

# 5

## Executive Control of Cognitive Search

Joshua W. Brown and Derek E. Nee

### Abstract

At a basic level, cognitive search involves several parameters: Under what circumstances should a search be initiated, and how should the goal be specified? What are the criteria by which the search is judged a success or failure? How are corrective actions implemented when search strategies are judged insufficient?

Studies of cognitive control have the potential to address each of these questions. In this chapter, a number of issues related to executive control of search are discussed, including the way in which hierarchical search goals are monitored and updated. A new theory of cognitive control is proposed to begin to answer these questions, and open questions that remain are highlighted for future enquiry.

### Initiating and Maintaining Searches

#### Initiation of Search

Searches are generally initiated on the basis of a goal and a lack of certainty about how best to achieve it. Goals may be anything, from finding a shape in a visual scene to remembering where the car keys are to finding a mate. Goals, and how to achieve them in an ever-changing environment, are the *raison d'être* of cognitive control. At the neural level, active goals are represented, in part, as a pattern of sustained activity across the dorsolateral prefrontal cortex (dlPFC) (Miller and Cohen 2001) and other regions, such as the intraparietal sulcus (IPS) (Chafee and Goldman-Rakic 2000). According to the biased competition model (Miller and Cohen 2001), sustained activity in dlPFC represents goals and working memory. Goal-related dlPFC activity interacts with posterior cortical regions to bias the flow of information across competing networks, much like a switch yard at a railroad station (Rogers and Monsell 1995), thereby enhancing activity in posterior regions that represent relevant information (Egner

and Hirsch 2005). Hence, activation in the dlPFC primes the cognitive system to encode and maintain information relevant to goals.

### **Goal Maintenance**

Goals are thought to be maintained in the dlPFC via sustained patterns of activity. This activity is modulated by neuromodulators, such as dopamine and norepinephrine, which influence the persistence of these goal representations and, in turn, influences how readily an animal will change goals as opposed to perseverate. In particular, a lower barrier to switching goals implies a lower barrier to either beginning or abandoning a search. Dopamine has been studied extensively as a principal mediator of reinforcement (Schultz 1998), but it also influences the stability of sustained activity patterns in dlPFC. Either too much or too little dopamine can reduce the stability of activity, thus making it easier for new working memory and goal representations to become active (Muly et al. 1998). The neural mechanisms underlying this “sweet spot” of stability have been modeled computationally (Brunel and Wang 2001; Durstewitz et al. 1999, 2000; Redish et al. 2007). Essentially, the optimal level of dopamine seems to deepen the attractor basins of the network state, which requires a stronger input to cause a change in the pattern of which units are active and inactive. At the behavioral level, as dopamine levels increase towards optimal stability, animals may perseverate on their current goal.

Perseveration on a goal constitutes the “exploitation” end of a spectrum between exploration and exploitation (Kaelbling et al. 1996). At the other end of the spectrum, a lack of stability in goal representations may lead to constant switching, which constitutes a process resembling a search except that it never terminates to allow consumption of what was found. This link among tonic dopamine levels, search, and the exploration/exploitation trade-off has been treated previously, and it appears that dopamine may bias behavior toward exploitation (Hills 2006). With regard to drug abuse, addictive substances typically cause a lasting release of dopamine (Grace 2000), which is associated with the recurring drug-taking behavior that characterizes addiction.

Norepinephrine has also been implicated in cognitive flexibility, although the neural mechanisms are somewhat less studied than those of dopamine. As discussed more fully by Cools (this volume), greater tonic norepinephrine seems to reduce cognitive flexibility, which corresponds with increased gain in the responsiveness of neurons to both excitation and inhibition (Hasselmo et al. 1997).

### **Internal versus External Search**

Searches may target internal cognitive processes in addition to the external environment. There is good evidence that many of the same neural mechanisms involved in searching the external environment are also recruited for

searching information in the mind held in working memory (Awh et al. 2006). For example, Nobre and colleagues (2004) demonstrated common recruitment of the IPS and the frontal eye fields when subjects either directed attention to an external location or to a location held in working memory. Nee and Jonides (2009) replicated this effect with more complex searches of external visual and internal memory information and demonstrated additional common dlPFC activation for both types of searches, presumably in the service of maintaining goals during search. Behaviorally, it has been demonstrated that attention is captured by externally presented objects that match objects held in working memory, indicating interactions between attentional and working memory systems (Downing 2000; Pashler and Shiu 1999). Moreover, holding information in working memory reduces filtering of distraction, consistent with the idea that both selective attention and working memory draw upon the same attentional resources (de Fockert et al. 2001). Taken together, attentional mechanisms that search the external world also appear to be necessary for searches of memory.

