

2

The Evolution of Cognitive Search

Thomas T. Hills and Reuven Dukas

Abstract

Search can be defined as an attempt to arrive at a goal at an unknown location in the physical environment, as well as in time, memory, or any other space. Search is necessary because the quantity and quality of resources essential to survival and reproduction vary in space and time. In addition to exploration through actual body movement in their environment, animals search their external information space through selective allocation of attention and their internal information space to retrieve relevant items from memory. This chapter integrates data on search in three distinct domains—physical movement, attention to external information, and locating items in memory—to highlight the remarkable similarities between these three domains. First, resources in all three domains are typically distributed in patches. Second, in each of the three domains, animals typically keep searching in patches where they have recently found resources and leave areas when none are found or where they have already depleted the resources. Third, the neurobiological mechanisms modulating the exploration for and exploitation of resources in all three domains involve dopamine as well as, in many vertebrates, regions of the prefrontal cortex and striatum. It is suggested that, throughout evolution, animals co-opted existing strategies and mechanisms used to search their physical space for exploring and exploiting internal and external information spaces. The cross-disciplinary integration of theory and data about search can be used to guide future research on the mechanisms underlying cognitive search.

Introduction

Search is one of the most fundamental of all organismal behaviors. Bacteria seek out essential nutrients and steer clear of noxious compounds (Koshland 1980; Eisenbach and Lengeler 2004), plant roots search for water and nutrients (Hutchings and de Kroon 1994; McNickle et al. 2009), and the protozoan *Paramecium* exhibits chemotaxis as well as thigmotaxis, geotaxis, and thigmotaxis (movement in response to touch) (Jennings 1906; Saimi and Kung 1987). In general, organisms that move are capable of searching for optimal abiotic

settings, such as temperature, humidity, and sunlight, as well as the best places for finding nutrients, avoiding danger, and securing sexual partners.

In addition to physically moving through the environment, animals may search within the information space for cues indicating relevant resources. The information space may be external (e.g., requiring the direction of attention in pursuit of cues that signal prey) or internal (e.g., requiring the directed activation of memory). Regardless of whether physical movement is involved, search entails navigating some space in pursuit of resources; that is, an individual has to decide whether to move (its body or its attention) or stay where it is, and, if it moves, where it should move to. In the domain of physical space, such search problems have been studied extensively in behavioral ecology (Stephens et al. 2007). Research on information search, in both external and internal environments, is developing rapidly (e.g., Fu and Gray 2006; Stephens and Krebs 1986; Hills and Hertwig 2010; Pirolli 2007; Wilke et al. 2009).

In this chapter, we are interested in drawing attention to the potential evolutionary parallels between search across external and internal domains. How might search in external and internal domains be related in an evolutionary sense? Three potential types of evidence can be used to address this question:

1. The neurobiological mechanisms that guide search in different animals may be functionally homologous, deriving from a common ancestral function that was also used to solve search-related problems.
2. Different environments may pose similar kinds of problems for search, generally involving navigating heterogeneous resource distributions to find locations containing resources that maximize fitness.
3. The underlying search strategies may share similar characteristics across different environments and domains.

We begin by providing a definition of search and then briefly review the three characteristics of environmental structure, search strategies, and neural mechanisms involved in search tasks in external and internal domains. The domain of physical movement of individuals in space is taken as a starting point, followed by allocation of attention to external cues and a closing discussion on search in memory.

What Do We Mean by Search?

Search can be defined as an attempt to arrive at a goal at an unknown location in the physical environment, time, memory, or any other space. Finding a resource typically involves at least two components: an *exploration phase* that investigates possible locations as to where the resource might be located and an *exploitation phase* that involves resource acquisition. Often, the exploration and exploitation phases are not mutually exclusive, as animals may sample and exploit during exploration and continue exploring while exploiting.

