

Search, Goals, and the Brain

Catharine A. Winstanley, Trevor W. Robbins,
Bernard W. Balleine, Joshua W. Brown, Christian Büchel,
Roshan Cools, Daniel Durstewitz, John P. O’Doherty,
Cyriel M. A. Pennartz, A. David Redish, and Jeremy K. Seamans

Abstract

The process of cognitive search invokes a purposeful and iterative process by which an organism considers information of a potentially diverse nature and selects a particular option that best matches the appropriate criteria. This chapter focuses on the neurobiological basis of such a goal-directed search by parsing the process into its main components, suggested here as initiation, identification of search space, deliberation, action selection, and evaluation and search termination. Unexpected uncertainty is suggested as a key trigger for the onset of the search process. Current data posit that this is represented in the anterior cingulate, parietal, and inferior frontal cortices, suggesting these areas could be particularly important in search initiation. A change in motivational state, likely signaled by a wide range of brain regions including the amygdala, can also play a role at this stage. The neural structures which represent the set of to-be-searched options may vary depending on the search domain (e.g., spatial, visual, linguistic). During deliberation, predictions regarding the consequences of selecting these options are generated and compared, implicating areas of frontal cortex as well as the hippocampus and striatum, which are known to play a role in different aspects of outcome evaluation. Action planning and selection likely involve an interplay between the prefrontal cortex and basal ganglia, whereas search termination could involve the specific neural networks implicated in response inhibition. The influence exerted over the search process by the major ascending neuromodulators (dopamine, norepinephrine/noradrenaline, serotonin, and acetylcholine) is also considered, and a particularly critical role suggested for dopamine and noradrenaline, given their ability to influence cognitive flexibility and arousal. Finally, pathologies of search processes are discussed, both with respect to brain damage and psychiatric illness.

Introduction and Overview

Search is defined as “movement in pursuit of a resource at an unknown location” (Hills and Dukas, this volume). This very general definition allows

search to be applied quite broadly from protozoa to humans. While laudable, this breadth could also motivate comparisons in process that are less desirable. The focus on superficial aspects of search, particularly the movements that collectively define the search response, provides a ready means of identifying search. However, it also implies that identity in response means identity in mechanism, and this may be problematic. It is tempting to argue that the occurrence of an organized set of responses associated with exploration (such as orienting, locomoting, pausing, turning, returning, and so on) always reflects a deliberated, goal-directed search process under cognitive control, whether nascent or explicit. Nevertheless, care should be taken with such assumptions. Considerable research has established that seemingly indistinguishable behavioral responses can, at different times and under different constraints, be controlled by quite distinct determinants. Take the case of lever pressing in rats as an example (see O'Doherty and Balleine, this volume). The behavior in which rats press a lever for food appears to be a quintessential goal-directed response mediated by both its relation to a goal (the specific food) and by the value of that goal; a movement in pursuit of a resource certainly qualifies as a search response. However, it is now well known that when the action is overtrained or goal access is placed under certain temporal constraints, the determinants of this response can change: it is no longer a flexible, deliberate goal-directed action; it becomes more routine, automatic, inflexible or habitual. Although it would still satisfy the broad behavioral definition of a search, such an automated process entails a reflexive movement elicited by antecedent stimuli, rather than its consequences. Hence, if we believe search to be essentially a goal-directed behavior, most exploratory behavior only looks like a search response; it utilizes different brain structures and depends on different computations within the mammalian brain. This leads us to reject it as a true cognitive search response.

As a consequence, it is necessary in all situations to establish whether a putative search response satisfies two conditions:

1. The performance of the search response is determined by the organism as being causal with respect to some specific resource or goal.
2. Its performance is sensitive to changes in the value of the goal.

There are, in fact, at least three kinds of search response which, by this definition, do not qualify as cognitive search. These responses reflect the operation of three different motivational constraints and can be referred to as "evaluative processes," "Pavlovian processes," and "habitual processes." Note first that sensory processing is common to each and is assumed to be more or less constant across all forms of search or search-like responses. In a novel or changing environment, sampling the sensory environment is critical, and search in this domain is likely to be general, constrained by a bottom-up attentional process sensitive to physical salience, regulated by motivational arousal, and subject to simple learning processes such as habituation.

Evaluative Processes

The first motivational constraint on search is the learning process by which stimuli become associated with specific, innate motivational processes, thereby conferring value on sensory events (e.g., contact with stimuli that provoke nutrient activity produces an association between those stimuli and the nutrient system resulting in what might be called the “representation of a specific food”). Increases in nutrient deprivation have long been reported to elicit an immediate increase in activity and orienting; food deprivation, for example, increases orienting to foods, as well as an increase in the production of vacuous consummatory/defensive reactions appropriate to those processes (e.g., food events will provoke consummatory responses—salivation, chewing, gastric motility, etc.). Thus although these appear to reflect search, they are actually reflexes elicited by internal states and not by their relationship to a specific resource (Changizi and Hall 2001).

Pavlovian Processes

A second motivational constraint is provided by the tendency of sensory events, or event representations, to become associated when they are paired in a manner that allows the activation of one representation to activate the other. Importantly, events that predict those sensory events that have been subject to evaluative conditioning provoke what is typically called Pavlovian conditioning. As a consequence, the former event (i.e., the “conditioning stimulus”) can produce (a) conditioned consummatory/defensive reactions and (b) conditioned preparatory reactions, such as behavioral approach/withdrawal. Whereas consummatory reactions are produced by activation of the specific sensory features of evaluative incentives, the preparatory reactions are produced by activation of either specific motivational states or affective states (e.g., appetitive and aversive states productive of general activity, and other conditioned responses like approach and withdrawal). As is well known, these responses are not determined by their relationship to the goal (or “unconditioned stimulus”; see Holland 1979; Holland and Straub 1979).

Habitual Processes

The third motivational constraint reflects the ability of environmental cues to become associated with responses and, under invariant conditions and by prolonged training, to elicit those responses irrespective of the value of the goal or the relationship between response and procuring the goal. These habit processes can transition between environmental states, like goal-directed search, but are not based on any knowledge of the structure of the environment. Instead, they are based on state-response associations. In the parlance of reinforcement

learning models, they are “model-free” rather than “model-based” responses (cf. Daw, this volume).

Goal-Directed Search

Now that we have considered processes that would *not* conform to cognitive search, let us consider the processes within the goal-directed system that would be considered representative of this class of search. Within this domain, search might be initiated to obtain information related to a number of different processes underpinning the goal-directed system: the perceptual level, the level of causal structure, the level of goal selection and the level of action selection. At the perceptual level, goal-directed search could be initiated over the perceptual environment in a deliberative sense to find and locate relevant stimuli for goal-directed action (such as visual search; see Wolfe, this volume). At the causal structure level, search might also be initiated through a set of internal hypothesis spaces to elucidate the likely causal structure of the decision problem (i.e., the rules governing the representation of states and transitions between states). This is necessary so that the appropriate decision structure is represented from which options can be selected. At the goal-selection level, search needs to be initiated to determine which goal from the multiple possibilities the animal wants to pursue. The final type of search is over the space of possible actions that might be selected to obtain a particular outcome.

We argue that the primary computational signal underpinning search in a motivated animal is the need to minimize uncertainty in the animal’s representation of information pertaining to each level of the goal-directed decision process. This uncertainty should be computed separately for each of the different component processes underpinning the goal-directed system.

The domains of cognitive search in Table 9.1 are those that we consider to be particularly amenable to an analysis in terms of brain mechanisms. From the outset, it should be made clear that it is unlikely that search can be reduced to a single process operating within a single neural system. Rather, searches may engage the articulation of different neural systems working in a combinatorial fashion, both in series and in parallel. However, it is also necessary to parse search processes further to investigate candidate neuronal mechanisms and to identify large-scale neural systems through which search is implemented. These neural systems may, for example, include regions in which relevant representations are held that may be accessed by other systems, for example, in a “top-down” manner. They may further be subject to modulation by ascending, diffuse neurochemical systems (e.g., the monoamines, dopamine, norepinephrine, and serotonin) mediating states of arousal, stress, and general motivation, which influence the fidelity of representations as well as the efficiency of search processes. It is highly likely that the search process itself utilizes some fundamental neuronal mechanisms common to many behavioral processes, such as prediction errors and outcome expectancies, though the exact

Table 9.1 Types of search process.

Non-goal-directed, search-like responses

1. Sensory/perceptual feature processes
 - Sensitive to levels of general arousal
 - Composed of reflexive orienting responses
 - Subject to habituation
2. Evaluative processes
 - Elicited by motivationally salient cues (stimuli associated with primary motivational states)
 - Composed of reflexive-orienting responses
 - Productive of arousal
 - Stimulus–motivation (S–M) associative structure
3. Pavlovian processes
 - Elicited by stimuli associated with evaluative incentives
 - Composed of consummatory/defensive reflexes (e.g., lick, chew, blink, freeze)
 - Preparatory responses (e.g., approach, withdrawal, restless activity)
 - Stimulus–stimulus (S–S) associative structure
4. Habitual processes
 - Elicited by antecedent stimuli with which the response has become associated through reinforcement
 - Model free
 - Stimulus–response (S–R) associative structure

Goal-directed search processes

1. Perceptual search: gathering information from the world
 2. Search over causal models: to identify (hidden) structures of environmental contingencies and define the search space; requires inference as well as perception
 3. Search over goals: internal, based on current motivational states and needs
 4. Searching over actions, exploring action–outcome (A–O) relationships
 5. Model based
-

mechanism by which these are integrated into the search structure is an open question, as theorized in the next section.

