


The “working” of working memory

Earl K. Miller, PhD



This review examines the evidence for a neurobiological explanation of executive functions of working memory. We suggest that executive control stems from information about task rules acquired by mixed selective, adaptive coding, multifunctional neurons in the prefrontal cortex. The output of these neurons dynamically links the cortex-wide networks needed to complete the task. The linking may occur via synchronizing of neural rhythms, which may explain why we have a limited capacity for simultaneous thought.

© 2013, AICH – Servier Research Group

Dialogues Clin Neurosci. 2013;15:411–418.

Keywords: basal ganglia; cognition; executive function; oscillation; prefrontal cortex; synchrony; working memory

Author affiliations: Picower Institute for Learning and Memory and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

Address for correspondence: Earl K. Miller, Picower Professor of Neuroscience, Picower Institute for Learning and Memory and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, MA 02139, USA (e-mail: ekmiller@mit.edu)

Working memory is, essentially, what we think of as thought. It is our mental sketchpad, where we hold information “in mind” and process it. Naturally, it has garnered much empirical interest, and this has yielded a robust and commonly reported neural correlate: sustained neural activity. Sustained activity can be seen when humans and animals are performing tasks thought to engage working memory. Delayed response tasks, for example, include a short gap in time (seconds) between a sensory cue and the opportunity to act based on that cue. Higher cortical areas, especially the prefrontal cortex (PFC), the putative “executive” cortex, show elevated levels of neural activity over that delay, as if the neurons are bridging the gap by sustaining their firing to cue. Because short-term buffering of information in an active “online” state is a keystone of working memory, sustained delay activity has become virtually synonymous with “working memory,” at least to neuroscientists. However, it is important to keep in mind that Baddeley’s original working memory model was meant to be a model of cognition and that there is more to thought than short-term memory (*Figure 1*).¹ Imagine planning a simple errand. You do not just hold elements of the plans in mind; you weigh alternatives, make decisions, and order the thoughts until you think the plan achieves your objectives. In short, there is more to working memory than “memory”; there is also the “working.” Indeed, what sets working memory apart from mere short-term storage and elevates it to a model of cognition writ large is the inclusion of a “central executive” (*Figure 1*), a set of mechanisms that together act to manage and regulate what we hold “in mind” (ie, contents of the short-term memory buffers). These executive functions are less well understood because they are less tractable than short-term buffering of information. But we have made

Translational research

progress. Here, we review work on the neural correlates of working memory and suggest candidate mechanisms.

Executive control, rules, and the prefrontal cortex

By definition, controlled thought and action are goal-directed and organized toward the completion of tasks. Consider a common cognitively demanding situation: navigating an airport. From the start, we know that we need a ticket, have to wait in line, board at the right gate, etc. We are not born knowing this; we have to learn the rules. As such, the neural substrates for executive control need to have access to the wide range of information needed to identify potential goals and the rules that can achieve them. This no doubt depends on many different brain areas. However, one cortical region is particularly necessary (but not sufficient): the PFC. It is this cortical area that reaches the greatest relative size in the human brain and is thus thought to be the neural instantiation of the mental qualities that we think of as “intelligent.”

The PFC is anatomically well situated to play a role as the brain’s executive. It receives information from, and sends projections to, forebrain systems that process information about the external world, motor system structures that produce voluntary movement, systems that consolidate long-term memories, and systems that process information about affect and motivational state.²⁻⁵ This anatomy has long suggested that the PFC may be important for synthesizing the external and internal information needed to produce complex behavior.

Neurophysiological studies suggest that this synthesis serves to form representations of task rules (for reviews see refs 6-8). This has been shown in studies that systematically vary task demands; subjects perform a different set of operations or make different decisions using

the same set of sensory inputs and motor options. For example, in one trial the subject may have to choose one of two pictures that matches one seen previously (a match rule); in another trial the subject has to choose the nonmatching picture (a nonmatch rule).⁹ These types of experiments have revealed that the PFC neural activity is highly sensitive to rule information. In fact, unlike sensory cortex, especially primary sensory cortex, it appears that task rules are more influential on how information is distributed across PFC neurons than bottom-up sensory information. More neurons reflect task demands than sensory information, indeed often at the expense of sensory information.¹⁰⁻¹¹ Interestingly, cognitively demanding tasks engage a very large proportion of PFC neurons; after training, as many as 30% to 40% of randomly selected PFC neurons show task-related activity.^{9,12-15}

