

Dynamic Coding for Flexible Cognitive Control

Mark G. Stokes, Timothy J. Buschman, and Earl K. Miller

Introduction—The Importance of Cognitive Control

A hallmark of cognitive control is its flexibility. Every situation we experience is essentially unique, and yet we are able to act appropriately in different contexts, integrating the current environmental contingencies, our internal goals, and the goals of those around us. Imagine sitting down to play a new card game—you are able to integrate the rules; form an initial, possibly rudimentary, strategy; and then play your hand appropriately.

On top of this incredible ability to coordinate complex rule-guided goal-directed behaviour, we can quickly change our behaviour in response to changes in the world around us or our own desires. Behavioural studies suggest this can happen extremely quickly—within a couple of hundred milliseconds (for a review, see Sakai, 2008). For example, you will likely adapt your strategy as you play the game, learn its nuances, and respond to the strategies of others—perhaps even before the next hand is played.

What are the neural mechanisms that support this flexibility? In this chapter, we suggest that cognitive flexibility depends on the capacity of the prefrontal cortex (PFC) to dynamically encode the task-relevant information. We begin by reviewing the evidence that the lateral PFC (IPFC) plays a particularly important role in flexible cognitive control. Next, we review the evidence for highly dynamic representations in IPFC. Finally, we consider how these representations can be used to establish task-relevant networks throughout the brain.

Cognitive Flexibility

Cognitive flexibility is a defining feature of animal intelligence. This ability has developed throughout evolution, reaching its pinnacle in humans. The size of the PFC has grown in parallel along the phylogenetic tree, from non-primates (e.g., <4% of the total cortical volume in cats) to primates (e.g., ~10% in macaques) to apes (e.g., ~17% in chimpanzees) to its apex in humans (~30%; for review, see Fuster, 2002). This evolutionary trajectory already suggests that the PFC plays a central role in cognitive flexibility, a hypothesis that is supported by the impact of damage to the IPFC on behaviour (see Chapter 29 by Fellows in this volume).

Systematic analysis of lesion extents and behavioural impairments demonstrate that the PFC is critically important for planning and executing sequences of actions, particularly when individuals are faced with abstract or open-ended tasks (Shallice & Burgess, 1991). A particularly powerful paradigm for investigating cognitive flexibility has been the Wisconsin Card-Sorting Task (WCST). The WCST requires individuals to apply a sorting rule to a multidimensional stimulus and then, after an unpredictable amount of time, flexibly switch to another rule. Patients with lesions involving the frontal lobe often display profound difficulties in performing the WCST, particularly at the point where they should switch between different sorting rules on the basis of negative feedback (Barcelo & Knight, 2002; causing perseverative errors, Milner, 1963). Subsequently, lesion studies in non-human primates has confirmed the importance of the PFC in performing a WCST analogue and further parcellated cognitive functions to different subregions in the PFC (Buckley et al., 2009). Such experiments are consistent with the view that the IPFC is particularly important for representing the current task set (Miller & Cohen, 2001).

Neurophysiological studies of the IPFC demonstrate that neurons have a set of properties that may make them uniquely well suited to acting as a cognitive controller. First, individual neurons in the IPFC are known to represent the contents of working memory (WM; Funahashi, Bruce, & Goldman-Rakic, 1989; Fuster & Alexander, 1971; Miller, Erickson, & Desimone, 1996). Although WM is distributed across many different brain regions, coding in the IPFC is more resistant to distraction than visual brain areas (Miller et al., 1996). More recent work also suggests that regions in the PFC have the longest time constants (Murray et al., 2014). In other words, IPFC neurons integrate information over long periods of time (Fuster, 2001). This is a prerequisite for intelligent behaviour—one must respond based on the situation and goals, even if these constraints are not immediately obvious. Second, neurons in the IPFC do not represent just any information from the environment; instead, it prioritises those stimuli that are relevant to the current task (Miller et al., 1996; Rainer, Asaad, & Miller, 1998). Such filtering ensures that only task-relevant information is used to guide behaviour. Third, neurons in the IPFC are able to represent a great diversity of such task-relevant material (Duncan & Miller, 2002). Importantly, this includes information about the current task (Wallis, Anderson, & Miller, 2001; White & Wise, 1999), recent events held in WM, and predictions of future events (Rainer, Rao, & Miller, 1999). In other words, no matter what the task, IPFC neurons are able to represent the relevant information (including what the task is), hold it ‘in mind’, and then use this task-relevant information to plan future action and guide current behaviour.

Flexible Representations in the PFC

To gain more precise insights into the computational role of the PFC, we need to look more closely at the neurophysiological principles of coding in frontal networks. As noted above, single-cell recordings in the monkey PFC provide extensive evidence that activity encodes the behavioural parameters that are currently relevant. However, it is implausible that the PFC is hardwired to perform all of these functions (Duncan, 2010). Indeed, in a typical experiment, about 30% of neurons represent any given task variable (Duncan & Miller, 2002), implying that the PFC is sufficiently flexible to allow task-relevant information to effectively dominate the available coding space in the PFC.

Duncan and Miller (2002) propose that PFC neurons represent information in a highly dynamic manner, adapting rapidly to the current task. According to the principle of *adaptive coding*, the PFC constitutes a flexible pool of neural resources that can be recruited ‘on the fly’ to represent whatever information is currently most relevant for achieving behavioural goals.

Real-Time Flexibility

Multivariate analysis tools allow us to characterise population-level neural dynamics that underpin cognitive flexibility. Applied to the PFC, these analytical approaches clearly show that neural coding is not fixed, but depends on the exact context of the current task (Meyers, Freedman, Kreiman, Miller, & Poggio, 2008; Sigala, Kusunoki, Nimmo-Smith, Gaffan, & Duncan, 2008; Stokes et al., 2013). For example, we previously observed that population coding adapts according to a flexibly defined behavioural context. In Stokes et al., monkeys were trained to perform a simple stimulus–stimulus association task. At the beginning of each trial, the animal was given a ‘cue’ stimulus that would indicate the current context. This context determined which stimulus in an upcoming stream of stimuli would be the current ‘target’, to which they must respond to receive a reward. Notably, recording neurons in the IPFC, we found the population response to a choice stimulus (target vs. non-target) critically depends on the behavioural context (Figure 13.1a). In context A, the response to stimulus ‘A’ followed a trajectory towards the ‘go’ activity state, whereas the response to stimuli ‘B’ or ‘C’ followed a different path to the ‘no-go’ activity state (Figure 13.1b, upper panel). Correspondingly, in context B/C, stimulus ‘B’/‘C’ followed the trajectory towards the ‘go’ state, whereas the response to stimuli ‘A’ or ‘C’/‘B’ followed respective path to ‘no-go’ states. This dynamic mapping can be rephrased as a temporary shift in the decision circuit for flexible accumulation of task-relevant information (see Figure 13.1b, lower panel). How is this dynamic shift in coding established and maintained in the IPFC?

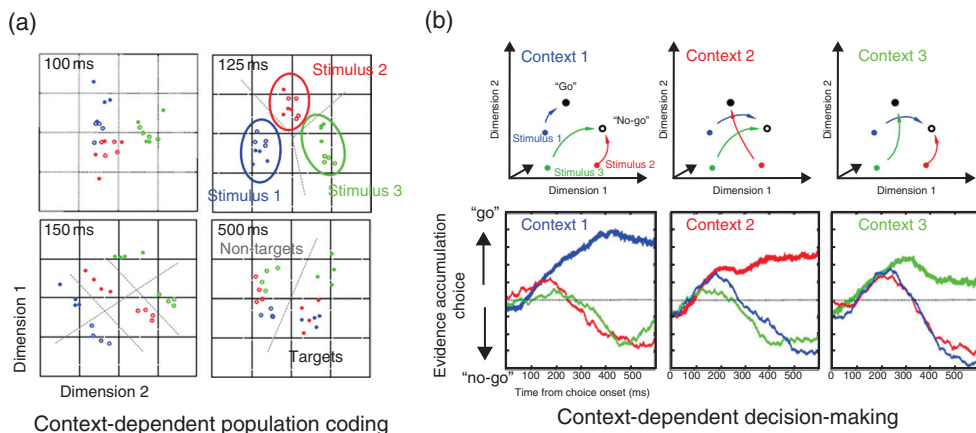


Figure 13.1 Context-dependent population coding in the prefrontal cortex. a. During WM-guided behaviour, population-level activity states quickly evolve from representing the physical properties of the presented stimuli (at 125 ms; stimuli are coloured coded) to the decision-relevant coding from ~150 ms (i.e., target vs. non-target; from Stokes et al. [2013], Open Access). b. This process can be schematised as a context-dependent path through the activity-state space. A functional shift in the response profile of the network effectively maps stimulus-specific activity states to the context-relevant position in the state space (upper panel; from Stokes et al., 2013). This flexible mapping could be framed as a context-dependent decision process. In the lower panel, the results from panel A are replotted as the accumulation of evidence for the ‘go’ or ‘no-go’ response as a function of the current rules (i.e., if context 1, then the evidence for stimulus 1 supports a ‘go’ response, but stimuli 2 and 3 provide counterevidence; upper panel). Source: Adapted from Stokes 2013. Reproduced with permission of Elsevier.

