6

The Position of the Prefrontal Cortex in the Cortical Hierarchy

John D. Murray and Christos Constantinidis

Abstract

The prefrontal cortex (PFC) exerts control on the flow of sensory information in cortical circuits, integrates current stimulus streams with stored memories, and plans motor action. Prefrontal neurons exhibit quantitatively distinct firing properties relative to its afferent inputs. These can be traced to unique anatomical morphology, neurotransmitter receptor composition, and relative distribution of different interneuron types. This evidence suggests a position of PFC on the top of the cortical hierarchy that processes sensory information and controls behavior. A functional specialization is also present within the PFC, as it comprises multiple areas that are hierarchically organized. Other brain structures exert influence on PFC activity critical for the control of behavior, including the thalamus and neuromodulator systems. In that sense, PFC is a critical node of the broader circuit that instantiates intelligent behavior.

Introduction

Our understanding of cortical function has been shaped by the hierarchical processing of sensory systems. The visual cortex provides a prototype of this organization (Felleman and Van Essen 1991): neurons in the primary visual cortex with small receptive fields respond to elementary properties of the visual scene. Inputs from multiple V1 neurons are integrated into higher cortical areas to extract progressively more complex properties of images, over larger parts of the visual field, ultimately allowing objects to be segmented, identified, and categorized. Other sensory systems (auditory, somatosensory) perform similar transformations along their own hierarchies. The output of all sensory systems is ultimately propagated to the prefrontal cortex (PFC) and in this sense the PFC sits atop the processing hierarchies. PFC is also connected with a number of subcortical structures, most importantly the thalamus (via the medio-dorsal nucleus) and the basal ganglia (Giguere and Goldman-Rakic 1988; Middleton and Strick 2002). Cortico-thalamic and cortico-striatal loops

From “The Frontal Cortex: Organization, Networks, and Function,” edited by Marie T. Banich,
Suzanne N. Haber, and Trevor W. Robbins. 2024. Strüngmann Forum Reports, vol. 34,
Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262549530

are essential for the maintenance of information in working memory, one of the critical functions of PFC (Jaffe and Constantinidis 2021). In turn, PFC broadcasts top-down signals to the rest of the brain influencing sensory processing, integrating current events with stored memories, and prompting motor action (Badre and D'Esposito 2009), commonly referred to as executive function.

The Baddeley and Hitch model of working memory, which has been tremendously influential, encompasses a central executive and three subsidiary systems: the phonological loop, the visuospatial sketchpad, and episodic buffer (Baddeley 2012). Although introduced as a conceptual model, parallels between the function of the PFC and the central executive, the module responsible for the control and regulation of the other components and the ability to switch between tasks, have been highlighted. In this view, PFC sits atop of the cortical hierarchy, whereas the subsidiary systems maintain the contents of working memory at the sensory cortices (D'Esposito and Postle 2015). This division of function, however, is tenuous, as strong evidence exists for sensory information being maintained within the PFC by the same neurons that implement top-down control, thus supporting the idea that PFC is the anatomical seat of both executive and subsidiary systems of working memory (Riley and Constantinidis 2016). The concept of "mixed selectivity," which has been popularized recently, provides a vivid illustration of how deeply maintenance and executive functions are entwined: individual neurons exhibit selectivity for stimuli that differ depending on what cognitive task a subject executes (Rigotti et al. 2013). At the population level, the "representational geometry" of stimuli also changes with task demands (Bernardi et al. 2020; Cueva et al. 2020; Minxha et al. 2020; Okazawa et al. 2021). This finding also suggests that neurons within the PFC instantiate executive function by flexibly altering the type of information they represent.

In this chapter, we consider experimental evidence of the PFC performing integrative functions, the neural substrates that allow the PFC to play such a role, and the organization of the PFC itself. We will emphasize experimental evidence primarily from nonhuman primate models, as these have allowed the most detailed experimental data pertinent to this question. However, when available, parallels with human neuroanatomy and imaging will be integrated.