So many PFC neurons (one third or more of the population) dedicated to a given rule might, at first blush, make it seem as if the PFC can only learn a few tasks. If the one third of PFC neurons represent the rules of one task, does that mean that only three tasks can be learned? In fact the opposite is true. Many PFC neurons are multitaskers with “mixed selectivity.”¹⁰ This mixed selectivity does not fit into the traditional view of brain function in which individual neurons have been thought to be specialized for single functions. Instead, in the PFC, neural specialization waters down in a mix of disparate information; there is no obvious function that unites the variety of information signaled by the individual neurons. Why this mixed selectivity, and why so many neurons? The answer is that large proportions of mixed selectivity neurons expand the brain’s computational power, increasing the complexity and number of task rules that can be learned, and speeding up their acquisition.^{16,17} The high dimensionality of the representational space they support allows learning algorithms to converge more quickly and reduces the plasticity mechanisms needed. Because mixed selectivity neurons already have a mixture of task-relevant information, only the readout neurons have to be modified during learning. In short, mixed selectivity amplifies our ability to quickly learn (and flexibly implement) complex rules.^{16,18}

Thus, the PFC seems to be a neural substrate ideal for absorbing the constellation of disparate information that forms rules. But how exactly does rule information exert control? Miller and Cohen⁸ suggested a possibility. Their




Figure 1. Baddeley’s working memory model. This includes short-term memory buffers (visuospatial and phonological loops) under the control of a central executive. Many experiments in neuroscience have focused on the short-term memory loops; here the executive is focused on.

central idea is that PFC rule representations are not esoteric descriptions of the logic of a task. Rather, the rules are represented in a particular format: as a map of the cortical pathways needed to perform the task (“rulemaps”—*Figure 2*). In other words, a task’s rules in the PFC are also maps of the neural pathways within and between other cortical regions that need to be engaged to solve the current task. In a given situation, cues about the current situation (context) and other external and internal cues activate and complete the PFC rulemap that includes that information as well as the course of action that has proven successful in the past. Rulemap activation (which can be sustained, if needed) sets up bias signals that feed back to other brain areas, affecting sensory systems as well as the systems responsible for response execution, memory retrieval, and emotional evaluation. The cumulative effect is the selection of the pattern of neural circuits that guide the flow of neural activity along the proper mappings between inputs, internal states, and outputs to reach the goal. It is as if the PFC is a conductor in a railroad yard

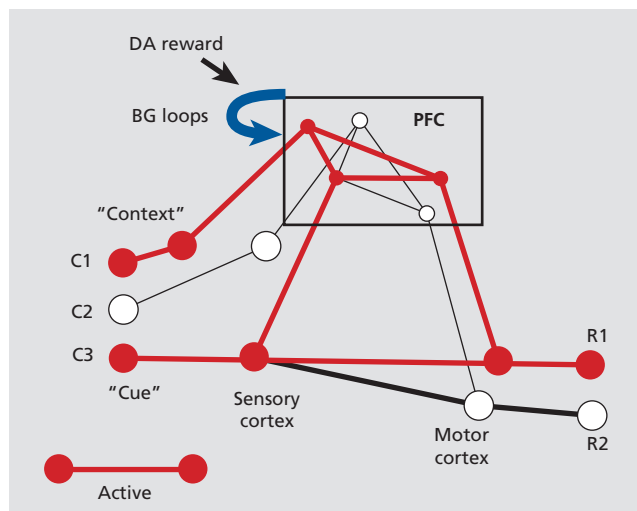


Figure 2. Miller and Cohen model of executive control. Shown are processing units representing cues such as sensory inputs, current motivational state, memories, etc (C1, C2, and C3), and those representing two voluntary actions (eg, “responses”, R1 and R2). Excitatory signals from the prefrontal cortex (PFC) feed back to other brain systems to enable task-relevant neural pathways. Thick lines indicate well-established pathways mediating a prepotent behavior that can be overcome by top-down PFC signals that activate an alternative pathway. Red indicates active units or pathways. BG, basal ganglia; DA, dopamine. Adapted from ref 8: Miller EK, Cohen JD. An integrative theory of prefrontal cortex function. *Annu Rev Neurosci.* 2001;24:167-202. Copyright © Annual Reviews Inc

and learns a map that it uses to guide trains (neural activity) along the right tracks (neural pathways). Next, we consider how these rulemaps are acquired.