To undertake a neural analysis of search, it is therefore necessary to identify its main components, and the following key elements are proposed:

1. Initiation of search
2. Identification of the set of to-be-searched options
3. Deliberation, including evaluation of the value of possible options and predicted outcomes
4. Action selection
5. Search termination, including evaluation of search success

The different subprocesses of cognitive search are now considered in turn, with their associated possible neural correlates. In general, we note that relevant neuroscientific investigations have generally relied on a rather limited number of species; namely rodents, nonhuman primates and humans, studied individually in rather artificial, laboratory-based environments. Nevertheless, we hope that at least some of what will be described has more general application to the situations and themes considered by this Forum. Some of the utility of this analysis will be considered in the context of various pathologies. For example, stress can have significant effects on aspects of the search process. Behavioral evidence of “search deficits” is also considered in human patients with discrete brain damage, or within functional cognitive deficits arising through neurological or neuropsychiatric disorders.

Components

Search Initiation

Search initiation can be thought of as a two-stage process, driven by both the onset of a motivational state (e.g., hunger, thirst, need for information) and an uncertainty regarding how to satisfy that need. There may well be competing goals to pursue, in which case there is also uncertainty as to which goal state should take precedence, and also a fundamental uncertainty about which action will best serve the organism in achieving its aim.

Prediction Errors

One key signal capable of triggering a search would be the occurrence of an unexpected event (i.e., if what is observed is inconsistent with what is expected). This process would also be key in other evaluative stages of the search process. Once organisms have become familiar with their environment and have learned about cues or subspaces that were previously associated with rewards, this knowledge can be used to generate predictions about the consequences of cues, events and actions. Computationally, the operation by which actual outcomes are compared to expectancies is cast as a calculation of prediction error. In basic form, a prediction error is computed by subtracting the expected outcome from the actual, observed outcome (Rescorla and Wagner 1972). Recent behavioral and neurophysiological studies have shed light on the neural systems involved in these computations. An important discovery (Schultz et al. 1992, 1997) was that the firing of dopaminergic neurons in the primate brain obeys a response pattern predicted by models of reinforcement learning based on temporal prediction errors (Sutton and Barto 1998). Before task acquisition, dopamine neurons transiently fire to rewards that are delivered unexpectedly, and also when rewards are preceded by a sensory cue

(conditioned stimulus). After the animal has learned that the reward is reliably preceded by a conditioned stimulus, dopamine neurons no longer increase firing when receiving a reward. In contrast, they still fire to reward delivery when this is unpredicted to the animal (i.e., when the sensory cue is omitted). When an expected reward is omitted, the firing rate of dopamine neurons transiently decreases. Overall, phasic increments in firing occur whenever a positive prediction error occurs (receiving more reward than predicted at that moment), and a decrement occurs when the error is negative (receiving less reward than expected). Importantly, once the animal is trained on a conditioning task and dopaminergic neurons stop firing in response to the now predicted reward, they will fire in response to stimuli or contexts that reliably predict reward in time. This backwards referral process transfers the dopaminergic signals from the end result (reward) to the environmental elements acting as the earliest predictors of reward.

In the context of search initiation, however, it is important to emphasize that dopamine probably serves more functions than just mediating an error in reward prediction. Dopamine neurons can also respond to novel stimuli as well as to generally salient stimuli, which may contribute to an animal's motivation to search novel spaces. Moreover, the tonic (sustained) component of dopaminergic signaling appears to be related to other processes, such as opportunity costs (Daw et al. 2006), vigor (Niv et al. 2007; Robbins and Everitt 1992), stability of representations (Durstewitz et al. 2000; Redish et al. 2007; Seamans and Yang 2004), uncertainty about future reward (Fiorillo et al. 2003), as well as to basic abilities of initiating motor actions and maintaining flexible posture and rhythmic movements, as is dramatically illustrated by Parkinson's disease. Salient, noxious stimuli and stress have also been described to enhance dopamine release (e.g., Matsumoto and Hikosaka 2009; Goto et al. 2007), and this may likewise have implications for search initiation and cessation. Finally, other brain systems have been shown to generate error- and surprise-related information; for example, the orbitofrontal cortex (Sul et al. 2010; van Duuren et al. 2009), anterior cingulate cortex (Gehring and Fencsik 2001), and habenula (Bromberg-Martin et al. 2010c). A more thorough consideration of the neurochemical regulation of search is provided later (see section on Deliberation and Evaluation).

The Nature and Importance of Uncertainty in the Initiation of the Search Process

Given the importance we have placed on prediction errors in mediating the search process, it follows that a guiding principle in the initiation and subsequent termination of search pertains to the degree of uncertainty present regarding aspects of the world. According to this idea, one of the main computational principles driving the search process is to minimize uncertainty in both the representation of relevant features of the environment and concerning the

nature of the interaction with that world. Building on the ideas outlined above, that there may be multiple components of search which differ in terms of the types of information being considered, it follows that there may be different types of uncertainty concerning information at different levels of the inference hierarchy, from perceptual features to action–outcome (A–O) relationships.

One useful way of thinking about the representation of these features and the consequent computation of uncertainty is through a Bayesian framework. Bayesian models are a class of simple models that build probabilistic representations that capture beliefs about the state of the world. Mathematically they use Bayes's theorem to update those belief representations (called priors), based on the difference between the actual observed outcomes and the expected representations (prediction errors). These models can represent inference processes about different features of the environment so that, for example, one inference process might encode beliefs about the perceptual environment (which stimuli are present), whereas another might capture beliefs about the hidden causal structure in the environment (e.g., which rules are in place, the context of the agent), and another inference process might compute beliefs about the relationship between particular actions and associated outcomes. Thus, the goal of minimizing uncertainty can operate for different types of inference process and motivate different goal-directed search strategies to minimize uncertainty for each type of inference process where necessary. Another feature of these types of models is that the inference structures can sometimes be arranged in a hierarchy where beliefs at one level of the hierarchy are used to inform and update beliefs at higher and lower levels. In this context, it may be useful to consider that inference over causal structure and inference over A–O relationships can usefully be considered to be part of a hierarchy, with causal structure at the higher level and A–O representations at the lower level; information about which action is currently rewarded (as computed at the lower level) will also be propagated up the causal structure hierarchy and used to update beliefs at that level (in Bayesian terminology the inferred A–O relationships can be used to construct the posterior beliefs). This is a bidirectional process because beliefs about causal structure can also inform priors about which action is currently rewarded.

Within each type of inference process, uncertainty can also be broken down into different components, only some of which are relevant to search. One proposal (Yu and Dayan 2005) is that there are at least two different types of uncertainty. The first is termed *expected uncertainty* and corresponds to the known variance in the world; for example, if an action gives reward only 50% of the time, compared to an action yielding reward 100% of the time, these actions would have different expected uncertainties over reward distributions. In the context of A–O relationships, this form of uncertainty corresponds to what is called risk in economics. Crucially, this form of uncertainty should *not* in principle instigate search, as it corresponds to intrinsic irreducible uncertainty

in the properties of the A–O contingencies; hence there is no way to minimize this through search.

The second type of uncertainty is *unexpected uncertainty*, which is proposed to correspond to features of the world that are unknown. For example, if a given action gives reward 80% of the time, and suddenly and unexpectedly shifts so that the probability of getting reward on that action is now only 20% of the time, this is a form of unexpected uncertainty. Unexpected uncertainty is likely to motivate search, because once an unexpected change is detected, the agent may need to resample the environment to update knowledge about its properties.

A third form of uncertainty described recently (Payzan-LeNestour and Bossaerts 2011) is *estimation uncertainty*. This form refers to the uncertainty in beliefs based on the fact that estimates of the true state of the world are noisy; if we have only sampled an A–O relationship a few times, we might have very high levels of estimation uncertainty about that A–O relationship, whereas if we sample that A–O relationship many times, our beliefs about that outcome will become more precise and our estimation uncertainty will be reduced. Estimation uncertainty is perhaps the most fundamental type of uncertainty that underpins search, as minimization of this kind of uncertainty is necessary to build an accurate picture of the decision problem for all types of representation (whether involving perceptual information, causal structure, or A–O structure). There is a complex relationship between estimation uncertainty and unexpected uncertainty; clearly, if there is a high level of volatility in the environment, unexpected uncertainty will be high and estimation uncertainty will also be high because the agent will constantly need to change its estimations as a function of the change in the underlying contingencies.

Given that the goal of this chapter is to focus on the neural correlates of search processes, we must consider *where* uncertainty is represented in the brain, particularly with respect to unexpected and estimation uncertainty, as these brain regions will be important in the initiation and termination of cognitive search according to the theoretical framework advanced here. In the economic literature, unexpected uncertainty is often described as “ambiguity,” and it has been studied in experimental situations where the precise odds of obtaining a reward outcome are hidden from the participant. Activity in parietal and inferior frontal cortex has been observed when participants are making choices over conditions of high ambiguity (when the probabilities are unknown) compared to low (when the probabilities are known) (Huettel et al. 2006). Other evidence for the representation of uncertainty in the brain comes from an fMRI study in which human subjects performed a simple bandit decision task (Behrens et al. 2007). Behrens et al. varied the “volatility” or rate of change of the reward contingencies at different times in the experiment; at some points, the probability of being rewarded on a particular action changed rapidly over time, whereas at other points the probability of being rewarded changed less rapidly. They used a Bayesian model that computed

a representation of uncertainty and correlated this with the fMRI data. They found that activity in the anterior cingulate cortex correlated with their uncertainty representation. Crucially, in their modeling, they did not distinguish between unexpected and estimation uncertainty, so it is unclear which of these signals is encoded in the anterior cingulate cortex.