WM States for Guiding Behaviour

WM provides a temporary but stable platform for flexible, context-dependent processing. Early neurophysiological recordings in the IPFC suggested that task-relevant information is maintained via persistent delay activity (e.g., Fuster & Alexander, 1971). Tonic activity states representing task-relevant parameters could, in principle, mix with new input for context-dependent processing (Mante, Sussillo, Shenoy, & Newsome, 2013). However, accumulating evidence suggests that the relationship between delay activity and WM maintenance is more complex. For example, content-specific delay activity has been observed prior to learning a specific WM task, and surprisingly, even reduced as the animals mastered the WM task (Qi, Meyer, Stanford, & Constantinidis, 2011). Even within a single trial, WM-specific activity seems to wax and wane during the delay period, depending on the current focus of attention (reviewed in Stokes, 2015). For example, when the duration of the maintenance period is fixed from trial to trial, robust delay activity may only emerge late during the delay period (Barak, Tsodyks, & Romo, 2010; Watanabe & Funahashi, 2007; Figure 13.2a). So-called ‘ramp-up’ activity implies that delay activity could optimise behaviour by preparing for the processing demands of the memory probe (e.g., temporal orienting of attention [Nobre, Correa, & Coull, 2007]). However, the relatively silent moments between encoding and response preparation also suggests that the continuity of vigorous delay activity is not always necessary for the continuity of the mental representation (Barak et al., 2010).

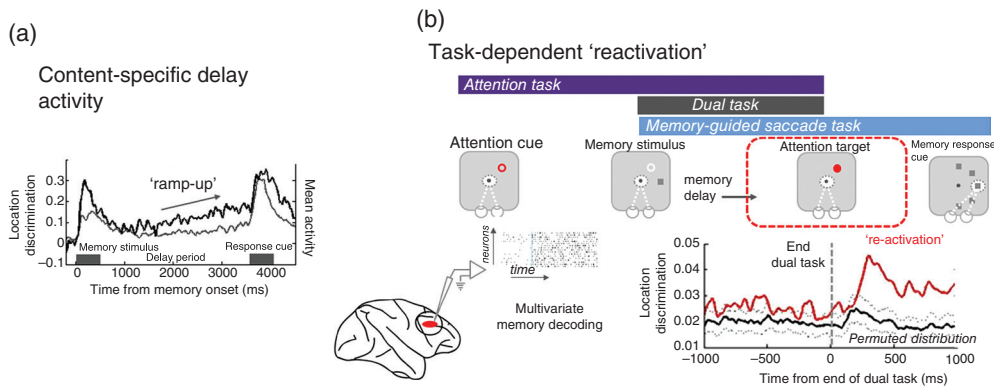


Figure 13.2 Delay activity is modulated by current task relevance. a. Neurophysiological recordings from the PFC during the delay period of a memory-guided saccade task reveal content-specific activity (location-specific activity is shown in black, and population mean activity is shown in grey for reference). However, content-specific delay activity also varies with the current task relevance of the memory item, resulting in a ‘ramp up’ in anticipation of the response (Source: Adapted from Watanabe 2007. Reproduced with permission of Oxford University Press). b. In a more recent dual-task experiment, an attention task partially overlaid the standard MGS design. Content-specific activity was effectively abolished by the dual task during the memory delay (red trace: from –1000 to onset of the attention target; black trace: permuted distribution). Critically, location-specific information was ‘reactivated’ at the end of the dual task (red trace, after the onset of the attention target), presumably reflecting a shift in task focus to the MGS (Source: Adapted from Watanabe 2014. Reproduced with permission of Nature Publishing Group). Such evidence suggests that delay activity reflects the task relevance of memoranda. Gaps in content-specific activity further suggest that WM could be maintained in an ‘activity-silent’ neural state.

Recent electrophysiological studies provide specific evidence that suggests persistent delay activity is not strictly necessary for accurate WM-guided behaviour. Warden and Miller (2007) trained monkeys to perform a two-item WM task: The animals were presented two stimuli in sequence, and then, after a memory delay, they were asked to report the identity and order of the remembered stimuli. Recordings in the IPFC demonstrated that neurons encoding the identity of the first stimulus in the sequence were suppressed during the presentation of the second stimulus. After the second stimulus, memory-specific delay activity for the first stimulus was 'reawakened', presumably in time for WM-guided behaviour. This was true, even when accounting for the high dimensionality of the representations (Rigotti et al., 2013). However, we should note that information was phase-locked to an ~30 Hz population oscillation (Siegel, Warden, & Miller, 2009), which might also be important for information coding (more on this later).

Watanabe and Funahashi recently reported similar results in a dual-task paradigm (Watanabe & Funahashi, 2014). As in the standard variant of a delayed saccade task (Funahashi et al., 1989), monkeys were required to encode and maintain the location of a saccade target for execution at the end of the trial. However, under a dual-task condition, they were also required to attend to a specific spatial location until part way through the WM delay period (Figure 13.2b). This cognitive manipulation almost abolished WM-specific delay activity during the dual-task period, even on trials in which WM performance was preserved (WM-correct trials). Moreover, when the competing task demands were completed, robust delay activity was again reinstated (see also Lebedev, Messinger, Kralik, & Wise, 2004). Together, these findings suggest that mnemonic delay activity in the PFC is not always critical for maintaining the continuity of WM, but can be dynamically re-established when attention is refocused to the task-relevant content. WM does not seem to depend on delay activity, but might be maintained in an 'activity-silent' format.

'Activity-Silent' WM Stored in Effective Connectivity

Mongillo and colleagues proposed that WM could be maintained during such activity-silent periods as a pattern of synaptic weights in the PFC (Mongillo, Barak, & Tsodyks, 2008), analogous to long-term memory (see also in IT; Sugase-Miyamoto, Liu, Wiener, Optican, & Richmond, 2008). In their computational model schematised in Figure 13.3, activity during encoding temporarily changes the synaptic efficacy within the neural network, leaving behind a temporary trace via short-term synaptic plasticity (STSP; Zucker & Regehr, 2002). In this particular model, calcium kinetics provide a window of approximately 2 s for STSP (Mongillo et al., 2008); however, different time constants inherent to other synaptic processes could underlie the diversity of mnemonic time scales observed in the PFC (Bernacchia, Seo, Lee, & Wang, 2011). The essential point is that task-relevant input changes the *effective connectivity* of the network to construct a temporary task-dependent circuit for WM-guided behaviour.

Mongillo et al. (2008) originally proposed that an 'activity-silent' coding scheme is more efficient than persistent firing models. A coding format that does not depend on an unbroken chain for persistent firing would clearly have a cost advantage (i.e., *sparsification* in time). However, here we focus on another possible application of such 'silent' memories: They can provide a neurobiologically plausible mechanistic account for WM-guided behaviour as generalised state-dependent processing (see below for further discussion). Memories are not stored as an active representation, but rather change the functional architecture of the neural network for future processing. WM is therefore expressed by the way the network responds to new input. State-dependent readout would avoid the need for explicit mechanisms for comparing the internal representation with a separate representation reflecting the new input (Mongillo et al., 2008; Sugase-Miyamoto et al., 2008). Moreover, because memories are

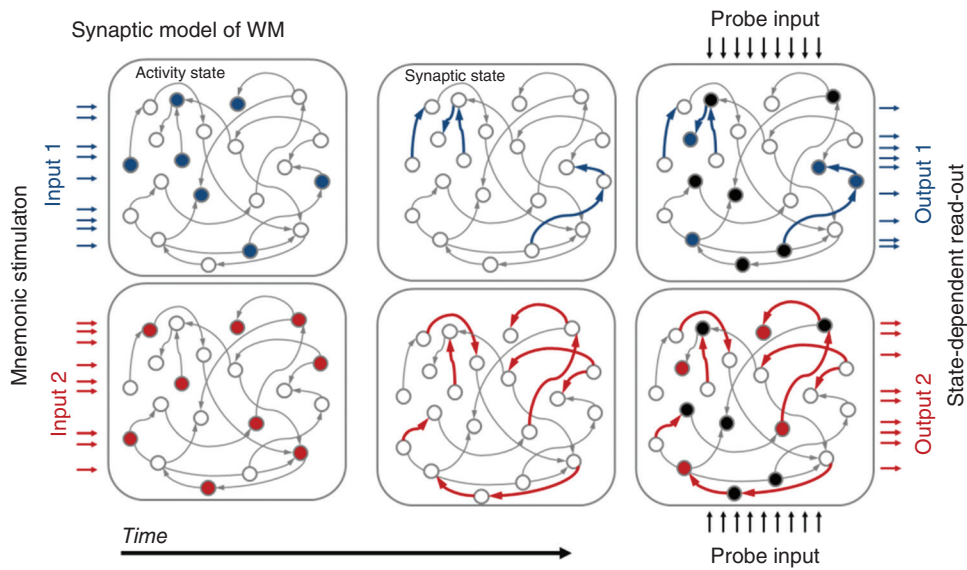


Figure 13.3 Maintaining ‘activity-silent’ working memory in effective connectivity. Schematic of the synaptic model of working memory described in Mongillo et al. (2008). Task-relevant input (left-side horizontal arrows, blue for ‘Memory A’ and red for ‘Memory B’) drives a stimulus-specific activity state (filled circles), which in turn triggers a specific pattern of short-term synaptic plasticity between cells (bold arrows). Memory is read out from this synaptic trace via the context-dependent response at retrieval (solid black circles, same for ‘Memory A’ and ‘Memory B’). The probe-driven response will be patterned by the hidden state of synaptic efficacy, resulting in a discriminable output pattern (right-side horizontal arrows). Source: Adapted from Stokes 2015.

stored in a format that is qualitatively different from more transient representations (e.g., ongoing perception and action), their informational content could be less confusable with other activity-based representations (Olivers, Peters, Houtkamp, & Roelfsema, 2011).