Teaching by dopamine

You can not learn rules unless you have some idea about the consequences of your actions. Simply put, the brain must strengthen neural connections that are successful at achieving a goal (rewarded), while breaking or weakening those that are ineffective. Dopamine (DA) neurons in the midbrain (ventral tegmental area and the substantia nigra, pars compacta) may provide this teaching signal.

At first, DA neurons activate to unexpected rewards, but then after a repeated pairing of a cue (eg, “bell”) with a reward (eg, “dinner”), they stop activating to the reward and activate to the cue as if the cue is a “stand-in” for the reward.¹⁹ Add another cue (eg, a light flash) that predicts the first cue (bell) and after a number of pairings the DA neurons will now activate to the light and no longer to the bell or dinner. Thus, DA neurons respond to the earliest unexpected event in a chain of events that are known to end in reward. They also pause their firing when an expected reward is withheld. Thus, their activity seems to correspond to prediction error signals in models of animal learning.²⁰ They are essential teaching signals that say “something good happened and you did not predict it, so remember what just happened so that you can predict it in the future.” As the organism learns and becomes an increasingly better predictor of what will lead to reward, DA neurons will activate progressively earlier, linking in the network of information needed to navigate toward that reward. The PFC is a main target of midbrain DA neurons.^{21,22}

Balancing different styles of learning

Normal learning has to find a balance between different demands. It is obvious that learning things quickly is often advantageous. You want to learn to get to the resources faster than your competitors. But there are also disadvantages in that fast learning: it is error-prone. If, for example, you have one-trial learning, you may mistake a coincidence for a real predictive relationship. Consider taste aversion. We often develop distaste for a food simply because we became ill after we ate it, even when that food had nothing to do with our illness. With

Translational research

slower learning rates, more experience can be taken into account, and this allows organisms to detect the regularities that indicate predictive relationships and leave behind spurious associations and coincidences. Further, slower, more deliberate learning also provides the opportunity to detect common structures across different experiences. It is these commonalities that form the abstractions, general principles, concepts, etc critical for sophisticated thought. We learn the concept of “fairness” from specific instances of being treated fairly or unfairly. Given the advantages and disadvantages associated with both forms of learning, the brain must balance the obvious pressure to learn as quickly as possible with the advantages of slower learning. The key to this may be balance and interactions between the PFC and the basal ganglia (BG).

The BG is a collection of subcortical nuclei that, similar to the PFC, have a high degree of convergence of cortical inputs. The frontal cortex receives the largest portion of BG outputs (via the thalamus), suggesting a close collaboration between the BG and frontal cortex.^{23,24} DA acts not only on the PFC but also on the BG. Importantly, the DA projections into the striatum (the input to the BG where cortical information converges) are much heavier than those to the PFC.²⁵ Thus, DA teaching signals may play a stronger role in gating plasticity in the striatum in contrast to the PFC, where DA influence may be more subtle—shading, not gating, plasticity. This may explain our observation that during operant learning, learning-related changes in the striatum appear sooner and faster than those in the PFC.²⁶ Thus, the trade-off between the advantages of slow plasticity versus fast plasticity may play out in interactions between the PFC and BG.²⁷

The idea is that during learning, specific associations between cues and immediate actions are quickly acquired by the striatum, by virtue of its heavy inputs from midbrain DA neurons. The output of the basal ganglia trains the PFC²⁶ where plasticity is “slower” (smaller changes in synaptic weights with each learning episode) because of the weaker DA influence. As a result, the PFC gradually builds up less error-prone, more elaborate, and generalized representations based on the patterns fed to it by the BG. This may explain why the PFC and BG seem to operate based on different types of representational schemes.²⁸ The fast striatal plasticity may be more suited for a quick stamping-in of immediate, direct associations. But, as a consequence, the striatum

learns complex behaviors in a piecemeal fashion, as a set of largely unconnected (cache) representations of which alternative was more successful at each decision point in the task.²⁸ By contrast, the slow PFC plasticity may be suitable for building elaborate rule representations that gradually link in more information (ie, tree-like representations).²⁸ The slow PFC plasticity may also find the commonalities and regularities among the simpler representations acquired by the striatum that are the basis for abstractions and general principles.²⁹ In other words, the striatum learns the pieces of the puzzle while the PFC puts the puzzle together.²⁷