Because exploration typically takes time away from exploitation, modulation between the two can be represented as an optimal control problem in which organisms attempt to minimize the time spent exploring for resources but still acquire sufficient information to maximize resource exploitation. When the search task involves a distinct individual target, the optimization problem is to choose the movement strategy that would minimize the time needed to find that target. Typically, however, biologically important resources show large variation in quality, and they vary over time and space. Thus an adaptive search usually involves a fitness-maximizing decision about the optimal balance between exploration and exploitation. More exploration can lead to finding better resources but to less time available for exploiting those resources. This trade-off between exploration and exploitation is common to both external and internal search problems.

External Search: Movement

The Structure of the External Resource Environment

All organisms encounter variation in the quantity and quality of resources. In terrestrial systems, physical factors (including the topography, soil types, winds, solar radiation, and precipitation) shape the spatial structure of temperature and availability of minerals and water. These, in turn, generate a variable spatial distribution of plant species and of the organisms associated with them. Such distribution may be either continuous or broken; the latter implies that distinct patches vary in the quality and quantity of a given resource, each surrounded by regions lacking that resource. Further diurnal and seasonal variation in abiotic factors adds temporal variation in organismal activity and productivity. This combination of spatial and temporal variation in essential abiotic and biotic resources means that an individual's exact location in time and space can dramatically affect its fitness. Hence individuals can be modeled as attempting to optimize their spatial position over time.

Search Strategies in External Space

Confined to the question of physical movement, the central issue concerning search in space is whether or not an organism should stay where it is or move elsewhere. Organisms should make this decision in response to heterogeneity in the density of resources in the surrounding environment—an area of study that has been extensively examined (Stephens et al. 2007; Stephens and Krebs 1986). One approach for examining adaptive “nonrandom” foraging behavior involves testing for area-restricted search, which refers to an individual's ability to restrict search to the local area where it has recently found resources before transitioning to more wide-ranging, global exploration (Kareiva and Odell 1987). Area-restricted search is related to patch-based models of foraging, like

