Cajal circuit shows the connection from photoreceptors to bipolar cells to ganglion cells. This is an essential step, but sequence of processing through neuronal elements provides a logic of the circuit and not the relative influence of the elements. The “quantitative circuit” captures the essential magnitude of the relations between

the elements. In the case of the retina, this would be the number of photoreceptors and horizontal cells that connect to a bipolar cell and the number of bipolar cells that connect to amacrine and to ganglion cells. In the retina, these numbers change greatly with distance from the area centralis or fovea, and so these changing numbers reflect essential properties of the receptive fields. For the cortex, the assumption is that the proportions of different types of neurons remain uniform across an area and that their patterns of connections are constant. However, the principle purpose of the quantitative circuit is to understand the functional consequences of the circuit, which are given by the details of the connectivity. No wonder then that there is currently a worldwide initiative to develop automated methods for solving the structure of brain circuits at synaptic resolution (Briggman and Denk, 2006).

The anatomist Tom Powell, a major modern proponent of cortical uniformity, had previously explored this quantitative aspect of cortical circuitry in a very imaginative way. His case for cortical uniformity rested on several observations, one of which was the stereotyped pattern of degeneration after a focal lesion, mentioned above (reviewed in Powell, 1981). Another was that at electron microscope level, he could distinguish three basic cell types, which were in the same proportions in all of the cortical areas that he studied in the monkey, cat, and rat (Sloper et al., 1978; Winfield et al., 1981; see below). A third, more controversial observation was that he simply counted the number of neurons in an arbitrary box ($25 \times 30 \mu\text{m}$) that extended from layer 1 to white matter. With one exception (and despite 3-fold variations in cortical thickness) the number of neurons in the box was astonishingly similar across the cortex of the mouse, rat, cat, old world monkey, and man. The exception was area 17 in primates, which had about half as many again the counts found for all other areas (Rockel et al., 1974, 1980). This latter result begs the other question of how the counts should be interpreted, for at face value it suggests that there are major differences in the way that area 17 is assembled compared to all other cortical areas, but there are not (for review see Dehay and Kennedy, 2007). Of course, the concept of a basic uniformity of cortex is at odds with the view from classical cortical cytoarchitectonics, where the nonuniformity of cortex is distinguished on the basis of various cellular staining methods.

At the very high magnifications of the electron microscope, these cytoarchitectonic differences vanish in the apparently relentless gray wallpaper of an electron micrograph. Nevertheless, skilled anatomists, like Gray, Szentágothai, Powell, Jones, and Peters, were able to discover in the welter of ultrastructural detail clear differences in synapse morphology and associated targets, as well as ultrastructural correlates of the smooth and spiny neurons that had first been seen by Cajal in Golgi-stained neocortex. When applied to the neocortex, the electron microscope also greatly encouraged quantification, and now, with modern stereological methods, unbiased estimates can be made of numbers of the structures of choice: neu-

rons, glia, or synapses. This at least removes the quirky selectivity inherent in Golgi stains, or the biases inherent in intracellular staining *in vivo* or *in vitro*. The quantitative aspects of the cortical circuits revealed clear consistencies across species and cortical areas (Douglas and Martin, 2004). About 80% of neurons are spiny and excitatory and form 85% of the synapses, while about 20% of the neurons are smooth and inhibitory, and they form about 15% of the synapses, because of their smaller axonal ramifications. One equally consistent finding is that in forming synapses, cortical neurons are promiscuous. Thus, a single neuron will receive its inputs from thousands of neurons and in turn it will provide output to thousands of target neurons.

These conversations between pairs of neurons have been monitored by cross-correlation or intracellular recording methods. However, after many years of work from many laboratories, we still do not yet have a comprehensive picture of the Cajal circuit for even one cortical area, let alone a quantitative circuit of all the connections. Thus, the first priority for neocortex is still to understand comprehensively what actually exists in a cortical area. For most of the 20th century, anatomists like Cajal and Lorente de Nó would be able to count the number of different types of cortical neurons on the fingers of both hands. These types were based only on their morphology. Now, with the advent of molecular markers and physiological methods for examining the biophysics and synaptic methods, the combinatorial possibilities of all the different parameters that can be measured are enormous, and the number of types presently hangs more on inclination than on any commonly agreed criteria, as Crick has long complained (Crick, 1999). In some hands, there is an immense diversity of neuron types, whether classified by morphology, cytochemistry, physiology, or gene expression (see reviews by Monyer and Markram, 2004; Nelson et al., 2006; Yuste, 2005). In others' hands, there were just two: simple and complex cells (Hubel and Wiesel, 1962). Perhaps a taxonomy of cortical neurons will only be properly resolved when we have tools for combining molecular tools with circuit analysis, as have been so effective for the retina.

Given that each neuron is polyeuronally innervated, what we badly need is to identify the neuron of origin of any particular synapse. However, all techniques in current use, whether physiological or anatomical, are simply unable to do this. Although significant steps have been taken to develop transsynaptic tracers for resolving the total input to a single cell, the successful application of this technique to quantitative problems remains for the future (Wickersham et al., 2007). Thus, in the absence of a direct means of tracing quantitatively the total circuit, alternative theory-based estimates have had to do. This process began with the statistical approaches of Sholl (1956) and Braitenberg and Luria (1960) and persists into the modern era.

Mapping the Matrix

The most determined attempt to assign a source to every synapse in a single cortical area was made by Binzegger et al. (2004) for the cat's area 17. Starting with high-fidelity

reconstructions of complete axons and dendrites of single neurons obtained from intracellular recordings *in vivo*, and continuing with a large database gleaned from the literature of cell counts and synaptic counts, Binzegger et al. applied “Peter’s rule” (Braitenberg and Schüz, 1991), which states that neurons interconnect in proportion to the contribution to the neuropil of their dendrites and axonal synaptic boutons. This of course is an approximation that does not hold universally, for wherever there is a selectivity in the connections between types of neurons, there is an exception to Peter’s rule. However, on average both spiny and smooth cell types follow Peter’s rule in their connections in the neuropil, as seen from the quantitative electron microscopic analysis described above. The analysis of Binzegger et al. (2004) for cat visual cortex showed that for some layers at least, it was possible to account for most of the synapses in a given layer on the basis of the known content of the local circuits (Figure 5).