Neural measures of uncertainty will be manifested either in direct measures of increased firing patterns with uncertainty or through measuring the internal self-consistency of neural representations. An important paradigm capturing changes in uncertainty in perception and decision making is the “diffusion model” of accumulating neural evidence, expressed by changes in firing rate (Churchland *et al.* 2008). As concerns self-consistency, representations are distributed across multiple cells, and the activity of a population of cells can either “agree” or “disagree” on a representation (Jackson and Redish 2003). These measures can be quantitatively identified through a three-step process from neural ensemble recordings, in which tuning curves are first derived from neural activity and behavior, then represented values are decoded from neural activity and those tuning curves, and finally, through a derivation of expected neural activity, from tuning curves and the decoded behavior (Johnson *et al.* 2009; Zhang *et al.* 1998). An important question is whether fMRI signals which correlate with uncertainty reflect the computation of uncertainty *per se*, or downstream processes associated with uncertainty, such as neural signals that reflect the generation or perception of increased autonomic arousal (i.e., changes in respiration and cardiovascular activity), or even direct effects on blood flow arising from such changes (Birn *et al.* 2006). The answer to this question remains to be empirically determined.

Comparators

We can identify comparison operations at three different stages of the search process. At initiation, a comparison needs to determine if there is unexpected uncertainty, hence leading to exploration and initiation of search. During the search, a continuing comparison process needs to continue to check whether or not the search has found the goal. Finally, after termination of the search, an evaluation process needs to compare the observed outcome from the expected outcomes; that is, did the search accomplish what was expected?

The identification of an environment as being novel can play an important role in the search process and is one of the most obvious examples of a huge rise in unexpected uncertainty. When rats are faced with a novel environment, their first priority is safety, and they run to a location within the environment that has some protection from potential predators (Chance and Mead 1955). This location forms what is called a “home base,” from which they then explore in small journeys with a distinctive pattern: rats leave the home base, exploring with a slowly variable path, until they suddenly turn toward the home base and run directly home (Chance and Mead 1955; Eilam and Golani 1989;

Redish 1999; Whishaw and Brooks 1999). The outbound journey and the return journey have very different behavioral characteristics; the outbound journey is slow and meandering, whereas the return journey is ballistic (Drai and Golani 2001). The length of subsequent outbound journeys increases with experience, suggesting that rats are exploring increasingly more of the environment. Following from the hypothesis laid out here that search entails the reduction of uncertainty, we can identify these outbound journeys as searches that reduce the uncertainty in the environment. Whether this termination of the exploration path occurs due to reaching a threshold of novelty stress or fear (Crusio 2001; Pardon et al. 2002) or due to recognition of unreliability in the spatial representation as a result of drift in dead-reckoning systems (Redish 1999) is as yet untested. It is possible that drift in dead-reckoning systems (measurable to the animal as uncertainty in its position) can drive stress and fear, leading to a threshold at which the rat decides that it must return to the home base to reset its dead-reckoning information from a known position (Redish 1999).

Hippocampal comparators. A number of researchers have suggested that the CA1 region of the hippocampus serves as a comparator (Vinogradova 2001), particularly for the detection of novelty (Lisman and Grace 2005; Lisman and Otmakhova 2001). These hypotheses were based, in part, on anatomical and neurophysiological studies of convergent inputs from entorhinal cortex and CA3 on individual CA1 neurons (Groenewegen et al. 1987; Witter and Amaral 1991) under the assumption that the recurrent connections in CA3 could provide a delay. While it is true that hippocampal lesions significantly reduce spatial exploration (Archer and Birke 1983; O’Keefe and Nadel 1978; Redish 1999), particularly through a reduction in recognition of changes in the environment (Clark et al. 2000; Thinus-Blanc 1996; Zola et al. 2000), single cellular activity purely reflecting novelty has not been found in hippocampus. However, mismatch-like signals have been found in this region when rats were swimming in an annular maze and searching around the location where they expected a hidden platform (Fyhn et al. 2002). Changes in novelty are also reflected in population activity within the hippocampus, in that more cells are active in novel environments due to a reduction in inhibitory activity (Wilson and McNaughton 1993). Place cells generally show activity in their place fields from the first experience through the field (Hill 1978); nevertheless, they change their activity over the course of several hours (Cheng and Frank 2008), through an NMDA-receptor-dependent mechanism (Austin et al. 1993; Kentros et al. 1998). This suggests that while the firing of individual CA1 cells primarily reflects information about the world, differences in activity—even correlations between cell firing patterns—can be used to provide additional signals such as novelty. While cross-trial reliability can be interpreted as reflecting uncertainty (Fenton and Muller 1998; Kelemen and Fenton 2010), it can also reflect unaccounted-for parameters, external or internal (Johnson et al. 2009). As animals familiarize themselves with an environment, the decoded

position from the neural population becomes more accurate due, in large part, to the stabilization of place fields with experience (Austin et al. 1993; Wilson and McNaughton 1993).

Anterior cingulate comparators. In terms of evaluating whether a search has successfully achieved the specified goal, the anterior cingulate cortex has been found to respond to errors (Gemba et al. 1986), and it was originally conceived of as a comparator between actual and intended outcomes (Falkenstein et al. 1991; Gehring and Fencsik 2001). Although some work has cast anterior cingulate cortex as a conflict detector (Carter et al. 1998), there is now evidence that the anterior cingulate cortex compares actual versus expected outcomes (Ito et al. 2003; Jessup et al. 2010), as distinct from actual versus intended outcomes. The anterior cingulate cortex is especially active when a search is initiated, and it shuts off once the object of the search has been found (Shima and Tanji 1998; Bush et al. 2002), or even once the uncertainty about the object of the search has been eliminated (Procyk et al. 2000). This suggests that the anterior cingulate cortex is active during search to compare expected findings (including, but not limited to, the object of the search) against the actual findings. As the anticipated successful completion of the search becomes nearer in space and time, the anterior cingulate cortex cells become progressively more active (Shidara and Richmond 2002). Overall, the anterior cingulate cortex may monitor an ongoing search in two ways: (a) it may continually anticipate the outcome of a search and (b) it may become active when a comparison between actual and expected outcomes yields a discrepancy, which in turn requires corrective action (Modirrousta and Fellows 2008). In this way, the anterior cingulate cortex may monitor and contribute to effective search.

Identification of the Set of To-Be-Searched Options

From an ethological perspective, search is usually seen in terms of progress through space to reach a goal; however, cognitive search can occur in both “spatial spaces” (e.g., a rat trying to find a food source in a maze) and “non-spatial spaces” (e.g., selection among different goals or among different actions available to the animal). An important issue, therefore, is whether search processes that occur within different domains are processed by different brain structures. It seems plausible that searches involving various types of information will involve different neural structures which specifically encode, retrieve, or store that type of information. Neural systems of imagination and planning often utilize the sensory systems involved in their sensory processing; thus, for example, visual imagination involves primary and secondary visual cortex (Kosslyn et al. 2001), and a similar pattern of activation holds for the auditory cortices during imagination of sounds (Daselaar et al. 2010; Zatorre and Halpern 2005).

Neural Representations Specific to the Domain of Information To-Be-Searches

Searching for semantic information. Retrieving information in verbal fluency (and naming) tasks can be understood as a mental search through an internal representation (“lexicon”). Whereas category fluency (e.g., naming all animals that come to mind) and letter fluency (e.g., naming all words that come to mind that begin with the letter “L”) both share the necessity to initiate and control search, they differ with respect to the information that is retrieved: category fluency requires access to semantic information, whereas letter fluency is related to orthographical and phonological information. Early functional neuroimaging studies implicated areas of the prefrontal, parietal, and temporal cortices in this task (Friston et al. 1991; Frith et al. 1991b). Subsequent studies have attempted to dissociate the functional roles of these structures with respect to specific subcomponents of the task, such as accessing semantic information. This was mainly inspired by the notion that objects are characterized by a variety of features and associations in multiple sensory domains (e.g., smell, taste, color, shape, sound) but also in the action domain (e.g., associated movement patterns). This led to the hypothesis that diverse attributes of an object are represented in cortical areas that are involved in processing each particular type of information.

This hypothesis has been investigated by asking volunteers to retrieve specific semantic associations of objects. For instance, if the color of an object was relevant to the search, this led to an activation of the ventral occipito-temporal junction, an area that is also activated in the context of color perception (Chao and Martin 1999). Certain objects are well characterized by their use. This implies that in the representation of tools, motor areas might play a role. In agreement with this notion, activation in the left ventral premotor cortex has been observed in tasks involving the retrieval of semantic information pertaining to tools, such as their names (Chao and Martin 2000; Martin and Chao 2001). All the examples mentioned above used univariate tests, in essence showing increased activation for certain object categories. This was then complemented by observations indicating that even distributed information in cortical areas can be “decoded” using multivariate pattern classification techniques of fMRI data (Haxby et al. 2001; Polyn et al. 2005).

