



From "The Neocortex," edited by W. Singer, T. J. Sejnowski and P. Rakic.
Strüngmann Forum Reports, vol. 27, J. R. Lupp, series editor.
Cambridge, MA: MIT Press. ISBN 978-0-262-04324-3

Functional Architecture of the Cerebral Cortex

David A. Leopold, Peter L. Strick, Danielle S. Bassett,
Randy M. Bruno, Hermann Cuntz, Kristen M. Harris,
Marcel Oberlaender, and Marcus E. Raichle

Abstract

Recent research in the neurosciences has revealed a wealth of new information about the structural organization and physiological operation of the cerebral cortex. These details span vast spatial scales and range from the expression, arrangement, and interaction of molecular gene products at the synapse to the organization of computational networks across the whole brain. This chapter highlights recent discoveries that have laid bare important aspects of the brain's *functional architecture*. It begins by describing the dynamic and contingent arrangement of subcellular elements in synaptic connections. Amid this complexity, several common neural circuit motifs, identified across multiple species and preparations, shape the electrophysiological signaling in the cortex. It then turns to the topic of network organization, spurred by routine capacity for noninvasive MRI in humans, where interdisciplinary tools are lending new insights into large-scale principles of brain organization. Discussion follows on one of the most important aspects of brain architecture; namely, the plasticity that affords an animal flexible behavior. In closing, reflections are put forth on the nature of the brain's complexity, and how its biological details might be best captured in computational models in the future.

Introduction

The human cerebral cortex consists of approximately 16 billion neurons (Herculano-Houzel 2009), whose integrated activity supports not only our

Group photos (top left to bottom right) David Leopold, Peter Strick, Randy Bruno, Kristen Harris, Marcus Raichle, Marcel Oberlaender, Hermann Cuntz, Danielle Bassett, Marcus Raichle, Marcel Oberlaender, Peter Strick, David Leopold, Danielle Bassett, Hermann Cuntz, Marcel Oberlaender, Kristen Harris, Danielle Bassett, David Leopold, Randy Bruno, Kristen Harris, Peter Strick

higher thoughts but also our sensory perceptions, verbal communication, and complex motor actions. While the number of neurons is clearly important, it is their organization and interconnections that determine the functional principles of a working brain. From one perspective, all mammalian brains have the same basic design, which includes a layered cerebral cortex governed by highly conserved developmental constraints (Workman et al. 2013). Mammalian species differ markedly, however, in brain size, peripheral sensory adaptations, and evolved ecological specializations. These differences strongly influence the brain organization of different taxa, with one pertinent example being the relatively dense packing of neurons into the primate cerebral cortex (Herculano-Houzel 2012). The principles of cerebral cortical architecture in primates and other mammals are simultaneously manifest at multiple scales: from the synaptic microenvironment, to local circuit motifs, to large-scale brain networks measured using methods such as functional magnetic resonance imaging (fMRI). At each scale, strong genetic determinism is complemented by modification through experience, which manifests both during early development and in the adult. In recent years, neuroscientists have learned a great deal about how flexibility in function can be superimposed upon an ostensibly fixed anatomical scaffolding.

Here, we take on the task of identifying key elements of cortical architecture that shape its basic functioning, plasticity, and capacity to drive flexible behavior. This review reflects our discussions at the 27th Ernst Strüngmann Forum—the third in a series of meetings spaced out over several decades. Given the amount of research that transpired since the initial meeting (Rakic and Singer 1988), we highlight new concepts and discoveries that have emerged pertaining to the cerebral cortex, its organization, and function. We focus on recent discoveries regarding synapse formation, local and long-range functional connections, and network organization at multiple scales. Where possible, we place findings in a historical context and direct readers to recent reviews on other important features of cortical architecture: prominent laminar organization of the cortex (Palomero-Gallagher and Zilles 2017), its columnar microcircuitry (Bastos et al. 2012), and its intimate and mysterious functional relationship to other prominent brain components, such as the thalamus (Sherman 2017).

We begin by reviewing the points of articulation between neurons, including biophysical and physiological features of the dendritic microenvironment that promote certain modes of information transmission. We then investigate the structural and functional connectivity between distant areas of the cerebral cortex, which is a field of study that has come to utilize the brain's spontaneous activity. Thereafter we highlight the inherent flexibility of the cerebral cortex, from experience-dependent changes during early development to adult learning and memory. We conclude by briefly considering the importance of computational and evolutionary frameworks in shaping our future conceptions of cortical functional architecture.

The Essence of a Neural Connection

Our current understanding of neural communication is grounded in the neuron doctrine, which is seen as the resolution of a nineteenth century debate between Ramón y Cajal and Golgi about whether neurons were interconnected through directed lines or as a broad syncytium (Bock 2013). Neurons are cells specialized to transmit information quickly through electrochemical signals that traverse a range of spatial scales. Individual neurons usually communicate through chemical cell-to-cell contacts, or synapses. The cartoon rendition of the neuron is familiar to all students of neuroscience: dendrites emerge from a cell body and a long axon makes synaptic connections with another neuron's dendrites. Unsurprisingly, the structure of real neurons is much more complex and variable than textbooks typically portray, and the physiology of synaptic connections is highly contingent on factors playing out over many spatial scales (Figure 9.1). Our knowledge about these details is growing at a rapid pace. Historically, our picture of neuronal physiology was strongly shaped by action potentials acquired in single-unit recordings. These clean and discrete pulses might suggest a brain that works by digital computation, perhaps reflecting the contemporary metaphor of the brain: the computer. However, the core of the brain's information processing is arguably its analog physiology, including the electrical, chemical, and genetic mechanisms that control the synaptic interconnections at a range of timescales.

Synaptic neuronal connections are diverse and commonly involve articulation between presynaptic axons and postsynaptic dendrites. Postsynaptic neurons integrate a massive and uneven array of axonal inputs, often stemming from diverse cell types. In the cortex, synapses onto dendritic spines have been the focus of much study, since the morphology of spines changes readily in the adult brain. These constant changes are thought to alter synaptic efficacy in the service of network plasticity and learning. Tracking the fate of individual spines can seem hopeless, given that there are estimated to be 10^{14} spine synapses in the human cerebral cortex (Matus 2009). Nonetheless, the principles governing their formation, retreat, or enlargement may be among the most important windows into how the adult brain remains adaptable and, in a sense, youthful: with spines, the brain's capacity for experience-dependent development seems endless.