Prefrontal Specializations

The position of a brain area in the cortical hierarchy can be assessed based on objective criteria related to anatomical circuits and properties of neuronal firing. Prefrontal neurons do not represent additional sensory attributes relative to those already represented at the top levels of the sensory pathways; instead, they exhibit distinct properties during the maintenance of information in working memory and modulation of neuronal activity during execution of different tasks. The generation of persistent activity is thus a critical property of PFC, which differs

From "The Frontal Cortex: Organization, Networks, and Function," edited by Marie T. Banich, Suzanne N. Haber, and Trevor W. Robbins. 2024. Strüngmann Forum Reports, vol. 34, Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262549530

at least quantitively from areas connected to it (Leavitt et al. 2017), though some controversy exists around this point (Christophel et al. 2017). Understanding which underlying specializations produce these unique prefrontal properties will be instructive as to the position of the PFC in the cortical hierarchy. Here, we review elements of cortical morphology and circuitry that differ systematically in the PFC and other areas, and specifically the anatomical morphology of pyramidal neurons, myelination of axonal fibers traversing the cortex, types of interneurons, and receptor composition for different neurotransmitters and neuromodulators.

The classical view of pyramidal neurons has been that they are essentially uniform across the cortex. This idea has been challenged by experimental findings that demonstrate a systematic difference across the cortical hierarchy, with the prefrontal pyramidal neurons exhibiting the most extensive dendritic trees and the largest number of spines among cortical neurons (Elston 2000, 2003). Functional correlates of this anatomical specialization are also reflected in the patterns of neuronal discharges at different areas. Prefrontal neurons receive a greater proportion of distal synaptic inputs compared to the neurons at other brain areas, with a substantial proportion of these inputs originating at distances greater than 1 mm. By contrast, the majority of inputs to posterior parietal neurons appear to originate from neurons at shorter distances, in the order of 0.2–0.5 mm (Hart and Huk 2020; Katsuki et al. 2014). It has long been speculated that prefrontal neurons with similar memory fields are grouped in clusters with reciprocal connections, often visualized in anatomical tracer studies (Goldman-Rakic 1984; Kritzer and Goldman-Rakic 1995; Levitt et al. 1993; Pucak et al. 1996). Modeling studies suggest that more extensive networks of interconnected neurons in the PFC account for the improved stability of prefrontal persistent activity during working memory and its ability to resist distracting stimuli, compared to equivalent neural circuits of other areas (Mejias and Wang 2022).

Independent evidence of prefrontal anatomical specialization has also been provided by anatomical studies of cortical myelin content. The MRI-based T1-weighted/T2-weighted ratio (T1w/T2w) is indicative of the extent of myelin presence within gray matter (Glasser and Van Essen 2011; Huntenburg et al. 2017). The cortical map of T1w/T2w exhibits a large-scale hierarchical gradient, with high values in primary sensory cortex and low values in association areas of the cortex. In the monkey, this was found to correlate with a measure of cortical hierarchy, based on the laminar patterns of feedforward and feedback inter-areal projections (Burt et al. 2018). As a proxy measure of cortical hierarchy applied to the human brain, the T1w/T2w map was found to align with the dominant spatial pattern of transcriptomic variation in human cortex, derived from the Allen Human Brain Atlas, which reflects multiple aspects of cellular specialization across cortex (Burt et al. 2018).

Another specialization that reveals the prefrontal position in the cortical hierarchy is the concentration and composition of NMDA receptors. These

From “The Frontal Cortex: Organization, Networks, and Function,” edited by Marie T. Banich, Suzanne N. Haber, and Trevor W. Robbins. 2024. Strüngmann Forum Reports, vol. 34, Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262549530

are glutamate-gated cation channels, critical for the generation of persistent activity, as they are capable of extending the duration of the postsynaptic depolarization by virtue of their slow decay time constant (Constantinidis and Wang 2004; Wang 2001). Thus, a circuit containing exclusively AMPA receptors, which produce synaptic currents with very fast decay time constants, would require unrealistically high firing rates to sustain neural activity during the delay period of a memory task (Wang 1999). Experimental results further support the role of NMDA receptors in the generation of persistent activity, as the systemic administration of ketamine, a nonspecific NMDA antagonist, seems to decrease the effective connectivity between prefrontal neurons, demonstrated by a decrease in the synchronous spiking between simultaneously recorded neurons (Zick et al. 2018). The area-specific expression of NMDA further underlies its role in facilitating the prevalence of persistent activity in the PFC. For example, GluN2B (the NMDA receptor subunit with the slowest decay time constant) is expressed in a gradient across the primate brain, with highest levels of expression observed in the PFC. A hierarchical cortical gradient of increasing expression of the GRIN2B gene is observed in human transcriptomics (Burt et al. 2018), and a higher ratio of NR2B/NR2A in pyramidal neuron EPSPs has been observed in rat medial frontal cortex compared to primary visual cortex (Wang et al. 2008).