Despite strong commonalities between external search and working memory, one consistent finding is that internal searches recruit ventrolateral prefrontal cortex (vlPFC)<sup>1</sup> to a greater degree, particularly in the left hemisphere (LaBar et al. 1999; Mayer et al. 2007; Nee and Jonides 2009; Nobre et al. 2004). Although the left vlPFC is often associated with the maintenance and manipulation of verbal information, one study has reported greater left vlPFC activation when selecting a spatial location from memory compared to selecting a spatial location in perception (Nobre et al. 2004). Moreover, left vlPFC involvement in memory search extends beyond working memory and includes searches of long-term memory as well (Cabeza et al. 2002). Hence, the left vlPFC may be generally involved in searching internal memory space (Zhang et al. 2004) in a way that is distinct from external searches.

### **Criteria for Search Success or Failure**

Searches end either in success or failure, but a key underlying question is: What criteria determine success versus failure? Suppose an animal is foraging for food, but finds only a little food and is still hungry. Should the search be considered a success or a failure? The answer to this question depends on prior expectation. If food is very scarce, then the expectation may be that virtually no food will usually be found, in which case finding even a little food may be considered a success. On the other hand, if food is typically plentiful, then

---

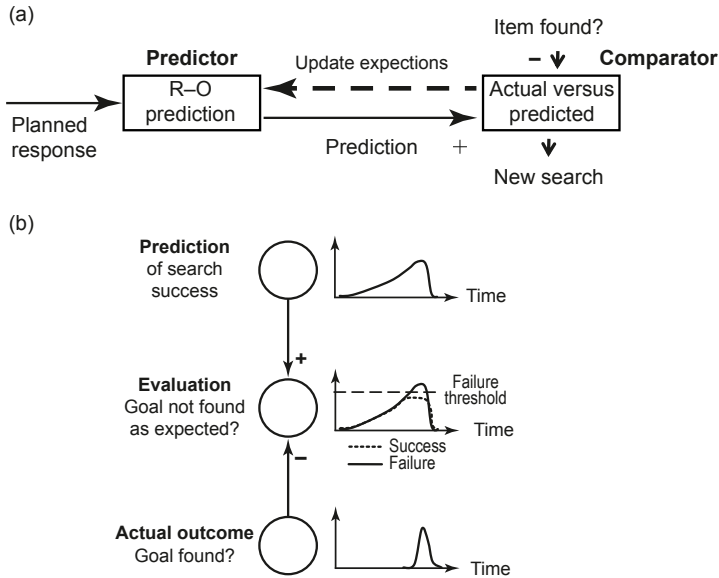
<sup>1</sup> We use the term “ventrolateral” prefrontal cortex (vlPFC) to distinguish these activations from the dorsolateral prefrontal cortex (dlPFC). Activations from the cited studies typically fall in and around pars triangularis (BA 45), which is the dorsal most aspect of the inferior frontal gyrus, and often also extends into the inferior frontal sulcus. Although there is some ambiguity as to where vlPFC ends and dlPFC begins, the activations reported here are ventral to activations we refer to as within dlPFC, which are on the middle frontal gyrus (BA 9 and 46).

finding only a little food may be considered a failure. This example illustrates two basic principles of evaluating search success or failure: (a) expectations are key to the evaluation and (b) expectations are formed on the basis of prior experience or information. Nonetheless, there is often no explicit environmental cue that a search has failed, so the evidence of absence must be inferred from the absence of evidence (Sagan 1996:213). The monitoring and evaluation functions required to infer search success or failure is a central concern of cognitive control.