Recent discoveries have emphasized the high specificity of synapse formation, the importance of neighboring synapses, and a number of commonly occurring synaptic and circuit "motifs" through which certain functional computations are achieved (Jiang et al. 2015). Multiple modes of plasticity are built into synaptic connections, and these adhere to complex learning rules and are subject to a wide range of cognitive and chemical contingencies. In recent years, researchers have attempted to gain a more holistic understanding of the dendritic microenvironment, its margins for plasticity, its modulation

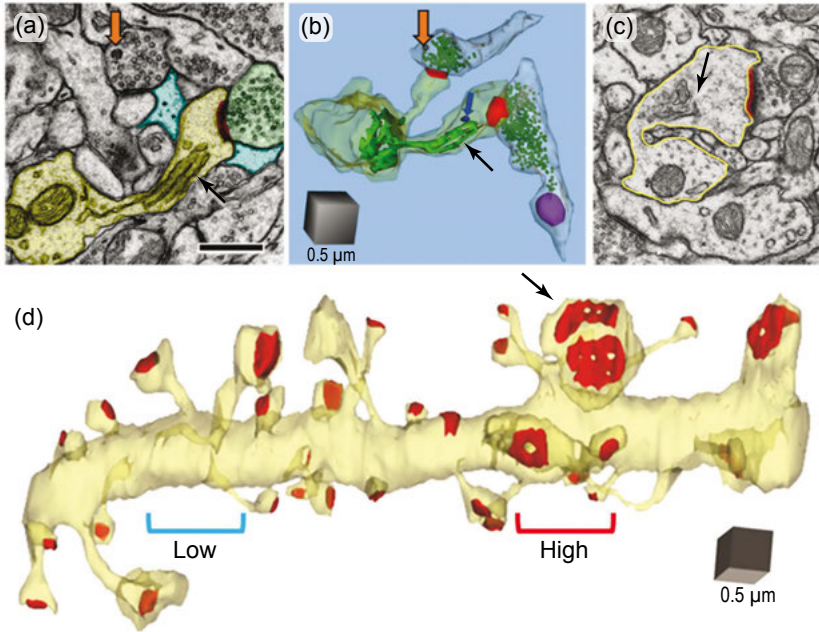


Figure 9.1 How do nonuniform distributions of subcellular resources in dendrites and axons influence where synaptic growth, maturation, and plasticity occur? (a) Single section electron micrograph and (b) three-dimensional reconstruction from serial section electron microscopy (3DEM) reveal extremes in the diversity of synapse size and composition between neighboring synapses on the same dendrite with different pre-synaptic partners. Dendrite (yellow), axon (translucent green), glia (blue), postsynaptic density (PSD, red), smooth endoplasmic reticulum forming a spine apparatus (black arrow), presynaptic dense core vesicle (orange arrow), large vesicle (blue arrow), and presynaptic mitochondrion (purple). (c) Single section through a spine with a spine apparatus (black arrow) and (d) 3DEM of the dendritic segment showing how synapses cluster even along short dendritic segments. The largest spine (arrow) along this segment contains a spine apparatus and the density of dendritic spine synapses surrounding it is high relative to other regions of the same length, where density is low. Scale bar is 0.5 μm .

by external factors, and the coordination of a dendritic tree's branches to issue action potentials in the parent cell.

Synaptic Specificity

The basic structure of dendritic spines and synaptic densities has been long known, as early electron microscopy studies unveiled the basic structural microcomponents of neural connections (Guillery and Ralston 1964). However, what has come as a surprise in recent years is the dynamic regulation and fine-tuning of these connections. Local processes, governed by complex genetic networks and shaped by electrochemical activity patterns, continually adjust

the positions, strengths, and types of synaptic connections. Ultimately, these adjustments determine, at any moment in time, how individual postsynaptic neurons integrate inputs from varied sources and ultimately issue action potentials. These transformations, in turn, define the temporally precise analog computations performed by a local patch of the cerebral cortex.

With the rise of genetic tagging of circuit elements in the mouse, researchers are moving quickly to understand the nature of synaptic specificity in a cortical column. Perhaps the best example of progress in this area pertains to the mechanistic role of various interneuron subtypes in the cortex. Inhibitory interneurons have long been recognized as important, morphologically diverse elements that serve to balance runaway pyramidal cell excitation. However, it is only in the past two decades that their molecular signatures have allowed for in-depth study. One particularly important finding has been that, unlike their pyramidal cell partners, they do not stem from a cortical origin but rather migrate tangentially into the cortical plate along multiple routes (Anderson et al. 1997). More recently, the pattern of genetically specified synaptic contacts of different interneuron subclasses has been elaborated in great detail. The spatial distribution of interneuron inputs along a cortical pyramidal cell is strongly specified by the interneuron subclass, and by extension its developmental origin and epigenetic state. Different classes of interneurons participate in blanket weak inhibition, targeted strong inhibition, and disinhibition (Kepecs and Fishell 2014). To a first approximation, these compartmentalized GABAergic synapses in the cortex are contributed by local interneurons (with some exceptions, such as the long-range GABAergic projection neurons originating in the basal forebrain). Much current work is attempting to establish when and how genetically specified interneuron subclasses find their cortical positions during early development as well as how synapse formation is regulated based on neural activity and experience (Wamsley and Fishell 2017).

By contrast, excitatory pyramidal cells receive hundreds or thousands of excitatory inputs from a varied combination of local and remote neurons. Local interconnectivity among cortical pyramidal neurons is relatively sparse, with only a tiny fraction of connections showing strong excitatory input (Lefort et al. 2009). It is important to note that neurons whose axonal and dendritic arborizations show a high degree of three-dimensional spatial overlap need not be interconnected, as the principles of pyramidal cell connectivity depend on more than spatial proximity (Mishchenko et al. 2010). While much has been learned about how cells find one another and make connections, the ultimate determinants of arborization for individual neurons or classes of neurons remain mysterious (Narayanan et al. 2017; Han et al. 2018). One fascinating observation is that growing axons appear to select their synaptic partners, in some cases, based on the projection *target* of a potential postsynaptic neuron. Such selective targeting has the interesting consequence that neurons with the same output targets tend to gather similar types of axonal inputs, and thus share their physiological response properties.

Neurons are very particular in their connections, with researchers gradually amassing a complex set of rules describing how, where, and when neural subclasses form synapses with one another in the cortical microcolumn. Let us thus take a closer look at the structure of the dendritic microenvironment, where recent experiments highlight the exquisite cellular mechanisms that enable dendrites to serve as the microscopic engines of computation and learning.

Smart Dendrites

Dendrites, like many of the brain's elemental structures, were initially misunderstood to be simpler than they are. Dendrites were long thought to act by computing the weighted sum of proximal and distal synaptic inputs through passive electrotonic conduction. The potentials detected at the cell body then determined the digital firing of action potentials of the neuron. However, research over the past years has demonstrated that this passive and capacitive view of dendrites is inaccurate. First, cortical dendrites are replete with active currents that propagate action potentials. Dendritic action potentials are sometimes generated locally and sometimes propagated backwards from the soma. Their discovery revealed a new dimension for how neural signals are integrated. Furthermore, the precise morphology and compartmentalization of the local dendritic microenvironment can strongly affect dendritic function, with recent work showing that these aspects of dendritic structure are constantly under renovation (Bourne and Harris 2012). Over time, the invisible hand of experience-dependent learning actively remodels local spine morphology: it adjusts synaptic strength and influences the postsynaptic neuron through multiple chemical and genetic pathways. Thus dendrites are now recognized as a bed of neural computation that far exceeds what was originally envisioned through the rules of electrotonic conduction, initially conceived by Rall (London and Häusser 2005).

One important principle of dendritic organization appears to be structural and functional optimization, which goes a long way in accounting for dendritic and axonal shapes and lengths (Chklovskii 2004). Dendrites optimize the amount of resources, such as cable length, and optimally enforce short conduction times and efficient current transfer from synaptic signals toward the soma (Cuntz et al. 2010). The close relationship between anatomical features and principles of connectivity serves as the basis for a large number of compartmental models that accurately account for neuronal electrophysiology (Hines et al. 2004).