We recently witnessed this hand-off from the striatum to PFC as animals transitioned from simple, specific learning to more generalized, abstract representation.³⁰ Each day, monkeys learned to associate two novel, abstract, dot-based categories with a right vs left saccade (*Figures 3A and 3B*). At first, monkeys only saw a few examples of each category, and they could learn specific stimulus-response (S-R) associations by rote. But as more and more new category exemplars were added, the capacity to memorize specific S-R associations was exceeded. To solve the task, the monkeys then had to learn the categories, and extract the common essence that united exemplars from the same group. As predicted, we saw neurophysiological evidence for a transfer of control of the task from the BG to the PFC over learning. Early, during S-R, learning striatum activity was an earlier predictor of the corresponding saccade (*Figure 3C*). However, as the number of exemplars exceeded the capacity of S-R learning and animals had to learn the categories, the PFC took over and began predicting the saccade before the striatum (*Figure 3C*). This dual-learner strategy allows the animal to perform optimally throughout the task; early on the striatum can learn associations quickly, while later in the task, when learning associations is no longer viable, PFC guides behavior.

Primates, especially humans, can learn a wide range of abstract categories like “peace.” But truly intelligent behavior depends on more than finding high-level structure across experiences. Humans can be creative and unique in finding new goals and strategies to pursue them. This means that the mechanisms that build the PFC rule representations must have a corresponding ability for open-ended growth. Another aspect of PFC-BG interactions may explain this. Anatomical loops between them may support recursive, bootstrapping interactions.

Recursivity and bootstrapping

The anatomical connections between the PFC and BG suggest “bootstrapping,” the process of building increasingly complex representations from simpler ones. The PFC and the BG form closed loops: channels within the BG return outputs, via the thalamus, into the same cortical areas that gave rise to their initial cortical input. This suggests some form of recursive processing.^{23,24} That

is, the neural representations that result from plasticity within and between the PFC and BG form cortical representations that can be fed back into the loop as fodder for further elaboration. In this manner, new experiences can be added onto previous ones, linking in more information to build more elaborate rule representations. It can allow the discovery of commonalities among more experiences and thus more high-level concepts and principles. Indeed, we often ground new concepts in familiar