In other layers, however, there was an embarrassment of riches: there were far more synapses than could be accounted for by the known neuronal composition of area 17. These synapses, the “dark matter” of the cortex, constituted a colossal 90% of all the synapses in layer 1 and 40% of all the synapses in layer 6. This latter number translates to about 3000 synapses per layer 6 pyramidal neuron. One may speculate at the sources of these synapses, but it is slightly disturbing to think that their very existence was unknown until the attempt was made to account for every one. Unfortunately, such an analysis has not yet been attempted for any other cortical area.

What the source is of the “dark matter” is not clear. One simple way of accounting for them is to suppose that they arise from an as yet unidentified type of local neuron. This seems unlikely, because many studies using a variety of different techniques have time and again turned up the same set of neuronal types. Thus, while there may be Yeti or Bigfoot neurons, large numbers of them are unlikely to be found. A more likely source of the dark matter synapses are interareal or subcortical sources, such as the claustrum or basal forebrain nuclei, which of course were not included in the original solution of Binzegger et al. (2004), who considered only the local neurons and the thalamic afferents. Synapses from these other sources must certainly account for some of the dark matter, particularly for those in layer 1, which is a target of the so-called “feedback” projections. Whether such large numbers of unaccounted synapses can be made up by the long-distance connections remains an open experimental question.

The reasons for the shadow of doubt about whether these interareal sources will be sufficient to soak up the dark matter are the quantitative experimental studies of Kennedy and cohorts on the monkey cortex. They have mapped by retrograde tracing the source of all the inputs to a point in a cortical area, such as V1 or V4 (reviewed in Vezoli et al., 2004; A. Falchier et al., 2006, Soc. Neurosci, abstract). They find consistently that over 80% of the neurons labeled by the retrograde tracer lay within the same

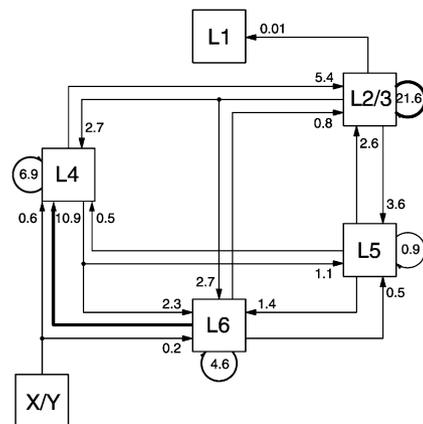


Figure 5. Quantitative Analysis of Excitatory Connections of Local Neurons and Thalamic Afferents in Cat Visual Cortex

Numbers are percentages of total excitatory synapses in area 17 of cat visual cortex contributed by the particular cell type. (See Binzegger et al., 2004.)

cortical area in which the tracer injection was made. The next largest group, about 10%, was in the neighboring cortical area. With progressive distances from the areas in which the tracer was laid, fewer and fewer labeled neurons were found, indicating progressively weaker connections between the two areas. For comparison, the lateral geniculate nucleus, which provides the major sensory drive to V1 and, relatively speaking, is a strong connection, contributed less than 1% of the total number of neurons that project to a point in V1. It is clear from these quantitative studies that if the dark matter is formed by neurons that lie outside the cortical area then each different projection contributes only a small fraction of the total additional synapses that are needed to account for the dark matter. Thus, long-distance projections from many different areas are needed to make up these numbers. This is a key implication in the light of the proposed functional roles of the long-distance projections on the neurons of the local circuits.

Modeling the Map of Synapses

Most modeling solutions to the cortical circuits are one-dimensional, because they reflect connectivity maps obtained from *in vitro* recordings or *in vivo* anatomical tracing studies. However, in the visual system, the two-dimensional receptive fields and cortical maps have given rise to more elaborate model circuits. Most of these have considered how properties like feature selectivity are generated, or how responses of cortical neurons become invariant. There are very large-scale simulations and often-ingenious expressions of these models, but these are almost without exception designed to accomplish their intended monistic goal, such as explaining orientation selectivity. Few have been generalized to encapsulate the generic computations that are carried out in any cortical area (Douglas et al., 1996; Douglas and Martin, 2007).

A big step forward in understanding the circuits of cortex came with the ingenious formulation of a processing hierarchy of thalamic relay cells, simple and then complex, devised by Hubel and Wiesel (1962). Now 45 years old, their back-of-the-envelope sketch of serial processing in the cat visual cortex has shown an admirably tenacious grip on life, to the extent that no alternative yet appears in the undergraduate textbooks, and it is still at the heart of much more elaborate computational models (see reviews by Ferster and Miller, 2000; Sompolinsky and Shapley, 1997). The enormous barriers to putting structure to function are illustrated by the stark fact that, although the connection between the LGN and layer 4 of cat visual cortex is one of the most studied connections in the central nervous system, it has taken decades to accumulate sufficiently detailed structural and physiological data to now be able to build biologically realistic models of the thalamic afferent input to spiny stellate cells (Banitt et al., 2007).