Despite considerable theoretical appeal, a synaptic model of WM is relatively difficult to test empirically. Extensive evidence confirms the prevalence of STSP (Zucker & Regehr, 2002), but its functional role has not yet been fully established. Extracellular recordings can only indirectly infer connectivity via interactions between simultaneous spike trains. The gold-standard evidence for a monosynaptic connection is correlated spiking between cell pairs. In practice, however, the probability of sampling any two neurons with a measurable monosynaptic connection is extremely low (~1–2% of all recorded pairs in Fujisawa, Amarasingham, Harrison, & Buzsaki, 2008). Such a poor yield has been effectively prohibitive for testing synaptic plasticity within a standard primate single neurophysiological study. Fujisawa and colleagues were able to collect sufficient samples of simultaneous recordings in the rat frontal cortex to show that the pattern of effective synaptic connections was dynamically modulated during a WM-dependent maze task (Fujisawa et al., 2008; Figure 13.4 left panel). Although this observation is compatible with a synaptic WM hypothesis, the rodent model is not ideal. The rat frontal cortex lacks key features of primate PFC cytoarchitecture (Wang et al., 2013). Moreover, the rodent WM analogue is difficult to compare with standard behavioural tasks used in primate WM experiments. Therefore, further research in the primate PFC is required to obtain more specific evidence for a synaptic theory of WM.

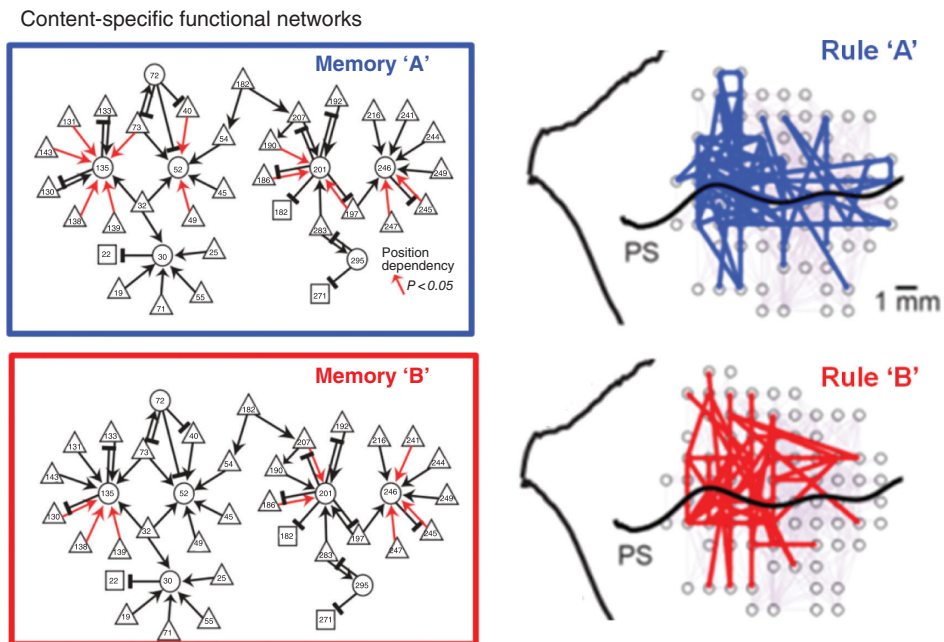


Figure 13.4 Evidence for content-specific functional networks. Simultaneous recordings in the rat frontal cortex revealed direction-specific patterns of synaptic efficacy (red arrows) between cells (putative pyramidal [triangles]; putative interneuron [circle]; unclassified [square]) coding for direction during a WM-based maze task (left panel, Source: Adapted from Fujisawa 2008. Reproduced with permission of Nature Publishing Group). This is consistent with a role for short-term synaptic plasticity in WM. In the monkey PFC (PS = principal sulcus), different task rules are associated with specific functional networks (synchronised electrodes for rule 'A' in blue; rule 'B' in red) coupled by synchrony at ~ 30 Hz (right panel, Source: Adapted from Buschman 2012. Reproduced with permission of Elsevier). This is consistent with the idea that coherence could also play a role in constructing functional networks for flexible behaviour.

Synaptic plasticity is not the only candidate mechanism for an effective connectivity-coding scheme for WM. For example, synchronous oscillations could provide an alternative, complementary, or even supportive, mechanism for rapid and temporary shifts in the effective connectivity between neurons. The idea is simple: If two neurons are rhythmically synchronised in phase with each other, then they will share the same periods of relative excitability and relative inhibition. This in turn will ensure that any action potentials from one neuron will have a maximal impact on the other neuron (Fries, 2005). Conversely, if two neurons are out of phase, the opportunity for successful information transfer between nodes will be reduced. In this way, functional pathways can be rapidly constructed (and dissolved) simply by modulating synchrony along pre-existing structural paths of the network (see also Chapter 14 by M. X Cohen in this volume).

Buschman et al. (2012) recently provided evidence for this kind of model. They trained monkeys to follow two different rules: Attend and respond to either the colour or orientation of an imperative stimulus. If synchrony acts to shape the effective connectivity in the PFC in support of the task rule, then one would expect the synchrony between populations

of neurons to change depending on the current rule. To test this prediction, Buschman et al. used about a dozen electrodes to simultaneously record the activity of a population of neurons in the IPFC. As predicted, they found two distinct patterns of synchrony—one for the colour rule and one for the orientation rule—within the IPFC at ‘beta’ frequencies (~25 Hz; Figure 13.4, right). Importantly, they also found that neurons carrying task-relevant information, such as the identity of the stimulus, were synchronised to the currently relevant task ensemble. When the animal switched from one rule to the next, these neurons changed their synchronous association, suggesting that synchronous ensembles pull together all of the relevant neurons to execute a task. A similar mechanism for the rapid configuration of content-specific network architectures could be used to keep other forms of task-relevant material in WM (Salazar, Dotson, Bressler, & Gray, 2012). As described above for STSP, such a rapid shift in the functional architecture of the PFC network could directly modulate the response dynamics, allowing for context-appropriate input-output behaviour (Fries, 2005).

Dynamic Coding

Activity-dependent dynamics in connectivity patterns (e.g., STSP) predict a cascading interaction between past and present input, resulting in a complex, but reproducible trajectory through the state space (Figure 13.5a). Consequently, the response profile of the network is in constant flux during active processing. While input drives a specific response to a network according to the current connectivity state, the resultant activity state in turn shifts the connectivity state. Therefore, the response to subsequent stimulation will trigger a unique response pattern according to the new connectivity state. Moreover, this new pattern will further modulate the new connectivity state of the system, thus determining the response to the next input, and so on (Buonomano & Maass, 2009). The reciprocal interaction between the activation state and the underlying connectivity states of the network should result in a complex spatiotemporal trajectory through the state space observed throughout different experimental contexts (Crowe, Averbeck, & Chafee, 2010; Meyers et al., 2008; Stokes et al., 2013; see Figure 13.5b). The full trajectory is reproducible across trials to the extent that the temporary connectivity state returns to the baseline connective state after some time period. The baseline state is determined by more stable connections established via long-term plasticity (i.e., long-term potentiation/depression).

By allowing activity and connectivity states to interact across time, dynamic coding dramatically expands the dimensionality of population coding. This can be thought of as a temporal extension of the mixed selectivity view outlined above (Rigotti et al., 2013). Neurons in the IPFC do not map neatly to distinct behavioural conditions, but tend to show some form of *mixed selectivity* (Rigotti et al., 2013). The usefulness of such mixed selectivity approaches is well established and has been used extensively in computer science (for example, it is directly analogous to the kernel trick used in support vector machines (Vapnik, 1998)). For example, the flexibility of mapping ‘A’ or ‘B’ to a ‘Go’ response is facilitated by having a high-dimensional representation that allows for space to be easily divided in different ways. By the same token, the information potential of a dynamic network representation is directly proportional to the statistical independence between time points (Stokes, 2011).

Perhaps more importantly, dynamic coding could condition the network to exhibit context-dependent response properties before the connectivity state relaxes back to the default configuration. This idea underpins probe-driven readout in synaptic models of WM (Mongillo et al., 2008; Sugase-Miyamoto et al., 2008), but could provide a mechanistic account for dynamic tuning profiles that could mediate adaptive coding in the IPFC (Duncan & Miller, 2002). Task-relevant input temporarily shifts the response profile of the IPFC (Stokes, 2015).