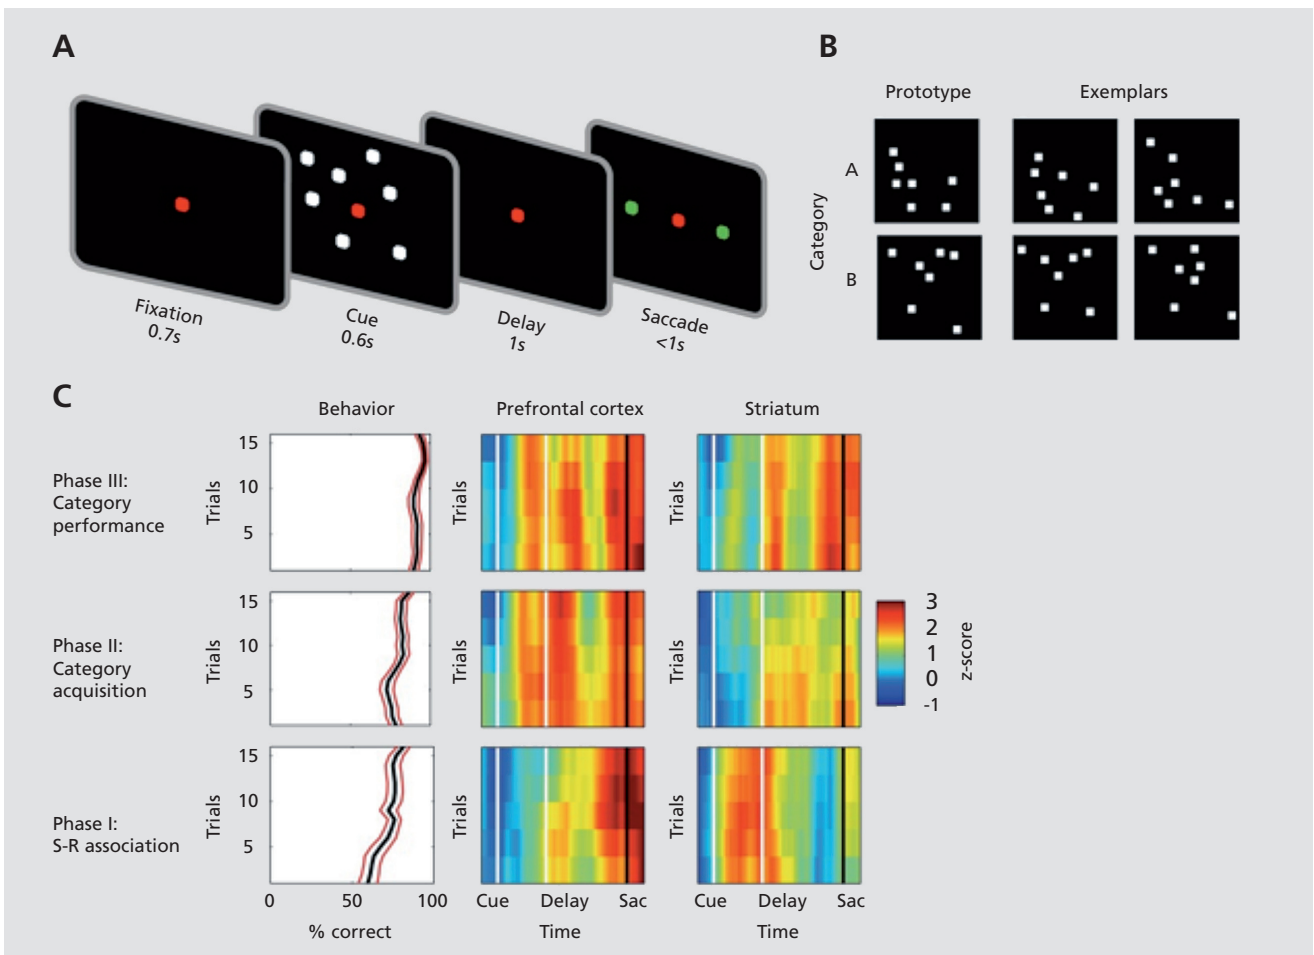


Figure 3. Specific vs generalized learning in basal ganglia vs prefrontal cortex. (A) Category-response association task. Animals were presented with a cue stimulus (an array of dots) that were exemplar stimuli created by morphing stimuli from one of two category prototypes. (B) The animals learned to associate each category with either a leftward or rightward saccade. After learning the associations for a given set of exemplars, new exemplars were added to the set, requiring the animal to generalize the learned association to these new stimuli. Earlier in training the animals were able to associate individual stimuli with a response (C, bottom row) but then, as the number of exemplars increased, they were forced to acquire the category (C, middle row), eventually generalizing to new stimuli (C, top row). Selectivity for the associated response was found in both prefrontal cortex and striatum. Early in training, during the stimulus-response phase, selectivity was seen earlier and stronger in striatum. However, prefrontal cortex took a lead role after generalized categories were learned (phase II and III). S-R, stimulus-response
Adapted from ref 30: Antzoulatos EG, Miller EK. Differences between neural activity in prefrontal cortex and striatum during learning of novel, abstract categories. *Neuron*. 2011;71:243-249. Copyright © Cell Press 2011

Translational research

ones because it seems to ease our understanding of novel ideas; we learn multiplication by serial addition, exponentiation by serial multiplication, etc.

The frontal cortex–BG loops also suggest an auto-associative type network, similar to that seen in the CA3 cell layer of the hippocampus. The looping back of outputs allow the network to learn to complete (ie, recall) previously learned patterns given a degraded version or a subset of the original inputs.³¹ Given the DA influence, the PFC-BG loops may be more goal-oriented (supervised) than hippocampal learning and memory. They could even explain the DA reward prediction signals. As previously described, midbrain DA neurons respond to earlier and earlier events in a predictive chain leading to a reward. Both the frontal cortex and the striatum send projections into the midbrain DA neurons, possibly underlying their ability to bootstrap to early predictors of reward.