Electron microscopy has offered deeper insights into the complex cell biology of dendrite remodeling (Bourne and Harris 2012). Creating spines, selecting axonal partners, and adjusting synaptic strengths all require a systematic redistribution of local subcellular resources, including plasma membrane,

structural and metabolic molecules, and cellular organelles such as mitochondria. Ultimately, it is through the parallel microscale adjustment of these elements within trillions of subcellular microcosms that the brain is able continually to tune and update its analog computations to support flexible cognitive and executive functions.

One surprise from recent years is that connections on the distal dendrites can be just as effective in driving the postsynaptic cell as those on proximal dendrites, contrary to conventional wisdom (Bromer et al. 2018). This may be for the simple reason that dendrites can actively adjust the strength of a synapse in any location, easily overriding the natural biases due to electrotonic conduction in a canonical dendrite model. In fact, investigations into this matter using electron microscopy (EM) suggest that synapses at distal dendrites are systematically larger than those near the soma. A look at a dendritic segment through EM reveals an uneven distribution of synapses of all sizes, which again is thought to reflect the specificity of synaptic connections (Figure 9.1). However, despite the nonuniformity of cellular resources, systematic investigation reveals a tight relationship between the size of a postsynaptic surface and the number of presynaptic vesicles and the presence or absence of a presynaptic mitochondrion (Figures 9.1a, b). The rules governing synaptic modification are not well understood, and those rules that have been well characterized pertain only to a subset of synapses. In the hippocampus, it has been observed that approximately 5% of spines are eligible to undergo changes in their synaptic size over time. These regional differences in plasticity speak further to the specificity of interconnections and may become important as we learn more about the learning principles that govern changes in circuit operation.

Finally, the spines themselves, which are abundant on many cortical pyramidal cells, are also highly structured, varied, and subject to morphological change (Bailey et al. 2015). Some large spines contain a so-called *spine apparatus*—an organelle involved in calcium regulation, protein and lipid trafficking, and posttranslational modification of proteins. Small spines lack this apparatus and have fewer of these resources. Some small or large spines contain a presynaptic dense core vesicle known to transport active zone proteins and vesicles between synapses. It is also important to point out that while spines have some degree of independence, the size/resource principle can extend beyond individual spines. For instance, on a given dendrite, the density of synapses is higher surrounding a large spine containing a spine apparatus (Figure 9.1c), compared to other regions where there is no proximal spine apparatus (Figure 9.1d). Future investigation should reveal the extent to which local subcellular resource allocation is a general principle that determines where synapses form, stabilize, and undergo plasticity across dendritic arbors, cell types, and brain networks. Ultimately, refined markers are needed to determine how synapses, when activated, actively redistribute resources during behavior and learning.

Circuit Motifs

With the collection of myriad observations about individual synapses, dendritic environments, and circuit contexts, the scientific community has come to discern recurring patterns or “motifs” which present themselves commonly. These motifs often refer to patterns of connectivity that specify certain physiological computations and, in some cases, extend to larger principles of cortical architecture and operation. These advances have been facilitated greatly, albeit not exclusively, through the advent of paired recording studies *in vitro* and transgenic techniques for targeting specific cell types, the latter driven primarily by the genetic tractability of circuits in the mouse. A now classic example, found throughout diverse parts of the nervous system, is the strong feedforward inhibitory circuit, which has been studied extensively *in vitro*, *in vivo*, and *in silico* (reviewed in Bruno 2011). The key ingredients of this motif are that a group of presynaptic neurons excite a downstream population of interconnected excitatory neurons and inhibitory neurons but provide greater drive to the inhibitory population (Figure 9.2a). Strong disynaptic inhibition then favors the propagation of signals encoded in the synchrony, rather than the absolute firing rates, of the presynaptic neurons. Another example that has gained much attention in recent years is a disinhibitory motif whereby excitatory synaptic input or nicotinic modulation of vasoactive intestinal peptide (VIP) inhibitory cells is effectively able to activate layer 2/3 pyramidal neurons via suppression of somatostatin inhibitory cells (Figure 9.2b). This motif has now been implicated in state-dependent modulation of sensation as well as in learning (Letzkus et al. 2011; Lee et al. 2013b; Pfeffer et al. 2013; Kepecs and Fishell 2014).

Consistent architectural features of columnar circuitry in mammalian neocortex have also become increasingly clear. These have been examined most closely in rodent sensory cortex (Figure 9.2c) but several aspects generalize across species and cortical areas. For instance, primary thalamic relay nuclei of at least the visual, somatosensory, auditory, and motor systems bifurcate to arborize in the middle (largely intracortical) layers and more sparsely at the border of layers 5 and 6 (both intracortically and subcortically projecting layers) in all mammals. This may allow some functional independence of the upper versus deep layers (Constantinople and Bruno 2013; Pluta et al. 2015). In contrast, secondary thalamic nuclei mainly innervate layer 5A and layer 1, which is also a target of other cortical regions and may exploit dendritic nonlinearity to enable top-down control (Larkum 2013). Layer 6 provides corticothalamic feedback but also targets excitatory and inhibitory cells of other layers and has been suggested as an important means to control circuit gain (Olsen et al. 2012; Vélez-Fort et al. 2014). Pyramidal neurons in layers 2/3 and 5 have extensive interconnections (omitted from Figure 9.2c for clarity). The degree to which these excitatory networks comprise motifs that depend on more than first-order connectivity statistics is a major area of active research.

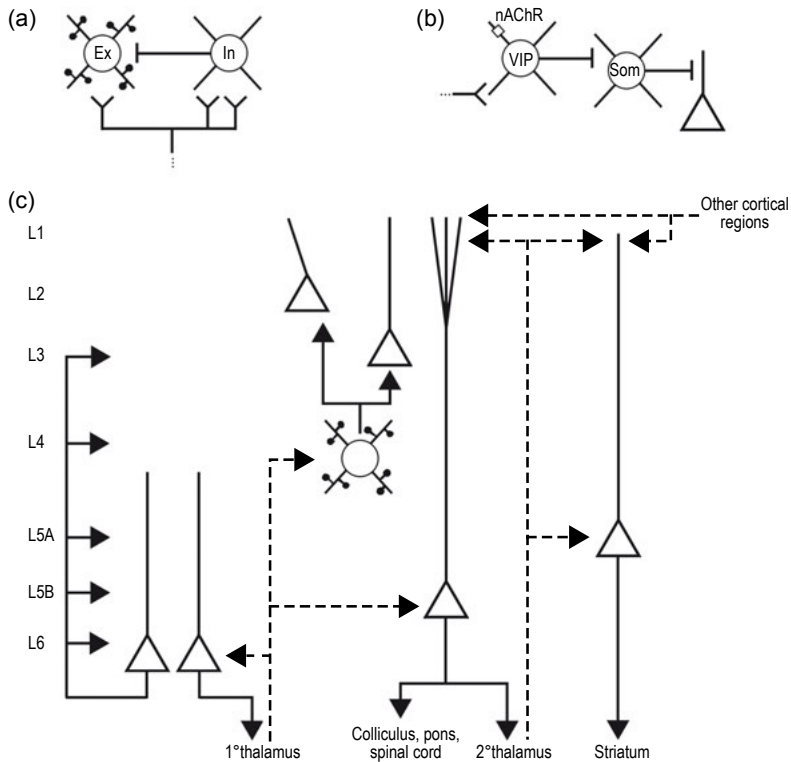


Figure 9.2 Examples of cortical circuit motifs. (a) Schematic of *feedforward inhibition*, where incoming information impinges in parallel on excitatory (Ex) neurons and local interneurons inhibiting the excitatory neurons (In). (b) Schematic of *disinhibition*, where excitatory input or neuromodulation stimulates one family of inhibitory interneurons (VIP), which inhibits another family of inhibitory interneurons (Som), synapsing on a pyramidal neuron, altogether resulting in the excitation of the pyramidal neuron. (c) Canonical patterns of input and output as well as arborization to different cortical layers.

