

Unpacking Cognitive Search

Mechanisms and Processes

Thorsten Pachur, Jeroen G. W. Raaijmakers,
Eddy J. Davelaar, Nathaniel D. Daw, Michael R. Dougherty,
Bernhard Hommel, Michael D. Lee, Sean M. Polyn,
K. Richard Ridderinkhof, Peter M. Todd, and Jeremy M. Wolfe

Abstract

This chapter discusses commonalities and differences in the cognitive mechanisms underlying different search tasks, such as spatial search, visual search, memory retrieval, action search, problem solving, and decision making. Three key issues relevant across all types of search are distinguished: (a) the initiation of search, (b) the maintenance and adaptive modification of the search process, and (c) the termination of search. As to search initiation, research is summarized concerning the effect of the number of cues on difficulty for executing search, and which factors structure the cue hierarchy. Discussion follows on how knowledge about metacognitive processes in memory might be used for better understanding the processes in maintenance of search, and heuristic principles for stopping search, possibly shared across different search tasks, are identified. Finally, consideration is given to how search processes might change as a function of experience and aging.

Introduction

In *The Disappearance of Lady Frances Carfax* (Doyle 1917), Sherlock Holmes is commissioned to track down a wealthy noblewomen who mysteriously vanished while traveling through Europe. To find Carfax, the detective and his ally, Dr. Watson, meticulously reconstruct the Lady's itinerary and visit the places where she had been seen prior to her disappearance. The investigation starts in Lausanne, where Watson is informed that Carfax has moved to Baden-Baden; there, he is sent to her long-term maid in Montpellier, who tells Watson that Carfax laid her off after making the acquaintance of a certain Dr. Shlessinger.

This information takes Holmes and Watson back to London, where they search for further clues as to the Lady's whereabouts. Although few of us engage in detective work regularly, Holmes' investigation resembles, in many respects, our more mundane search activities. In particular, it seems fair to say that most of our cognitive activities involve search of some kind, whether for a name to go with a face, a word to describe how we are feeling, an object hidden somewhere in the scene before us, or a solution to a problem encountered on the job.

But how do we search? As a starting point, consider how animals go about searching for resources in space (Bell 1991). In general, animals attempt to find as much resource in as short a time as possible. If there are cues to locations of resources that can be sensed from afar (e.g., seeing prey, chemically sensing conspecifics already at a resource), then these should govern the search; this is similar to visual search being guided to areas of interest detected in peripheral vision. Otherwise, in cases of uncertain resource location, organisms should tend to search in a way that brings them to new locations without going over recently visited locations again (akin to sampling without replacement). A random search (e.g., Brownian motion) does not accomplish this well (as evident, e.g., in the protest among iPod users against Apple's original random shuffle algorithm, which brought up recently played songs too frequently). Therefore, animals often use search strategies that move across the environment on a roughly straight course, or use a more systematic "space-filling" path (e.g., spiraling outward from a starting point).

Spatial foraging is but one example of a task that involves search. Search is also a key factor in memory retrieval, visual search, action search, problem solving, and decision making. In this chapter we discuss both the commonalities and differences between these different types of search. To structure our discussion, we distinguish three basic issues in search: (a) how search is initiated, (b) how search is maintained and adaptively modified, and (c) when and how search is terminated. In addition, we discuss individual differences in search that may arise due to developmental changes, due to prior experience in these or similar tasks, or due to preexisting (possibly genetic) differences in information processing.

How Is Search Initiated?

Cue Selection

The first step to get the search process going is to establish a set of features (or cues) that define the object of the search. To illustrate, consider a visual search task where you set out to look for a turquoise ring in the bedroom: What are the prerequisites for initiating the search? Search for the ring will not proceed randomly; rather, attention will be guided to items that share basic features with the target. These basic features (e.g., size, shape, color of the desired ring)

serve as a template for comparing items encountered as search proceeds, and the goal is to find an adequate match between this abstract representation of the target item and the visual image encountered during search. Interestingly, guiding features are not the same as perceptual features. Guidance seems to be based on a coarse and categorical representation of a set of basic features. To illustrate: although the ring may clearly be a particular shade of “turquoise,” your ability to use color for guidance is limited to directing your attention to items that are broadly, categorically “blue” (Daoutis et al. 2006). Moreover, as rings occur in some places (e.g., on dressers) and not in others (e.g., mid-air), search will be guided by scene-based properties, or context features (Biederman 1972).