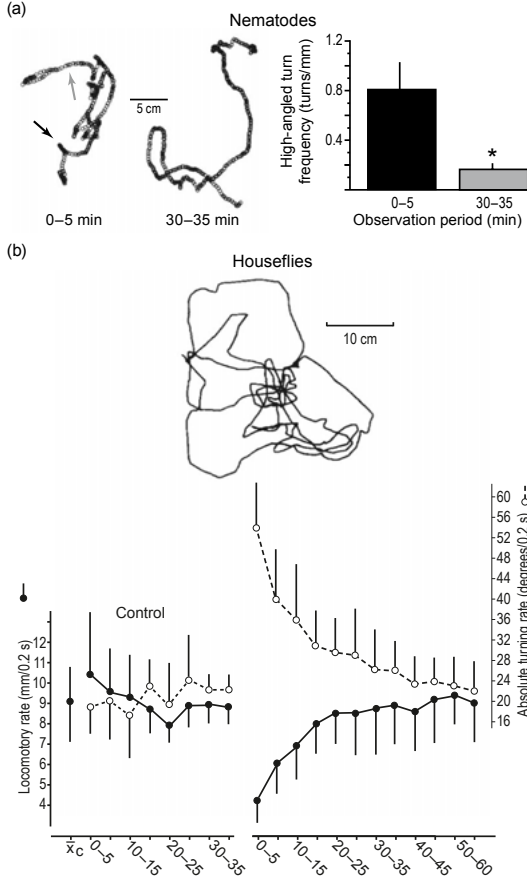


Figure 2.1 Evidence of resource-contingent foraging in (a) nematodes (*Caenorhabditis elegans*), (b) houseflies (*Musca domestica*), (c) bumblebees (*Bombus bimaculatus*), and (d) humans (*Homo sapiens*). (a) The left panel shows the foraging paths for *C. elegans* 0–5 min after encountering food and 30–35 min later. The black arrow indicates a high-angled turn; the gray arrow denotes a region of the path with no turning. The right panel shows that high-angled turns are significantly more likely to occur for the interval more recently associated with food (Hills et al. 2004). (b) The top panel shows a 69 s path for *M. domestica* immediately after it encounters food (at the central dot). The lower panel shows the quantitative comparison of turning angle (open circles) and locomotory rate (closed circles) for control flies (on the left) and flies immediately after encountering food (on right) (redrawn from White et al. 1984). (c) The top panel shows a significantly decreasing flight distance to the next flower following sequences of one, two, or three rewarding flowers for *B. bimaculatus*. The lower panel shows a significantly increasing flight distance after a series of one, two, or three nonrewarding flowers (data from Dukas and Real 1993). (d) The top panel shows typical paths for humans foraging in a three-dimensional environment with invisible resources arrayed in distributed or clustered arrangements. The lower panel illustrates that humans show significantly increased turning after encounters with resources in clustered environments than in distributed environments (Kalff et al. 2010).

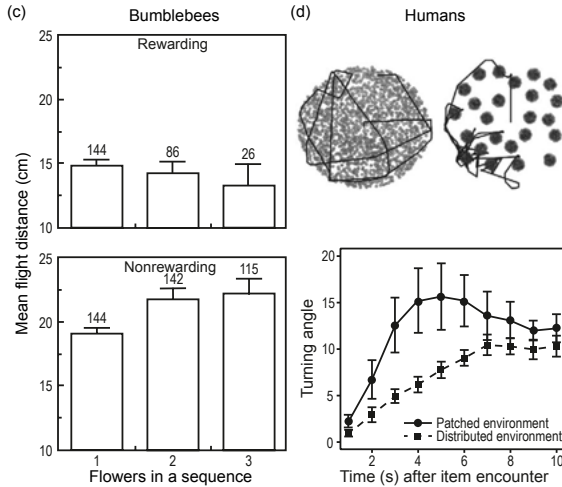


Figure 2.1 (cont'd)

the marginal value theorem (Charnov 1976), but is often employed when patch boundaries are difficult to detect or are otherwise “fuzzy” (Benhamou 1992; Adler and Kotar 1999).

One of the most primitive forms of search transitions between local and global foraging is the run-and-tumble behavior of bacteria such as *Escheria coli*. *E. coli* exhibit a change in behavior upon detecting increasing or decreasing food concentration gradients (Eisenbach and Lengeler 2004; Koshland 1980). When *E. coli* encounter increasing resources as they move, they engage in directed “runs” of swimming behavior using their flagellar motor. When they experience decreasing resources, the direction of the flagellar motor changes and this causes the bacteria to tumble randomly before engaging in another directed swim. This behavior appears to serve as a method for moving toward high concentration gradients and away from low concentration gradients. Thus, bacteria show evidence of area-restricted search by attempting to stay in areas with higher resource density, but move away from areas with lower resource density.

Figure 2.1 illustrates patterns of area-restricted search observed for several classes of animal species: nematodes (*Caenorhabditis elegans*), houseflies (*Musca domestica*), bumblebees (*Bombus bimaculatus*), and humans (*Homo sapiens*). In each case, the central result is that the animal responds to low resource densities by traveling away from them and to high resource densities by staying near them. In nematodes (*C. elegans*), individuals engage in high-angled turns (or pirouettes) following recent encounters with resources, but reduce their number of pirouettes as the time since the last encounter increases (Hills et al. 2004). Similar patterns of increased turning in response to resource encounters have been observed in flies (White et al. 1984), bumblebees (Dukas

and Real 1993), and humans (Kalff et al. 2010). This pattern of density-contingent foraging in space is ubiquitous across metazoans (Bell 1990; Hills 2006).

Neural Mechanisms of Search in External Space

What are the neural modulators of spatial search? Despite the abundance of evidence that animals can respond to changing resource densities in space, the neural mechanisms that control this ability are not well understood. Here we focus primarily on dopamine, because other neuromodulators (e.g., norepinephrine and serotonin) are less well understood from a comparative perspective, though they are potentially critical to search and other reward-seeking behaviors (Barron et al. 2010; Cools, this volume).