Summary

The connections among neurons are at the heart of understanding information processing in the cerebral cortex. The twentieth-century metaphor of a *wiring diagram* fails to take into account the remarkable complexity surrounding synapses, dendrites, and functional specificity among genetically specified cells comprising circuit motifs. In some ways, synaptic microenvironments are more like living ecosystems, in which a panoply of neurites, organelles, and genetic instructions cooperate and sometimes compete for resources. At the same time, an understanding of the cerebral cortex cannot

rely solely on microscopic structure and function, since much of the brain's architecture hinges on evolved cortical areas, long-range connections, and interplay with more primitive subcortical structures. Together, their operation as a coordinated, single entity is at the heart of brain function. We next discuss the concept of large-scale functional connections, specifically how they are investigated and summarized as organized networks distributed across the cerebral cortex.

Functional Connections in a Restless Brain

Investigating the large-scale organization of the brain involves a different set of questions and tools, and dates back to the nineteenth century. Most of the early neuroanatomists assessed cortical organization through postmortem methods, revealing histological subdivisions (Brodman 1909), gross fiber bundles (Curran 1909), and specific patterns of fiber degeneration following a lesion (Nauta and Gygax 1954). Through painstaking work, a portrait of the brain's anatomical connective skeleton gradually took form and, for the cortex, is perhaps best summarized in the diagrams constructed by Felleman and Van Essen (1991). These diagrams put forth a hierarchy of cortical areas based in part on laminar patterns of inter-areal projections.

A handful of early physiological studies also appropriated the brain's signaling capacity to study its large-scale organization. For example, in chemical neuronography, small amounts of the neuroactive agent strychnine were applied to a given cortical site in an experimental animal instrumented with large electrocorticography arrays. The depolarizing action of the strychnine led to voltage deflections in a subset of cortical surface electrodes, thus revealing which areas received axonal connections from the stimulated site (Pribram and MacLean 1953). Chemical stimulation was gradually replaced by "electroanatomy," where the effects of electrical stimulation at one location were assessed at other locations across the brain (Miller and Bloomfield 1983). Together with the systematic investigation of brain circuitry through lesions and electrophysiological recordings, the concept of *functional anatomy* gradually emerged, placing emphasis on the large-scale organizational principles of networks in the brain, and particularly the cerebral cortex.

Structured Spontaneous Activity

An unexpectedly fruitful source of neural signals with which to study brain organization has been spontaneous activity. Traditionally viewed as a nuisance background signal, little attention was paid to the spatial organization of spontaneous activity until the 1990s. Then, two different brain imaging methods abruptly increased neuroscientists' respect for this ongoing background activity: brain imaging in anesthetized animals and in awake humans. Optical

imaging methods in the anesthetized cat showed that spatial patterns of spontaneous activity in the visual cortex followed the pattern of orientation preferences present in the local functional architecture (Arieli et al. 1995; Kenet et al. 2003). Around the same time, human fMRI studies demonstrated that spontaneous hemodynamic fluctuations in subjects, in the absence of any task, showed correlated activity within established functional networks (Biswal et al. 1995). The brain's ongoing signals, it appeared, could be harnessed as a tool for studying the layout of its functional networks. In the two decades that followed, this approach contributed significantly to, and in some ways even came to dominate, the study of the human brain.

Analyzing spontaneous activity forces researchers to depart from conventional experimental paradigms, in which brain responses are typically locked in time to stimuli or actions. During the resting state, brain organization is characterized in terms of the internal statistical dependencies of neurons or voxels at different spatial positions. In the simplest case, this involves computing the temporal correlation of a signal measured at one location with all other simultaneously recorded locations, rendering a brain-wide map. Tools to formalize terms and concepts related to neural interactions were initially developed in the context of single unit electrophysiology (Gerstein and Aertsen 1985). This formalism, summarized and expanded by Friston et al. (1995), set the stage for thousands of future neuroimagers to study what is now termed *functional connectivity*. Broadly defined, functional connectivity is the statistical relationship between the dynamic neural activity measured in two or more parts of the brain. This statistic is sometimes computed between pairs of points, but can also be evaluated for many areas in parallel using data-driven methods, such as independent component analysis (Smith et al. 2013). It is important to note that the relationship between functional and anatomical connectivity is complex and often underdetermined, particularly when the functional signal is assessed through an indirect measure such as blood-based hemodynamic responses. Nonetheless, the emergence of functional connectivity in the fMRI field has revolutionized the study of the human brain, by first establishing the basic correlations between related areas and then offering a new way to visualize and study brain networks. These methods and descriptions currently play a fundamental role in research into the human brain, including its dysfunction in psychiatric and neurological disorders.

A fascinating aspect of spontaneous activity, one that has drawn additional attention, is its potential *effects* on normal brain operation. At the microscopic level, ongoing activity provokes neurons to vary their responses from trial to trial. In an early observation by Bishop (1932), electrical stimulation of the optic nerve led to neural responses in the visual cortex that varied with each stimulation, an effect attributed at the time to the state of the cortex. In more recent studies in experimental animals, spatially coherent waves of ongoing activity have been shown to explain a large proportion of the response variance, even in primary sensory areas (Arieli et al. 1995; Fukushima et al. 2012).

Within the domain of human fMRI, ongoing fluctuations have been shown to alter task-based responses and directly impact perception and behavior. For example, subjects detecting faint visual stimuli are likely to be fooled into a false percept on trials in which fMRI activity in the visual cortex is high in the absence of a stimulus (Ress et al. 2000). Similarly, the reaction time of button presses is shortened during trials in which activity is high in sensorimotor areas. Large-scale fluctuations are sometimes directed by expectations or the structure of a specific task. For example, Sirotin and Das (2009) demonstrated a robust task-entrained change in hemodynamic signals in the primary visual cortex amid an unchanging visual stimulus. These results demonstrate that hemodynamic responses in sensory cortical areas can be subject to cognitive influences, such as the expectation of a stimulus. Such anticipatory modulation might be directed through long-range connections from the basal forebrain (Turchi et al. 2017), the adrenergic system (Reimer et al. 2016), the frontal cortex (Noudoost and Moore 2011), or the amygdala (Hadj-Bouziane et al. 2012), all of which have the capacity to alter responses of neurons to sensory stimuli.