The loops may also explain another important aspect of goal-directed behavior: the sequencing of thought and action. A key feature of auto-associative networks is their ability to learn temporal sequences of patterns and thus make predictions. This relies on the activity pattern being fed back into the network with a temporal lag, allowing the next pattern in the sequence to arrive as the previous pattern is fed back, building an association.^{32,33} Inhibitory synapses in the pathways through the BG may add the temporal delay needed, as they have a slower time constant than excitatory synapses. Another way to add lag is via a memory buffer. As highlighted earlier, the PFC is well-known for this type of property; its neurons can sustain their activity to act as a bridge for learning contingencies across several seconds, even minutes. The introduction of lag into the recursive loop through either mechanism (or both) may be enough to allow sequencing and prediction. This would seem to be key to the development of tree-like rule representations that describe an entire sequence of goal-directed actions (discussed above).

A brief word about capacity limitations

Despite the remarkable power and flexibility of human cognition, our working memory—the “online” workspace that most cognitive mechanisms depend upon—is surprisingly limited. An average adult human has a capacity to retain only four items at a given time.^{34,36} Why this limited capacity? The answer may lie in the mixed selectiv-

ity neurons that amplify the brain’s computational power (as previously discussed). Mixed selectivity neurons that participate in many different functions would seem to create problems. Don’t downstream neurons sometimes receive signals that are irrelevant or counterproductive? A solution is oscillatory brain rhythms.

It has long been known that brain waves (coordinated oscillations among many neurons) vary their frequency with cognitive focus. Oscillations create synchronous spikes that can have a greater impact than unsynchronized spikes, as they all arrive simultaneously at downstream neurons. They could allow neurons to communicate different messages to different targets depending on those with which they are synchronized (and how, eg, phase, frequency).

Evidence for this comes from a variety of studies. Different frequency synchronization between human cortical areas supports recollection of spatial vs temporal information.³⁷ Different phases of cortical oscillations preferentially signal different pictures simultaneously held in short-term memory.³⁸ Monkey frontal and parietal cortices synchronize more strongly at lower vs higher frequency for top-down vs bottom-up attention, respectively.³⁹ Entraining the human frontal cortex at those frequencies produces the predicted top-down vs bottom-up effects on behavior.⁴⁰ Thus, activity from the same neurons has different functional outcomes depending on their rhythmic dynamics.

This suggests that our brain does not operate continuously, but rather discretely, with pulses of activity routing packets of information.⁴¹ Such discrete cycles would provide a backbone for coordinating computations (and their results) across disparate networks. They can provide a substrate via which the PFC can “direct traffic,” guiding the flow of neural activity along pathways that establish the proper mappings between inputs, internal states, and outputs needed to perform a given task. However, it comes at a cost: oscillations are naturally limited in bandwidth; only so many things can be computed or carried in a single oscillatory cycle. This can explain the most fundamental property of consciousness, the limited capacity for simultaneous thought. Interestingly, Duncan and colleagues have linked individual differences in fluid intelligence to each person’s working memory capacity for task rules.⁴² This suggests that fluid intelligence may depend on how much rule information from mixed selectivity neurons can be packed into an oscillatory cycle.

Summary

Here we have reviewed evidence and suggested mechanisms and substrates to help provide a neurobiological explanation for executive functions—that is, neurobiological rather than homuncular. We have discussed how interactions and balance between different styles of plasticity in the PFC and BG acquire the rules of the game needed to organize goal-directed thought and action. The computational power to quickly learn, store, and flexibly implement the large number of complex rules may be provided by large proportions of mixed selectivity, adaptive multifunction neurons (and other higher cortical areas). Synchronization of oscillatory rhythms between

neurons in local and global networks may disambiguate the output of the mixed selectivity neurons, allowing them to selectively participate in different networks with different functions by virtue of synchrony at different frequencies, phases, etc. Executive control may result when rule information in the PFC dynamically establishes networks that link together the corresponding information throughout the cortex. If oscillatory synchrony indeed plays this role, it could explain why conscious thought is so limited in bandwidth. Any oscillatory signal has a natural bandwidth limit; only so much information can be packed into a cycle. And with a limited bandwidth, it is critical to have executive functions that can single-mindedly focus those limited resources on the task at hand. □