In nematodes (*C. elegans*), the modulation between local area-restricted perseveration and wider-ranging exploration is governed, at least in part, by a relationship between presynaptic dopaminergic neurons modulating downstream glutamatergic locomotory interneurons. Higher levels of dopamine increase turning angles, whereas lower levels reduce turning angles. Selectively killing dopaminergic neurons or applying a dopaminergic antagonist (raclopride) removes the capacity for area-restricted search (Hills et al. 2004). Dopaminergic mechanisms also facilitate the increased turning that fruit flies (*Drosophila melanogaster*) show under the influence of cocaine (Bainton et al. 2000), and this has even been found to extend to associative learning for places in the flatworm, *Dugesia japonica* (Kusayama and Watanabe 2000). In rats (*Rattus norvegicus*), turning increases in response to agonists for dopaminergic receptors (Robertson and Robertson 1986), and modulation between explorative and exploitative behaviors is mediated by midbrain dopaminergic neurons (Fink and Smith 1980). In random foraging experiments, injection of a specific antagonist for the dopaminergic receptor subtype D1 into the nucleus accumbens of rats significantly impaired performance, measured by an increase in wrong entries into maze arms (Floresco and Phillips 1999).

Perseveration in response to resources is known to involve a significant dopaminergic component across animal phyla (for a recent review, see Barron et al. 2010). In part, this may be due to the relationship between dopaminergic processing and reward sensitivity. Numerous observations of dopaminergic activity in response to rewards as well as novel and aversive stimuli have been made and given rise to terms like “reward detector” and “novelty detector” (Salamone et al. 1997). Critically, dopaminergic neurons adjust their firing rates in response to unpredicted stimuli that are associated with fitness, such as appetitive and aversive stimuli (Salamone et al. 1997). Dopaminergic neurons are also involved in learning to predict outcomes associated with conditioned stimuli (Ljungberg et al. 1992; Kusayama and Watanabe 2000). In vertebrates, the dopaminergic neurons most often associated with goal-directed behaviors are located in the thalamus, striatum, and frontal cortex. These appear to

work together to control goal-directed movement in physical space and the focus of attention.

Attentional Search for External Information

The Structure of the External Information Space

As noted above, animals encounter nonrandom distributions of abiotic and biotic resources as they move through their physical environment. This means that the cues indicating the availability and quality of relevant resources (including food, predation, potential mates, and competitors) also show nonrandom distribution in time and space. Hence individuals can rely on the spatial and temporal structure of certain information for locating resources.

It is obvious that, in many species, search involves movement in physical space but the issue of search within the external information space is less apparent. Intuitively, one might argue that individuals should just process all incoming relevant information. It is indeed possible for some organisms with very limited perceptual ability to adopt such an inclusive strategy. In animals with extensive perceptual ability, it is clearly optimal to tune out all irrelevant information. Often, however, even the flow of relevant information exceeds the information processing rate of both the sensory organs and the brain (Dukas 2002, 2009). In humans, for example, only the fovea, which occupies about 0.01% of the retina and 1.7° of the visual field, transmits high-quality visual information. In primates, in general, the optic nerve transmits only approximately 2% of the information captured by the retina, and only about 1% of that information is processed by the visual cortex (Van Essen and Anderson 1995). In short, an individual's sensory organs can capture only a small proportion of the incoming information flow, and the rate of information capture by the sensory organs far exceeds the brain's rate of information processing. This necessitates a strategy for allocating attention to the most relevant cues in the information space at any given time.