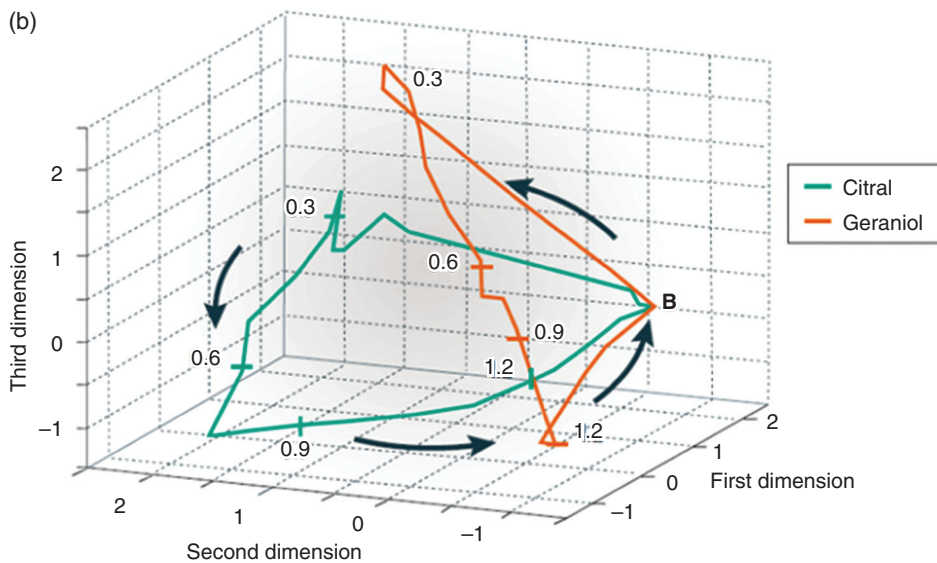
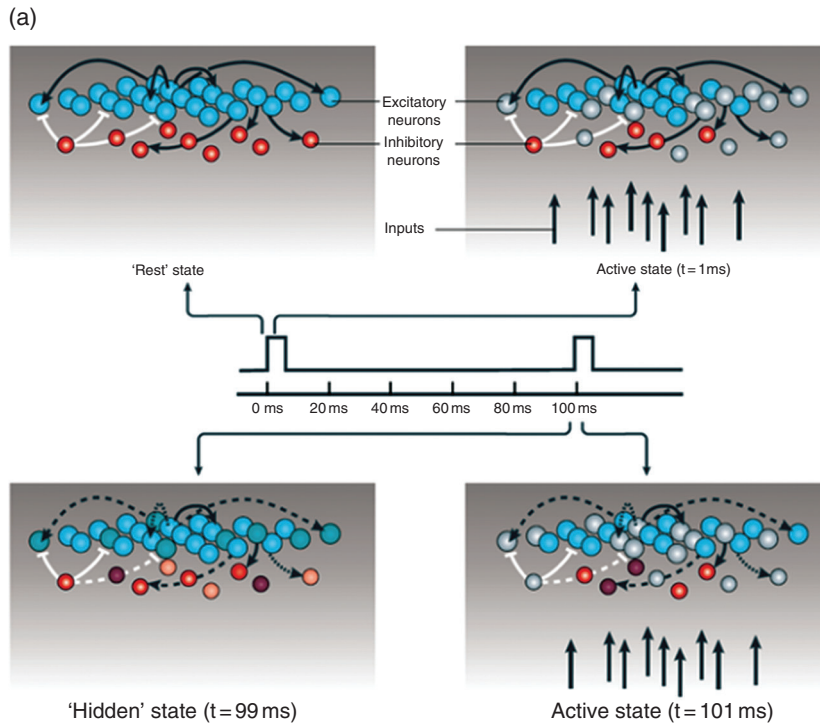


Figure 13.5 Dynamic coding predicts non-stationary activity states. a. Activity-dependent changes in effective connectivity (e.g., STSP) will result in a cascading interaction between past and present activity states. The activity state of the system depends on the input drive and the current connectivity of the network. At each moment, activity changes the connectivity profile, which in turn alters the subsequent response of the network, which in turn results in a different response to the same input drive, and so on (with permission, from Buonomano & Maass, 2009). b. This mutual interaction between activity states and the underlying connectivity states means that even a fixed input to the system will drive a dynamic trajectory through activity-state space, as illustrated by the complex trajectories observed during odour processing in the locust olfactory system for two specific smells: citral and geraniol. Source: Buonomano 2009. Reproduced with permission of Nature Publishing Group.

Although it is difficult to measure connectivity states with standard extracellular recordings, it is possible to infer changes in network architecture from systematic changes in network behaviour. As already mentioned, the complex spatiotemporal trajectory observed during cue processing is a characteristic feature of dynamic changes in network behaviour associated with activity-dependent effective connectivity. However, it is difficult to rule out other kinds of dynamics, including different cell latencies (Harvey, Coen, & Tank, 2012) or nonlinear interaction between activity states (Mante et al., 2013). A specific hidden state should also emit a signature hum as a result of ambient spontaneous activity (Sugase-Miyamoto et al., 2008); however it would be difficult to differentiate such patterned activity from the classic persistent activity states previously ascribed to WM (Shafi et al., 2007).

One possible way to explore otherwise hidden states is to ‘ping’ the network with a neutral input and see how it responds. As an illustration, consider echolocation (e.g., active sonar), where a simple impulse (e.g., ‘ping’) is used to probe the hidden contours of an unseen structure (Figure 13.6a). Analogously, the impulse response to neural perturbation should co-depend on the pattern of input activity and the current neural state (activity and connectivity state). If the input pattern is held constant, then any response-related difference could be attributed to underlying changes in network behaviour. Exploiting this logic in Stokes et al. (2013), we found that a neutral task-irrelevant stimulus presented during a WM delay period generated WM-specific patterns of activity in the IPFC. We suggested that this context-dependent response pattern could reflect differences in the hidden neural state, such as effective connectivity. We have recently extended this impulse–response approach for non-invasive EEG (Figure 13.6b), providing new opportunities to study hidden states in the human brain (Wolff, Ding, Myers, & Stokes, 2015).

To summarise, a dynamic coding framework proposes that context-defining input configures a temporary functional architecture in the IPFC, perhaps via effective connectivity (schematic Figure 13.7). Although this kind of state could be ‘activity silent’, such hidden neural states can influence subsequent processing, including dynamic trajectories observed during encoding (see Figure 13.7). Such a shift in the underlying hidden state would also emit a signature hum from ambient spontaneous activity, and condition the response to an impulse perturbation, and/or establish specific activity states when attended. Most importantly, however, the temporary shift in network behaviour could provide a mechanistic account for WM-guided behaviour specifically (Stokes, 2015), and perhaps context-dependent processing more generally (Stokes et al., 2013). The flow of information is effectively re-routed through the prefrontal network. This can be reformulated as a ‘dynamic decision circuit’, where evidence accumulation is flexibly determined by the current task context (Deco, Rolls, & Romo, 2010).

Flexibility in Networks

So far we have considered the physiological properties of neurons in the IPFC and how their dynamic nature can support cognitive flexibility. However, cognitive flexibility relies on using these representations to guide neural activity throughout the brain. Indeed, what may give the IPFC its special role in cognitive flexibility may be its connections throughout the brain. Here, we will discuss anatomical and physiological evidence that the IPFC plays a key role in biasing representations throughout the brain.

The PFC is especially well situated anatomically to play a controlling role over cognition. First, subregions within the PFC are densely interconnected, both within and across

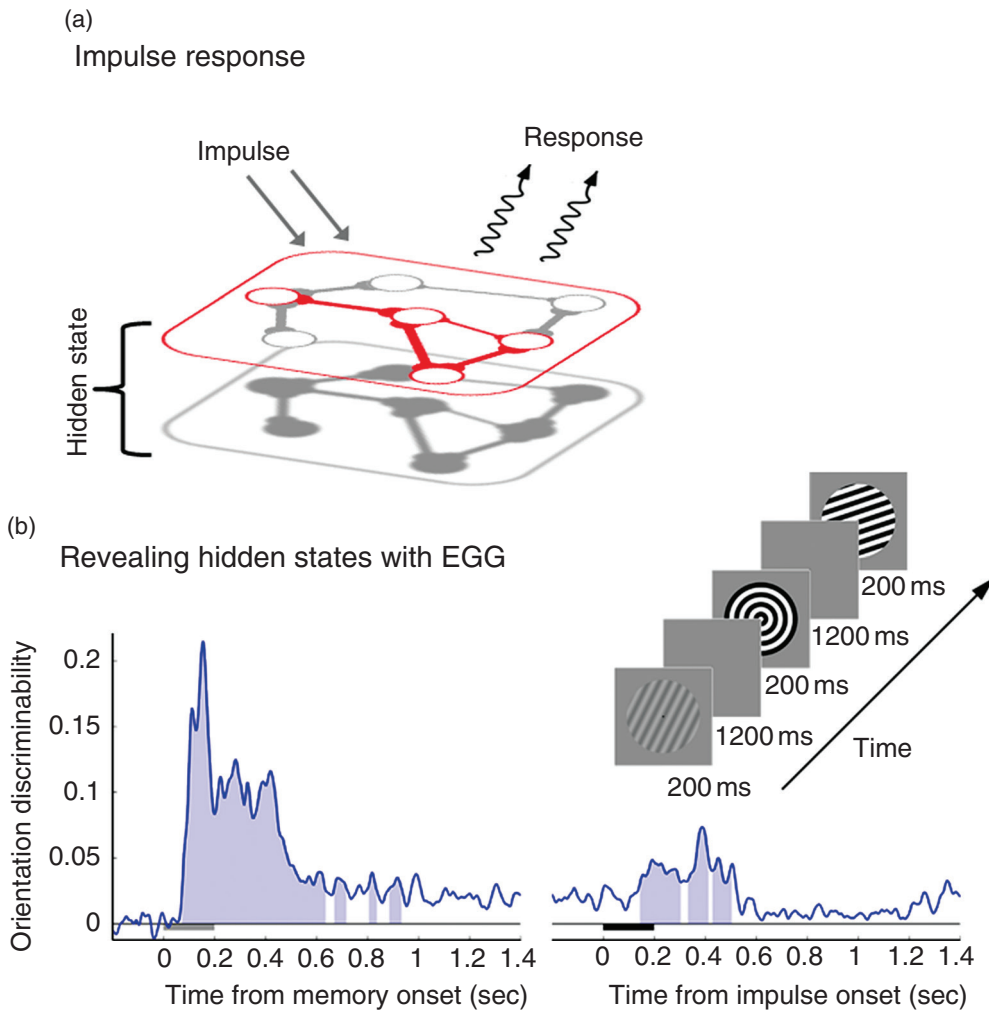


Figure 13.6 Revealing hidden brain states. a. Synaptic working memory is effectively ‘activity silent’; however, the information content can be inferred from a shift in the behaviour of the network. This can be probed using an impulse of activity to drive the network. Source: Stokes 2015, Reproduced with permission of Elsevier. b. Source: Wolff et al. (2015) *Frontiers in Neuroscience* recently developed an approach for measuring impulse-driven reactivation of WM using non-invasive EEG in human participants. Using multivariate statistics, we could show that a neutral driving stimulus triggers a distinct pattern of activity that reflects the content of WM (blue trace), even though the driving stimulus was exactly the same for all WM conditions.