Similar principles hold for search in memory, where it is assumed that a template representation consisting of a set of “retrieval cues” is used to constrain the output of the memory system during search. As in visual search, the retrieval cues contain information that distinguishes the to-be-retrieved item from all of the other traces that may reside in the memory system. Examples of these constraining retrieval cues are semantic characteristics (e.g., animals) or the temporal circumstances of the item’s occurrence (e.g., recall the items that were on list A). There is also evidence that people can use multiple cues to constrain simultaneously what is being retrieved during memory search (e.g., recall all the animals that were on list A; Polyn et al. 2011).

In some search tasks, the set of features guiding search may arise directly from the task. In visual search tasks, for instance, the description of the target (e.g., turquoise ring) readily provides the features that will lead to the target (e.g., round object, has a hole). In other search tasks, however, the set of features has to be actively generated by the participant (below we discuss factors that can affect the construction of the feature set, such as the predicted effectiveness of the cues). In memory search, the set of target features might also be defined by the specific recall strategy used by the participant (e.g., search for items in alphabetical order). Whereas there is considerable research on memory and visual search, we still know relatively little about cue selection in nonvisual search, such as auditory (e.g., speaker identification) or tactile search. As discussed below, given that several principles in search are shared across different search types, it is likely that aspects of feature or cue selection described above also generalize to these types of search.

How Do Multiple Cues Affect Search?

Often, a target is defined by a set of multiple features, such as when one searches for an item that is both a ring *and* blue. How does the complexity of the target affect the difficulty of search? Interestingly, the answer could vary across different types of search. For instance, visual search based on conjunctive rules (e.g., find the red X) seems to be more difficult than when based on a single cue (e.g., find the X). Similarly, it is assumed that search prior to probability

judgments is more difficult for judgments of conditional probability (e.g., the probability of breast cancer given a positive mammogram), which are based on multiple cues, than search with only one cue (because cues are processed sequentially; e.g., Dougherty et al. 1999). Memory models such as search of associative memory or SAM (Raaijmakers and Shiffrin 1981), by contrast, do not necessarily assume that combining multiple cues complicates search.

Which Factors Guide the Selection of Cues?

The selection of retrieval cues used in standard memory-retrieval paradigms is relatively well understood, at least compared to cue-selection processes in real-world tasks (e.g., medical diagnosis). For example, in laboratory tasks, each item (or in some cases an entire list) is generally associated with a single and unique cue (e.g., in paired-associates learning; Calkins 1894). In contrast, in many real-world retrieval tasks, cues are shared across items or “lists” and are thus only probabilistically related to the target item. To illustrate, in medicine a retrieval cue such as “high white-blood cell count” is associated with several different pathologies, ranging from bacterial infection to disorders of bone marrow (e.g., leukemia). These pathologies are often organized hierarchically, such that there are many specific examples of the general class of bacterial infection and many specific examples of the general class of bone-marrow disorders. Within each class of pathologies, individual examples (which could be called “hypotheses”) might be associated with specific symptoms (“data”). Given a representation that can be expressed in terms of hypotheses and data, we can now ask the question: What is the probability of the data (a symptom) given a particular hypothesis (a disease)? The answer is the diagnosticity or validity of that symptom cue. Though not a deterministic cue, the presence of a high white-blood cell count may still be a diagnostic piece of information. In the context of memory retrieval tasks, one can imagine that the diagnosticity (or validity) of a particular memory retrieval cue can be exploited to help guide the retrieval of potential hypotheses from long-term memory (Thomas et al. 2008; Dougherty et al. 2010). For example, in a simplified environment, imagine that the diagnosticity of a particular cue (symptom) for discriminating between two mutually exclusive and exhaustive categories of diseases is 2:1. This would imply that twice as many hypotheses from disease category 1 are related to the observed symptom compared to disease category 2. Such cue diagnosticity provides valuable information that can and should be used in determining how to search through memory in diagnosis tasks; namely, search using the most diagnostic cues available. Put more generally, the statistical properties of a cue likely inform basic memory search processes. Unfortunately, there is little work on how statistical properties of the retrieval cues affect cue selection in memory search (see, however, Anderson 1991).