A further physiological signature of prefrontal cortical specialization has to do with the intrinsic timescales of neural activity. At the level of single-neuron recordings, the autocorrelation of spontaneous neuronal firing exhibits a characteristic timescale, which increases across the cortical hierarchy, with faster dynamics in sensory areas and slower dynamics in prefrontal areas (Murray et al. 2014). These timescale differences presumably reflect different properties of the local microcircuit operating regime across areas, with contributions from cellular and synaptic properties (e.g., NR2B/NR2A ratio due to different time constants associated with those NMDA subunits) and differences in net recurrent strength (e.g., stronger recurrence in PFC, as hypothesized to play a role in generation of persistent activity for working memory).

Another line of evidence suggestive of prefrontal specialization is related to inhibitory interneurons. Prefrontal interneurons exhibit persistent activity with higher baseline firing rates and broader tuning than pyramidal neurons (Constantinidis and Goldman-Rakic 2002). Their action thus serves to “sculpt” the spatial and temporal tuning of prefrontal neurons (Constantinidis et al. 2002), without which stimulus-specific persistent activity is much less viable in computational models (Compte et al. 2000). Multiple types of cortical interneurons are hypothesized to form specialized networks for the purpose of facilitating stimulus-specific persistent activity (Wang et al. 2004b). Three broad types account for the vast majority of interneurons in the cortex: those that express (a) parvalbumin (PV), (b) vasoactive intestinal peptide (VIP), which tends to co-localize with calretinin (Gabbott and Bacon 1997), and (c) somatostatin (SST), which tends to co-localize with calbindin. PV interneurons target the cell

From “The Frontal Cortex: Organization, Networks, and Function,” edited by Marie T. Banich, Suzanne N. Haber, and Trevor W. Robbins. 2024. Strüngmann Forum Reports, vol. 34, Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262549530

bodies of pyramidal neurons and, when activated by their preferred stimulus, they would tend to suppress the activation of pyramidal neurons with different spatial tuning than their own and sharpen the tuning function of those with similar tuning (Li et al. 2020b). Without feedback inhibition, recurrent excitation may shift the excitatory/inhibitory balance and bring the network into an unstable, hyperexcited state, which would be deleterious for the maintenance of working memory (Constantinidis and Wang 2004).

VIP/calretinin interneurons are thought to inhibit other types of interneurons, including SST/calbindin ones (Fish et al. 2018; Melchitzky and Lewis 2008; Meskenaite 1997). Furthermore, interneuron-targeting cells are more abundant in association cortices, particularly in the PFC, compared to the sensory cortex whereas soma-targeting PV interneurons exhibit the opposite trend (Defelipe et al. 1999; Elston and Gonzalez-Albo 2003). SST interneurons, on the other hand, are peridendritic-targeting cells and are thought to exhibit high spontaneous firing rates that may tonically inhibit all pyramidal neurons during baseline, prior to any stimulus presentation. SST neurons would lift inhibition on the pyramidal neurons that are excited by a stimulus maintained in working memory, whereas other SST neurons, not recruited by the maintained stimulus, would continue to inhibit nonactivated pyramidal neurons, thus suppressing both background noise and any potential activation by subsequent, distracting stimuli (Wang et al. 2004b).

Anatomical and physiological evidence supports the greater prominence of the disinhibiting circuit in PFC compared to other areas. Calretinin-positive interneurons are more numerous in PFC compared to the medial temporal and medial superior temporal visual cortical areas (Torres-Gomez et al. 2020). Moreover, interneurons with high baseline firing rate and inverted tuning (consistent with the profile of disinhibiting neurons) have also been found to be more numerous in the PFC than in the posterior parietal cortex (Zhou et al. 2012). The importance of inhibitory-to-inhibitory connections has been confirmed by neural network modeling studies (Kim and Sejnowski 2021). Such connections emerge in the network as training of synaptic weights progresses, and they play a critical role in maintaining working memory activity. Thus, these circuits underlie the prefrontal specialization toward persistent activity.