Searching through space. The hippocampus has been long identified as a key component of spatial navigation (Morris et al. 1982; O’Keefe and Nadel 1978; Olton and Papas 1979; Redish 1999), particularly in the context of spatial search processes (Johnson and Redish 2007; Morris 1981; Tse et al. 2007). There is also ample evidence to suggest that the hippocampus encodes more than just spatial representations, but may likewise be important for complex temporal information (Fortin et al. 2002). Recently, hippocampal cells have been shown to divide up temporal sequences when animals must run on a

treadmill during a delay to a goal (Macdonald *et al.* 2010; Pastalkova *et al.* 2008; Takahashi *et al.* 2009a). These firing patterns appear to act much like spatially encoding place cells, each of which fires only in a small portion of an environment, or along a small portion of a repeated journey (Levy 1996). Amnesic patients with medial temporal lobe lesions have been observed to be impaired in trace eyeblink conditioning, in which a temporal gap is introduced between the conditioned and unconditioned stimuli (Clark and Squire 1998; McGlinchey-Berroth *et al.* 1997). fMRI studies have also revealed activation of the hippocampus in Pavlovian trace conditioning (Buchel *et al.* 1999) that was not seen in a similar cue conditioning paradigm (Buchel *et al.* 1998; LaBar *et al.* 1998).

Contextual conditioning probes the association of a large set of multisensory stimulus features, including spatial information. Such learning has been shown to involve the hippocampus in rodents (Bouton 2004; Kim and Fanselow 1992) and humans (Cohen and Eichenbaum 1993), particularly in the face of contextual changes (Rawlins 1985; Redish 1999). Similar observations have been made in human functional neuroimaging, showing activation in the hippocampus in contextual fear conditioning (Lang *et al.* 2009; Marschner *et al.* 2008). Although many fMRI studies have highlighted the role of the hippocampus in establishing “maps” that include the dimensions of space and time, it is important to note that current functional imaging cannot provide enough detail about the underlying mechanisms of how the hippocampus integrates these features into such a map.

The hippocampus is not necessary, however, for simple one-step representations of causal structure in the world: hippocampal lesions do not interfere with either acquisition or performance of a lever press for food task (Corbit and Balleine 2000); animals with hippocampal lesions remain sensitive to devaluation, indicating that even without a hippocampus, animals remain knowledgeable about the consequences of their actions. Whether the hippocampus is necessary for deeper searches through causal structure is still unknown. Even in spatial tasks, the hippocampus is primarily necessary for the development of a world schema (cognitive map) on which expected outcomes can be placed; once the schema is learned, even new outcomes can be learned in the environment. Lesion data suggest the existence of nonhippocampal representations of such schematic, causal structure (Tse *et al.* 2007).

Nevertheless, as noted above, hippocampal lesions have profound effects on exploration and on the ability to use knowledge about the spatial world to find goals and targets, particularly when there is uncertainty (Kesner and Rogers 2004; Morris 1981; Redish 1999; Sutherland *et al.* 2011). As a classic example, the hippocampus is necessary to learn the location of a platform within a cloudy pool of water (the “Morris water maze”; Morris 1981; Sutherland *et al.* 2011), particularly when animals are started from many locations within the pool (e.g., with uncertainty in the starting point), and during early learning (e.g., with uncertainty in the location of the platform). The hippocampus is no

longer necessary if animals have a single cue they can approach (Eichenbaum et al. 1990), or if they are overtrained (Day et al. 1999), both of which reduce the uncertainty in the location of the platform. In these cases, other nonhippocampal systems are capable of guiding the rat to the platform, including systems for stimulus–response (S–R)-based, egocentric navigation (McDonald and White 1993; Packard and McGaugh 1992).

Neural Representations Independent of the Domain of Information To-Be-Searched

In humans, the anterior cingulate cortex is activated across a wide array of seemingly unrelated cognitive tasks involving very different cues and responses (Duncan and Owen 2000). In rats, anterior cingulate cortex and medial prefrontal cortex neurons appear to encode virtually all relevant aspects of any task the animal is required to perform, including cues and choices as well as reward magnitude, reward probability, action sequences, and abstract task rules (Hyman et al. 2005; Jung et al. 1998; Lapish et al. 2008; Narayanan and Laubach 2009). Furthermore, if the rules of a task change, there is a tightly correlated change in the way the same stimuli and responses are represented by the anterior cingulate cortex, both at the level of single neurons and ensembles (Durstewitz et al. 2010; Jung et al. 1998; Rich and Shapiro 2009). Therefore, the anterior cingulate cortex represents actions and stimuli with reference to the task being performed (Hoshi et al. 2005). Accordingly, it has been proposed that the anterior cingulate cortex, forming a continuum with adjacent medial prefrontal areas, is an integral part of a network that formulates task sets; that is, the dynamic configuration of perceptual, attentional, mnemonic, and motor processes necessary to accomplish a particular task (Dosenbach et al. 2006; Sakai 2008; Weissman et al. 2005).

As reviewed by Ridderinkoff and Harsay (this volume), the idea that the anterior cingulate cortex formulates task sets has been expanded to suggest a more general role of the region as part of a salience network that tracks all homeostatically relevant (salient) stimuli and events. In support of this idea, it has recently been observed that in the absence of an overt task situation, ensembles of anterior cingulate neurons formed highly distinct representations of novel environments which became less distinct as the environments became more familiar. However, the manner in which anterior cingulate cortex ensembles represented environments changed when tasks were performed. For instance, if rats had learned to perform a specific action in a specific environment, the action and the corresponding environment was represented by similar activity state patterns. In contrast, when rats were required to perform the same task continuously across different environments, the ensembles consistently represented only the task elements, and the representation of the environment was much less evident. While these data support the idea that the anterior cingulate cortex represents whatever is currently salient, they suggest that tasks are the

key factor in organizing these representations. Therefore, once a goal has been selected, the anterior cingulate cortex may formulate a task set or a representation of the set of stimuli and actions that are relevant to attaining the goal. It then tracks the progress of the animal within the task space (Lapish *et al.* 2008; Shidara *et al.* 2005).

Deliberation and Evaluation

Evaluation of stimuli and outcomes is important to the selection of intermediate or final targets of search and is equally important for computing prediction errors (discussed above). The assessment of outcome can occur in many ways, and we need to distinguish various modalities in the representation of outcomes: by value (understood with reference to homeostatic brain mechanisms defining the animal's needs) as well as by sensory properties defining the identity and quality of the outcome. For instance, a monkey searching for bananas can be said to have successfully completed its search once it finds a banana; however, in some cases this outcome is more valuable than in others. In the case of sensory-specific satiety (where the monkey has had its fill of a particular food, in this case bananas), the banana will be less valuable than if the monkey had not encountered a banana in some time. The specific taste and consistency of the banana define qualities other than its reward value. For example, an apple may be equal in reward value as compared to the banana, but yet have a different behavioral significance to the animal, potentially affecting its future search. Therefore, an important aspect of search is the determination of stimuli for which to search.

Whereas the gustatory cortex codes specific tastes of food rewards, the orbitofrontal cortex is important because it represents the value of outcomes (Padoa-Schioppa 2009) and it contains neurons that code the expected value of stimuli and actions (Baxter *et al.* 2000; Schoenbaum *et al.* 1998). Orbitofrontal neurons are activated by both primary rewards and conditioned reinforcers and may become activated before, during, or after a reward delivery. Neurons in this region can also discriminate between different rewards, largely irrespective of the actual features of reward-predicting stimuli or the responses used to obtain them (Padoa-Schioppa and Assad 2006). Neurons in rat orbitofrontal cortex are sensitive to different parameters of reward outcome (e.g., magnitude and probability of an upcoming reward; van Duuren *et al.* 2007, 2009). Perhaps most importantly, the responses of orbitofrontal neurons discriminate rewards based on their relative preference or value to the animal (Tremblay and Schultz 1999). Accordingly, the neural coding of food reward is subject to satiety (Rolls *et al.* 1999), confirming that neuronal activity is related to value coding. In addition, however, signaling within the orbitofrontal region also appears to reflect the sensory-specific qualities of the outcome regardless of value (e.g., a banana versus an apple, when valued equally; McDannald *et al.* 2011). Furthermore, orbitofrontal neurons also respond to aversively

predicting stimuli; here the response is again related to the relative preference of one aversive outcome versus another (Hosokawa et al. 2007; Morrison and Salzman 2009). Therefore, orbitofrontal neurons and networks might weigh the relative preference of different rewards as well as factor in whether the search for reward is offset by the potential harms involved.

The results of these calculations might then bias striatal activity so as to guide the appropriate actions to be taken (Simmons et al. 2007). Although the rodent ventral striatum receives little direct input from the orbitofrontal cortex (Schilman et al. 2008), it has also been strongly implicated in coding the value of outcomes as well as expected values. This structure receives strong inputs from the hippocampal formation and basolateral amygdala, which are important in forming stimulus–outcome (S–O) associations and relaying these to downstream areas, such as the ventral pallidum, to affect motor behavior (Parkinson et al. 2000). Many ventral striatal cells generate “ramps” in firing rate when animals are expecting a reward, with the firing becoming more intense as the animal gets temporally or spatially closer to reward delivery (Lansink et al. 2008; Lavoie and Mizumori 1994; Schultz et al. 1992; van der Meer and Redish 2011). Distinct subsets of ventral striatal cells code expected value at different task phases in advance of reward, or distinctly respond upon reward delivery. A possible difference between orbitofrontal and ventral striatal coding may reflect differences in representation of value and identity: while the ventral striatum is necessary for rodents to recognize any change in value, whether it be due to changes in amount of food delivered or in identity of food delivered, the orbitofrontal cortex was only necessary for rodents to recognize changes in identity (McDannald et al. 2011).