The idea that search for cues is guided by their usefulness is, by contrast, a common one in other realms of judgment and decision making. For instance,

the take-the-best heuristic (Gigerenzer and Goldstein 1996) assumes that cues are inspected in sequential order according to their validity (defined as the probability that the cue leads to a correct response given that it discriminates between two options). The large literature on multiple-cue probability learning has examined the processes by which people acquire knowledge about the validity of cues (e.g., Klayman 1988; for other approaches, see Dieckmann and Todd 2012). Alternatively, people may use cues in an order, based on how likely they are to lead to any decision (i.e., their “discrimination rate”; cf. Rakow et al. 2005), or a combination of validity and discrimination rate (i.e., their “success”; Martignon and Hoffrage 1999), on intuitive causal beliefs about the cues’ importance (Chapman and Chapman 1969).

Search Initiation in Action Selection

While both visual search and memory search usually have a clearly specified target, other types of search are more open-ended and, as a consequence, may be guided in a rather different fashion. For instance, consider exploring which out of many possible actions will yield desired outcomes (or will avoid undesirable outcomes). Animals and humans are often confronted with a variety of opportunities for action in a particular situation. Some of these options might be more alluring or potentiating than others; some might be more risky or more effort-consuming.

Action selection in these circumstances implicates search in at least two senses: external (i.e., traversing the environment and exploring the results of actions so as to learn action-outcome relationships, reward contingencies, or cognitive maps) and internal (i.e., the use of these learned representations to evaluate candidate courses of action to guide subsequent action selection toward those most likely to maximize reward). (Here reward may be determined by a cost-benefit analysis of potential outcomes vis-à-vis current motivational states.) The search for those actions that have maximal (subjective) expected utility (another way of talking about reward) is well captured by reinforcement learning models, which describe how regularities among action-outcome contingencies are extracted from experience. Broadly, this type of action search will be initiated, constrained, guided, and terminated by an agent’s current concerns, intentions, and prior experience, as well as by its present motivational state (e.g., fatigue, satiation). In these respects, search based on reinforcement learning may differ from visual and memory search, in which initiation, guidance, and termination are generally influenced more explicitly by instructions and cues.

Although it is traditionally assumed that search occurs over actions, computing their values by averaging over the possible reward outcomes to which they might lead, a recent alternative proposes that agents might first choose between outcomes, then search over action plans for how best to obtain the desired outcome (Padoa-Schioppa and Assad 2006; Krajbich et al. 2010). The

initial *goal choice* in this case may occur in a similar fashion as initiation of memory or visual search (e.g., based on feature templates).

Construction of the Search Space

In addition to coming up with a set of cues that define *what* to look for, some types of search also require the definition of the search space; that is, *where* it is possible to look. In visual search, it has been hypothesized that individuals need to construct the search space by “parsing” the scene into *proto-objects*, regions that will be selected by attention (Rensink 2000a). Moreover, in searches that are extended in time, there may be an initial plan for a search path, which is refined during search as a function of what is found (in the next section, we elaborate on such search maintenance processes). For instance, when asked, in a verbal fluency task, to retrieve all movies seen over the last six months, one might first search among the films seen in a particular movie theater and then move on to search among films of a particular genre, rather than probing memory for movies in general. In memory search, the refinement of search can thus consist of switching between different sets of retrieval cues.

A similar construction of a search space is relevant in problem solving and in multi-attribute choice, where a set of possible options needs to be generated from which a final choice can be made (Marewski et al. 2010; Tversky 1972). Sometimes, such a consideration set might be generated more or less automatically—and efficiently. In a study that examined action selection in sports, Johnson and Raab (2003) found that options which are quickly generated tend to be of higher quality than options generated more slowly (see also Dougherty et al. 1997; Gettys and Fisher 1979).

The construction of the search space can have a considerable effect on the efficiency of search. For instance, in a verbal fluency task, search becomes more difficult the larger the category from which objects are recalled (though the effect can depend on the retrieval strategy; Indow and Togano 1970; see also Murdock and Okada 1970).