However, the experimental difficulties of determining quantitatively the actual numbers that should be put in the models cannot be exaggerated. One simple example is illustrated by the brief history of attempts to determine the fraction of thalamic synapses in layer 4 of cat visual cortex. Estimates of the proportion of thalamic synapses in layer 4 of cat visual cortex have varied over a factor of ten, depending on the experimental method used. The highest estimate was 28%, which was made by LeVay and Gilbert (1976), who used tritiated proline as the tracer and made their counts of autoradiographs of thin EM sections. LeVay (1986) later repeated the experiment with wheat-germ agglutinin as the tracer and obtained a markedly lower value of 5%, which, however, was much more in agreement with the estimates of Garey and Powell (1971), who used degeneration methods, and with those of Ahmed et al. (1994), who mapped the entire synaptic input to the dendrites of layer 4 spiny stellate cells (Anderson et al., 1994) and then used ultrastructural criteria to identify which were the LGN synapses. It seems likely that technical issues, such as background activity, generated the inflated figures obtained with tritiated proline. The consensus figure from the anatomy is that the thalamus provides around 5% of the excitatory synapses in its main target layer. This estimate is supported by electrophysiological studies using cross-correlation, where about 30 geniculate neurons are estimated to converge on a single simple cell in layer 4 (Usrey et al., 2000). Similar estimates have been made for barrel cortex (De Kock et al., 2007; see White, 1978). The majority of the remaining synapses in layer 4 of the cat come from other neurons in the local cortical circuit, such as the layer 6 pyramidal cells and spiny stellate cells. The question then is why cortex is connected to its principal thalamic nucleus by such a fragile long-distance link? The answer, we believe, lies in understanding the spatial and temporal relations between the local circuit, with its strongly recurrent excitatory and inhibitory subcircuits, and the long-distance connections that arise from many sources, including thalamus and cortex.

That's Just Enough!

The small number of thalamic synapses in relation to the excitatory synapses from intracortical sources may seem puzzling to the generations who have relied on the textbook model of the simple cell, which shows no other inputs than those from the thalamus. However, this number raises a key question that has never really been seriously asked: how many synapses should be made for a given projection? Because it is so well studied, the thalamic projection offers our best case study for establishing a likely number. If the thalamic afferents are the dominant excitatory input to the cortical simple cells, they clearly cannot be too dominant for the following reason: the synaptic mapping experiments showed that the just over 4000 excitatory synapses are formed with the dendritic tree of a spiny stellate cell (Ahmed et al., 1994). The simulations of a very detailed biophysical model show that the synchronous activation of about 100 thalamocortical synapses can drive the voltage of the spiny stellate cell membrane from resting potential through the spike threshold, assuming no inhibitory inputs are coactivated and there is no strong synaptic depression (Banitt et al., 2007). If the thalamic synapses are a substantial fraction of the excitatory synapses of a spiny stellate neuron, then it is obvious that even the activity of a small fraction of these would be sufficient to excite the postsynaptic cell. This means that even at the nonoptimal orientations, more than sufficient excitation would arrive from the thalamus to drive the spiny stellate cell through threshold, and so degrade the observed selective response of these neurons.

The classical solution to the problem of strong excitation is always: use strong inhibition. Thus, one argument has been that it is intracortical inhibition that prevents the cell from firing at the nonoptimal orientations. However, the case for so-called "cross-orientation inhibition" is weak, because direct measurements show that cortical inhibition is tuned to the same orientation as the excitation and is weak or absent at the nonoptimal orientations (Ferster, 1986; Douglas et al., 1988; Monier et al., 2003; Anderson et al., 2000). Thus, the fact that the layer 4 neurons do not fire at nonoptimal orientations must be because the relatively small numbers of thalamocortical synapses are not firing coherently enough to provide the spatial summation needed to drive the spiny stellate membrane through its spike threshold. If the thalamic drive is insufficient, then the recurrent connections of the local circuit will not be engaged. But when the thalamic relay cells drive the spiny stellate cell, then the recurrent circuits of the canonical circuit amplify to the synaptic drive initiated by the thalamic afferents (Douglas et al., 1989, 1996). In fact, the simulations of Banitt et al. (2007) indicate that some sort of background activity is required to help the thalamic afferents to drive their postsynaptic neurons. Thus, we are led inexorably to the conclusion that the thalamus contributes "just-enough" excitation to the spiny stellate cell to drive it when the appropriate stimulus configuration is presented. Similar mechanisms have been proposed for the thalamocortical link in the rodent barrel cortex, which is

similarly modest in its contribution to a single layer 4 neuron (White, 1978; Bruno and Sakmann, 2006).

Of course, it is trivially true that a threshold-crossing stimulus must provide just enough stimulation to evoke a response. But notice that our use of “just-enough” means that the total effect of input projections is also capped by the small number of synapses that are available to drive the target neuron from any particular source, such as thalamus or another cortical area. That is, the input projections are critically small. This feature has particular significance in relation to the pathways between cortical areas, which are thought to provide the essential “top-down modulation” or “bottom-up driving” of their targets. The significant effects of these connections are now taken for granted, for example in the celebrated distributed hierarchical processing network of the primate cerebral cortex (Felleman and Van Essen, 1991). But how is all this driving and modulation to be achieved with projections that provide at most a few percent of synapses to their target areas (Anderson et al., 1998; Vezoli et al., 2004)? It appears that these tiny long-distance projections must provide “just-enough” to guarantee a significant effect on their target under some spatiotemporal configurations of their input spikes, but on the other hand, their maximum effect is also capped by design. Thus, the “just-enough” principle may be true for all cortical circuits, because these long projections are the only ones that can provide the biases for the recurrent excitatory operations of the local canonical circuits (Douglas and Martin, 2007). The functional maps of cortex should then be seen not as the static units suggested by Hubel and Wiesel’s ice-cube model, but dynamic ensembles (Basole et al., 2003, 2006; Tsodyks et al., 1999; Kenet et al., 2003) whose state reflects not only the spatial interaction of local and long-distance circuits, but also their spatiotemporal interactions.