The final specialization informative about the prefrontal position in the cortical hierarchy has to do with dopamine. Dopamine innervation is concentrated in the frontal lobe (Levitt et al. 1984), and the D1 receptor has, in particular, been implicated in the generation of persistent activity. Iontophoretic application of D1 receptor antagonists, at least in large doses, compromise working memory function and erode persistent activity in the oculomotor delayed response task (Sawaguchi and Goldman-Rakic 1994; Williams and Goldman-Rakic 1995). In contrast, D1 agonists increase activity for preferred stimuli and suppress nonpreferred responses (Ott et al. 2014; Vijayraghavan et al. 2007). However, the effects of dopamine receptors are complex and depend on dosage (Vijayraghavan et al. 2007; Williams and Goldman-Rakic 1995),

From "The Frontal Cortex: Organization, Networks, and Function," edited by Marie T. Banich,
Suzanne N. Haber, and Trevor W. Robbins. 2024. Strüngmann Forum Reports, vol. 34,
Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262549530

with differential effects on pyramidal neurons and interneurons (Jacob et al. 2016). D2/D3 antagonists also suppress persistent activity, though their action primarily modulates motor responses of prefrontal neurons (Wang et al. 2004a). Computational and experimental studies suggest that the overall effect of dopamine is to improve the signal-to-noise ratio of persistent activity (Chen et al. 2004; Durstewitz et al. 2000; Seamans et al. 2001; Yang and Seamans 1996). Thus, dopamine innervation in the frontal lobe endows PFC with properties that are distinct from its afferent pathways.

Prefrontal Intrinsic Organization

Thus far we have referred to the PFC as a single brain region; however, embedded within this collective term is considerable heterogeneity of structure and functional specialization. The PFC can be subdivided into a medial, a lateral, and an orbital aspect, each consisting of several cytoarchitectonic areas. By itself, the lateral PFC comprises at least 12 functionally distinct areas, defined by unique cytoarchitectonic patterns and interconnected with a different set of brain areas (Petrides 2005). Many of the aspects of prefrontal specialization, delineating differences relative to other non-prefrontal areas in the cortical hierarchy, are also present in different areas *within* PFC: feedforward/feedback projection patterns, intracortical myelination variation, transcriptomic and cell-type gradients. An anterior-posterior hierarchical specialization has been suggested within the lateral PFC based on anatomical and imaging studies, with more abstract operations localized anteriorly in the prefrontal surface (Baird et al. 2013; Cole et al. 2015a; Koechlin et al. 2003; Ramnani and Owen 2004; Strange et al. 2001). A newer, more nuanced view suggests that the organization of function along the rostro-caudal axis is not an absolute gradient serving a unitary model of frontal control function; instead, separate frontal networks interact within an overall hierarchical structure to support task demands (Badre and Nee 2018). This idea is further developed by Badre (this volume). Neurophysiological evidence from animal models supports the idea of a rostro-caudal hierarchy in two respects: more neurons are selective for stimulus properties in posterior PFC, and more abstract qualities (e.g., task variables and rules) in anterior areas (Figure 6.1). Depending on task demands, plasticity of responses to the same stimuli is also more prominent in anterior areas.

Neuronal selectivity for the spatial location of stimuli was thus found to decrease along the anterior-posterior axis, so that the most highly selective neurons for stimulus properties were located more posteriorly in the PFC (Riley et al. 2017). Conversely, neurons in more anterior areas exhibited little selectivity to stimuli per se but were more likely to represent task variables. Neuronal selectivity for nonspatial stimulus attributes, such as shape and color, was also found to decrease along the anterior-posterior axis. The most highly

From "The Frontal Cortex: Organization, Networks, and Function," edited by Marie T. Banich, Suzanne N. Haber, and Trevor W. Robbins. 2024. Strüngmann Forum Reports, vol. 34, Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262549530