Open Questions

While some of the principles guiding search initiation seem to be similar across different types of search, there are also some differences. What is currently unclear, however, is the extent to which the observed differences between various types of search may be due to the experimental paradigms used to study the different types of search. Natural environments may provide a much richer context than the rather artificial settings used in the laboratory. Consequently, navigation through search spaces in the real world may be much easier, due to the constraints imposed. Furthermore, the selection of cues and the construction of the search space are likely to arise from a dynamic interplay of divergent (i.e., global) and convergent (i.e., local) search strategies, possibly

applied sequentially; methods must be developed to investigate such dynamic processes, to which we turn to next.

How Is Search Maintained and Dynamically Modified?

After initiation, how is search maintained? The answer might depend strongly on whether a search targeted just one single thing (e.g., a nest or partner) or whether search is ongoing (e.g., for food). Whereas in the first case search is (often) stopped after the target has been found, in the latter case, search may continue after finding a target, seeking other targets.

Global versus Local Search Strategies

In many situations, search may be characterized as switching between exploration (or divergent search) and exploitation (or convergent search). The respective contributions of exploration and exploitation are influenced by the structure of the environment, in particular by whether the desired resource occurs in patches or not. If resources are patchy (i.e., distributed in clumps with relatively empty regions between them), then finding one resource indicates that others may be nearby. Here, the organism can benefit from switching from exploration between patches to exploitation of the discovered patch. Because the resources within a patch are themselves often not immediately detectable, and thus also require search (e.g., a berry bush is a patch in which berries must be sought by looking underneath leaves), this switching can also be thought of as going from global to local search. Local within-patch search can be implemented by taking smaller steps or making smaller movements to stay within the patch, turning more to stay in the same vicinity, and turning back if the edge of the patch is detected (Bell 1991).

A popular way to study the dynamic interplay between exploration and exploitation is with so-called *bandit problems*, in which there are M choices you can make for a sequence of N trials (e.g., Gittins 1979; Kaelbling et al. 1996). Each choice has some fixed, but unknown, rate of providing a binary reward. The goal is to maximize the total number of rewards obtained, and the search problem is then to explore the M choices sufficiently to determine which one(s) to exploit on further trials (for experimental studies, see Daw et al. 2006; Lee et al. 2011).

When to Leave the Patch?

When resources are distributed in patches and one is currently being exploited, another problem arises: as the resources in the patch are increasingly being depleted, the benefit of staying in the patch decreases, and at some point the organism must decide to leave that patch and return to exploring for other

patches. The general principle guiding animals in this situation is captured in optimal foraging theory by the *marginal value theorem* (Charnov 1976). Accordingly, the highest rate of return of resources can be achieved if the patch is left as soon as the rate of finding things in that patch falls below the expected mean rate of finding things across the environment as a whole when the optimal strategy is followed.

Although a useful benchmark, the marginal value theorem makes strong and often unrealistic assumptions about the organism's knowledge and computational capacities with respect to determining instantaneous and expected rates of return. Consequently, a variety of heuristics for patch-leaving decisions have been proposed that are based on simple, easy-to-compute cues: how much time one has already spent in the patch, how many items one has found in the patch, how long it has been since the previous item was found in the patch, and how long it took to get to this patch in the first place. The effectiveness of specific patch-leaving heuristics depends in part on how resources are distributed across patches. For instance, if resources are aggregated such that there are some very good patches along with many middling ones, then it is appropriate to leave the patch after some giving-up-time has passed since the previous item was found. Humans seem to use such a rule in some spatial and memory search tasks (Wilke et al. 2009; Hutchinson et al. 2008; see also Payne et al. 2007).

Another way to conceptualize the dynamic transition between global and local search is area-restricted search, in which an organism performs more high-angle turns when resources are encountered and so stays in a local area, gradually returning to low-angle turns when resources are not encountered for some time. Area-restricted search can yield more continuous transitions back and forth between local and global search over time, and may be more appropriate where patch boundaries are fuzzy (for an overview, see Hills 2006). In addition to factors such as the current and expected rate of return and the time spent in a patch, it has also been shown that animals sometimes take into account the variability of the patches they seek. For instance, when they must reach a threshold amount of food to survive the night, they might prefer a patch with greater variability but lower overall mean return rate over a less variable but higher mean one if the former, but not the latter, has a chance of providing enough food for survival (as described in risk-sensitive foraging theory; for an overview, see McNamara and Houston 1992).

How Is Search Monitored?