hemispheres (Barbas, Hilgetag, Saha, Dermon, & Suski, 2005; Barbas & Pandya, 1989). This means that even if different subregions within the PFC show an initial preference for different cognitive variables, this information will be rapidly shared throughout the circuit, ensuring that all regions have the necessary task-relevant information.

Second, the PFC is widely interconnected with many other cortical regions, with especially prominent connections to the temporal, parietal, and cingulate cortex (Kunzle & Akert,

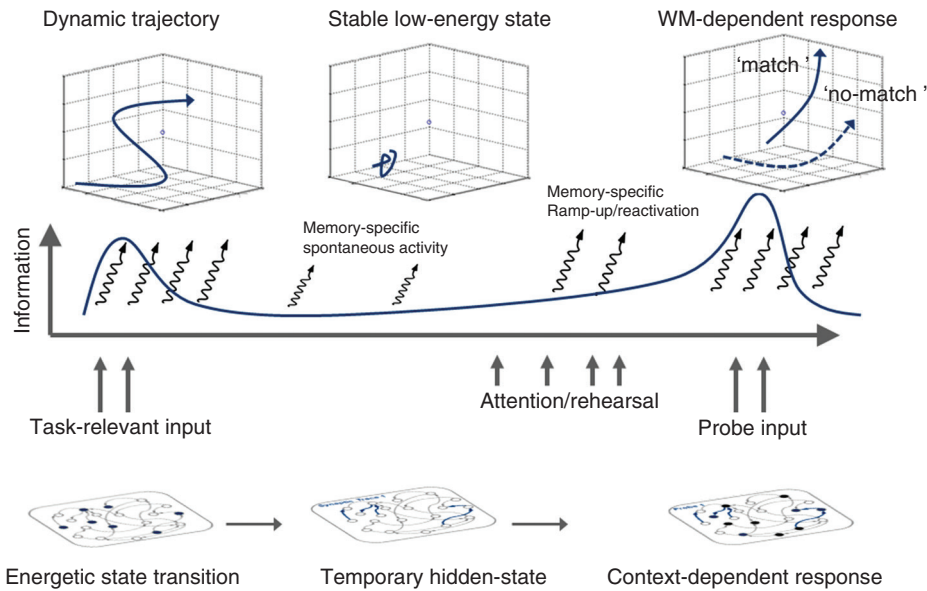


Figure 13.7 Summary schematic of dynamic coding for WM-dependent behaviour. The initial input triggers a specific pattern in the activity state, which in turn alters the underlying hidden state of the network via a temporary shift in effective connectivity (e.g., short-term synaptic plasticity or coherence). Activity-dependent changes in the hidden state drive a dynamic trajectory during the initial high-energy phase via the reciprocal interaction between hidden states and the activity states that modulate them (Meyers et al., 2008; Stokes et al., 2013). After activity has relaxed to baseline levels, the hidden state remains patterned according to the WM item. Although, in principle this temporary hidden state could be ‘activity silent’, spontaneous activity in the network will be patterned according to the WM context, resulting in a WM-specific activity state during spontaneous firing (Stokes et al., 2013). This kind of ‘baseline emission’ could help explain content-specific delay activity observed under some circumstances (Sugase-Miyamoto et al., 2008). Increasing the level of network activity, via attention/rehearsal mechanisms, could increase the discriminability of the activity state, resulting in ‘ramp-up’ delay activity (Barak et al., 2010; Watanabe & Funahashi, 2007) or task-dependent ‘reactivations’ (Watanabe & Funahashi, 2014). Finally, when the critical memory probe is presented, the context-dependent response maps activity states for WM-guided behaviour (e.g., match/non-match decision; see Stokes et al., 2013). Source: Adapted from Stokes 2015.

1977; Petrides & Pandya, 1999). Graph analysis of a large database of connections between brain regions shows how the PFC is centrally located in the large neural network that is the brain (Markov et al., 2013; Figure 13.8).

Finally, the PFC is closely interconnected with several subcortical regions. In particular, the PFC is the largest source of projections into the basal ganglia via the striatum (Haber, Kunishio, Mizobuchi, & Lynd-Balta, 1995; Jarbo & Verstynen, 2015; McFarland & Haber, 2002), forming cortico-ganglio-thalamocortical loops (Haber, 2003). We have previously suggested that these loops may allow the brain to ‘bootstrap’ from simple representations to more complex ones, allowing for the complexity of cognition to be built through learning (Buschman & Miller, 2014). The basal ganglia is a ‘fast’ learner, for example, quickly capturing associations between a stimulus and response (e.g., a bullmastiff barks). In contrast, the ‘slower’ learning of the PFC allows it to integrate over many experiences. Importantly,

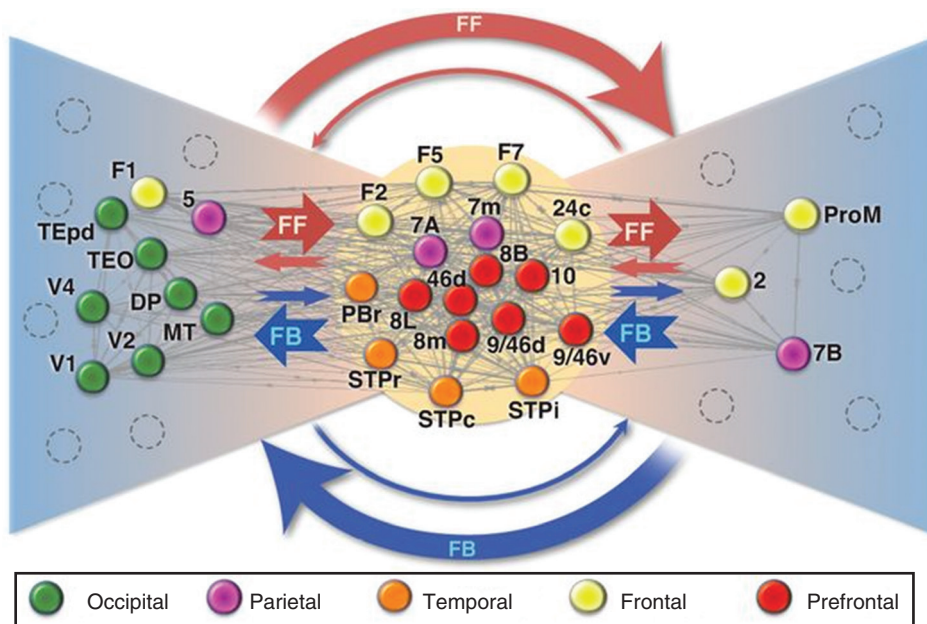


Figure 13.8 PFC is a ‘hub’ in the brain. Detailed anatomical work has shown that the brain is a highly interconnected network with both feedforward (FF) and feedback (FB) connections between regions. Detailed graph analysis of these connections suggests a core ‘hub’ of regions is more densely interconnected with other brain regions. Prefrontal regions (shown in red) are a prominent member of this hub. Source: Adapted from Markov 2013; region labels follows. (Source: Markov 2013, reproduced with permission of AAAS.)

the representations learned by the basal ganglia are available to the PFC during learning (given the strong connections between the regions). This allows the PFC to integrate these associations when forming a more generalised representation (e.g., a bullmastiff is a type of dog; and all dogs bark). This generalised representation then becomes available to further fast learning in the basal ganglia, allowing the brain to continually bootstrap increasingly complex and abstract associations.

Evidence for Top-Down Control Over Other Brain Areas

Given the anatomical position of the PFC within the broader network of brain structures, it seems natural that it plays an important role in cognitive control and, thus, cognitive flexibility. However, structural anatomy is necessary but not sufficient to determine what role the PFC plays in cognition. Critically, we must also consider how these network properties are modulated for different task contexts. Again, we need to consider the role of effective connectivity for cognitive flexibility.

Detailed studies of visual attention provide rich insights into the role of the PFC in cognitive control. Attention is the ability to selectively enhance task-relevant representations. Typically, this is studied in the domain of visual perception, where decades of research have demonstrated that attention can improve the perception of a selected visual stimulus while simultaneously suppressing distracting stimuli (for a recent review, see Carrasco, 2011). In this way, visual attention is a case study for cognitive control: Contextual/task variables are used to

guide what information is represented in the brain. Just as in cognitive control more broadly, lesions of the PFC cause deficits in controlling attention (Rossi, Bichot, Desimone, & Ungerleider, 2007). In particular, deficits can be characterised as causing a reduction in the ability to shift attention—much like in the WCST, patients with PFC lesions tend to get stuck on attending to a single location or object (Knight, 1997; Knight, Grabowecky, & Scabini, 1995; see also Chapter 29 by Fellows in this volume).