Search Strategies for External Information

External information can be envisioned as a multidimensional space generated by the information flow from all sense organs. At any given time, animals must choose what information to attend to. This is analogous to the location choices that animals make in their physical space (discussed above). In the information space, animals should attend to the portion of information flow that would have the greatest effect on fitness (Dukas and Ellner 1993). For example, when human subjects were more likely to find targets at certain angles of the visual field, they devoted more attention to and had higher detection rates at these angles than subjects searching for randomly distributed targets (Shaw and Shaw

1977). Similarly, human subjects tend to focus their visual attention in the vicinity of a recently detected target but switch their attention to other spatial locations if no target is found at this area within a short giving-up time. This behavior, which is reminiscent of area-restricted search, is called *inhibition of return* (Klein 2000; Posner and Cohen 1984). In general, animals foraging in natural settings should focus their attention on the sensory cues associated with the most profitable food and most likely danger (Dukas 2002). Whereas much of the research on attention has been done in the visual domain, auditory and olfactory studies have revealed similar patterns of animals focusing on the most relevant cues at any given time (Skals et al. 2005; Fritz et al. 2007; Cross and Jackson 2010).

Animals searching for resources in the physical environment must often choose the search rate (distance moved per unit time) that would maximize their rate of finding resources (Dukas 2002; Gendron and Staddon 1983). Similarly, animals have to choose their range of information processing, which should be negatively related to the difficulty of processing certain information (Dukas and Ellner 1993). That is, animals can distribute attention broadly (e.g., devote little attention per unit area) when handling easy information but must adopt a narrow focus of attention when handling difficult information. Consider, for example, blue jays (*Cyanocitta cristata*) that were trained to search for two prey types: a caterpillar, which could appear in the center of the visual field at a probability of 0.5, and a moth, which could appear in either right or left peripheries of the visual field at a probability of 0.25 per side. Jays were three times more likely to detect the peripheral moth targets when the central caterpillar was conspicuous (i.e., easy to detect) than when it was cryptic and hence difficult to detect. This result is consistent with the prediction that the jays would process information from the whole visual field when the primary task is easy, but would narrow down their focus of attention to the center field when the primary task is difficult (Dukas and Kamil 2000). Jays modulated their focus of attention, reducing the area from which they processed information when the task became more difficult (see also Wolfe, this volume).

Neural Mechanisms Controlling Attention to External Information

Exactly as dopamine is a key neuromodulator of search in physical space, it plays an important role in search within the external information space. In general, dopamine is involved in subjects' ability to focus and sustain attention on relevant cues. For example, mice (*Mus musculus*) that were genetically manipulated to eliminate selectively phasic firing of dopaminergic neurons showed selective impairment in using relevant cues for learning. This suggests that the phasic firing of dopaminergic neurons modulates selective attention to relevant information (Caron and Wightman 2009; Zweifel et al. 2009). In humans, subjects with a subtype of the dopamine transporter gene associated with higher dopamine levels in the striatum (a region of the brain associated with attention)

show a different pattern of inhibition of return than control subjects (Colzato et al. 2010a). This suggests involvement of dopamine in the spatial allocation of attention over time.

Dopamine deficit is currently the leading theory for explaining attention deficit hyperactivity disorder (ADHD), a mental disorder characterized by a reduced ability to focus and sustain attention and by an excessive level of activity. Brain imaging studies indicate smaller sizes and lesser activation of brain regions related to dopamine in individuals with ADHD. Allelic variation in two genes, the dopamine receptor D4 and the dopamine transporter, has been linked to ADHD, and the principal drug for treating ADHD, methylphenidate (Ritalin®), increases synaptically released dopamine (Iversen and Iversen 2007; Swanson et al. 2007). Together, these examples provide strong evidence that dopamine modulates the focus of attention to external information similarly to the way it modulates perseverative local foraging in external space.

Internal Information Search

Having focused on search via physical movement in the environment as well as through selective tuning to external information, we now explore search for information in memory or for solutions to problems that require internal manipulation of information.

The Structure of Internal Information

As demonstrated above, external stimuli often present themselves in a nonrandom, spatially autocorrelated fashion—with rewards associated with a specific location likely to signal rewards close to that location in the near future. Does the structure of relationships between items in memory also implicate an autocorrelated structure, and do we see evidence of this structure in recall from memory?