Models of cognitive control typically have two main components: one for monitoring and one for control. For example, one model (Botvinick et al. 2001) casts the monitor as a response conflict detector, while the controller implements attentional focus or increased caution. Response conflict occurs when cues in the environment are associated with two different responses that are mutually incompatible. Computationally, response conflict can be detected by multiplying the activities associated with the mutually incompatible responses (Botvinick et al. 2001). While conflict models predominate (Yeung et al. 2004, 2005; Yeung and Nieuwenhuis 2009), others have cast the monitor as detecting errors; that is, a failure to achieve a desired goal (Holroyd and Coles 2002) or the likelihood of errors (Brown and Braver 2005, 2007). There is ongoing debate on whether such dedicated monitoring and control pathways are necessary to account for cognitive control phenomena. Some argue for the existence of such mechanisms (Monsell 2003; Rogers and Monsell 1995), while others argue that simpler mechanisms (e.g., priming) are sufficient (Altmann 2003; Altmann and Gray 2002; Mayr et al. 2003). With respect to this debate, we propose that dedicated control structures can provide useful contributions to the control of search processes.

### **The Predicted Response–Outcome Model**

Recently, we proposed a new model of performance monitoring and cognitive control functions in the medial prefrontal cortex (mPFC, including anterior cingulate cortex), which we refer to as the predicted response–outcome (PRO) model (Alexander and Brown 2011). The PRO model can detect when searches fail, and provides a monitoring and evaluation function with two interacting components (Figure 5.1). The first component, the predictor, generates a prediction of the expected outcomes of an action, which in the context of search would correspond to successfully finding the object (and in the expected quantities). The neural activity representing this expected outcome increases as time elapses, such that if the object of the search is available, then it ought to be found within a certain time frame. In other words, not finding sufficient quantities early on would not be considered failure, but failure will be signaled if sufficient quantities are not found after a longer period of time. This kind of representation can be thought of as qualitatively similar to a hazard function of the probability of finding the searched-for object, given that it exists (Ghose and



**Figure 5.1** Predicted response–outcome (PRO) model: (a) Planned searches activate learned response–outcome (R–O) predictions. These predicted outcome signals indicate the expected findings of the search. (b) The Comparator unit receives a timed prediction from the Predictor unit that signals when the search should yield a finding. The actual findings (the outcome) are compared against the expected findings, and failure to find the searched-for item leads to both an update of the search outcome predictions and a possible initiation of a new search.

Maunsell 2002). The second component, the comparator, subtracts the actual outcome from the expected outcome. The net result is that when a searched-for object is found, a signal of the actual finding suppresses the expectation activity in the comparator. Conversely, when the object of the search is not found, the predictor activity increases unopposed and signals search failure in the comparator. Of note, failure can be detected at any point in time, whenever the difference of the prediction activity minus the actual success outcome exceeds a specified threshold. In addition, finding greater than expected amounts of the goal, or finding it sooner than expected, would not be evaluated as a failure, although other aspects of the PRO model not discussed here would signal it as a surprising event. Neurophysiological findings in monkey mPFC are consistent with the PRO model, as described below.

The PRO model differs from existing models of mPFC in that it does not compute response conflict, as do some other models (Botvinick et al. 2001). Our simulations suggest that the PRO model can simulate virtually all of the known effects in mPFC such as error, response conflict, and error likelihood, among others (Alexander and Brown 2011). The PRO model derives from our error likelihood model (Brown and Braver 2005), but it differs in two ways. First,

the PRO model predicts various outcomes, including possible rewards, and is not restricted to predicting only errors. These prediction signals may be formed by mechanisms within mPFC, or they may instead be formed elsewhere and sent to the mPFC. We are actively investigating this question. Second, the PRO model adds a mechanism that signals any discrepancies between the outcome predictions and the actual outcomes. These discrepancy signals resemble a dopaminergic temporal difference signal that has been proposed as an alternative account in the earlier RL-ERN models (Holroyd and Coles 2002; Holroyd et al. 2005). Nonetheless, the PRO model posits different mechanisms to account for the signals. In one earlier model, the dopaminergic error signal from the midbrain disinhibits the mPFC (Holroyd and Coles 2002). In contrast, the PRO model suggests that such error signals are computed internally by the mPFC. In another model, the dopamine signals train mPFC to recognize conjunctions of events that constitute errors (Holroyd et al. 2005). In contrast, the PRO model signals not conjunctions but comparisons of actual versus expected events. Furthermore, while the PRO model may influence dopaminergic signaling, it does not depend on external dopamine signals to function per se.