Memory Search

Managing the search process effectively requires keeping track of the contents of the current as well as past search space. This monitoring probably relies heavily on what are termed metacognitive processes; that is, processes which

keep track of how aspects of cognition are proceeding (Koriat et al. 2000; Metcalfe and Shimamura 1994; Nelson 1996). Although we still know relatively little about the exact nature of the metacognitive processes involved in search, it is possible that they are similar to those used in other domains, such as (a) monitoring with respect to the contents of memory and (b) monitoring with respect to the acquisition or learnability of study material. Several paradigms have been developed to investigate these processes (Nelson 1996; Nelson et al. 2004). Monitoring of the contents of memory is often studied by asking people to assess their confidence in the accuracy of a retrieved piece of information, or by asking them to assess the likelihood that they will be able to retrieve a particular piece of information in the future, given that it cannot be retrieved at the present time (*feeling-of-knowing judgments*). The monitoring of the acquisition of information is studied by asking people to assess how well learned a piece of material is (*judgments-of-learning*). People's assessments in these tasks are often rather accurate and have been shown to predict future recallability. A second method to study monitoring of information acquisition is by asking people to assess how easily they will be able to learn a newly experienced piece of information (*ease-of-learning*); for instance, when estimating how much study time to allocate to studying for a test based on the difficulty of the material. Ease-of-learning judgments may play out in cognitive search by influencing how long one spends in an exploration mode, assuming that one goal of exploration is to discover or learn environmental or statistical contingencies (see also Metcalfe and Jacobs 2010; for a review of different types of metacognitive judgments, see Nelson 1996).

Although some work has examined how metacognitive monitoring limits or informs search behavior in single-item recall tasks (e.g., Dougherty et al. 2005; Nelson et al. 1986), there is no work on more complex search tasks. For instance, in verbal fluency tasks (e.g., name all animals you can think of) it is necessary to monitor how much of a semantic space has already been exploited, or to estimate the size of the remaining unused portion of the "information patch." Another important gap in the understanding of metacognitive monitoring of memory is how it might relate to error monitoring and detection, as carried out by functions localized to the prefrontal cortex. Shimamura (2008) proposed a neurocognitive model of metacognition that postulates a fundamental role of cognitive control for regulating and monitoring metacognitive representations.

Decision Making and Problem Solving

Metacognitive monitoring processes in other search tasks are even less well studied. For example, in decision-making tasks, which are often assumed to be based on sequential search of cues (e.g., Payne et al. 1993; Gigerenzer et al. 1999), how does one monitor which cues have been previously accessed in the course of a decision? Does a physician use a similar process to monitor

symptoms that have already been checked or evaluated while generating a diagnosis? Finally, monitoring processes might also be important for searching and navigating the solution space in problem-solving tasks, where one has to monitor one's current location in the solution space and which locations have already been visited (cf. hill climbing; Newell and Simon 1972).

Open Questions

Although the categorical distinction between exploration and exploitation describes some search processes quite well, in other cases it may be more appropriate to use a continuous approach. In addition, it is still relatively unclear what mediates the switch between exploration and exploitation. One possibility is that switching is based on a form of conflict signaling, indicating that there is a mismatch between the encountered stimuli and the target (see Hommel, this volume). Specifically, mild conflict might lead to increased top-down control (exploitation), whereas stronger conflict might lead to stress and a change in the search strategy (exploration).

Another issue concerns the metaphors and analogies we use to conceptualize search. We often liken internal search in memory to external search in a spatially laid out environment. Might this spatial metaphor critically constrain the way we think about and understand search? Clearly, there are alternative conceptualizations, such as distributed, symbolic, or temporal representations, which might highlight different aspects of the search process rather than portray search in spatial terms (see Schooler et al., this volume). For instance, the importance of navigation costs may be less important if search occurs within a distributed representation.

How Are Search Processes Controlled?

As mentioned above, effective search often requires maintenance and control processes (e.g., to switch dynamically between exploration and exploitation). What are the proximate psychological capacities that are tapped by these control processes? A general assumption in influential models of cognitive control (inspired by the cybernetic approach; Wiener 1948) is that information is sampled and matched against a goal representation until a reasonable fit is achieved (e.g., Botvinick et al. 2001; Miller et al. 1960). Top-down control over cognitive search might be achieved in a similar manner. In a visual search task, for example, this would suggest that a representation of the target stored in working memory is matched against stimuli encountered during search until the target is identified. In tasks requiring action search, it has been proposed that conflict—that is, when there is a mismatch between the target and the stimuli encountered—leads to an increase of top-down control (e.g., Botvinick et al. 2001). Control, however, can sometimes also be governed by local priming (i.e., arising from the stimuli) rather than managed in a top-down fashion.