El “trabajo” de la memoria de trabajo

Este artículo revisa la evidencia para la explicación neurobiológica de las funciones ejecutivas de la memoria de trabajo. Se sugiere que el control ejecutivo deriva de la información sobre las reglas de tareas adquiridas a través de neuronas multifuncionales, con selectividad combinada y con adaptación a códigos de la corteza prefrontal. El resultado de estas neuronas es vincular dinámicamente las redes de toda la corteza necesarias para completar la tarea. La interacción puede ocurrir mediante la sincronización de los ritmos neuronales, lo que puede explicar el porqué tenemos una capacidad limitada para el pensamiento simultáneo.

Le « travail » de la mémoire de travail

Cet article analyse les arguments d’une explication neurobiologique des fonctions exécutives de la mémoire de travail. Nous pensons que le contrôle exécutif naît de l’information des consignes acquise par des neurones multifonctionnels, sélectifs, hétérogènes et de codage adaptatif présents dans le cortex préfrontal. Ces neurones ont des connexions efférentes dynamiques avec des réseaux présents dans tout le cortex et nécessaires à l’accomplissement de la tâche. Ces connexions apparaîtraient par synchronisation des rythmes neuronaux, ce qui expliquerait notre capacité limitée à avoir des pensées simultanées.

REFERENCES

1. Baddeley A. Working memory. *Science*. 1992;255:556-559.
2. Barbas H, de Olmos J. Projections from the amygdala to basoventral and mediodorsal prefrontal regions in the rhesus monkey. *J Comp Neurol*. 1990;300:549-571.
3. Croxson PL, Johansen-Berg H, Behrens TEJ, et al. Quantitative investigation of connections of the prefrontal cortex in the human and macaque using probabilistic diffusion tractography. *J Neurosci*. 2005;25:8854-8866.
4. Porrino LJ, Crane AM, Goldman-Rakic PS. Direct and indirect pathways from the amygdala to the frontal lobe in rhesus monkeys. *J Comp Neurol*. 1981;198:121-136.
5. Selemon LD, Goldman-Rakic PS. Common cortical and subcortical targets of the dorsolateral prefrontal and posterior parietal cortices in the rhesus monkey: evidence for a distributed neural network subserving spatially guided behavior. *J Neurosci*. 1988;8:4049-4068.
6. Fuster JM. Frontal lobes. *Curr Opin Neurobiol*. 1993;3:160-165.
7. Fuster J. *The Prefrontal Cortex*. 4th ed. London, UK: Elsevier; 2008.
8. Miller E, Cohen J. An integrative theory of prefrontal cortex function. *Annu Rev Neurosci*. 2001;167-202.
9. Wallis JD, Anderson KC, Miller EK. Single neurons in prefrontal cortex encode abstract rules. *Nature*. 2001;411:953-956.
10. Cromer JA, Roy JE, Miller EK. Representation of multiple, independent categories in the primate prefrontal cortex. *Neuron*. 2010;66:796-807.
11. Roy JE, Riesenhuber M, Poggio T, Miller EK. Prefrontal cortex activity during flexible categorization. *J Neurosci*. 2010;30:8519-8528.
12. Asaad WF, Rainer G, Miller EK. Task-specific neural activity in the primate prefrontal cortex. *J Neurophysiol*. 2000;84:451-459.
13. Mansouri FA, Matsumoto K, Tanaka K. Prefrontal cell activities related to monkeys’ success and failure in adapting to rule changes in a Wisconsin card sorting test analog. *J Neurosci*. 2006;26:2745-2756.
14. White IM, Wise SP. Rule-dependent neuronal activity in the prefrontal cortex. *Exp Brain Res*. 1999;126:315-335.
15. Bunge SA, Kahn I, Wallis JD, Miller EK, Wagner AD. Neural circuits subserving the retrieval and maintenance of abstract rules. *J Neurophysiol*. 2003;90:3419-3428.
16. Rigotti M, Barak O, Warden MR, et al. The importance of mixed selectivity in complex cognitive tasks. *Nature*. 2013;497:585-590.
17. Rigotti M, Rubin DBD, Wang XJ, Fusi S. Internal representation of task rules by recurrent dynamics: the importance of the diversity of neural responses. *Front Comput Neurosci*. 2010;4:24.