The coding of outcome value by ventral striatal cells may have two important functions. First, given the strong projection from the ventral striatum to the ventral tegmentum (the primary source of dopaminergic projections to cortical and limbic regions), the ventral striatum may provide expectancy and/or outcome signals that are used in the computation of prediction errors at the level of the dopamine cells. However, which brain areas are needed to compute dopamine error signals is not precisely known. Recent discoveries characterized the firing of habenula cells as an inverse signal, with high firing during disappointment (unexpected losses) and decreases in firing during surprising rewards (unexpected gains) (Bromberg-Martin et al. 2010a; Matsumoto and Hikosaka 2009). The habenula has an inhibitory influence on the dopamine cells through an inhibitory nucleus called the tail of the ventral tegmental area (VTA) (AKA rmTG) (Jhou et al. 2009). In addition to the contribution made by a ventral striatum to VTA projection, the orbitofrontal cortex may have an important role because contralateral orbitofrontal cortex–VTA inactivations have been reported to disrupt learning from unexpected outcomes (Takahashi et al. 2009b). Second, the ventral striatum projects to downstream structures, such as the ventral pallidum, and from there on to lower downstream structures in the brain stem, or up to the thalamus to complete an anatomical loop back to

prefrontal cortex. These output pathways are thought to convey motivational influences on patterns of motor behavior and cognitive processing.

Regarding the evaluation of different options in the spatial realm, hippocampal representations have been found to reflect future paths ahead of the rat, both during foraging behavior on mazes and in open situations (Lisman and Redish 2009), and specifically during decision making when rats are faced with explicit choices (Johnson and Redish 2007). During normal navigation, some hippocampal neurons fire when the animal is located in a particular place. These place cells fire at specific phases relative to an internal 7–10 Hz hippocampal local field potential called “theta” (Maurer and McNaughton 2007; O’Keefe and Recce 1993; Skaggs *et al.* 1996). It has been suggested that these phase precession phenomena represent a prediction of future paths that could be taken by the animal (Jensen and Lisman 1996). The discovery that the phases of firing in bidirectional place fields only converge in the two directions at the end of the field suggests that these place fields are, in fact, representing distance to a goal (Lisman and Redish 2009; Battaglia *et al.* 2004; Huxter *et al.* 2008). When rats are forced to make explicit choices, they sometimes pause and look back and forth between options, as if confused (or searching) between those options (Tolman 1938). During this pause-and-look behavior, termed “vicarious trial and error” (Muenzinger and Gentry 1931), hippocampal representations in area CA3 serially represent the potential options ahead of the rat (Johnson and Redish 2007). In downstream evaluative structures, such as ventral striatum and orbitofrontal cortex, cells that normally respond to reward also respond during these vicarious trial-and-error events (Steiner and Redish 2010; van der Meer and Redish 2009), suggesting a covert search-and-evaluation process (van der Meer and Redish 2010). These forward sweeps may represent cued memory retrieval, given the functions attributed to CA3 in this process (Marr 1971; McNaughton and Morris 1987; O’Reilly and McClelland 1994; Redish 1999), and may subserve prospective search.

One of the unsolved questions in this field concerns how the various evaluation systems interact. It is important to note that the orbitofrontal cortex most densely projects to the dorsomedial striatum (Price 2007; Schilman *et al.* 2008). Given that this striatal sector has been implied in mediating A–O associations (Yin *et al.* 2005), it is reasonable to hypothesize that orbitofrontal cortex may provide information about the outcome component of this associative process. However, caution should be exercised in this context; recent anatomical studies suggest five divisions of orbital cortex, from medial to lateral, with only the medial orbital and the most medial portion of ventral orbital projecting to medial striatum (Schilman *et al.* 2008). In fact, more lateral regions project largely to lateral and ventral regions of striatum and appear to play a role in stimulus-based rather than action-based predictions of outcomes (Ostlund and Balleine 2007). Certainly, as noted above, the ventral striatal system has generally been strongly implicated in mediating the motivational effects of Pavlovian cues and contexts on behavior. Finally, it should be emphasized that

there are many more brain structures where information about outcome expectancy is expressed. For instance, areas implied in visual processing, visuospatial behavior, and visual search express reward value information (parietal cortex: Platt and Glimcher 1999; visual cortex: Shuler and Bear 2006). These data reveal how widespread the effects of reward expectancy are across the brain, whereas the causal generation of evaluative signals driving action selection is likely to be primarily restricted to frontal-basal ganglia circuitry.

Working memory is also critical at this stage for strategy development as it allows the organism to consider multiple options online. The effect of strategy representations in working memory is essentially to narrow down the initial pool of candidate actions that may be employed in the search. There is general agreement that working memory involves the network interactions among lateral prefrontal cortex and parietal cortex (Chafee and Goldman-Rakic 2000), although subcortical regions such as the striatum and hippocampus can also contribute. Overlapping regions of ventrolateral prefrontal cortex may therefore provide top-down cognitive control of cognitive search from this perspective (Nobre et al. 2004).

Action Selection

At the conclusion of this deliberation and evaluative stage, it can be presumed that an action is needed to test the predictions of the search and obtain the goal in question. This is true regardless of whether the goal was abstract, such as information (a test of the individual's newly updated representation of the world), or physically substantive, such as food reward. There are times when this will involve complex action planning, thereby requiring almost a separate subsearch in which a set of possible actions must be identified and evaluated, versus simpler engagement of a well-known motor movement.

Analyses of action systems in human or animal subjects usually depend on experimentally highly constrained situations, such as voluntary, as distinct from stimulus-elicited, limb or eye movements receiving rewarding feedback. In functional imaging tasks in humans, these engage regions of the prefrontal cortex, including dorsolateral regions, as well as the premotor and supplementary motor cortex, which project into the so-called parallel loops of the corticostriatal systems (Alexander et al. 1990; Berendse et al. 1992; Voorn et al. 2004; Zahm and Brog 1992). Again, it is important to realize that these activations can also be produced (generally to a lesser degree of activation) by instructions to imagine a given action, such as serving the ball in tennis, or even thinking of the meaning of a hammer (Martin and Chao 2001). The human functional imaging observations have been paralleled by experimental observations from electrophysiological recordings in nonhuman primates. Thus, there is a cortical representation of many voluntary actions in premotor regions that may also be sensitive to observations of others performing similar actions (so-called mirror neurons: Cattaneo and Rizzolatti 2009; Rizzolatti et al. 2009). It has

been presumed that such “ideo-motor” representations may be important in inferring intentions in social situations. Electrophysiological observations linking action representations and outcome representations in nonhuman primates suggest that there is a distribution of reward-related activity throughout the entire prefrontal cortex but that it is only in certain regions coincident with representations of action information (Kennerley and Wallis 2009; Wallis and Miller 2003); this notably includes the anterior cingulate cortex. Furthermore, human neuroimaging and nonhuman primate lesion data highlight an important role for this region in representing A–O information (Walton et al. 2005).

Role of Prefrontal Cortex in Action Planning

Action planning is likely to involve the prefrontal cortex and supplementary/premotor areas. For instance, in rodents and nonhuman primates, the medial prefrontal cortex contains a large fraction of neurons sensitive to the setting of task rules (Birrell and Brown 2000; Durstewitz and Seamans 2002; Mulder et al. 2003; Peyrache et al. 2009; Rich and Shapiro 2009; Wallis et al. 2001). Upon a switch of strategy and adoption of another task rule, subsets of prefrontal cortex ensembles that were previously active now become inactive, and previously silent ensembles are activated. Further evidence from primates has implied prefrontal and premotor/supplementary motor area structures in planning and executing complex action sequences (Averbeck et al. 2006; Wise et al. 1996), and the most rostral components of prefrontal cortex appear to be involved in the hierarchical organization of behavior and of complex cognitive operations (Koechlin et al. 2000, 2003). It is not yet known whether the rapid alterations in the temporal organization of frontal activity correspond to internal, generative search processes themselves or to the execution of planned actions and application of task rules. However, if we assume that the information retrieved during forward sweeps in the hippocampus is of a generative nature (see earlier section on Deliberation and Evaluation), and that it is coupled in time to similar processes in connected brain areas, then it is straightforward to hypothesize that internal search for future actions involves medial prefrontal cortex—which receives strong hippocampal input (Jay and Witter 1991) that produces firing time-locked to the hippocampal theta rhythm, particularly during decision making and attentive tasks (Hyman et al. 2010; Jones and Wilson 2005)—and related structures for planning and action selection.

Passingham (1993) has reviewed evidence that the medial premotor cortex is required to retrieve the appropriate movement in the absence of external cues or prompts. However, he also concludes that the dorsolateral prefrontal cortex is required for self-directed sequences of actions that often make up goal-directed behavior. Damage to Brodmann area 46 impairs self-ordered visual search behavior in monkeys (Passingham 1985); analogous results have been found following dorsolateral prefrontal lesions in humans (Manes et al.

2002; Owen et al. 1990). It is, however, not yet clear whether the deficits arise from working memory or response selection impairments (or both). Frith and colleagues have provided evidence that self-generated sequences (of “willed action”) activate areas 9/46 within the dorsolateral prefrontal cortex, when there is no obvious working memory component (Frith et al. 1991a).