Neural activity and hemodynamic fluctuations have a notoriously complex relationship (Logothetis 2008), part of which can be seen in relationship to cellular metabolism. One straightforward explanation of neurovascular coupling is that neural responses spend local energy; this, in turn, causes local metabolic increases, which summons more regional blood perfusion (Magistretti 2000; Raichle and Gusnard 2002). There are many known examples, however, in which this linear explanation fails, particularly since fluctuations can stem from physiological signals other than from neurons, including metabolic changes (Goldbeter 1996) and even hemodynamic changes themselves (Moore and Cao 2008). These multiple levels of biological complexity pose significant challenges for pinpointing the neural processes that underlie the commonly observed whole-brain correlation patterns in humans, and for understanding the causal chain by which ongoing fluctuations might influence functional responses and behavior.

Nonetheless, through the mapping of fMRI temporal correlations, the resting human brain offers an array of at least a dozen interleaved networks, many of which are straightforward to identify in the majority of subjects and appear to be in a mature form in early childhood (Damoiseaux et al. 2006; Thornburgh et al. 2017). While less studied in animals, the basic features of many of the resting-state networks appear similar (Hutchison and Everling 2012; Belcher et al. 2013). These networks have provided a new and extremely useful approach to study brain organization and physiological processes in healthy subjects and patients. For example, it is now possible to use spontaneous signals to establish the functional layout of the brain, and to use this information to define regions of interest for analysis in subsequent fMRI experiments. It is also possible to compare the functional integrity of such networks between patient groups and control subjects through straightforward resting-state scans collected in just a few minutes.

Conceptualizing and Analyzing Neural Networks

A network, broadly defined, is a complex set of interacting elements. Above we highlighted the whole-brain networks identified using fMRI. However, neuroscientists also speak of networks defined by the interactions among genes, cell types, interconnected areas, or even individual animals. At some level, each network is embedded within a superordinate network that may be impacted by its perturbation. In terms used by network theoreticians, if a given node is disrupted (e.g., by suppressing activity of a cell type within a circuit) or a given edge removed (e.g., by suppressing one particular set of connections), other network components are likely to be affected. Such “diaschisis” or “off-target effects” present a potential pitfall for scientists attempting to draw cause-and-effect conclusions through manipulations of a biological system. The remedy might be for a researcher to focus on a separate perturbation of multiple nodes and edges, perhaps under multiple conditions. To approach the problem this way, however, one needs a considerable understanding of the network structure, and this is often only achievable through an integrated computational framework. At the same time, well-developed network approaches promise to facilitate unity *across* levels of analyses, not just among components at one level. The feasibility of studies which cross levels (genetic, cellular, circuit, system, behavior) has increased tremendously, and much effort is currently being directed toward network models that link two or more levels.

As mentioned above, functional connections are sometimes characterized by their pairwise interactions, primarily because the computational methods involved are straightforward. However, interactions in real brains are much more complex, may be more difficult to assess, and are ultimately tuned at a systems level to achieve certain behaviors. Encapsulating and analyzing these dynamics is of great interest as we attempt to integrate isolated descriptions of anatomical connections or physiological measurements into principles of brain function. As in many fields of biology, fitting the components of a complex system into a rigorous mathematical framework is an immense challenge. Nonetheless, this approach has already led to fruitful insights into modes of cortical dynamics.

Because whole-brain wiring diagrams capture only particular aspects of brain connectivity, quantitative models of brain architecture have gained popularity. Quantitative models require a mathematical language capable of capturing different types of interaction in a highly interconnected network. Built on fundamental mathematics in the form of graph theory and fundamental physics in the form of statistical mechanics, network science provides exactly such a language (Albert and Barabasi 2002) (Figure 9.3). In its simplest form, network science can be used to model intricate connectivity patterns as graphs, where computational units (neurons, ensembles, cortical columns, brain areas) are represented as nodes, and connections between them (functional relations or physical links) as edges (Bassett and Sporns 2017). By modeling the system

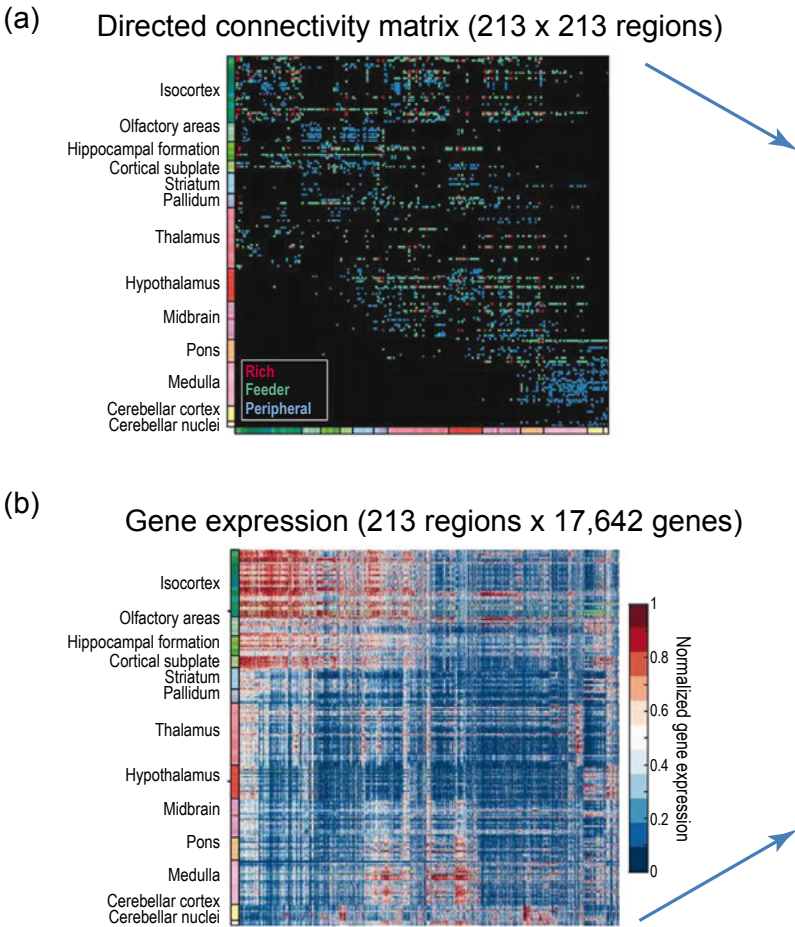


Figure 9.3 Relations among anatomical connectivity and gene co-expression networks. (a) Matrix of anatomical connections among 213 mouse brain regions: regions (nodes) with more than 44 distinct connections were considered hubs, and connections were classified as hub→hub (rich), hub→nonhub (feeder) or nonhub→nonhub (peripheral). (b) Normalized expression levels of 17,642 genes across 213 brain regions: genes with highly correlated expression profiles are placed near each other.

as a graph, one can apply computational tools to characterize quantitatively the architecture of the graph with various metrics, and then compare those metrics across measurement modalities, spatial scales, temporal scales, individuals, and species (van den Heuvel et al. 2016).