For instance, in memory search, the search for the next item to be considered will be influenced by both top-down constraints (the target representation) and the similarity between the current and the previous items (and the priming the previous item generates).

Given that controlling the search process requires continuous updating of the information currently in the focus of attention, processes in working memory are likely to play a key role. Within cognitive psychology, the construct of working memory has been defined as the ability to maintain focus of attention on goal-relevant information in the face of distraction (Kane et al. 2001). Factor analysis and experimental work have revealed that working memory capacity (as measured by operation-related tasks) is correlated with performance in a number of laboratory and nonlaboratory tasks, including response inhibition tasks (anti-saccade, Stroop), auditory tasks (dichotic listening tasks), resistance to proactive interference (Brown-Peterson task), measures of general fluid abilities, note taking, and planning (Engle 2002).

Increasing evidence indicates that key characteristics of cognitive control during search are indeed correlated with working memory capacity. For instance, Hills et al. (2010b) have proposed that a higher working memory capacity is associated with a lower frequency of switching between patches. Currently it is unclear how exactly working memory capacity affects the switching behavior. For instance, working memory could affect the signal-to-noise ratio in information processing (i.e., the ability to discriminate between targets and distractors), which might help focusing on the current task. Alternatively, a higher working memory capacity could lead to better conflict resolution (Bäckman et al. 2010; Li et al. 2001), for instance, by facilitating the identification of the actual signal within the noise or by suppressing task-irrelevant information. Moreover, it is likely that not all subcomponents of working memory affect control processes equally during search (Friedman et al. 2008; Miyake et al. 2000). Thus, further investigation is needed to distinguish more precisely the relevant components.

In light of the current evidence for the influence of cognitive control on switching behavior, it might be useful to distinguish between switching which results from a strategic decision and switching that occurs due to unsystematic factors (i.e., distraction). On one hand, higher working memory is assumed to help individuals stay focused on searching within a patch (while a patch still yields successful outcomes), thus decreasing the switching rate, as shown by Hills and Pachur (2012) and Hills et al. (2011). On the other hand, to the extent that strategic patch switching (e.g., when disengaging from a patch and switching to exploration once the current patch has been depleted) involves task-switching costs, higher cognitive control might be associated with an increased switching rate. For instance, Mayr (2001) found that older adults (who are likely to have a reduced working memory capacity) display higher switching costs than younger adults in a task-switching paradigm. Given the potentially multiple roles of working memory in exploitation and exploration,

future research should delineate more clearly the factors that moderate the relationship between working memory and search.

Modes of Control versus Levels of Control

There is agreement that agents can change their strategy in search-related tasks, but how can we describe this change between strategies? One approach is to group the strategies in terms of binary dimensions (e.g., exploitation vs. exploration), so that changes between them can be seen as moving agents toward one or the other pole of each dimension. For instance, starting to search within a patch (spatial foraging) or a visual group (visual search) or item category (memory search) can be described as moving out of exploration toward exploitation. Alternatively, changing between strategies can be viewed as up- or down-movements in a goal hierarchy (Miller et al. 1960). Accordingly, the same strategic choice can be considered as moving down one level in a hierarchy of possible representations of search targets, from a more general level, which includes all available patches, groups, or categories as possible target locations, to a more specific level that restricts the search space to one patch, group, or category. The advantage of viewing the dynamics in search maintenance in terms of different levels of control is that it allows further levels that are more concrete (lower) or abstract (higher) to be added without giving up the general theoretical scheme.

How Is Search Stopped?

Earlier, in our discussion of the maintenance of search, we addressed the issue of how to decide when to modify the current search behavior (e.g., leaving a patch to move on to the next one). Similar principles apply to decide when to terminate the search process altogether. Although the decision to stop search is relevant for most search tasks, relatively little is known about the extent to which similar principles govern people's stopping behavior across these tasks. In any case, for search to be effective, stopping rules need to be sensitive to the characteristics of the task. In some tasks, for instance, it might be crucial to find at least one object (e.g., in food search or mate choice), whereas in others one can be more selective and stop search if a threshold is not met (e.g., in consumer product search or information foraging), irrespective of whether anything has been found at all.