...and Just-In-Time

The flip side of the “just-enough” coin is the problem of punctuality. In conventional computation, time enters in a very different way to that used by the nervous system. Conventional computing is organized according to two requirements (Mead and Conway, 1980). The first is the topological ordering that describes the sequence of transformations of the process. The second is the temporal property of the underlying physical circuit that implements the communications and transformations. The values of the input data must be held steady for some minimum time before and after the clock signal that activates the transformation is applied. Beyond this technical requirement of the transformation, there is also the question of how the topological sequence of the algorithm is bound to the physical time of the underlying machine. There are two basic models: systems that operate synchronously and those that operate asynchronously. In the first case, the communication and transformation actions are synchronized by a global clock, while in the second case, each transformation element times only its local transformation. It simply waits until it has received all its necessary

input before emitting its output. This means that all transformation elements that take input from the waiting element must also wait until that output is available. In both cases, the physical time at which transitions occur have no role in the representation of the data being processed. It is only the sequence number of the transformation that is relevant, and physical time is simply the medium that permits data exchange and transformations to be coordinated at each step for the process (on either the global or the local scale). In contrast to these styles of machine computation, it appears that spiking neurons follow a rather different strategy, in which physical time is intrinsic to the data processing and representation.

Generally, neurons are quiescent, providing an output only on demand when driven by suitable input. Estimates based on power requirements for neuronal signaling and the power actually consumed by the brain suggest that only a few percent of neurons are significantly active at any time, an estimate that supports notions of sparse coding and on-demand (event-driven) processing (Lennie, 2003). In the cerebral cortex, the output of neurons is not synchronized by a global clock, although synchronization to a local oscillating field potential such as hippocampal theta and cortical gamma oscillations has been interpreted as providing a local clock (Hopfield and Brody, 2001). Unlike their silicon counterparts, neurons are ready to process at any time and will emit asynchronous output provided that some spatiotemporal input conditions on its dendrites are approximately met within some sliding time window of output excitability.

Because neurons are not waiting for a clock-tick, they are always prepared to process, and because neuronal processing is everywhere intimately linked to physical time, there is no need for a global clock to provide for synchronous processing steps. All that is required is that just enough signals rendezvous at a given location just-in-time to perform the transformation at that location, otherwise processing along that branch will die. Moreover, the processing in the brain is then intrinsically synchronized with the natural dynamics of the events in the external world that it strives to characterize and predict. How many spikes should rendezvous? The simulations of Banitt et al. (2007) indicate that a quiescent soma requires between 100–200 synaptic events within a window determined by the synaptic time constant in order to honor an output computation. That is, a few percent (a few hundred) of its excitatory inputs should be active during a window whose time constant is about 10–20 ms.

Is it possible for neurons to compute on such patterns? In principle, yes. Hopfield and others have described such mechanisms (Hopfield, 1995; Hopfield and Brody, 2001; Gütig and Sompolinsky, 2006). The “Tempotron” of Gütig and Sompolinsky (2006) is a model integrate-and-fire neuron that is able to learn decisions based on the spatiotemporal pattern of the spikes input to its simple dendrite. The Tempotron can learn to detect the presence of a particular spatiotemporal pattern anywhere in a time interval that is significantly longer than the duration of the patterns

themselves. Their model performs best for synaptic time constants of about 10 ms, provided the number of patterns that must be detected is no more than about twice the number of input synapses. A Tempotron signals its recognition of an input pattern by generating an output spike. So a model composed of parallel Tempotrons with some common inputs would offer a simple circuit element that detects and transforms temporal patterns. An important property of the Tempotron is that it is robust against small temporal jitter and even missing spikes in the input patterns. So the signal is restored, and signals that do not arrive in time can be neglected.

Permitted Sets

When just-enough and just-in-time are combined in the cortical circuits, the result is a rich matrix of possible outcomes. Their manner of interaction is illustrated schematically in Figure 6, which shows how a computation evolves in time in a network of neurons, represented here for only a single spatial dimension. The network illustrated consists of three spatiotemporal “rendezvous nodes” (Rab, Rcd, Rbd). Each node represents a temporal window over a clusters of neurons within the same orientation domain or within the same whisker barrel, for example.

The rendezvous nodes receive input from distant sites, such as thalamus or another cortical area. These distant sites are represented by nodes Ra and Rb, which emit messages composed of spatiotemporal spike patterns along the communication edges indicated as red arrows. These messages are similar to the synfire chains of Abeles (Abeles, 1982). The arrows represent the axonal projections between the nodes, while the blue-gray paths flanking each arrowed edge represent the spatial dimension (y axis) of the anatomical projection (just-enough) and the temporal window (x axis) dimensions of the effective spike pattern being transmitted. Messages A and B pass through the rendezvous node Rab, where as a result of a local interaction in the node, Rab emits spatiotemporal messages B' and C. C, in this example, consists of a small number of neurons that hold their outputs steady for some while until the message D arrives at rendezvous node Rcd. The interaction of C and D leads to dynamic switching within node Rcd, which results in the emission of message E and the quenching of a potential progress of D to D'. As a result of this dynamic switching in node Rcd, the interaction of B' and D' in node Rbd cannot occur with the consequence that the potential emission B'' from node Rbd will not be honored. This scheme obviously places emphasis on the role of populations of neurons in collectively transforming multineuronal input messages, rather than on the input-output relations of single neurons.