Action planning has also been studied in humans in terms of the Tower of London problems, which involve sequencing a set of actions to obtain a single specified goal (Shallice 1982). Note that this sequence can also be an imagined sequence (Owen et al. 1995). To solve such tasks, subjects have to search through a number of possible sequences in a finite problem space, a process that may correspond to “deliberation.” These sequences can include various key “intermediate positions,” or subgoals, which can serve as aids to a solution when it begins to exceed working memory capacity. Performance on such tasks is known to depend on a fronto-parietal-striatal system (Baker et al. 1996; Owen et al. 1990; Shallice 1982), notably involving the dorsolateral and dorsomedial prefrontal cortex. The presumption is that the anterior cingulate cortex may represent the general task set, as reviewed above, whereas the set of visuospatial options may be encoded by the parietal cortex and the execution of the selected sequence in the basal ganglia. Finally, the dorsolateral prefrontal cortex may be especially involved in response selection (Frith et al. 1991a).

Solution of the Tower of London problems is not conventionally related to reward outcomes unless a specific payoff matrix is devised, in which case these action sequences are more likely to engage reward representations in the neural systems encoding value, such as within the orbitofrontal cortex (Wallis et al. 2001). In the conventional task, however, a successful outcome is symbolized simply by correct feedback for the solution. For this reason, such tasks are often labeled as exemplifying “cold” cognitive processes. Planning can, however, involve more complex A–O searches, for example, in selecting actions that anticipate future long-term motivational needs. In addition, planning can involve the scheduling of actions to obtain multiple goals (as in shopping), a task exemplified by the so-called “six elements test,” which is especially sensitive to damage of the anterior frontal prefrontal cortex (Burgess et al. 2000).

Action Planning within the Basal Ganglia

Although areas of prefrontal cortex are no doubt involved in action planning, action selection itself is thought to depend critically on activity within the basal ganglia. Action selection initiates a process of action evaluation through the response–outcome (R–O) association; that is, the value of the action is estimated on the basis of the predicted reward value of the outcome which is contingent on that action. Finally, the action selection and evaluation processes combine to initiate an action (see also Balleine and Ostlund 2007). Of the current theories of how this is achieved, perhaps the most plausible is the

associative-cybernetic model, which posits that action selection is largely controlled by stimulus(outcome)–response learning, S(O)–R, and hence by outcome retrieval (Dickinson and Balleine 1993). Although the learning of A–O associations has been associated with discrete structures (prelimbic cortex) within the rat medial prefrontal cortex, such A–O learning also implicates the dorsomedial striatum, to which this region of the rodent frontal cortex projects (see O’Doherty and Balleine, this volume). Unlike the prelimbic cortex, the dorsomedial striatum is also apparently necessary for the expression of goal-directed actions, and so the dorsomedial striatum is presumably responsible for response selection in situations where search is required between different established options. Other regions of the basal ganglia, such as the dorsal putamen (or dorsolateral striatum in rodents), are implicated in the learning and expression of S–R habits where no outcome or goal is represented and which therefore is not considered to require a cognitive search.

The fact that the dorsomedial region of the striatum mediates the encoding of R–O associations, whereas the dorsolateral region mediates S–R learning, poses some problems for the associative-cybernetic model: the critical associative and S–R memory systems that contribute to instrumental performance course through corticostriatal circuits localized to adjacent regions of the dorsal striatum, but it is unclear how these two pathways interact to permit the formation of S(O)–R associations identified as critical for action selection. The generally accepted architecture of the basal ganglia emphasizes the operation of functionally distinct, closed parallel loops connecting prefrontal cortex, dorsal striatum, pallidum/substantia nigra, thalamus, and feeding back onto the originating area of prefrontal cortex (Alexander *et al.* 1986). According to this view, there is considerable vertical integration within loops but less clearly lateral integration across loops. As a consequence, various theories have had to be developed to account for lateral integration: the split loop (Joel and Weiner 2000) or spiraling midbrain-striatal integration (Haber *et al.* 2000; Haruno and Kawato 2006). These models have not yet found wide acceptance. In contrast, older theories of striato-pallido-nigral integration proposed that, rather than being discrete, corticostriatal connections converge onto common target regions particularly in the globus pallidus—a view that allows naturally for integration between various corticostriatal circuits (Bar-Gad *et al.* 2003). Although anatomical studies challenge this view, recent evidence has emerged supporting a hybrid version; in addition to the segregated loops, there may also be integration through collateral projections from caudate (or dorsomedial striatum) converging with projections from the putamen (or dorsolateral striatum) onto common regions in both the internal and external globus pallidus (Nadjar *et al.* 2006). Whether these converging projections underlie the integration of the O–R and R–O associations, which the associative-cybernetic model identifies as critical for the initiation of instrumental performance, remains an open question.

Search Termination

Once a particular search action has been executed, the outcome of the action must be evaluated in terms of whether or not it led to successfully achieving the anticipated goal of the search. As outlined earlier in our discussion, this process involves comparator operations which likely take place within the anterior cingulate cortex. If the comparator output indicates a discrepancy between the actual versus expected outcome, this signals two items of information. The first is that a corrective action may need to be taken (Modirrousta and Fellows 2008). In the case of search, the corrective action is to terminate the current strategy and initiate a new search. There is evidence from monkey neurophysiology that anterior cingulate cortex is active especially during the time of a search; that is, from the time when an unsuccessful strategy is rejected until a new strategy is found (Procyk et al. 2000; Shima and Tanji 1998). In experienced animals, such searches may not be random but instead near optimal (Procyk and Joseph 1996), such that at least in certain tasks, experienced animals do not often try an unsuccessful option twice during a search. This suggests a kind of inhibition of return in higher-level cognitive search, similar to that found in lower-level visual search (Wolfe 2003). A second piece of information signaled by discrepancies is that the A–O prediction was potentially incorrect and should therefore be updated (Matsumoto et al. 2007). When the environment is nonstationary or highly volatile, such predictions will be continually out of date and will therefore yield ongoing discrepancy with the actual outcomes, as has been observed with fMRI studies (Behrens et al. 2007).

Hence, just as search is initiated by rising uncertainty or enhanced motivational drive, search termination can be triggered by a reduction in uncertainty or the recognition that the uncertainty is irreducible, implying that the uncertainty is expected, rather than being a form of unexpected or estimation uncertainty. As discussed above, detection of changes in uncertainty will again involve comparators. Search termination may also result from a shift in motivational state, either appetitive or aversive. For example, in the aversive case, termination of open space exploration in rodents might be triggered by increases in anxiety and stress upon departure from the home base, as indexed, for example, by increases in heart rate variability (Aubert et al. 1999; Behrens et al. 2007). Increases in danger signs (e.g., suddenly bright lights or the addition of predator odor) will send an exploring rat scurrying back to its home base. Brain regions involved in aversion-induced processes include the amygdala and the prefrontal cortex, suggesting that they may well play a role in search termination. Similarly, parts of the frontal cortex are also likely to play a role in the case of shifts in appetitive motivation (e.g., through detection of satiety signals). Thus, the orbitofrontal cortex may play a central role as evidenced by the existence of satiety-responsive neurons in the medial orbitofrontal cortex of the macaque (see section on Deliberation and Evaluation).

This observation highlights a role for reductions in hedonic value in search termination (Pritchard *et al.* 2008).

The final stage of search termination is the process of action stopping or response inhibition, which involves fronto-basal ganglia networks (Eagle and Baunez 2010). Two forms may be distinguished, with nonselective stopping (or “clamping”) of already initiated actions recruiting primarily a network connecting the inferior frontal cortex with the supplementary motor area and the subthalamic nucleus (Aron *et al.* 2007; Schall *et al.* 2002; Stuphorn 2006; Stuphorn and Schall 2006). However, a form of stopping that might have wider validity in the context of search is selective stopping, which involves a plan to stop only a select set of actions (Aron 2010). This latter, more proactive form of inhibitory control is generated according to the goals of the subject rather than by an external signal, and has hypothetically been claimed to involve the striatum and its modulation by dopamine.

Neurochemical Regulation of Search

The major ascending neuromodulatory systems, dopamine, norepinephrine (noradrenaline), serotonin, and acetylcholine, play a critical role in many, if not all, of the subcomponent processes of search that we have outlined above (see also Cools *et al.*, this volume). Dopamine, for example, is well known to alter performance on high-level cognitive tasks, such as the Tower of London forward planning test, probabilistic reversal learning, and self-ordered search in spatial working memory (Robbins 2007; Cools 2006). Although much work has highlighted the role of dopamine in reinforcement-based habit learning associated with the dorsolateral striatum, these high-level cognitive effects likely reflect modulation of goal-directed search processes associated with the prefrontal cortex and dorsomedial parts of the striatum (Cools *et al.*, this volume). These dopamine-sensitive processes may include search initiation, option identification, search evaluation, option selection, or search termination. To illustrate the importance of neurochemical modulation in search, we highlight in this section some data evidencing its implication in search initiation.