With this different approach, the PRO model can account for data that other models cannot. For example, whereas the response conflict model of the anterior cingulate cortex (ACC) may account for greater activity during search, it cannot account for feedback-related ACC responses (Holroyd et al. 2004). The PRO model accounts for activation at the time of feedback as a discrepancy between actual and expected outcomes. Of note, when a search is expected to fail or is rarely successful, then the PRO model would predict that activity related to search failure should be weaker. In fact, weaker error signals have been found when errors are more likely (Brown and Braver 2005), and error signals even reverse when success occurs unexpectedly (Jessup et al. 2010). Such reverse reward effects are difficult to reconcile with dopamine-based models of the ACC that compute signed differences in reward expectation (Holroyd and Coles 2002; Holroyd et al. 2005). By contrast, the PRO model interprets this latter finding with complementary mechanisms that detect surprising occurrences as well as surprising nonoccurrences.

With the PRO model framework, the threshold at which a search failure is signaled is a product of two parameters. The first parameter is the strength of the prediction. Prediction strength in the PRO model is proportional to the previously experienced probability of finding the searched-for object and the quantity of the searched-for object that is expected to be found (Amiez et al. 2005). The second parameter is the strength of what was found. The greater the quantity found, the greater the suppression of the expectation signal in the comparator and the less likely it is that a failure will be signaled.

Predictions about the success of a search are learned from past experience. When a searched-for object is not found, the resulting error signal from the comparator drives learning in the predictor unit to reduce the predictions of success in similar searches in the future. This is tantamount to raising the

threshold for signaling failure. In this way, predictions about what constitutes success for a given search are dynamically updated in nonstationary environments. The greater the nonstationarity, the greater the ongoing discrepancy signals and resulting mPFC activity, as has been found with fMRI results in humans (Behrens et al. 2007).

Certain kinds of task-switching tasks can be thought of as cognitive searches, similar to foraging. For example, some tasks require subjects to choose a certain option or strategy to gain reward, but after some trials, the reward is depleted. The subjects must then detect the depletion and search for a new strategy that yields reward, similar to patch-leaving in animal foraging. Tasks of this kind include the Wisconsin card sort task (Grant and Berg 1948) as well as searches for a correct sequence of button presses (Bush et al. 2002; Procyk et al. 2000) or lever manipulations (Shima and Tanji 1998). Notably, these tasks differ somewhat from much of the traditional task-switching literature, which involve either explicit cues or unambiguous patterns that cue a task switch (Altmann and Gray 2002; Rogers and Monsell 1995). In cases where the new task is explicitly cued, searching for the appropriate task is not required. In contrast, task switches due to the absence (or surprising reduction) of an expected reward may involve different neural mechanisms. When reward is reduced, ACC is active prior to a switch, but it is not active for explicitly cued switches (Bush et al. 2002; Shima and Tanji 1998). Thus the task switches cued implicitly by reduced or absent reward can be thought of as a disconfirmation of the current strategy, which may in turn lead to renewed or continued searches for a more effective strategy. In a broader sense, task switches due to reward omission and exploration of the environment (as opposed to exploitation) may be thought of as two sides of the same coin: reduced reward may be a cue to switch or a current reward level may still be deemed insufficient if a possibly greater reward may be found elsewhere. The common question is: What constitutes a sufficient level of reward, below which a search for better reward will be initiated?

The PRO model framework, as applied to cognitive search, can account for a variety of findings regarding implicit task-switching paradigms. The Wisconsin card sort task (Grant and Berg 1948) yields activation in the mPFC for negative feedback that leads to a search for the new correct strategy (Monchi et al. 2001). In monkeys, negative feedback also leads to greater activity in ACC during search (Procyk et al. 2000; Shima and Tanji 1998), as is also the case in humans (Bush et al. 2002). Similarly, ACC is more active when monkeys are actively searching than when behavior is routine (Procyk et al. 2000). More broadly, monkey supplementary eye fields in the mPFC have distinct subpopulations of cells with activity profiles that apparently anticipate the outcome of actions and shut off when expected outcomes occur, and other subpopulations of cells that signal the actual occurrence of an expected outcome, such as a reward (Amador et al. 2000; Ito et al. 2003).

## **Controlling and Correcting Search Strategies**

Searches can fail for many reasons. A primary concern of cognitive control is to minimize the possibility of failure, while at the same time detecting failure when it does occur and driving corrective action. A central question then is: How are these functions accomplished? First, a signal is needed to indicate when failure is likely. Such a prediction signal can, in turn, drive greater attention and effort to maximize the chance of success with the existing strategy, or it can drive a change in strategy to find another search tactic that is more likely to succeed. Previous models cast the prediction of an error as driving increased caution by slowing down response processes (Botvinick et al. 2001; Brown and Braver 2005) or by increasing attentional focus (Botvinick et al. 2001; MacDonald et al. 2000; Posner and DiGirolamo 1998). In this case, slower and more careful processing of the environment may lead to detection of the searched-for object when environmental cues are otherwise weak and easy to miss (Clark and Dukas 2003). A second issue involves how failures are detected and corrected. There is evidence that the mPFC is involved in error avoidance (Magno et al. 2006) as well as error correction (Modirrousta and Fellows 2008).