Translational research

18. Duncan J, Miller EK. Cognitive focus through adaptive neural coding in the primate prefrontal cortex. In: Stuss DT, Knight RT, eds. *Principles of Frontal Lobe Function*. New York, NY: Oxford University Press; 2002;278-291.
19. Schultz W, Apicella P, Ljungberg T. Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. *J Neurosci*. 1993;13:900-913.
20. Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science*. 1997;275:1593-1599.
21. Goldman-Rakic PS, Leranth C, Williams SM, Mons N, Geffard M. Dopamine synaptic complex with pyramidal neurons in primate cerebral cortex. *Proc Natl Acad Sci U S A*. 1989;86:9015-9019.
22. Thierry AM, Blanc G, Sobel A, Stinus L, Glowinski J. Dopaminergic terminals in the rat cortex. *Science*. 1973;182:499-501.
23. Middleton FA, Strick PL. Basal ganglia output and cognition: evidence from anatomical, behavioral, and clinical studies. *Brain Cogn*. 2000;42:183-200.
24. Middleton FA, Strick PL. Basal-ganglia "projections" to the prefrontal cortex of the primate. *Cereb Cortex*. 2002;12:926-935.
25. Lynd-Balta E, Haber SN. The organization of midbrain projections to the ventral striatum in the primate. *Neuroscience*. 1994;59:609-623.
26. Pasupathy A, Miller EK. Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature*. 2005;433:873-876.
27. Miller E, Buschman T. Rules through recursion: how interactions between the frontal cortex and basal ganglia may build abstract, complex rules from concrete, simple ones. In: Bunge SA, Wallis JD, eds. *Neuroscience of Rule-Guided Behavior*. New York, NY: Oxford University Press; 2007;419-440.
28. Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ. Cortical substrates for exploratory decisions in humans. *Nature*. 2006;441:876-879.
29. Seger CA, Miller EK. Category learning in the brain. *Annu Rev Neurosci*. 2010;33:203-219.
30. Antzoulatos EG, Miller EK. Differences between neural activity in prefrontal cortex and striatum during learning of novel abstract categories. *Neuron*. 2011;71:243-249.
31. Hopfield JJ. Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci U S A*. 1982;79:2554-2558.
32. Kleinfeld D. Sequential state generation by model neural networks. *Proc Natl Acad Sci U S A*. 1986;83:9469-9473.
33. Sompolinsky H, Kanter I. Temporal association in asymmetric neural networks. *Phys Rev Lett*. 1986;57:2861-2864.
34. Alvarez GA, Cavanagh P. The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychol Sci*. 2004;15:106-111.
35. Luck SJ, Vogel EK. The capacity of visual working memory for features and conjunctions. *Nature*. 1997;390:279-281.
36. Vogel EK, Machizawa MG. Neural activity predicts individual differences in visual working memory capacity. *Nature*. 2004;428:748-751.
37. Watrous AJ, Tandon N, Conner CR, Pieters T, Ekstrom AD. Frequency-specific network connectivity increases underlie accurate spatiotemporal memory retrieval. *Nat Neurosci*. 2013;16:349-356.
38. Siegel M, Warden M, Miller E. Phase-dependent neuronal coding of objects in short-term memory. *Proc Natl Acad Sci U S A*. 2009;106:21341-21346.
39. Buschman TJ, Miller EK. Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science*. 2007;315:1860-1862.
40. Chanes L, Quentin R, Tallon-Baudry C, Valero-Cabr e A. Causal frequency-specific contributions of frontal spatiotemporal patterns induced by non-invasive neurostimulation to human visual performance. *J Neurosci*. 2013;33:5000-5005.
41. Buschman TJ, Miller EK. Shifting the spotlight of attention: evidence for discrete computations in cognition. *Front Hum Neurosci*. 2010;4:194.
42. Duncan J, Schramm M, Thompson R, Dumontheil I. Task rules, working memory, and fluid intelligence. *Psychon Bull Rev*. 2012;19:864-870.