Several concepts from network science have proven useful in our understanding of neural systems. Locally, processing units tend to display strong clustering with neighboring regions, and these clusters combine to create

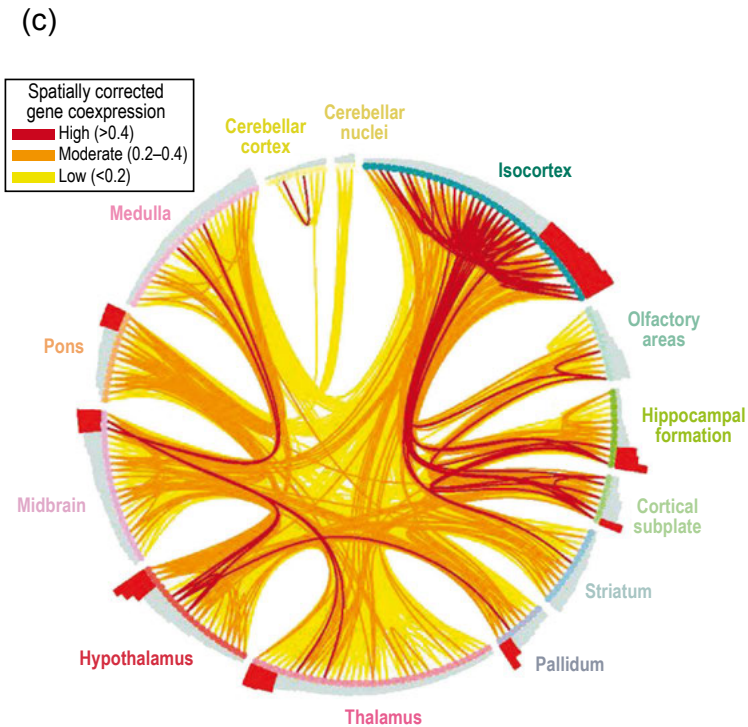


Figure 9.3 (continued) (c) Brain regions have been arranged around a circle, ordered by number of connections (bars) in each anatomical subdivision. Hubs are marked by red bars. The connection diagram traces anatomical connections between pairs of brain regions, color-coded by the corresponding gene co-expression value, after applying a correction for spatial distance. Statistical analysis revealed strongest gene co-expression among pairs of regions linked by reciprocal connections (as compared with unidirectional or unconnected pairs), as well as for rich connections linking hubs (as compared with feeder and peripheral connections). Genes driving correlations in expression in connections involving hub regions are functionally enriched in oxidative energy metabolism. Connectivity data derived from Bassett and Sporns (2017).

modules (Sporns and Betzel 2016). Each module is composed of a set of nodes that are more densely interconnected to one another than they are to nodes in other areas of the network. Complementing the local clustering, strong long-distance connections exist to link diverse areas of the network and enhance the complexity of functional dynamics (Betzel and Bassett 2018). The combination of local clustering and a few long-distance connections produces a small-world network topology, which intuitively can support segregated processing

of information in combination with information transmission to spatially disparate units (Bassett and Bullmore 2017). Within this modular, small-world topology, one also observes network hubs—nodes that display an unexpectedly high number of edges—that are often connected to one another, and which are thought to be capable of exerting a particularly salient influence on the system (van den Heuvel and Sporns 2013). Recent work bridging network neuroscience and control theory has identified additional node types thought to be capable of enacting diverse control strategies, altering system dynamics to support cognitive function (Gu et al. 2015; Kim et al. 2018). Other complementary work that bridges network neuroscience and applied algebraic topology has identified additional network motifs that capture higher-order interactions, which may also be particularly important for neural computations (Giusti et al. 2016). These and related efforts have demonstrated promise in understanding (and quantitatively characterizing) alterations in network organization that accompany neurological disease and psychiatric disorders (Stam 2014; Fornito et al. 2017).

One of the advances over the past three decades has been in how we think about neural networks—progress that has resulted, in part, from improvements in computational methods and speed. What has changed at a conceptual level is the linearity, or perhaps seriality, in thinking about brain operation. This change is manifest at the level of the synapse, with highly selective, multidirectional, and dynamic analog interactions in the neuropil. This has also altered how we look at whole-brain networks more broadly, characterizing them not within the framework of a climbing hierarchy but rather as a syncytium of interconnected areas. There may be some historical irony here, in that the emerging network conceptualization of the brain may resonate somewhat more with Golgi's original syncytium view of the brain's organization than Cajal's revered neuron doctrine.

The Far Reach of Cerebral Control

The exquisite spatial organization of spontaneous signals in the cortex came as a surprise to many systems neuroscientists. As described above, the discovery was made only after stepping away from conventional paradigms designed to tap into the presumed sensory, cognitive, and motor functions of the cerebral cortex. In some ways, this finding fits with the long-known fact that much of the brain's activity is concerned with internal regulation and homeostasis rather than interaction with the external environment. Beyond homeostasis, internal signaling can also shape a range of behaviors by adjusting state parameters that require action. For example, small groups of cells in the hypothalamus and elsewhere act through the endocrine and descending autonomic systems, utilizing bidirectional communication with visceral organs in the regulation of feeding, sexual behavior, and other actions critical for survival.

Subcortical brain structures have been shown to regulate aspects of immune functions (Wrona 2006), body metabolism (Morton et al. 2006), and perhaps the gut microbiome (Foster et al. 2017).

The surprising thing that we have learned in recent years is that the *cerebral cortex* has a hand in controlling autonomic and visceral function. The polysynaptic cortical control over these structures was determined using injections of rabies virus into the end organs. Rabies is a retrograde virus that crosses synaptic connections and can thus be used to identify higher-order upstream neurons, including those in the cerebral cortex (Dum and Strick 2013).

Several studies have demonstrated that restricted regions of the motor cortex and a handful of other cortical areas hold reign over stations in the peripheral nervous system as well as over visceral organs (e.g., kidney, adrenal medulla, stomach, heart) and likely many other corporeal structures (Levinthal and Strick 2012; Dum et al. 2016). After several synaptic crossings, typically four, through sympathetic or parasympathetic chains, the retrogradely transported rabies virus reaches layer 5 neurons of highly specific and circumscribed cortical regions. In the motor cortex, these regions occupy portions of the homunculus map of the body that are commensurate with their possible roles. For example, one of the sites projecting to the adrenal medulla is present in the face area of the motor cortex, which may be related to the regulation of sympathetic responses elicited in concert with facial emotions. These findings offer an entirely new perspective on cortical anatomy and may have potential implications for clinical work, where it is known, for example, that stimulation or transection of the vagus nerve can relieve symptoms that might originate from cortical dysfunction.

Summary

Over the last three decades, the appropriation of spontaneous functional signals in the study of brain organization has opened new doors for the study of cortical organization, particularly in humans. It has also launched multidisciplinary approaches for the analysis of cortical activity, gaining insights from disciplines accustomed to making sense of complex interactions in large-scale networks. At the same time, neuroscientists remain aware that the cerebral cortex is not a computer in isolation, but is integrated within a biological system that includes the body itself, and can exert descending control over endocrine, immune, and even visceral functions.