The PRO model yields two relevant signals in this regard from the predictor and comparator components (Figure 5.1). The predictor provides a prediction of what will be the outcome of a search, including possible failure. These prediction signals would be sufficient to provide a greater level of control toward the goal of avoiding failure, whether by searching more carefully or by trying a different strategy, according to the predicted outcome. For example, if the model predicts a likely failure to detect some event in a certain situation, then an increase in attention is most likely to lead to reward and consequent reinforcement. If instead the model predicts that the resources are likely depleted such that no amount of greater attention will succeed, then a change in strategy is most likely to be rewarded. In the PRO model, a second control signal derives from the comparator. As described above, this signals when a failure has in fact occurred, in that the searched-for object has not been found. This signal is exactly what is needed to drive a change in strategy, which is essentially a task switch that is cued implicitly by reduced reward, as described above.

### **Hierarchies of Strategy**

In a foraging task, the appropriate change in strategy may involve giving up exploiting the current patch or environment and returning to an exploratory set to search for new resources. Of note, this foraging example highlights the hierarchical nature of search goals. At the lowest level of the hierarchy, foraging in a given part of an environment may involve searching a limited region for a particular resource, and many individual resources may be found. In this case, a visual search may be conducted, and if the searched-for object is not



found in the fovea, then the search “fails” at the lowest level and a new location is searched in the immediate vicinity. We have suggested that the mPFC may yield evaluations of failure in general, but it is an open question as to whether the mPFC may detect failure of lower-level visual search.

As resources are depleted and become scarcer, more careful and attentive processing may be needed to find the resources, expressed as greater attentional focus and longer processing times (Botvinick et al. 2001; Brown et al. 2007b). When resources are depleted beyond a certain level, this modulation of the lower-level strategy is no longer sufficient. In such a case, it is time to change the higher-level strategy and switch from exploiting the current location and instead explore for a new location. The PRO model comparator would provide the signal necessary to drive the change in strategy. This proposal is consistent with ACC activation due to reduced available reward in both humans and monkeys (Bush et al. 2002; Shima and Tanji 1998).

The hierarchical nature of search leads to a credit assignment problem, which may be seen in an expanded variant of the explore versus exploit foraging task. We may suppose, for example, that foraging strategies may be for a more- versus less-preferred food and that there is a choice of continuing to forage versus waiting and conserving energy until more resources become available. In this scenario, we might suppose that an animal will search for a preferred food in a limited region until the preferred food is depleted in that region. As food becomes scarcer, the animal might implement increased attentional control to find less salient food items. Once the local region is depleted, the animal will qualitatively switch control strategies from greater attention in the current region to exploration of a different region for the preferred food instead. Once the preferred food is depleted, the animal might again switch strategies to forage for nonpreferred food, subsequently switching strategies between exploration versus exploitation for the nonpreferred food. Once even the nonpreferred food is depleted, the animal may switch strategies between foraging and resting or waiting for more resources. This example of a hierarchical goal structure for search leads to an important question: If a failure occurs, at which level of the goal hierarchy should failure be ascribed? For example, when a food item is not found, does this mean the animal should pay more attention to the local region and look in another nearby location (lowest-level search failure)? Or does it mean that the animal should explore for new regions (mid-level search failure)? Or does it mean that the animal should switch to foraging for another, less-preferred food (higher-level failure)? Or does it mean that the animal should give up searching entirely and conserve energy until new resources arrive (highest-level failure)?

The PRO model suggests an answer to the hierarchical goal credit assignment problem. The answer begins with the assumption that just as there are multiple levels of goals, there are also multiple levels of outcome predictions. At the lowest level, a visual search may involve the expectation of a particular object in the fovea. This is not to suggest that the evaluation of the visual scene

is necessarily carried out in the ACC. Instead, there is evidence that visual cortex may carry similar temporally structured expectation signals in anticipation of particular visual cues, which appear as attentional signals (Ghose and Maunsell 2002).