Dopamine likely contributes to search initiation by signaling a reward prediction error (Schultz 2007). However, it also contributes to search initiation via mediating changes in the motivational state. For example, increases in anxiety and stress are known to be accompanied by changes in neurochemical state, such as supra-optimal increases in dopamine, norepinephrine, and corticosteroids (Arnsten 2009), which in turn have been demonstrated to disrupt the optimal functioning of the prefrontal cortex (Seamans and Yang 2004). Accordingly, anxiety or stress might mediate search termination by inducing supra-optimal levels of dopamine and norepinephrine in the prefrontal cortex. The importance of neurochemical state changes are also likely to play a role in the case of appetitive motivational shifts, such as satiety, which involves

modulation of the orbitofrontal cortex and the striatum by the mesolimbic dopamine system (and its interactions with the hypothalamus) and appetite-regulating hormonal systems (Farooqi et al. 2007; Kringsbach and Stein 2010). Search initiation might also depend on noradrenergic activity, which has long been known to affect attention, particularly in the face of different levels of uncertainty (Robbins 1997). Thus elevated tonic norepinephrine levels, arising from activity within the locus coeruleus, might play an important role in initiating search by serving a network reset function; such a function enables the interruption of ongoing activity (Sara 2009) and the revision of internal representations based on new sensory input (Aston-Jones and Cohen 2005b; Yu and Dayan 2005). Specifically, it has been suggested that norepinephrine is particularly crucial when changes in the environment are unexpected (as opposed to expected; Yu and Dayan 2005). Consistent with this hypothesis are observations that manipulations of norepinephrine affect performance on paradigms in which behavioral change is driven by unexpected uncertainty, such as those measuring extra-dimensional set shifting (Robbins and Roberts 2007). Extra-dimensional set shifting requires subjects to shift attention from one dimension of multidimensional stimuli (e.g., shape) to another (e.g., color), and might be particularly appropriate for modeling search processes due to the requirement to identify and evaluate different response strategies based on alternate sets of stimulus features.

Conversely, it has been argued that behavioral change signaled by expected uncertainty is mediated by acetylcholine, a hypothesis that is consistent with observations that cholinergic changes are associated with attentional shifts in Posner-like attention-orienting paradigms where subjects are aware of cue invalidity (Hasselmo and Sarter 2011). In contrast, cholinergic manipulations generally leave extra-dimensional set shifting unaffected. This general distinction between the norepinephrine and acetylcholine systems fits with observations on intra-dimensional reversal learning in response to changes in reward contingencies: late, but not early, reversal learning (i.e., when changes are expected) is sensitive to acetylcholine, but not norepinephrine (Robbins and Roberts 2007). Accordingly, both increases in (tonic) norepinephrine and acetylcholine may align attention with a source of sensory input, perhaps by enhancing sensory thalamic input to the prefrontal cortex and by shutting down current top-down internal models held online by the prefrontal cortex (Hasselmo and Sarter 2011; Chamberlain et al. 2006; Yu and Dayan 2005). However, the signals that trigger this norepinephrine- and acetylcholine-mediated shifting might differ, with only the former type of shifting (i.e., that triggered by unexpected uncertainty) being relevant for search as defined here.

In addition to dopamine and norepinephrine, serotonin is also critical for search initiation, as evidenced by its implication in behavioral extinction (Walker et al. 2009), which depends on behavioral change in response to unexpected uncertainty. Specifically, Walker et al (2009) have shown that depletion of both dopamine and serotonin in the orbitofrontal cortex of nonhuman

primates (marmosets) impaired extinction of previously rewarded behaviors, albeit in different ways, with serotonin depletion specifically suppressing the exploration of the never-rewarded option, though allowing overall extinction to proceed normally. By contrast, depletion of orbitofrontal dopamine allowed normal exploration of alternatives to occur but greatly retarded extinction.

Finally, we note that the relationship between effects of neurotransmitters and search is complex and nonlinear. In the case of dopamine, it is well established that there is an inverted U-shaped relationship between levels of dopamine receptor stimulation and performance on working memory tasks, whereby both too little and too much dopamine are associated with poor performance (Arnsten 1998). The implication of this observation is that increases in dopamine (e.g., through administration of dopamine-enhancing drugs) will improve performance of individuals with suboptimal baseline levels of dopamine, while impairing performance of individuals with already optimized baseline levels of dopamine. Similar nonlinear functions have been established for Tower of London planning (Williams-Gray *et al.* 2008), cognitive switching (Cools and D'Esposito 2011), and probabilistic reversal learning (Clatworthy *et al.* 2009), all involving cognitive search. Furthermore, there are multiple inverted U-shaped functions, so that effects of dopamine depend not only on the baseline neurochemical state of the system, but also on task demands (Cools and D'Esposito 2011; Cools and Robbins 2004). Thus, administration of dopaminergic drugs have been shown to improve performance on one type of cognitive search (i.e., probabilistic reversal learning), while simultaneously impairing performance on another type of cognitive search (i.e., spatial working memory), even within the same individual (Clatworthy *et al.* 2009; Cools *et al.* 2001).

Interim Summary

To summarize, we suggest that cognitive search is a goal-directed behavior that can exist across multiple domains (spatial, causal structure, goals, actions) and that a fundamental aim of a cognitive search is to reduce the unexpected (or estimation) uncertainty present at any of these levels. The search process itself can be compartmentalized into five general stages: initiation; outlining of the to-be-searched options; deliberation and evaluation; action planning/selection; termination. A theoretical depiction of how search could be structured based on what we know of neural function and specialization, is outlined in Figures 9.1 and 9.2.

A rise in unexpected uncertainty, represented in the brain within the anterior cingulate, parietal and inferior frontal cortices, can provide the trigger for search initiation. Comparator computations, such as those performed within the hippocampus and anterior cingulate, may make a critical contribution in terms of detecting outcomes which deviate from what was expected. Dopamine

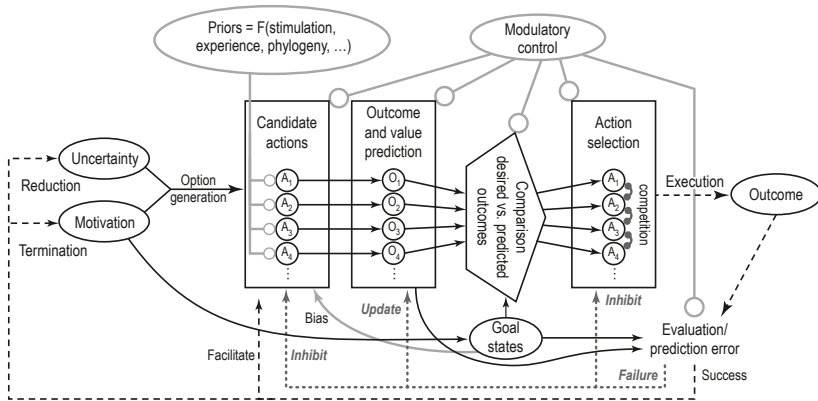


Figure 9.1 Functional circuit for cognitive search. The search process is initiated by a basic motivation or need, such as hunger, combined with an uncertainty about how to attain this goal. This leads to the generation of a number of candidate strategies for how to resolve this situation. The candidate options generated depend on the currently present sensory input, prior experience, biological biases, etc. A prediction for the outcome and value attained by choosing any particular option is generated, which is then compared to the desired goal state. This will narrow down the pool of candidate options. In a competitive process, the option is selected that is most strongly favored by biases, prior experience, proximity of predicted outcome to the goal, etc. The selection of the appropriate actions and thereby the execution of this option will lead to an actual outcome, which is then evaluated with respect to the desired goal state, yielding a prediction error signal. Depending on the sign of this signal, action, value, and outcome representations will be updated. In case of failure, the action will be inhibited for subsequent selection, whereas in the event of success, uncertainty will be reduced and the need may be resolved. At any stage, these processes may be modulated to widen or narrow, for example, the scope of the search.

signals, which are thought to carry prediction error information, can likewise play an important role at this stage. Motivational states, such as hunger or fear, can also stimulate an environmental search, and these are signaled by a broad range of brain systems, including the amygdala. The set of to-be-searched options is identified, and the neural structures involved may vary by the type of information under scrutiny; the hippocampus, for example, is involved in representing searchable spatial or temporal representations. However, the anterior cingulate may play a relatively unique role at this stage in that it appears to represent diverse sets of information that are relevant for the task at hand.

During the deliberation phase, predictions regarding the outcomes and values of these options are generated and evaluated in terms of whether they are likely to meet the goals of the search. Key areas of frontal cortex, such as the dorsolateral and orbitofrontal cortex as well as the hippocampus, ventral striatum and caudate putamen, play distinct roles in this process. Again, the anterior cingulate is implemented due to its involvement in conflict and prediction error monitoring. Once the most promising option has been identified, the

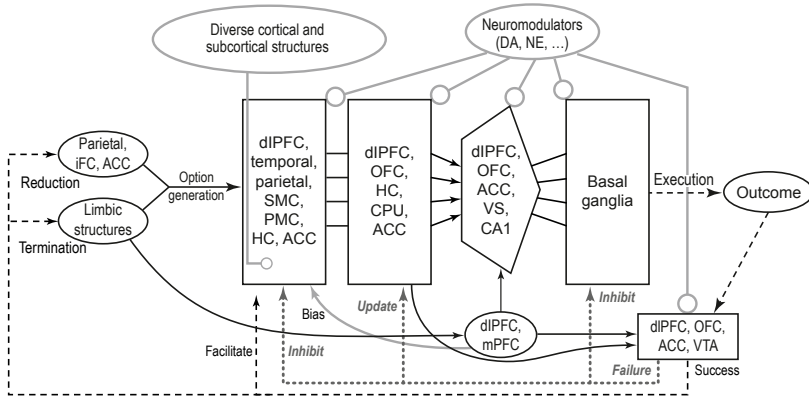


Figure 9.2 Anatomical circuit for cognitive search. This schema illustrates how the functional processes described in Figure 9.1 may map onto anatomical structures. ACC: anterior cingulate cortex; CA1: region of hippocampus; CPU: caudate putamen; DA: dopamine; dIPFC: dorsolateral prefrontal cortex; HC: hippocampus; iFC: inferior frontal cortex; mPFC: medial prefrontal cortex; NE: norepinephrine; OFC: orbitofrontal cortex; PMC: premotor cortex; SMC: supplementary motor cortex; VS: ventral striatum; VTA: ventral tegmental area.

actions required to execute this strategy are determined largely through competition within the basal ganglia, but also via input from the dorsolateral frontal cortices, as well as premotor and supplementary motor cortex, if the motor sequence is complex. The final outcome of the search is then evaluated, a phase which recruits similar regions involved in comparison processes at other stages of the search, and the search process terminated if either the motivational state is resolved or uncertainty reduced. The neuromodulators dopamine and norepinephrine may be particularly important in multiple aspects of the search process due to their ability to influence cognitive flexibility and arousal, but other neurotransmitters such as serotonin may also play a role. Optimal levels of these neurotransmitters may vary for different types of search depending on the precise cognitive processes involved.