Whether by single neuron or by collective computation, how is it that inputs which are barely enough and nearly in time can be used to compute a reliable output? It is here, we believe, that the recurrent connections of the cortical circuits play their role, providing the signal gain necessary for actively reorganizing the relatively small contributions of long-range afferents. It is this collective processing of

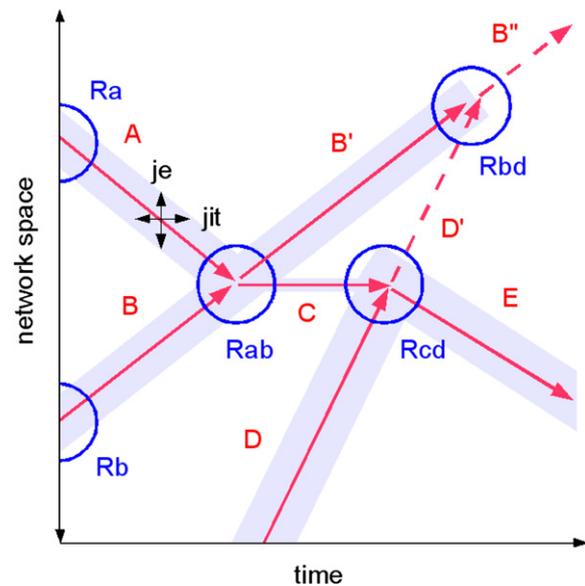


Figure 6. Schematic Representation of Just-Enough and Just-in-Time Computation in a Cortical Network

The computation is shown evolving in time in a network of neurons, represented here spatially in one dimension. Two local populations of neurons located at computational “rendezvous” nodes Ra and Rb emit messages composed of spatiotemporal spike patterns along the communication edges (“axons”) indicated as red arrows. The “widths” of the connection pathways (the number of connections) and the temporal window during which they can be effective are indicated by the blue-gray paths flanking each arrowed edge (widths “je” and “jit”). Messages A and B are shown passing through a rendezvous node Rab. As a result of that local interaction, Rab emits messages B' and C. C in this case consists of a small number of neurons that hold their outputs steady for some interval until the arrival of message D at rendezvous node Rcd. The interaction of C and D then leads to the emission of message E from Rcd. As a result of this interaction, another possible output, D', from node Rcd does not occur, which means that a possible interaction between B' and D' in rendezvous node Rbd does not occur, and B'' is not emitted.

transformations that could provide the signal restoration properties sought by von Neumann, who pondered how the brain was able to transform data through many successive stages without degradation of significant signal (von Neumann, 2000). The positive feedback between neurons of the cooperating population can be used to enhance the features of the input that match patterns embedded in the weights of the excitatory feedback connections. At the same time, the overall strength of the excitatory response of the population is used to suppress outliers via the dynamical inhibitory threshold imposed by inhibitory neurons. In this sense, the network can actively impose an interpretation on an incomplete or noisy input signal by restoring it toward some activity distribution across a “permitted set” of neurons (Hahnloser et al., 2002; Douglas and Martin, 2007).

Thus, the potential instability of the strong positive feedback, which is a cardinal feature of the recurrent cortical circuits, is exploited during the transient behavior of the networks, because the network can use this instability to explore new interpretations (partitions of active neurons)

until a suitable (stable) interpretation, which is consistent with the input pattern, is found (Hahnloser et al., 2000). The computational primitives found in the neocortex include linear operations such as summation, division, and sign inversion, and also nonlinear operations such as winner-take-all, invariance, and multistability, in which lateral interactions of excitation and inhibition play a key role (Douglas et al., 1996; Douglas and Martin, 2007). These primitives can form the basis for many of the computations that have been observed in different cortical areas. It is perhaps important to point out that these computational primitives arise through the collective action of the whole circuit and are not carried out within the dendritic tree of single neurons. As a consequence of this collective action of many thousands of neurons, the neural avalanche of Cajal, which begins at the peripheral sense organs and rolls on through the matrix of canonical cortical circuits, is not out of control, but is constrained such that it generates a coherent activity consistent with the context in which the animal finds itself. The cortical output along the final common path of Sherrington then generates an adaptive motor action.

Lost (and Found) in Translation

Our understanding of the organization of specific neocortical circuits derives in large part from electrophysiological work carried out over the past 40 years. This is exemplified in the work in visual cortex where the results from single-electrode recordings were the basis for inferring a functional architecture for the visual cortex. Details of this functional architecture were expressed in circuits, most famously those for simple and complex cells in area 17 of the cat (Hubel and Wiesel, 1962). Because the specificity of the connections inferred was on a very fine grain, it has proved impossible to translate any of these functional models into an anatomically proven circuit. Thus, the translation of structure into function remains elusive for neocortex, and Francis Crick's dictum that, "if you do not make headway understanding a complex system, study its structure and knowledge of the function will follow automatically" has yet to be realized for the neocortex.

As is by now a familiar pattern, new floating signifiers have been coined to soak up the excess of existence, as in the "dark matter" of cat visual cortex, coined by Binzegger et al. (2004) to explain the existence of large numbers of synapses of no fixed address. Nevertheless, our current knowledge of the structural richness of cortical neurons and the happy reappearance of familiar neuronal morphologies and patterns of connection in many different cortical areas and in different species challenges us to provide a far more comprehensive translation of structure into function. This challenge is now being addressed by the development of methods for large-scale solutions of neocortical circuits at synaptic resolution. But in translating structure into function we need more than the weight matrix of the connectivity or even the complete wiring diagram. To understand the syntax of cortex is going to be a major advance, but we also need to understand

deeply the semantics of the many areas of neocortex and the conversations they have with each other. As a wise person once remarked: the great difference between the telephone directory and a play of Shakespeare is that, while both have a grand cast of characters, only the play has a plot.

ACKNOWLEDGMENTS

Preparation of this review was supported by European Union Daisy Project grant FP6-2005-015803. We thank John Anderson, Tom Binzegger, Nuno da Costa and Dylan Muir for assistance.