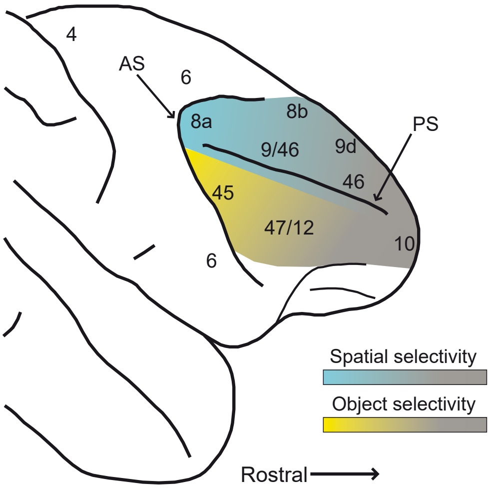

Figure 6.1 Diagram of the macaque monkey brain, with the lateral PFC highlighted. Labels denote anatomical areas: AS: arcuate sulcus; PS: principal sulcus. From Constantinidis and Qi (2018).

selective neurons for stimulus properties were located more posteriorly in the PFC (Riley et al. 2017).

Similar to the anterior-posterior axis of specialization, ventral areas are also more sensitive to task variables and cognitive factors rather than stimulus properties, so that robust selectivity to the location of stimuli may emerge as a result of training in task that requires tracking of reward. Ventral PFC has greater sensitivity to the learning of new, rewarded conditions, and this is likely due to the action of dopamine D1R receptors (Puig and Miller 2012).

A hierarchical prefrontal organization with respect to abstract functions implies that areas at the apex of this hierarchy are not activated solely by sensory stimuli, but by certain stimuli only in a specific context. In turn, such a context-depending property implies a greater capacity for plasticity of neural responses according to task demands. Direct evidence of systematic variation of plasticity markers between limbic and eulaminate areas, roughly organized across the anterior-posterior axis of the PFC, has been documented (Garcia-Cabezas et al. 2017). Calcium/calmodulin-dependent protein kinase II (CaMKII), which is essential for plasticity, is more impoverished in area 46d compared to anterior limbic areas. By contrast, makers of cortical stability, including intracortical myelin, perineuronal nets, and parvalbumin, show the reverse pattern. Changes in neuronal morphology, molecular profiles of the synaptic apparatus, and the influence of neuromodulator systems have also been implicated in long-term prefrontal plasticity (Laroche et al. 2000; McEwen and Morrison 2013), and these differ as well from posterior to anterior PFC.

From “The Frontal Cortex: Organization, Networks, and Function,” edited by Marie T. Banich, Suzanne N. Haber, and Trevor W. Robbins. 2024. Strüngmann Forum Reports, vol. 34, Julia R. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 9780262549530

Conclusions and Future Directions

Over the past decade, substantial progress has been made in our understanding of the unique functional properties of the PFC and the neural circuit substrates responsible for this specialization. This work has made it possible to account for the role of PFC, at least in some tasks, in terms of elemental neural circuits. A great deal of progress has been made in uncovering the functional organization of the PFC itself, as well as the hierarchy within its borders.

Nonetheless, some questions and controversies remain. Although we emphasize the hierarchical organization of areas within the PFC, comparative studies that would provide a comparison of their properties are lacking, particularly for the most anterior ones (e.g., area 10). These are critical regions for understanding the function of PFC, and additional experimental data will be valuable to inform our understanding.

The concept of executive control of other circuits is still poorly understood, though several chapters in this volume provide a good overview of the state of knowledge. To what extent functions of the PFC, such as working memory maintenance, are centralized in the PFC or more distributed is also a matter of controversy. Finally, how generalizable the conclusions drawn from animal models are on the human brain is still relatively unexplored. Attending to these questions over the upcoming decade will allow us to determine definitively not only the position of the PFC in the cortical hierarchy but also the mechanisms of executive cognitive functions.

Acknowledgments

This work was supported by NIH grant award R01 EY017077 (CC). We wish to thank Mark D'Esposito for helpful comments on an earlier version of this manuscript.

From "The Frontal Cortex: Organization, Networks, and Function," edited by Marie T. Banich,