Pathologies of Search Processes

Examination of deficits in patients with focal lesions (e.g., in the prefrontal cortex) begins to inform us about the mediation of specific neural components of the search process. In general, patients with prefrontal cortex lesions are impaired in search-like situations that benefit from the application of strategy or structure to the problem, such as the Tower of London (Shallice 1982). For example, impairments in the application of strategy in a self-ordered spatial search task have been observed in patients with focal lesions in the lateral

prefrontal cortex, but not in temporal lobe lesion patients (who nevertheless exhibit mnemonic deficits on the task; Owen et al. 1990). Further examples of the role of the prefrontal cortex in search focus on retrieval or generation strategies as exemplified by verbal fluency and alternate use (divergent-thinking) tasks (Eslinger and Grattan 1993). Frontal patients are impaired in imposing a strategy on category and letter retrieval, even though their semantic lexicon is relatively intact (Baldo et al. 2006). Moreover, functional neuroimaging studies have strongly implicated the lateral prefrontal cortex in memory retrieval processes involving recall and/or selection of either verbal or nonverbal material (Badre and Wagner 2002; Cabeza and Nyberg 2000; Thompson-Schill et al. 1997). Frontal patients have difficulty not only with searching the past, but also with “searching the future,” and the left frontal cortex is thought to underlie such “mental time travel” (Nyberg et al. 2010).

By contrast, patients with brain damage in the parietal cortex can exhibit deficits in search-like processes (“neglect”), not because of a problem with imposing structure or strategy, but rather because of a basic spatial representational deficit, leading to a restricted set of options available for search (Vossel et al. 2010). Moreover, some patients with predominantly posterior cortical lesions in the left hemisphere experience forms of apraxia that may resemble search deficits, but can be understood in terms of difficulty with retrieving semantic representations of actions.

Patients with medial temporal lobe lesions exhibit difficulties in cognitive search (Hassabis et al. 2007) and tend to use action-selection systems that do not depend on search processes. Some have suggested that this deficit occurs due to deficits in stored memory representations (Squire 1987; Buckner and Carroll 2007). Others have suggested that this deficit arises from a problem in the construction of novel conjunctions of representations, particularly of episodic representations of the potential future options (Hassabis et al. 2007; Buckner and Carroll 2007). Both suggestions are controversial (Atance and O’Neill 2001; Sutherland et al. 2011; Holland and Smulders 2011; Nadel and Moscovitch 1997).

Problems with search-like processes also surface in a wide variety of neuropsychiatric and neurological disorders, which are characterized by a more diffuse pattern of neuropathology but striking functional deficits. For example, certain symptoms of obsessive-compulsive disorder, depression, Parkinson’s disease, schizophrenia, addiction, and attention deficit hyperactivity disorder can be interpreted within the current theoretical framework. In the case of depression, for example, the search space might be restricted as a result of negative and affective biases that limit the capacity to recall information or generate future options (Beck et al. 1979; Sutherland et al. 2011; Lloyd and Lishman 1975; Murphy et al. 1999). By contrast, such affective biases may be required for normal socio-emotional decision making (Damasio 1994), including moral judgments. These may go awry in proactive aggressive disorders, like psychopathy (Blair 2008; Blair and Mitchell 2009). Thus psychopaths may search

an abnormally wide range of options in a way that is not constrained by social rules of affective biases, such as disgust, empathy, and fear.

Obsessive-compulsive disorder provides some particularly interesting potential applications of the current framework. On one hand, obsessive-compulsive disorder can be characterized as a failure to complete a search, particularly in the domain of obsessions, leading to excessive checking or monitoring behavior and “worrying,” possibly as a consequence of anterior cingulate dysfunction (Schlosser et al. 2010). On the other, there is some evidence that the normal balance between A–O knowledge and habitual knowledge is biased toward the latter (Gillan et al. 2011), likely reflecting the known orbitofrontal-striatal dysfunction present in obsessive-compulsive disorder (Menzies et al. 2008).

Addiction can also be characterized in terms of a narrowing of effective goal states. Thus the search for drugs occludes that for other goals that drive adaptive behavior, such as food and social interaction (Hyman and Malenka 2001). Whether this is due to motivational deficits, search process deficits, or other problems is still unknown and controversial (Altman et al. 1996; Redish et al. 2008). Finally, delusional symptoms in psychosis, including schizophrenia, can also be cast in terms of the current framework. Specifically, these symptoms of “abnormal beliefs” have been argued to reflect a search-like disturbance in constructing causal models of the world, which can lead to inappropriate “jumps to conclusions” (Fletcher and Frith 2009). The anterior cingulate cortex shows reduced error signaling in schizophrenia (Carter et al. 2001), and subsequent work showed that these reduced error effects stem from an underlying deficit in the ability to predict the consequences of an action in schizophrenia (Krawitz et al. 2011).

Concluding Remarks

Our aim was to consider cognitive search in such a way that would allow some hypotheses to be generated regarding its underlying neural and neurochemical bases. As is often seen, when evaluating the contribution of behavioral neuroscience to the larger field of psychology, consideration of the biological underpinnings of search helped to critically inform the discussion as to the nature of the search process itself. At the outset of this discussion, we defined search as a goal-directed behavior which could be parsed into five key stages. Although few studies have addressed the biological basis of cognitive search per se, careful consideration of the psychological constructs implicated at each stage has allowed for the creation of a model that reflects the neural circuitry so far identified in mediating these subprocesses.

When exploring the rationale for this model, data were considered from a range of experimental paradigms, including human imaging studies, neuropsychological assessment of brain-damaged patients, lesions, and

electrophysiological studies in animals. The evidence from these disparate fields largely overlapped in pinpointing which brain areas may be responsible for performing the specialized functions we identified as inherent in the search process. The ability of researchers to record ongoing neural activity while animals are performing certain goal-directed behaviors clearly allows advances to be made in determining how particular computational functions (such as the calculation of prediction errors or the generation of comparisons) may be accomplished at a neuronal level. Computational modeling theories continue to evolve, and their ability to approximate, decode, and predict both single cell, network, and population activity is constantly improving. Our understanding of how our brains are capable of implementing complex processes, such as a cognitive search, will certainly benefit from this growing field.

While the anterior cingulate cortex appears to be crucial to so many of the stages of search identified here, particularly with respect to evaluation of ongoing behavior, questions still remain as to how expectancies are generated and interpreted within this and other brain regions. Although it seems fairly well established that dopaminergic firing can signal prediction errors, which area(s) provide(s) the critical inputs that drive those predictions? How does neuronal activity within the anterior cingulate shape the prediction error signal, or change as a result of its detection? If the anterior cingulate is already crucial for many phases of search, what are the additional functions of the striatum?

Current data also suggests that the dopamine system does much more than carry prediction errors, yet this signal has proved particularly amenable to investigation at both the neuronal and behavioral levels of analysis. Understanding how drugs and chemicals can influence, and are influenced by, neuronal and cognitive function remains an important goal of neuroscience research, particularly with regards to improving treatment options for psychiatric illness. Models which capture how neuronal circuits are modified by the tonic and phasic firing patterns generated not just by dopamine neurons, but by neurons that produce norepinephrine and serotonin and other neurotransmitters, may be heuristically useful in guiding experimental design in this field. Such models depend on continuing evaluation of drug effects on behavior and brain function, experiments which are highly informative in their own right.

We have been proscriptive in specifying how cognitive search might operate in the mammalian brain; namely, in the context of goal-directed action. We have indeed eschewed what might turn out to be only superficial comparisons with behavior in many species that is ostensibly goal-directed, but which has not been subjected to rigorous experimental tests of its goal-directed nature. It is nevertheless possible that the physical basis of search processes postulated here as contributing to goal-directed search might be related to more general biological processes. Only very high-level descriptions of what search processes entail will ultimately be able to address this issue. In the interim, a useful strategy will be to compare the nature of the search processes for the various components we have defined as contributing to goal-directed search, most

of which depend on distinct neural networks. Such comparisons will determine whether similar neurocomputational principles are implicated, and hence whether there are fundamental aspects of search mechanisms in the brain held in common.