REFERENCES

- Abeles, M. (1982). *Local Cortical Circuits: An Electrophysiological Study* (Berlin: Springer).
- Ahmed, B., Anderson, J.C., Douglas, R.J., Martin, K.A., and Nelson, J.C. (1994). Polyeuronal innervation of spiny stellate neurons in cat visual cortex. *J. Comp. Neurol.* **347**, 39–49.
- Anderson, J.C., Douglas, R.J., Martin, K.A., and Nelson, J.C. (1994). Map of the synapses formed with the dendrites of spiny stellate neurons of cat visual cortex. *J. Comp. Neurol.* **347**, 25–38.
- Anderson, J.C., Binzegger, T., Martin, K.A., and Rockland, K.S. (1998). The connection from cortical area V1 to V5: a light and electron microscopic study. *J. Neurosci.* **18**, 10525–10540.
- Anderson, J.S., Carandini, M., and Ferster, D. (2000). Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. *J. Neurophysiol.* **84**, 909–926.
- Banitt, Y., Martin, K.A.C., and Segev, I. (2007). A biologically realistic model of contrast invariant orientation tuning by thalamocortical synaptic depression. *J. Neurosci.* **27**, 10230–10239.
- Barkow, J.H., Cosmides, L., and Tooby, J., eds. (1992). *The Adapted Mind: Evolutionary Psychology and the Generation of Culture* (New York: Oxford University Press).
- Basole, A., White, L.E., and Fitzpatrick, D. (2003). Mapping multiple features in the population response of visual cortex. *Nature* **423**, 986–990.
- Basole, A., Kreft-Kerekes, V., White, L.E., and Fitzpatrick, D. (2006). Cortical cartography revisited: A frequency perspective on the functional architecture of visual cortex. *Prog. Brain Res.* **154**, 121–134.
- Binzegger, T., Douglas, R.J., and Martin, K.A.C. (2004). A quantitative map of the circuit of cat primary visual cortex. *J. Neurosci.* **24**, 8441–8453.
- Blasdel, G.G., Lund, J.S., and Fitzpatrick, D. (1985). Intrinsic connections of macaque striate cortex: axonal projections of cells outside lamina 4C. *J. Neurosci.* **5**, 3350–3369.
- Braitenberg, V., and Luria, F. (1960). Toward a mathematical description of the grey substance of the nervous system. Supplemental Volume XVIII Serie X Dei Nuovo Cimento, 149–165.
- Braitenberg, V., and Schüz, A. (1991). *Anatomy of the Cortex* (Berlin: Springer-Verlag).
- Briggman, K.L., and Denk, W. (2006). Towards neural circuit reconstruction with volume electron microscopy techniques. *Curr. Opin. Neurobiol.* **16**, 562–570.
- Bruno, R.M., and Sakmann, B. (2006). Cortex is driven by weak but synchronously active thalamocortical synapses. *Science* **312**, 1622–1627.
- Cajal, S.R. (1937). *Recollections of My Life* [translated by E.H. Craigie and J. Cano, 1989] (Philadelphia, PA: American Philosophical Society).