Studies of written language—presumably reflecting the internal structure of cognitive information—find evidence for a strongly clustered environment. With nodes representing words and links representing relations between words, these language networks often reveal a small-world structure, indicating that words are much more likely to appear together in small clusters of related items than one would expect by chance (Cancho and Solé 2001). A similar small-world structure has also been identified in internal search when people are asked to say the first word that comes to mind after hearing another word (i.e., free association) (Steyvers and Tenenbaum 2005). Moreover, this structure of language and free association networks is well correlated with the order in which children learn about language (Hills et al. 2010a). This indicates that the patchy internal structure of memory may be tightly linked with the patchy external structure of information.

Search Strategies for Internal Information

Research on free recall from natural categories and list learning consistently finds that groups of semantically similar words are produced together (Bousfield 1953; Romney et al. 1993). This clustering in output is often considered to be the result of a dynamic search process that modulates between local and global search policies. One of the most prominent and successful memory search models, the search of associative memory model, employs this dynamic local-to-global search policy (Raaijmakers and Shiffrin 1981). Local search is assumed to occur via item level similarity, with recently recalled items in memory activating other related items in memory. Global search activates items in relation to the overarching category and context such as according to their typicality or frequency of occurrence in that category. For example, in the animal fluency task—“say all the animals you can think of”—a person might search globally and produce “dog” and then search locally for similar items, like “wolf” and “fox,” before transitioning to a global search and producing “cow.” In the model, transitions from local to global search occur when local resources become depleted, such as when there is nothing similar to “fox” that has not already been produced. Interestingly, this model of memory search was developed in cognitive psychology independent of models in behavioral ecology, but it shares the signature behavioral pattern associated with area-restricted search in physical space: modulating between exploration and exploitation in response to recent experience with the resource environment.

Similar evidence for local perseveration due to memory activation has been found in experiments based on word priming. In these experiments, a person is first shown a word prime (e.g., BIRD) and then asked to determine whether a second shown word target is a true word or a nonword (e.g., ROBIN or ROLIN, respectively). Relative to an uninformative word prime, Neely (1977) demonstrated both facilitation (faster response times) and inhibition (slower response times) in people’s ability to determine the identity of the word target by manipulating whether the word target was expected or unexpected following the word prime. This elegantly demonstrates that expectations create local activation in memory following the presentation of a prime, and that this can both reduce the time it takes to recognize objects associated with those memories and also increase the time it takes to recognize objects that are not associated with those memories.

Research on sequential solutions in problem-solving tasks also demonstrates that people show local perseveration in internal search environments. For example, people tend to produce solutions that are more clustered together (i.e., similar) than one would expect by random generation; for example, in math search tasks (Hills 2010) and anagram search tasks (Hills et al. 2010b). In one case, Hills et al. (2010b) had participants search within scrambled sets of letters for multiple words. Participants would see a letter set, like BLNTAO, and they could find “BOAT,” “BOLT,” etc. An analysis of the string similarity