Design Principles Promoting Flexible Behavior

The capacity to learn, adjust, and flexibly direct behavior constitutes a very important architectural design consideration of the cerebral cortex. These

capacities derive from a multiplicity of mechanisms that are manifest simultaneously at multiple scales, from the synapse to the system. Many of the mechanisms for neural plasticity found in the adult resemble those that were involved in the initial building of the brain and may actually extend developmental windows that permit lifelong modification of neural circuits.

Essential Early-Life Training of the Cerebral Cortex

The cerebral cortex has a protracted period of growth and maturation compared to many brain structures outside the telecephalon, which reflects one of its fundamental design principles: it needs training. The genetic specification of cortical development is staggering and strongly conserved across mammals (Workman et al. 2013), with the most prominent difference related to factors such as brain size and peripheral sensory specialization. Built into that genetic program, however, are explicit mechanisms for the postnatal refinement of cortical connections guided by activity and experience. These developmental steps are critical if we are to understand how the cerebral cortex of a particular mammal (e.g., a human) takes form. Early-life plasticity, through typical sensory experiences, ecological constraints, and parental relationships, has the effect of shaping species-specific sensory inputs, cortical connections, and processing domains.

Interestingly, some of the relevant behaviors that drive cortical training originate not in the cortex itself, but rather in the early-developing control centers of behavior which reside in the hypothalamus and midbrain (Swanson 2000). Thus, early-life plasticity may stem from one part of the brain training another, with innately programmed subcortical structures driving experiences needed for the normal development and maturation of the cerebral cortex. Among primates, whose extended childhood offers much time for experience to mold the adult patterns of cortical connections, experience-dependent learning is particularly obvious in the domain of complex social interaction.

The postnatal shaping of the cerebral cortex involves multiple mechanisms, and builds in an initial exaggeration, or exuberance, of axonal projections followed by a gradual pruning and restriction to their adult target locations. One measurable example of this exuberance in primates is the initial overproduction of interhemispheric fibers passing through the corpus callosum. In rhesus monkeys, for example, less than one third of the interhemispheric fibers present at birth persist into adulthood (LaMantia and Rakic 1990). For other corticocortical connections, less is known about the principles that underlie the elimination of axons and synapses during early life, particularly in primates. Recent work, however, has tracked the age-dependent expression of other developmental markers; for example, those related to myelination and structural molecules such as neurofilament protein in the primate visual cortex (Mundinano et al. 2015). These studies illustrate that the dorsal, parietal areas

mature more quickly than the ventral, temporal areas, and that early maturation of the cerebral cortex may be driven by a transient thalamic connection that is present only during early development (Mundinano et al. 2018).

The important message that we wish to stress is that the cerebral cortex is, from the beginning, critically shaped by early sensory experiences and behaviors. Many of these behaviors are primitive or innate, driven by specialized subcortical circuits. Subject to the sensory and social consequences of these innate actions, the cerebral cortex prunes and refines its connections and steadily takes control over many overt and internal behaviors. The cortex learns to interpret complex sensory signals, to establish contingencies, to withhold reactions and reflexes, and to plan goal-oriented action sequences critical for survival. This learning is an essential element of cerebral cortex design and may ultimately be as important as the adult wiring diagram for understanding its core principles.

Learning in the Adult: Plasticity of Synapses and Systems

Following the extreme plasticity evident during development, the adult brain retains a critical ability to alter itself in response to the environment; this serves as the basis for learning and memory. Above, we provided an overview of the complex and interacting cellular substructures involved in the tuning of neural connections. Over the last decade, researchers have discovered myriad modes of synaptic regulation stemming from mechanisms that govern changes to the chemical, genetic, and structural composition of synaptic connections (Alberini 2009; Holtmaat and Svoboda 2009; Bailey et al. 2015).

At one level of description, synapses can be seen as independent actors, regulating their potency independently of their neighbors. The phenomenon of long-term potentiation has long been known (Bliss and Lomo 1973), and the list of factors regulating the strength of individual synapses is ever growing (Malenka and Nicoll 1999). At the same time, a broader view of the synaptic microcosm, highlighted in an earlier section, indicates that efficacy of a given synapse is highly dependent on extrasynaptic factors. In other words, altering the morphology or synaptic strength at one location can affect the contribution of neighboring synapses. This partial dependency of nearby sites on the dendritic shaft, together with other active features of dendritic physiology, provides many degrees of freedom for fine-tuning circuit plasticity. While learning principles are still being discovered, recent work suggests that even extrinsic cells, other than presynaptic and postsynaptic neurons, play important roles in shaping the dendritic microenvironment. These cells range from microglia, which contribute to synaptic modification by actively and aggressively removing spines, to particular subtypes of inhibitory interneurons, whose level of input may be important to shift neurons into a mode that is receptive to synaptic plasticity (Bavelier et al. 2010).

Under some conditions, opening the window to synaptic modification can resemble a local reversion to a more immature state of brain development. While this is a somewhat new field of study, the genes thought to be involved in this process are sometimes called *neotenus*, referring to their prominent role in early life. Such developmental genes have been found to co-localize with other markers of brain activity, such as cortical areas undergoing aerobic glycolysis (Goyal et al. 2014). As we learn more about the regulation of plasticity through physiological, chemical, and genetic mechanisms, opportunities may arrive to improve plasticity for a range of cognitive and neurological disorders.

At a more holistic level, the cerebral cortex participates in multiple systems that appear designed to adapt behavior through various types of learning. This also involves cortical connections to external brain structures. Although this is a very broad topic, brief mention is warranted here since so much of the cerebral cortex participates in these adaptive systems. The two most prominent regions working with the cortex to facilitate behavioral adaptation are the basal ganglia and cerebellum (Bostan and Strick 2018). These structures are anatomically interconnected to form an integrated network capable of adaptation over different timescales and under different contingencies. Often overlooked, this network is topographically organized so that motor, cognitive, and affective territories at each node are interconnected with the corresponding territory of another node. During a particular task, the interlinked nodes are coactivated, and during learning there is an orderly shift in the progression of learning between the different structures as the performance level changes. Within each structure, activation in cognitive territories often predominates when a task is first performed. Thereafter, learning is expressed in motor territories, where changes emerge simultaneously with improvements in motor performance. Figure 9.4 illustrates the temporal evolution of the involvement of different networks in behavior. It is typical of the changing involvement of corticocortical networks (e.g., default mode network, dorsal attention network) in various tasks. The conserved connectivity of the cerebral cortex with external elements such as the basal ganglia and cerebellum allow the adult brain of mammals to learn and adapt their behaviors. This general mammalian pattern has been expanded in primates and is particularly prominent in the massive human brain, conferring a remarkable level of flexibility in adult behavior. Unlike any other animal, a human being can participate in learned activities as varied as acrobatics, music, politics, and typing on a keyboard. This flexibility stems from the multiscale organization of learning mechanisms in the brain, from plasticity in the dendritic microenvironment to whole-brain circuits specialized for learning. Through its interconnection with areas such as the basal ganglia and cerebellum, the cerebral cortex learns to control myriad aspects of behavior, conferring a centrality to our actions and thoughts, the expression of which, through consciousness, is one of the most fascinating and elusive aspects of brain science.