- Creutzfeldt, O.D., Garey, L.J., Kuroda, R., and Wolff, J.-R. (1977). The distribution of degenerating axons after small lesions in the intact and isolated visual cortex of the cat. *Exp. Brain Res.* 27, 419–440.
- Crick, F. (1999). The impact of molecular biology on neuroscience. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 354, 2021–2025.
- Daniel, P.M., and Whitteridge, D. (1961). The representation of the visual field on the cerebral cortex in monkeys. *J. Physiol.* 159, 203–221.
- De Kock, C.P., Bruno, R.M., Spors, H., and Sakmann, B. (2007). Layer- and cell-type-specific suprathreshold stimulus representation in rat primary somatosensory cortex. *J. Physiol.* 581, 139–154.
- Dehay, C., and Kennedy, H. (2007). Cell-cycle control and cortical development. *Nat. Rev. Neurosci.* 8, 438–450.
- Douglas, R.J., and Martin, K.A.C. (1991). A functional microcircuit for cat visual cortex. *J. Physiol.* 440, 735–769.
- Douglas, R.D., and Martin, K.A.C. (2004). Circuits of neocortex. *Annu. Rev. Neurosci.* 27, 419–451.
- Douglas, R.J., and Martin, K.A.C. (2007). Recurrent neuronal circuits in the neocortex. *Curr. Biol.* 17, R496–R500.
- Douglas, R.J., Martin, K.A.C., and Whitteridge, D. (1988). Selective responses of visual cortical cells do not depend on shunting inhibition. *Nature* 332, 642–644.
- Douglas, R.J., Martin, K.A.C., and Whitteridge, D. (1989). A canonical microcircuit for neocortex. *Neural Comput.* 1, 480–488.
- Douglas, R.J., Mahowald, M., Martin, K.A., and Stratford, K.J. (1996). The role of synapses in cortical computation. *J. Neurocytol.* 25, 893–911.
- Elston, G.N. (2002). Cortical heterogeneity: Implications for visual processing and polysensory integration. *J. Neurocytol.* 31, 317–335.
- Felleman, D.J., and Van Essen, D.C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cereb. Cortex* 1, 1–47.
- Ferster, D. (1986). Orientation selectivity of synaptic potentials in neurons of cat primary visual cortex. *J. Neurosci.* 6, 1284–1301.
- Ferster, D., and Miller, K.D. (2000). Neural mechanisms of orientation selectivity in the visual cortex. *Annu. Rev. Neurosci.* 23, 441–471.
- Fisken, R.A., Garey, L.J., and Powell, T.P.S. (1975). The intrinsic, association and commissural connections of area 17 of the visual cortex. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 272, 487–536.
- Fodor, J.A. (1983). *The Modularity of Mind* (Cambridge, MA: MIT Press).
- Garey, L.J., and Powell, T.P.S. (1971). An experimental study of the termination of the lateral geniculo-cortical pathway in the cat and monkey. *Proc. R. Soc. Lond. B Biol. Sci.* 179, 21–40.
- Gatter, K.C., and Powell, T.P.S. (1978). The intrinsic connections of the cortex area 4 of the monkey. *Brain* 101, 513–541.
- Gilbert, C.D., and Wiesel, T.N. (1979). Morphology and intracortical projections of functionally characterised neurones in the cat visual cortex. *Nature* 280, 120–125.
- Gilbert, C.D., and Wiesel, T.N. (1983). Clustered intrinsic connections in cat visual cortex. *J. Neurosci.* 3, 1116–1133.
- Gütig, R., and Sompolinsky, H. (2006). The tempotron: a neuron that learns spike timing-based decisions. *Nat. Neurosci.* 9, 420–428.
- Hahnloser, R., Sarpeshkar, R., Mahowald, M., Douglas, R., and Seung, S. (2000). Digital selection and analog amplification co-exist in an electronic circuit inspired by neocortex. *Nature* 405, 947–951.
- Hahnloser, R.H., Douglas, R.J., and Hepp, K. (2002). Attentional recruitment of inter-areal recurrent networks for selective gain control. *Neural Comput.* 14, 1669–1689.
- Hamilton Paterson, J. (1993). *Seven-Tenths: The Sea and Its Thresholds* (London: Vintage).
- Heinzle, J., Hepp, K., and Martin, K.A.C. (2007). A microcircuit model of the frontal eye fields. *J. Neurosci.* 27, 9341–9353.
- Helmstaedter, M., de Kock, C.P., Feldmeyer, D., Bruno, R.M., and Sakmann, B. (2007). Reconstruction of an average cortical column in silico. *Brain Res. Rev.* 55, 193–203.
- Hopfield, J.J. (1995). Pattern recognition computation using action potential timing for stimulus representation. *Nature* 376, 33–36.
- Hopfield, J., and Brody, C. (2001). What is a moment? Transient synchrony as a collective mechanism for spatiotemporal integration. *Proc. Natl. Acad. Sci. USA* 98, 1282–1287.
- Horton, J.C., and Adams, D.L. (2005). The cortical column: a structure without a function. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 360, 837–862.
- Hubel, D.H., and Wiesel, T.N. (1959). Receptive fields of single neurones in the cat's striate cortex. *J. Physiol.* 148, 574–591.
- Hubel, D.H., and Wiesel, T.N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* 160, 106–154.
- Hubel, D.H., and Wiesel, T.N. (1963). Shape and arrangement of columns in cat's striate cortex. *J. Physiol.* 165, 559–568.
- Hubel, D.H., and Wiesel, T.N. (1968). Receptive fields and functional architecture of monkey striate cortex. *J. Physiol.* 195, 215–243.
- Hubel, D.H., and Wiesel, T.N. (1972). Laminar and columnar distribution of geniculocortical fibers in the macaque monkey. *J. Comp. Neurol.* 146, 421–450.
- Hubel, D., and Wiesel, T. (1974a). Uniformity of monkey striate cortex: A parallel relationship between field size, scatter, and magnification factor. *J. Comp. Neurol.* 158, 295–306.
- Hubel, D.H., and Wiesel, T.N. (1974b). Sequence regularity and geometry of orientation columns in the monkey striate cortex. *J. Comp. Neurol.* 158, 267–293.
- Hubel, D.H., and Wiesel, T.N. (1977). Ferrier lecture. Functional architecture of macaque monkey visual cortex. *Proc. R. Soc. Lond. B Biol. Sci.* 198, 1–59.
- Hübener, M., Shoham, D., Grinvald, A., and Bonhoeffer, T. (1997). Spatial relationships among three columnar systems in cat area 17. *J. Neurosci.* 17, 9270–9284.
- Jeanmonod, D., Rice, F.L., and Van der Loos, H. (1981). Mouse somatosensory cortex: alterations in the barrelfield following receptor injury at different early postnatal ages. *Neuroscience* 6, 1503–1535.
- Kalisman, N., Silberberg, G., and Markram, H. (2005). The neocortical microcircuit as a tabula rasa. *Proc. Natl. Acad. Sci. USA* 102, 880–885.
- Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A., and Arieli, A. (2003). Spontaneously emerging cortical representations of visual attributes. *Nature* 425, 954–956.
- Lashley, K.S. (1930). Basic neural mechanisms in behavior. *Psychol. Rev.* 37, 1–24.
- Lashley, K.S. (1950). In search of the engram. *Symp. Soc. Exp. Biol.* 4, 454–482.
- Le Be, J.V., and Markram, H. (2006). Spontaneous and evoked synaptic rewiring in the neonatal neocortex. *Proc. Natl. Acad. Sci. USA* 103, 13214–13219.
- Lennie, P. (2003). The cost of cortical computation. *Curr. Biol.* 13, 493–497.
- LeVay, S. (1986). Synaptic organization of claustral and geniculate afferents to the visual cortex of the cat. *J. Neurosci.* 150, 53–86.