However, it is important to note that the PFC does not act alone in controlling attention. In particular, lesions in the parietal cortex will also lead to problems in controlling attention that are almost identical to the impact of a PFC lesion (Eglin, Robertson, & Knight, 1991). Electrophysiological recordings in the PFC and the parietal cortex find that neurons in both regions reflect shifts in attention (for a detailed review, see Miller and Buschman, 2013). However, by using multiple electrodes to record simultaneously from both regions, Buschman and Miller (2007) found that the PFC plays a leading role when attention is wilfully directed, using internal knowledge. Once these signals are observed in the IPFC, they supply feedback in a ‘top-down’ manner to other prefrontal regions (the frontal eye fields, FEF) and the parietal cortex (the lateral intraparietal region, LIP). Recent work has extended these findings to humans (Li, Gratton, Yao, & Knight, 2010). These results provide evidence that the IPFC is the source of top-down control, using internal knowledge to control other brain regions.

Causal evidence that the PFC acts to bias posterior representations comes from a highly influential series of experiments using electrical microstimulation to directly stimulate neurons in the FEF (Moore, 2006). Microstimulation of FEF cells tuned to visuospatial coordinates will boost the animal’s behavioural discriminability at the corresponding spatial location (Moore & Fallah, 2001, 2004). In other words, microstimulation of the FEF is functionally equivalent to shifting attention to a location. FEF microstimulation was also found to have attention-like effects on V4 neurons (Moore & Armstrong, 2003), providing direct evidence that the FEF modulates activity in the visual cortex to bias sensory processing. Similar results have been observed in humans using a combination of transcranial magnetic stimulation (TMS) and electrophysiology (Taylor, Nobre, & Rushworth, 2007).

More recently, Siegel et al. (2015) found electrophysiological evidence that the PFC plays a leading role in cognitive control during a decision-making task. Two monkeys were trained to perform a decision-making task that required the task-dependent integration of sensory information to make a choice. To trace sensory, task, and choice information across the brain, Siegel and colleagues used multiple electrodes to record simultaneously from six brain regions: the sensory cortex (V4, MT, and IT), associative cortex (IPFC and LIP), and premotor cortex (FEF) cortex. As expected, sensory information flowed in a bottom-up manner, starting in the sensory regions (V4, MT, and IT) and then propagating into the parietal and frontal regions (Figure 13.9a). In contrast, sustained task information and choice information flowed in a top-down direction, starting in the frontal and parietal cortex and then feeding back into the sensory regions (Figures 13.9b and 13.9c). The timing of the signals suggests that task information represented in fronto-parietal networks is used to selectively integrate the currently relevant sensory information and transform this signal into an action. These results complement the results from the prefrontal and parietal cortex in the control of attention, demonstrating the PFC represents context information in order to support task-relevant computations throughout the brain.

PFC as the Central Executive

Taken together, all of these results suggest that the PFC is ideally suited to play the role of a ‘central executive’. As reviewed above, IPFC neurons have the ability to adaptively and dynamically encode the currently task-relevant information. Furthermore, because of its

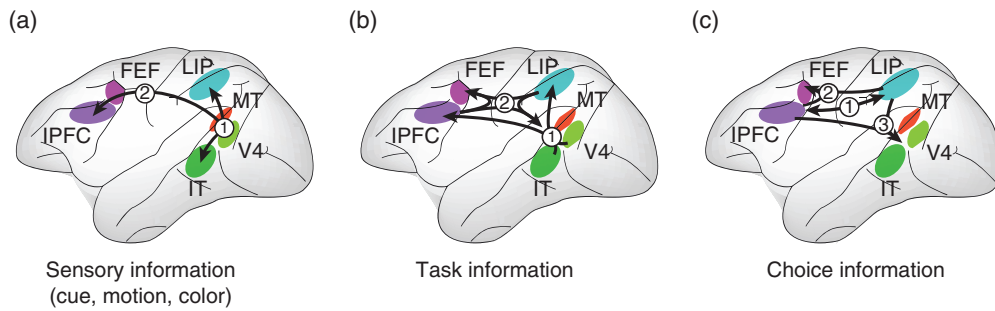


Figure 13.9 Dynamics of task information across six different brain regions. (a) Sensory information flows from the sensory cortex into the frontal-parietal associative areas. In contrast, (b) sustained task information and (c) decision information are seen first in the frontal-parietal network. Source: Siegel 2015. Reproduced with permission of AAAS.

anatomical connections, it is well positioned to use this activity to influence much of the brain. In this way, the PFC can act as a ‘central executive’, guiding activity throughout the brain in a way that accomplishes the current task.

But what are the neural mechanisms that would allow the PFC to act as a central executive? When considering possible mechanisms, it is important to remember that they must meet two competing criteria. First, any mechanism of cognitive control must be able to selectively boost those representations that are relevant to accomplishing the current task (while also suppressing undesired representations). However, this selection cannot be absolute—behaviour in the real world is highly dynamic, changing from moment to moment. Therefore, the selection mechanism must be flexible enough to allow for rapid changes between behaviours.

The first requirement can be met via experience-dependent changes in network connections. Learning could establish new anatomical pathways that are optimised for repeated scenarios and behavioural routines. However, the second requirement means that we need to consider more rapid changes in network behaviour. The brain simply cannot rewire on a moment-to-moment basis. Instead, flexibility demands that some changes must be rapid and temporary. In other words, anatomy acts as the substrate that defines the set of all possible behaviours, and temporary changes in effective connectivity ride on top to select the currently relevant behaviour.

Miller and Cohen (2001) outline one biophysically plausible mechanism through which this might occur. They propose that, at any given moment in time, we are faced with a plethora of stimulus inputs, all of which can lead to a large set of possible responses. As we reviewed above, there is strong evidence that the PFC is the source of internally directed attention and that attention selects a specific desired stimulus from competing stimuli. Miller and Cohen propose that this same selection mechanism can be generalised from stimulus inputs to all possible stimulus–response mappings. In other words, the model suggests that neural activity in the PFC is distributed across the brain in a manner that boosts the representations of the contextually appropriate behaviour. Competition between representations leads to suppression of inappropriate behaviours. Finally, because control depends only on the pattern of activity in the PFC, it can be highly flexible, changing as rapidly as the patterns of activity in the PFC.

However, the biasing mechanism underlying the Miller and Cohen model may not be the only mechanism supporting cognitive control. We suggest that cognitive control (and its

flexibility) is the result of changes in effective connectivity. In particular, we suggest that the same changes in effective connectivity that act to support dynamics within the PFC can be generalised to the whole-brain network in order to support cognitive flexibility.

Effective Connectivity as a General Mechanism for Dynamic Coding Within and Across Brain Areas

As reviewed above, changes in effective connectivity between neurons can alter the dynamics of a population of neurons within a cortical region. This could occur through either STSP and/or through changes in the synchrony between populations of neurons. These same mechanisms could also apply at the scale of whole-brain networks, with changes in the effective connectivity between regions acting to bias connections in a way that supports the current behaviour.

This is exemplified in many of our everyday experiences. For example, when walking home from work, you follow a typical path that has been learned over many experiences: You know to go straight when approaching a familiar intersection (Figure 13.10a). Once well learned, such behaviours are reliable and operate effortlessly, reflected in the fact that one's mind often wanders as one traverses the route. However, if you suddenly remember you have to get milk on the way home, you need to override this default path, deviating towards the corner store. Now the same stimulus (the familiar intersection) must lead to a different response (turning left; Figure 13.10b). In this case, the PFC becomes engaged and represents the current rule in a pattern of activity across a population of neurons. Synchrony within the PFC couples the task-relevant neurons into an ensemble, whereas synchrony between regions may act to bias currently relevant associations while suppressing alternatives.

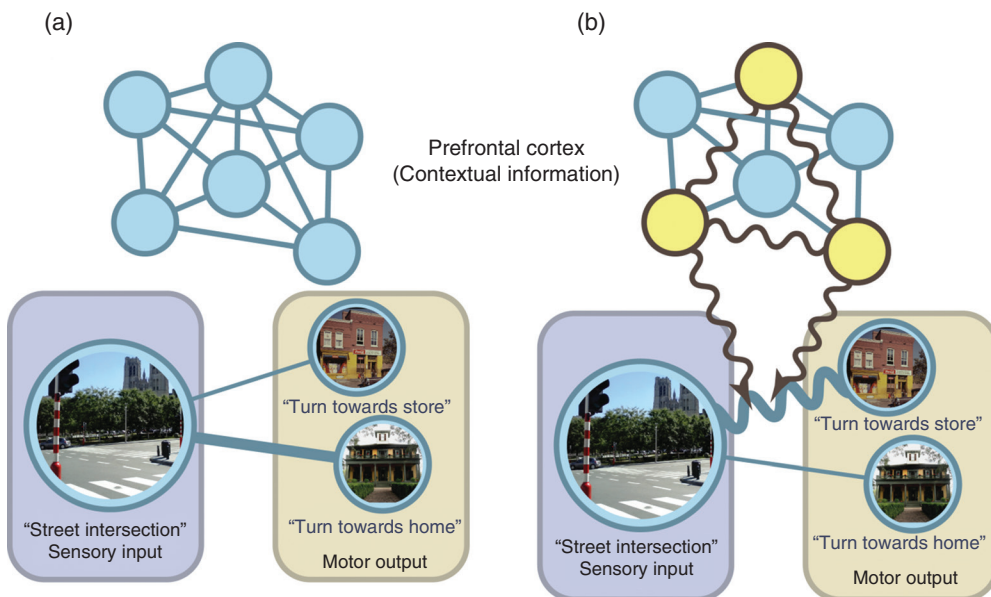


Figure 13.10 Model of cognitive control. (a) At baseline, stimuli lead to habitual responses. (b) If one is in a situation where the context requires alternative responses, then the PFC represents the current task and uses this to bias representations throughout the brain.