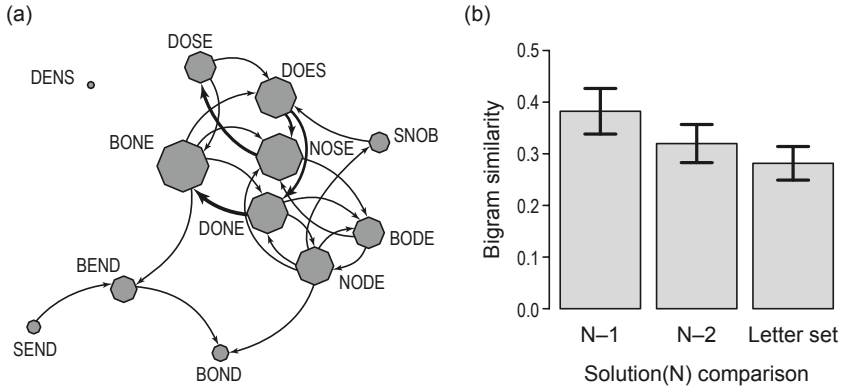


Figure 2.2 Behavior in an anagram search task. (a) Visual depiction of the between-word transitions produced by all participants in the letter set NSBDOE. Participants looked for words they could make from letters in the letter set (using four or more letters). Nodes represent solutions and links between nodes represent transitions between words, with the arrow showing which word came second. Node size is proportional to the number of participants who provided that solution for this letter set. Link thickness is proportional to the number of participants who made that transition. For visual clarity, only transitions that took place more than twice are represented with a link. (b) The bigram similarity of the present word solution to previous ($N-1$) and two-back ($N-2$) solutions and to the original letter set, showing that solutions tended to have the highest string similarity to solutions produced nearby. Error bars are standard error of the mean. Reprinted with permission from Hills et al. (2010b).

(e.g., bigram similarity comparing the number of shared letter pairs: “BO,” “OA,” etc.) between subsequent solutions determined that participants tended to produce solutions that were most similar to their last solution. This was true even though previous solutions were not visible. Results indicate that participants were searching locally around previous solutions, before transitioning to a global search strategy (Figure 2.2).

Neural Mechanisms in Internal Information Search

Several studies have found that the trajectories taken through long-term memory are related to working memory span (Rosen and Engle 1997), which is well known to be tightly connected with dopaminergic processing (Cools and D’Esposito 2009). Rosen and Engle (1997) found that participants with higher working memory spans tend to produce longer sequences of clustered items in a category fluency task than individuals with lower working memory spans. Hills and Pachur (2012) used a social fluency task (“say all the people that you know”) and had participants reconstruct the social network over which they were searching. Using semantic memory models, they found that participants with lower working memory spans transitioned more frequently between global and local cues in memory than individuals with higher working memory

spans. This passage is similar to the transition between exploratory and exploitative behavior described above for spatial and attentional foraging.

Cools and D'Esposito (2009) suggest that a proper balance between prefrontal and striatal dopamine levels is the key modulator of cognitive stability and cognitive flexibility and that this proper balance is also related to working memory. This is similar to Kane and Engle's (2002) interpretation that the cognitive control of attention (i.e., the ability to focus on one subgoal to the exclusion of other, distracting stimuli) is the underlying factor that determines working memory span. Furthermore, they suggest that this ability is mediated by prefrontal cortex modulation of activity in other areas of the brain. In other words, individuals with higher working memory spans are better at exploiting local information in internal search, whereas individuals with lower working memory spans tend to leave patches of local information more readily.

Prospects

The data we have presented above indicate three central points about external and internal search:

1. The environments in which organisms search both externally and internally share similar structural properties, and resources tend to be patchily distributed.
2. Various search strategies often rely on this patchiness to focus search around areas where resources have been recently found, and thus to facilitate resource acquisition based on their nonrandom distribution.
3. The neural mechanisms that control search—especially those involving dopamine, the prefrontal cortex, and the striatum—are often shared across species and search environments.

Although the data help us integrate information about the structure, strategies, and mechanisms of search in external and internal environments, we still lack substantial knowledge about the cognitive ecology of search. Below we highlight key issues that require further research.

Physical Search as an Evolutionary Precursor of Cognitive Search

Might the similarity between external physical search and internal information search indicate an origin for goal-directed cognition (i.e., cognitive control) from an evolutionary precursor devoted to spatial foraging and feeding related behaviors? Across metazoans (i.e., vertebrates and invertebrates), we find similar mechanisms modulating physical search for resources (Barron et al. 2010). As outlined above, in vertebrates (especially mammals) we find roughly the same mechanisms modulating search for information. This suggests a potential evolutionary homology between search in physical space and cognitive

search, with the derived form broadening the domains of search to information (Hills 2006). What other evidence would provide support for or against this hypothesis?