- LeVay, S., and Gilbert, C.D. (1976). Laminar patterns of geniculocortical projections in the cat. *Brain Res.* 113, 1–19.
- Lorente de Nó, R. (1949). The cerebral cortex: architecture, intracortical connections, motor projections. In *Physiology of the Nervous System* (New York: Oxford University Press), pp. 288–330.
- Lund, J.S., Yoshioka, T., and Levitt, J.B. (1993). Comparison of intrinsic connectivity in different areas of macaque monkey cerebral cortex. *Cereb. Cortex* 3, 148–162.
- Markram, H. (2006). The blue brain project. *Nat. Rev. Neurosci.* 7, 153–160.
- Martin, K.A.C., and Whitteridge, D. (1984). Form, function and intracortical projections of spiny neurones in the striate visual cortex of the cat. *J. Physiol.* 353, 463–504.
- Mead, C., and Conway, L. (1980). *Introduction to VLSI Systems* (Reading, MA: Addison-Wesley).
- Monier, C., Chavane, F., Baudot, P., Graham, L.J., and Fregnac, Y. (2003). Orientation and direction selectivity of synaptic inputs in visual cortical neurons: a diversity of combinations produces spike tuning. *Neuron* 37, 663–680.
- Monyer, H., and Markram, H. (2004). Interneuron diversity series: Molecular and genetic tools to study GABAergic interneuron diversity and function. *Trends Neurosci.* 27, 90–97.
- Mountcastle, V. (1957). Modality and topographic properties of single neurons of cat somatosensory cortex. *J. Neurophysiol.* 20, 408–434.
- Mountcastle, V.B. (1997). The columnar organization of the neocortex. *Brain* 120, 701–722.
- Nelson, S.B., Hempel, C., and Sugino, K. (2006). Probing the transcriptome of neuronal cell types. *Curr. Opin. Neurobiol.* 16, 571–576.
- Ohki, K., Chung, S., Ch'ng, Y.H., Kara, P., and Reid, R.C. (2005). Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature* 433, 597–603.
- Ohki, K., Chung, S., Kara, P., Hubener, M., Bonhoeffer, T., and Reid, R.C. (2006). Highly ordered arrangement of single neurons in orientation pinwheels. *Nature* 442, 925–928.
- Peters, A., and Yilmaz, E. (1993). Neuronal organization in area 17 of cat visual cortex. *Cereb. Cortex* 3, 49–68.
- Powell, T.P.S. (1981). Certain aspects of the intrinsic organization of the cerebral cortex. In *Brain Mechanisms and Perceptual Awareness*, O. Popeiano and C. Ajmone Marsan, eds. (New York: Raven Press), pp. 1–19.
- Powell, T.P., and Mountcastle, V.B. (1959). Some aspects of the functional organization of the cortex of the postcentral gyrus of the monkey: a correlation of findings obtained in a single unit analysis with cytoarchitecture. *Bull. Johns Hopkins Hosp.* 105, 133–162.
- Pucak, M.L., Levitt, J.B., Lund, J.S., and Lewis, D.A. (1996). Patterns of intrinsic and associational circuitry in monkey prefrontal cortex. *J. Comp. Neurol.* 376, 614–630.
- Purves, D., Riddle, D.R., and LaMantia, A.S. (1992). Iterated patterns of brain circuitry (or how the cortex gets its spots). *Trends Neurosci.* 10, 362–368.
- Rockel, A.J., Hiorns, R.W., and Powell, T.P.S. (1974). Numbers of neurons through full depth of neocortex. *J. Anat. London* 118, 371.
- Rockel, A.J., Hiorns, R.W., and Powell, T.P.S. (1980). The basic uniformity in structure of the neocortex. *Brain* 103, 221–244.
- Rockland, K.S., and Lund, J.S. (1982). Widespread periodic intrinsic connections in the tree shrew visual cortex. *Science* 215, 1532–1534.
- Rockland, K.S., and Ichinohe, N. (2004). Some thoughts on cortical minicolumns. *Exp. Brain Res.* 158, 265–277.
- Sholl, D.A. (1956). *The Organization of the Cerebral Cortex* (London: Methuen).
- Sloper, J.J., Hiorns, R.W., and Powell, T.P.S. (1978). A qualitative and quantitative electron microscopic study of the neurons in the primate motor and somatic sensory cortices. *Philos. Trans. Roy. Soc. Lond. B. Biol. Sci.* 285, 351–359.
- Sompolinsky, H., and Shapley, R. (1997). New perspectives on the mechanisms for orientation selectivity. *Curr. Opin. Neurobiol.* 7, 514–522.
- Sperry, R.W. (1947). Cerebral regulation of motor coordination in monkeys following multiple transection of sensorimotor cortex. *J. Neurophysiol.* 10, 275–294.
- Sperry, R.W., Miner, N., and Myers, R.E. (1955). Visual pattern perception following subpial slicing and tantalum wire implantations in the visual cortex. *J. Compar Physiol. Psych.* 48, 50–58.
- Swindale, N.V. (1998). Cortical organization: modules, polymaps and mosaics. *Curr. Biol.* 8, R270–R273.
- Szentágothai, J. (1978). Specificity versus (quasi-) randomness in cortical connectivity. In *Architectonics of the Cerebral Cortex*, M.A.B. Brazier and H. Petsche, eds. (New York: Raven Press), pp. 77–97.
- Tsodyks, M., Kenet, T., Grinvald, A., and Arieli, A. (1999). Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* 286, 1943–1946.
- Usrey, W.M., Alonso, J.M., and Reid, R.C. (2000). Synaptic interactions between thalamic inputs to simple cells in cat visual cortex. *J. Neurosci.* 20, 5461–5467.
- Vezoli, J., Falchier, A., Jouvé, B., Knoblauch, K., Young, M., and Kennedy, H. (2004). Quantitative analysis of connectivity in the visual cortex: extracting function from structure. *Neuroscientist* 10, 476–482.
- von Neumann, J. (2000). *The Computer and the Brain*, Second Edition (New Haven, CT: Yale University Press).
- White, E.L. (1978). Identified neurons in mouse Sml cortex which are postsynaptic to thalamocortical axon terminals: a combined Golgi-electron microscopic and degeneration study. *J. Comp. Neurol.* 181, 627–661.
- Wickersham, I.R., Lyon, D.C., Barnard, R.J., Mori, T., Finke, S., Conzelmann, K.K., Young, J.A., and Callaway, E.M. (2007). Monosynaptic restriction of transsynaptic tracing from single, genetically targeted neurons. *Neuron* 53, 639–647.
- Winfield, D.A., Brooke, R.N., Sloper, J.J., and Powell, T.P.S. (1981). An electron microscopic study of the types and proportions of neurons in the cortex of the motor and visual areas of the cat and rat. *Brain* 103, 245–258.
- Yuste, R. (2005). Origin and classification of neocortical interneurons. *Neuron* 48, 524–527.
- Zeki, S. (1993). *A Vision of the Brain* (Oxford: Blackwell Scientific Publications).