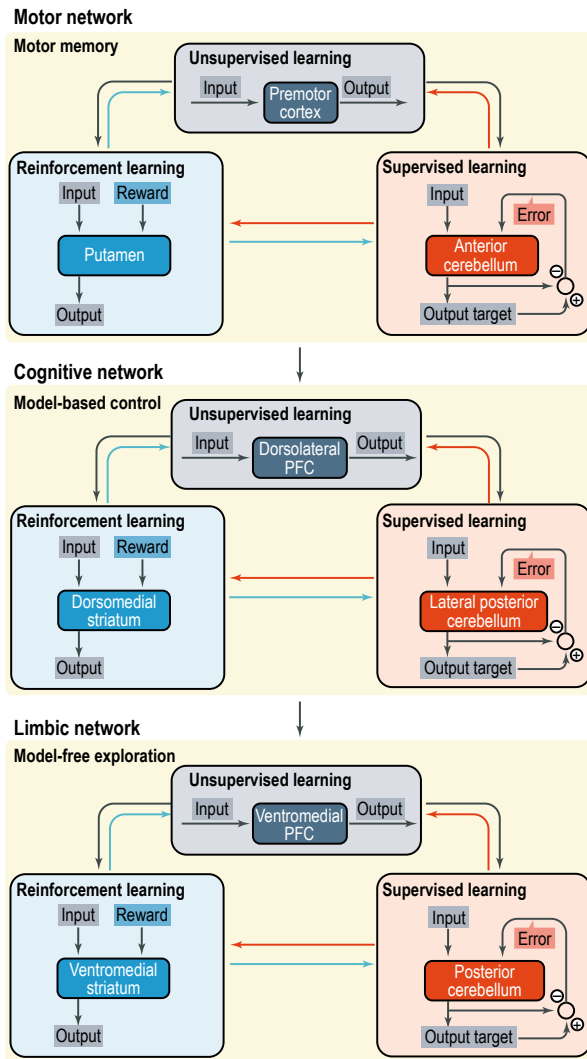


Figure 9.4 Functionally related cortical, basal ganglia, and cerebellar sites within interconnected networks participate in progressive stages of action planning. On the basis of these results, learning through exploration involves a limbic network, including the ventromedial prefrontal cortex (PFC), ventromedial striatum, and posterior cerebellum. Model-based learning involves an associative (cognitive) network, including the dorsolateral PFC, dorsomedial striatum, and lateral posterior cerebellum. Performance based on motor memory involves a motor network, including the supplementary motor areas, putamen, and anterior cerebellum. The authors’ imaging data suggest that as learning progresses, the sites of activation shift in a topographically organized fashion. Our interpretation of these data is that each stage of the learning process involves a different set of interconnected basal ganglia, cerebellar, and cerebral cortical regions. Reprinted with permission from Bostan and Strick (2018).

Summary

Without the capacity to learn and adapt its behavior, a brain would be unsuccessful in a complex and contingent world. Many if not most interactions among animals, including predation, foraging, social interactions, and navigation, require continual learning and modification of behavior. The cerebral cortex is at the heart of this flexibility. In early life, its neural circuits undergo major experience-dependent changes as they tune themselves to the basic statistics and behavioral requirements of the environment. In adulthood, the brain retains the capacity to learn and adapt, facilitated both by the capacity to reconfigure cortical synaptic connections, as well as the utilization of cortical connections to subcortical structures specialized to support behavioral flexibility.

Conclusions

The functional architecture of the cerebral cortex is a formidable topic and its details could fill multiple volumes. In this chapter we have highlighted three areas that have seen particular conceptual advancement over the past decades: the synaptic microenvironment, large-scale cortical networks, and plasticity and adaptation in cortical circuits. Many other important features of cortical architecture include the prominent laminar organization of cells in cortical microcircuits, the growth and initial wiring of the brain, the functional consequences of its columnar organization, and its principles of connectivity with subcortical structures, such as the thalamus, striatum, claustrum, and superior colliculus. A growing body of exciting findings has investigated functional specificity through genetically modified mice, as well as specificity of interneurons and circuit motifs. Details from this burgeoning field will continue to shape our understanding of functional architecture at both microscopic and macroscopic scales. One of the great challenges in the study of the brain is to synthesize a large number of details into principles for understanding. While many details are known about the cerebral cortex, our level of understanding about its overarching architectural and functional principles remains, arguably, primitive. For the optimist, this is a situation of great opportunity, where any scientist who is able to identify and integrate the most relevant of these details will reap the benefits of fundamentally new insights into brain function.

There is no denying that remarkable conceptual progress has been made over the last decades. Reflecting back to the first meeting in this series in 1987, the community was largely familiar with the basic connective anatomy of the cerebral cortex. Cortical anatomists and physiologists pictured the brain as a step-by-step sequence of hierarchical computations, in some ways captured by the newly available information codified in the VLSI diagram of

the cerebral cortex (Felleman and Van Essen 1991). There was also the general awareness that perhaps superimposed on that hierarchy were two distinct streams for visual information processing: one concerned with objects, the other with locations (Ungerleider and Mishkin 1982). However, there was not yet a community of brain imagers, outside of a few specialists in positron emission tomography. At the top of the cortical hierarchy sat the hippocampus alone, which, perhaps coincidentally, was the target of study for most synaptic physiologists.

Our current shared understanding of the cerebral cortex has changed considerably, particularly in reference to the richness of synaptic interconnections and the large-scale organization of networks. The view of a serial processing hierarchy has been complemented by the concept of interacting and dynamic functional networks. As is often the case, the scientific conquests have had the effect of raising new questions at a rate that seems to exceed our steps forward in understanding. The staggering biological complexity of the brain is not for the weak willed, and neuroscientists struggle to gain traction on basic architectural and functional principles.

One critical element for the future is almost certainly the continued development and refinement of computational and theoretical tools. Within the realm of computation, there is a spectrum of detail built into different models, ranging from realistic, elemental models of the brain, to mathematical encapsulations of brain function, to deep learning approaches that only marginally reflect brain operation. An important question for the future is how closely to tune modeling efforts to the empirical details of brain anatomy and physiology as they are continually revealed. Perhaps even more important is the philosophical question of where the essence of brain function, cognition, and behavior lies. Is there a primacy of computational descriptions, as one might conclude from visionary thinkers such as David Marr (1982)? Or does focusing on computation while ignoring biological details come at a cost that is too heavy to bear? Theorists will have to grapple with these questions as both computation power and the capacity to collect empirical data about the brain accelerate.

A complementary perspective on cognition and behavior is that structural and functional principles of the brain can only really be understood through genetics, development, and evolution. From this perspective, the human brain is a fundamentally composite structure, with evolved layers of control subsuming and overriding more primitive control systems. A truly mechanistic understanding will therefore depend on disentangling ancestral versus derived features of the brain, including its multiscale structural and functional details, its layers of genetic and molecular control, its growth and refinement during embryonic and postnatal development, and much more. While this pursuit will also draw upon computational tools, it pushes neuroscientists to embrace the natural biological complexity that has continually shaped the human brain, and its massive and powerful cerebral cortex, over hundreds of millions of years.

Acknowledgment

This research was supported in part by the Intramural Research Program of the NIMH (ZIA MH002838).