The comparative evolutionary approach to search raises several other questions. Are different forms of cognitive search domain-specific or domain-general? Recent research demonstrated priming in humans from external to internal search (Hills et al. 2008), based on empirical data indicating that prior experience in spatial foraging influenced a subsequent search in an “internal” problem-solving task. In this experiment, participants who first searched in a visuospatial task for clustered or diffuse resources subsequently searched for word solutions in anagrams as if those solutions were also more or less clustered, respectively. This may indicate a domain-general search process, consistent with our understanding of executive processing in cognition as a method for navigating hierarchical subgoals (Hills et al. 2010b). Which other forms of search are guided by such a domain-general process, or by different domain-specific processes (e.g., mate search)?

Do flexible cognitive capacities rely on balancing neuromodulation, similar to the cognitive search trade-off between exploration and exploitation outlined above? Many pathologies of goal-directed behavior (e.g., ADHD, Parkinson’s, stereotypies in autism, drug addiction) involve dopamine in a way that would be predicted from the neural control of animal foraging behavior, with more (or less) synaptic dopamine leading to higher (or lower) levels of perseveration and attentional focus (Hills 2006). Cools and Robbins (2004) argue that there is a balance between too-high and too-low dopamine levels and that this generates the “optimal state of mind”; patterns of behavior associated with too much or too little dopamine are consistently inflexible, often being too compulsive or impulsive for the demands of the environment. Flexibility is potentially one of the guiding selective forces in the evolution of the brain, as relatively larger brains appear to confer greater flexibility—an observation called the *cognitive buffer hypothesis* (Sol 2009). Can we better operationalize what flexibility means, in terms of searching for information? What might be the various evolutionary origins of this flexibility?

What Are the Biological Mechanisms of Cognitive Search?

In our analysis of the neural mechanisms underlying search, we focused on the common denominator of neuromodulation by dopamine, which, in vertebrates, is localized principally in the prefrontal cortex and striatum. Whereas this shared characteristic of neuromodulation by dopamine is intriguing and deserves further exploration, a fuller examination must also include more specific details about other brain regions, neuromodulators, and patterns of neuronal firings involved in search within each of the distinct spaces discussed here. Are there additional common mechanisms at this deeper level of analysis? Can existing knowledge about biological mechanisms of search within one domain,

such as selective attention in external space (Knudsen 2007; Salamone et al. 1997), help us understand mechanisms of search in another area (e.g., retrieval from an internal information space) (see also Winstanley et al., this volume)?

The Organization of Internal Information

Thus far we have focused on similarities across search environments and search mechanisms; however, important differences do exist. Perhaps the most significant distinction between external and internal search environments is that searchers typically cannot control the distribution of targets in the external environment but may affect the way they store their own information. That is, natural selection may have shaped the architecture of internally stored information to maximize some utility, such as the speed of recall or the numbers of items recalled. Existing models and data on search in external space may be able to help us understand the selective pressures and constraints operating on the structure of internal search environments.

How Are Algorithms for Search Shared across Domains?

What other dimensions can be used to characterize search? Part of the power of search as a paradigm is our ability to use search algorithms in one domain to inform research in other domains. In this discussion we highlighted the trade-off between exploitation and exploration, which is closely aligned with models of patch foraging. Similar search strategies borrowed from behavioral ecology have recently been applied to human information processing, for example, in terms of giving-up rules in problem solving (Wilke et al. 2009) and information-foraging strategies that capitalize on the structure of linked pages in the World Wide Web (Fu and Pirolli 2007). There are, however, other ways to implement search policies and many dimensions along which they may be defined. Given that some characterizations of search (e.g., exploitation versus exploration) lend themselves better to comparative analysis—both across organisms and algorithms—understanding how we define the dimensions of search and characterize different search policies may help us integrate our understanding of search and cognitive abilities more effectively.

Acknowledgments

We thank Lauren Dukas, Luc-Alain Giraldeau, David Redish, Peter Todd, and an anonymous referee for comments on the manuscript, and participants of the Ernst Strüngmann Forum for helpful suggestions. Our work has been supported by a grant from the Swiss National Research Foundation (100014 130397/1) (TH) and by the Natural Sciences and Engineering Research Council of Canada, Canada Foundation for Innovation and Ontario Innovation Trust (RD).