Recent neuroimaging experiments provide general support for such a model. For example, Rowe et al. (2005) found that the effective connectivity between the PFC and the posterior cortex changed as subjects switched between performing different tasks. Subjects were asked to perform two ‘free-choice’ tasks: one where they selected a colour of their choosing and one where they selected an action of their choosing. Using structural equation modelling, they measured the effective connectivity between the PFC, motor cortex, and visual cortex during both tasks. Interestingly, they found PFC–motor connectivity during the free-choice action task and greater PFC–visual connectivity during the free-choice colour task. When subjects were instructed as to which colour/action to choose, this was no longer the case, suggesting that effective connectivity is only seen when cognitive control is needed to select among competing representations. Similar results have also been seen during visual discrimination tasks (Chadick & Gazzaley, 2011).

In addition, there has been growing evidence that changes in synchrony within the IPFC may support flexible cognitive control. As reviewed above, Buschman et al. (2012) found that synchrony within the PFC acts to carve task-relevant ensembles of neurons from the PFC population. This could also act to manage the highly dynamic representations within the PFC—if neurons are multidimensional, then understanding their current meaning requires integrating their activity with a larger population (in order to disambiguate between multiple preferred cases). Synchrony could act to bring such a population together, allowing downstream neurons to accurately decode the current state (see also ‘synapsemsbles’; Buzsaki, 2010)

Such changes in synchrony have also been observed across different brain regions. For example, Pesaran et al. (2008) found that synchrony between the PFC and the parietal cortex was increased when animals were making a free choice regarding what response to make in comparison to when they were being instructed as to their choice. Similarly, Buschman and Miller (2007) showed that synchrony between the PFC and the parietal cortex changes depending on whether attention is internally directed or externally captured. Furthermore, Buschman and Miller found that the frequency of synchrony between regions changed: When the PFC was internally directing attention, synchrony increased at ‘beta’ frequencies (~25 Hz). When attention was externally captured, synchrony between brain regions was at a higher frequency (‘gamma’ at ~50 Hz). These results suggest that synchrony at different frequency bands may reflect different types of interactions between regions, a model that has been recently confirmed in the visual and auditory cortex (Bastos et al., 2015; Fontolan, Morillon, Liegeois-Chauvel, & Giraud, 2014).

Together, these results suggest that changes in synchronous oscillations could directly alter the effective connectivity within and between brain regions. Furthermore, changes at different frequency bands could impact the direction in which information flows. For example, if the brain needed to boost bottom-up information flow from the visual cortex to the parietal cortex, then it would synchronise these two regions at gamma frequencies. Similarly, if the brain needed to increase top-down information flow from the frontal cortex to the parietal cortex, it would increase beta-band synchrony between these regions. It will be important to consider how STSP could play a parallel, or even mechanistically supporting, role for inter-regional coherence.

Discussion Topics/Future Directions

In this chapter, we have reviewed evidence that the PFC plays a central role in flexible cognitive control. We have highlighted the dynamic nature of PFC representations and how such representations may support the flexibility needed in high-level cognition. Finally, we

reviewed evidence that the PFC can use these dynamic representations to bias activity throughout the brain in order to achieve the task at hand. In particular, we highlighted the importance of flexible networks, suggesting that either short-term changes in synaptic plasticity or changes in synchrony could be effective mechanisms for temporarily altering the connectivity between and within populations of neurons.

However, there are several remaining questions. First, it is not clear how the dynamics observed within and between brain regions relate to the specific underlying anatomy. In this chapter, we have assumed that changes in effective connectivity act within the framework of the available anatomical connections. However, it remains to be seen to what degree this is true—what proportion of connections are malleable in this way? It is also unclear what neural mechanisms drive changes in effective connectivity. For example, if changes in synchrony are important for guiding the flow of information within and between networks, what are the neural mechanisms that alter this synchrony? How are specific patterns of effective connectivity established and maintained?

Cognitive neuroscience has come a long way since the first PFC recordings in awake, behaving monkeys (e.g., Fuster & Alexander, 1971). In this chapter, we have attempted to highlight the dynamic nature of PFC representations, and consider potential mechanisms that may allow for prefrontal neurons to alter responses throughout the brain to support the current behaviour. Together, these ideas could help shed some light on the most important aspects of intelligent behaviour: cognitive flexibility.

References

- Barak, O., Tsodyks, M., & Romo, R. (2010). Neuronal population coding of parametric working memory. *Journal of Neuroscience*, *30*(28), 9424–9430.
- Barbas, H., Hilgetag, C. C., Saha, S., Dermon, C. R., & Suski, J. L. (2005). Parallel organization of contralateral and ipsilateral prefrontal cortical projections in the rhesus monkey. *BMC Neuroscience*, *6*, 32.
- Barbas, H., & Pandya, D. N. (1989). Architecture and intrinsic connections of the prefrontal cortex in the rhesus monkey. *Journal of Comparative Neurology*, *286*(3), 353–375.
- Barcelo, F., & Knight, R. T. (2002). Both random and perseverative errors underlie WCST deficits in prefrontal patients. *Neuropsychologia*, *40*(3), 349–356.
- Bastos, A. M., Vezoli, J., Bosman, C. A., Schoffelen, J. M., Oostenveld, R., Dowdall, J. R.,...Fries, P. (2015). Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron*, *85*(2), 390–401.
- Bernacchia, A., Seo, H., Lee, D., & Wang, X. J. (2011). A reservoir of time constants for memory traces in cortical neurons. *Nature Neuroscience*, *14*(3), 366–372.
- Buckley, M. J., Mansouri, F. A., Hoda, H., Mahboubi, M., Browning, P. G., Kwok, S. C.,...Tanaka, K. (2009). Dissociable components of rule-guided behavior depend on distinct medial and prefrontal regions. *Science*, *325*(5936), 52–58.
- Buonomano, D. V., & Maass, W. (2009). State-dependent computations: Spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, *10*(2), 113–125.
- Buschman, T. J., Denovellis, E. L., Diogo, C., Bullock, D., & Miller, E. K. (2012). Synchronous oscillatory neural ensembles for rules in the prefrontal cortex. *Neuron*, *76*(4), 838–846.
- Buschman, T. J., & Miller, E. K. (2007). Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science*, *315*(5820), 1860–1862.
- Buschman, T. J., & Miller, E. K. (2014). Goal-direction and top-down control. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *369*(1655).
- Buzsaki, G. (2010). Neural syntax: Cell assemblies, synapses, and readers. *Neuron*, *68*(3), 362–385.
- Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, *51*(13), 1484–1525.