At the next level, the ACC evaluates reward, or the lack thereof, and may drive a corresponding change in mid-level strategy aimed at successfully searching for the same reward (Bush et al. 2002; Procyk and Joseph 2001; Shima and Tanji 1998). Exploration-related activity in cognitive search is also associated with anterior prefrontal activation (Daw et al. 2006). When the searched-for reward is not found, then, in the framework of the PRO model, it may be that a prediction signal of finding a certain quantity of reward over a longer timescale (many trials) eventually goes unmet by a longer-term measure of actual successful trials, and this could lead to a higher-level switch in strategy to search for other kinds of reward. In the same way, an even longer timescale prediction of total reward aggregated across multiple reward types may develop, and if it is not met by successful search across a variety of reward types, then a highest-level switch in strategy may be made to give up the search and switch to a strategy of waiting until new resources become available.

The key point, and the proposed solution to the hierarchical credit assignment problem, is that outcome predictions are associated with a corresponding action. If a low-level visual foveation action fails to yield the searched-for object in the fovea within a few hundred milliseconds, then the failure violates the expectation of the eye movement foveating an object within a short time. It does not necessarily violate the higher-level expectation associated with the overall search strategy, which is that the object will eventually be found, perhaps after some longer time period of minutes. The PRO model's ability to specify not only the nature of expected outcomes but also their timing allows for a short-term failure to be signaled without necessarily signaling the failure of a higher-level goal that is expected to take more time to achieve.

The template of hierarchical goals in search is ubiquitous, with examples ranging from animals' search for food to cognitive search of memory to humans searching for employment or mates. In the end, the hierarchical monitoring and control of search goals may be carried out by a corresponding hierarchical structure of evaluating shorter and longer timescale predictions about the outcomes of one's own actions.

## **Conclusion and Open Questions**

Our aim in this chapter has been as much to raise questions as to propose answers. In the course of exploring the topic of cognitive control in search, several potentially controversial or at least unresolved issues may be highlighted. First is the question of whether and to what extent there are distinct structures in the brain that provide executive control of search. We have outlined the PRO

model as a possible mechanism of executive control, but undoubtedly there are other possibilities. It may be that what appear to be effects of executive function are in fact properties that derive from the nature of regions that drive the search, so that no additional control structures are necessary. This question parallels the debate over whether executive control is necessary to account for effects associated with explicitly cued task switching.

If there are indeed neural mechanisms dedicated to executive control of search, then the next question is what distinct brain regions are involved in monitoring and evaluating the different kinds and hierarchical levels of search, and whether or to what extent there is overlap. There is evidence that mPFC monitors the outcome of actions and drives changes in at least higher-level strategies, but it is less clear whether these same regions are involved in lower levels of search (e.g., visual search). This leaves open the question of what possible distinct brain regions are involved in detecting failure at different levels in the hierarchy of search goals, and whether those distinct regions share nevertheless a common neural architecture related to prediction and evaluation.

Correspondingly, there is a question of how and where in the brain the search goals are represented. We have generally referred to working memory and goal representation in the dlPFC, but this is a relatively large region. Some have argued that the hierarchy of lower- to higher-level goals is represented along a posterior-to-anterior gradient within the lateral PFC (Koechlin et al. 2003; Kounieher et al. 2009). This leaves open the possibility that a similar gradient exists in the medial PFC that interacts with lateral PFC, although this has yet to be explored.

Another open question is the degree to which search has memory or not. There is evidence that visual search has no memory (Horowitz and Wolfe 1998), but it is not clear how this finding can be reconciled with effects showing inhibition of return (Klein and MacInnes 1999), which would imply memory. With respect to higher levels of cognitive search, monkeys seem to have a strong memory and ability to infer which search spaces remain as plausible resources, as they show near optimally short and successful searches at the cognitive level (Procyk and Joseph 1996). If higher cognitive search has memory, then the interaction between dlPFC and mPFC may be reciprocal. More specifically, mPFC may drive changes in the strategy represented by dlPFC and IPS, but dlPFC may, in turn, constrain how error signals are generated in mPFC and what kinds of new strategies may be implemented in response to search failure. The effect of working memory context on performance monitoring is suggested by recent studies of individuals with schizophrenia (Krawitz et al. 2011).

### **Acknowledgments**

This work was supported by R01 DA026457.