- Chadick, J. Z., & Gazzaley, A. (2011). Differential coupling of visual cortex with default or frontal-parietal network based on goals. *Nature Neuroscience*, *14*(7), 830–832.
- Crowe, D. A., Averbeck, B. B., & Chafee, M. V. (2010). Rapid sequences of population activity patterns dynamically encode task-critical spatial information in parietal cortex. *Journal of Neuroscience*, *30*(35), 11640–11653.
- Deco, G., Rolls, E. T., & Romo, R. (2010). Synaptic dynamics and decision making. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(16), 7545–7549.
- Duncan, J. (2010). The multiple-demand (MD) system of the primate brain: Mental programs for intelligent behaviour. *Trends in Cognitive Sciences*, *14*(4), 172–179.
- Duncan, J., & Miller, E. K. (2002). Cognitive focusing through adaptive neural coding in the primate prefrontal cortex. In D. Stuss & R. T. Knight (Eds.), *Principles of frontal lobe function* (pp. 278–291). New York: Oxford University Press.
- Eglin, M., Robertson, L. C., & Knight, R. T. (1991). Cortical substrates supporting visual search in humans. *Cerebral Cortex*, *1*(3), 262–272.
- Fontolan, L., Morillon, B., Liegeois-Chauvel, C., & Giraud, A. L. (2014). The contribution of frequency-specific activity to hierarchical information processing in the human auditory cortex. *Nature Communications*, *5*, 4694.
- Fries, P. (2005). A mechanism for cognitive dynamics: Neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, *9*(10), 474–480.
- Fujisawa, S., Amarasingham, A., Harrison, M. T., & Buzsaki, G. (2008). Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nature Neuroscience*, *11*(7), 823–833.
- Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology*, *61*(2), 331–349.
- Fuster, J. M. (2001). The prefrontal cortex—An update: Time is of the essence. *Neuron*, *30*(2), 319–333.
- Fuster, J. M. (2002). Frontal lobe and cognitive development. *Journal of Neurocytology*, *31*(3–5), 373–385.
- Fuster, J. M., & Alexander, G. E. (1971). Neuron activity related to short-term memory. *Science*, *173*(3997), 652–654.
- Haber, S. N. (2003). The primate basal ganglia: Parallel and integrative networks. *Journal of Chemical Neuroanatomy*, *26*(4), 317–330.
- Haber, S. N., Kunishio, K., Mizobuchi, M., & Lynd-Balta, E. (1995). The orbital and medial prefrontal circuit through the primate basal ganglia. *Journal of Neuroscience*, *15*(7 Pt 1), 4851–4867.
- Harvey, C. D., Coen, P., & Tank, D. W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature*, *484*(7392), 62–68.
- Jarbo, K., & Verstynen, T. D. (2015). Converging structural and functional connectivity of orbitofrontal, dorsolateral prefrontal, and posterior parietal cortex in the human striatum. *Journal of Neuroscience*, *35*(9), 3865–3878.
- Knight, R. T. (1997). Distributed cortical network for visual attention. *Journal of Cognitive Neuroscience*, *9*(1), 75–91.
- Knight, R. T., Grabowecky, M. F., & Scabini, D. (1995). Role of human prefrontal cortex in attention control. *Advances in Neurology*, *66*, 21–34; discussion 34–26.
- Kunzle, H., & Akert, K. (1977). Efferent connections of cortical, area 8 (frontal eye field) in Macaca fascicularis: A reinvestigation using the autoradiographic technique. *Journal of Comparative Neurology*, *173*(1), 147–164.
- Lebedev, M. A., Messinger, A., Kralik, J. D., & Wise, S. P. (2004). Representation of attended versus remembered locations in prefrontal cortex. *PLoS Biology*, *2*(11), e365.
- Li, L., Gratton, C., Yao, D., & Knight, R. T. (2010). Role of frontal and parietal cortices in the control of bottom-up and top-down attention in humans. *Brain Research*, *1344*, 173–184.
- Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, *503*(7474), 78–84.
- Markov, N. T., Ercsey-Ravasz, M., Van Essen, D. C., Knoblauch, K., Toroczkai, Z., & Kennedy, H. (2013). Cortical high-density counterstream architectures. *Science*, *342*(6158), 1238406.
- McFarland, N. R., & Haber, S. N. (2002). Thalamic relay nuclei of the basal ganglia form both reciprocal and nonreciprocal cortical connections, linking multiple frontal cortical areas. *Journal of Neuroscience*, *22*(18), 8117–8132.

- Meyers, E. M., Freedman, D. J., Kreiman, G., Miller, E. K., & Poggio, T. (2008). Dynamic population coding of category information in inferior temporal and prefrontal cortex. *Journal of Neurophysiology*, *100*(3), 1407–1419.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*, 167–202.
- Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of Neuroscience*, *16*(16), 5154–5167.
- Mongillo, G., Barak, O., & Tsodyks, M. (2008). Synaptic theory of working memory. *Science*, *319*(5869), 1543–1546.
- Moore, T. (2006). The neurobiology of visual attention: Finding sources. *Current Opinion in Neurobiology*, *16*(2), 159–165.
- Moore, T., & Armstrong, K. M. (2003). Selective gating of visual signals by microstimulation of frontal cortex. *Nature*, *421*(6921), 370–373.
- Moore, T., & Fallah, M. (2001). Control of eye movements and spatial attention. *Proceedings of the National Academy of Sciences of the United States of America*, *98*(3), 1273–1276.
- Moore, T., & Fallah, M. (2004). Microstimulation of the frontal eye field and its effects on covert spatial attention. *Journal of Neurophysiology*, *91*(1), 152–162.
- Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X.,...Wang, X. J. (2014). A hierarchy of intrinsic timescales across primate cortex. *Nature Neuroscience*, *17*(12), 1661–1663.
- Nobre, A. C., Correa, A., & Coull, J. (2007). The hazards of time. *Current Opinion in Neurobiology*, *17*(4), 465–470.
- Olivers, C. N., Peters, J., Houtkamp, R., & Roelfsema, P. R. (2011). Different states in visual working memory: When it guides attention and when it does not. *Trends in Cognitive Sciences*, *15*(7), 327–334.
- Pesaran, B., Nelson, M. J., & Andersen, R. A. (2008). Free choice activates a decision circuit between frontal and parietal cortex. *Nature*, *453*(7193), 406–409.
- Petrides, M., & Pandya, D. N. (1999). Dorsolateral prefrontal cortex: Comparative cytoarchitectonic analysis in the human and the macaque brain and corticocortical connection patterns. *European Journal of Neuroscience*, *11*(3), 1011–1036.
- Qi, X. L., Meyer, T., Stanford, T. R., & Constantinidis, C. (2011). Changes in prefrontal neuronal activity after learning to perform a spatial working memory task. *Cerebral Cortex*, *21*(12), 2722–2732.
- Rainer, G., Asaad, W. F., & Miller, E. K. (1998). Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature*, *393*(6685), 577–579.
- Rainer, G., Rao, S. C., & Miller, E. K. (1999). Prospective coding for objects in primate prefrontal cortex. *Journal of Neuroscience*, *19*(13), 5493–5505.
- Rigotti, M., Barak, O., Warden, M. R., Wang, X. J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, *497*(7451), 585–590.
- Rossi, A. F., Bichot, N. P., Desimone, R., & Ungerleider, L. G. (2007). Top-down attentional deficits in Macaques with lesions of lateral prefrontal cortex. *Journal of Neuroscience*, *27*(42), 11306–11314.
- Rowe, J. B., Stephan, K. E., Friston, K., Frackowiak, R. S., & Passingham, R. E. (2005). The prefrontal cortex shows context-specific changes in effective connectivity to motor or visual cortex during the selection of action or colour. *Cerebral Cortex*, *15*(1), 85–95.
- Sakai, K. (2008). Task set and prefrontal cortex. *Annual Review of Neuroscience*, *31*, 219–245.
- Salazar, R. F., Dotson, N. M., Bressler, S. L., & Gray, C. M. (2012). Content-specific fronto-parietal synchronization during visual working memory. *Science*, *338*(6110), 1097–1100.
- Shafi, M., Zhou, Y., Quintana, J., Chow, C., Fuster, J., & Bodner, M. (2007). Variability in neuronal activity in primate cortex during working memory tasks. *Neuroscience*, *146*(3), 1082–1108.
- Shallice, T., & Burgess, P. W. (1991). Deficits in strategy application following frontal lobe damage in man. *Brain*, *114*(Pt. 2), 727–741.
- Siegel, M., Buschman, T. J., & Miller, E. K. (2015). Cortical information flow during flexible sensori-motor decisions. *Science*, *348*(6241), 1352–1355.

- Siegel, M., Warden, M. R., & Miller, E. K. (2009). Phase-dependent neuronal coding of objects in short-term memory. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(50), 21341–21346.
- Sigala, N., Kusunoki, M., Nimmo-Smith, I., Gaffan, D., & Duncan, J. (2008). Hierarchical coding for sequential task events in the monkey prefrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *105*(33), 11969–11974.
- Stokes, M. G. (2011). The spatiotemporal structure of population coding in monkey parietal cortex. *Journal of Neuroscience*, *31*(4), 1167–1169.
- Stokes, M. G. (2015). ‘Activity-silent’ working memory in prefrontal cortex: A dynamic coding framework. *Trends in Cognitive Sciences*, *19*(7), 394–405.
- Stokes, M. G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., & Duncan, J. (2013). Dynamic coding for cognitive control in prefrontal cortex. *Neuron*, *78*(2), 364–375.
- Sugase-Miyamoto, Y., Liu, Z., Wiener, M. C., Optican, L. M., & Richmond, B. J. (2008). Short-term memory trace in rapidly adapting synapses of inferior temporal cortex. *PLoS Computational Biology*, *4*(5), e1000073.
- Taylor, P. C., Nobre, A. C., & Rushworth, M. F. (2007). FEF TMS affects visual cortical activity. *Cerebral Cortex*, *17*(2), 391–399.
- Vapnik, V. N. (1998). *Statistical learning theory*. New York: Wiley.
- Wallis, J. D., Anderson, K. C., & Miller, E. K. (2001). Single neurons in prefrontal cortex encode abstract rules. *Nature*, *411*(6840), 953–956.
- Wang, M., Yang, Y., Wang, C. J., Gamo, N. J., Jin, L. E., Mazer, J. A.,...Arnsten, A. F. (2013). NMDA receptors subserve persistent neuronal firing during working memory in dorsolateral prefrontal cortex. *Neuron*, *77*(4), 736–749.
- Watanabe, K., & Funahashi, S. (2007). Prefrontal delay-period activity reflects the decision process of a saccade direction during a free-choice ODR task. *Cerebral Cortex*, *17*(Suppl. 1), i88–100.
- Watanabe, K., & Funahashi, S. (2014). Neural mechanisms of dual-task interference and cognitive capacity limitation in the prefrontal cortex. *Nature Neuroscience*, *17*(4), 601–611.
- White, I. M., & Wise, S. P. (1999). Rule-dependent neuronal activity in the prefrontal cortex. *Experimental Brain Research*, *126*(3), 315–335.
- Wolff, M., Ding, J., Myers, N., & Stokes, M. G. (2015). Revealing hidden states in working memory using electroencephalography. *Frontiers in Neuroscience*.
- Zucker, R. S., & Regehr, W. G. (2002). Short-term synaptic plasticity. *Annual Review of Physiology*, *64*, 355–405.