Several empirical tasks have been developed to investigate the effectiveness of people's stopping rules in sequential choice. A prominent approach uses *optimal stopping problems*, for which optimal points to end search can, in principle, be determined. In one type of optimal stopping problem, known as the *secretary* or *dowry problem* (Ferguson 1989; Gilbert and Mosteller 1966), there is a sequence of N numbers distributed in some unknown way and independently

sampled. The searcher's goal is to choose the maximum number in the sequence, under the constraint that only the current number can be chosen at the time it is presented and that one cannot return to a previous number. In this version of the task, the searcher only learns the rank of the current number relative to all those previously seen. In other versions, the numbers themselves are seen, N may be unknown, the distribution may be known, and the utility function may differ (e.g., with continuous payoff rather than just success or failure; Bearden 2006; Smith et al. 2007). Investigations of people's performance in the secretary problem have been conducted, for instance, by Dudey and Todd (2001) and Lee (2006). Finally, in *deferred decision tasks*, searchers have to decide whether to continue information search (e.g., conduct another test) or to stop search and make a diagnosis about a situation (e.g., which of two diseases a patient has). Models of stopping rules to describe people's search behavior in such a task have been tested, for instance, by Busemeyer and Rapoport (1988; see also Browne et al. 2007).

Is There Evidence for Similar Stopping Rules across Different Types of Search?

As mentioned above, in many situations the determination of optimal stopping rules will exceed the cognitive capacities of an organism. In such situations, decisions to stop will need to be based on heuristic principles which can, under some circumstances, approximate the optimal solutions. Given that the need to decide when to stop search is relevant across many different tasks, we must ask whether similar heuristics for stopping search may be used across various domains. Although only very few studies have directly compared stopping behavior in different search tasks, the existing evidence hints at some commonalities. Comparing patch-leaving rules—akin to stopping rules at the patch level—in spatial and memory search, an interval-based rule (specifically, time since the last encountered item) accounted in both tasks for the data best (Hutchinson et al. 2008; Wilke et al. 2009). Similarly, findings suggest that people's decision to terminate retrieval from memory is a function of the number of retrieval failures, which is usually highly correlated with the temporal interval since the last retrieval (Harbison et al. 2008).

Process Tests of Stopping Rules in Decision Making

How can people's stopping behavior be studied and measured? Whereas search in memory is usually not directly observable, decision-making paradigms have been developed that enable tracking of external information search (for an overview, see Schulte-Mecklenbeck et al. 2011). In multi-attribute decision making, where people have to search for attributes to evaluate the alternatives, process tracing methodologies such as Mouselab (Payne et al. 1993) or eye-tracking have been used to test how people stop search. For instance, according

to the take-the-best heuristic (Gigerenzer and Goldstein 1996), alternatives are compared by sequentially inspecting their attributes (in the order of their validity or importance for this decision) and stopping that search and inspection as soon as the alternatives differ on an attribute. Thus, to infer which of two cities has more inhabitants, take-the-best starts by comparing the cities on the highest validity attribute (e.g., whether it is a state capital): if both cities have the same value on that attribute (e.g., if neither is a state capital) then the second most valid attribute is inspected (e.g., whether the city has an international airport). If this attribute discriminates (e.g., if only one of the two cities has an international airport), search is stopped and no further attribute is inspected. Using the Mouselab experimental tool, several studies have shown that people's stopping behavior indeed follows such a simple rule when information costs are high (Bröder 2003), there is time pressure (Rieskamp and Hoffrage 2008), cognitive resources are limited (Mata et al. 2007), or the number of alternatives is high (Ford et al. 1989). More recently, Khader et al. (2011) developed a neuroimaging paradigm that allows tracking the neural correlates of information search in memory-based decision making. The authors obtained evidence that people using take-the-best show reduced retrieval activity in the brain areas representing attribute knowledge when the heuristic stops search early as compared to when the heuristic searches more extensively. Pachur and Scheibehenne (2012) used a sequential sampling paradigm to show that when pricing a lottery, people stopped information search about the lotteries differently depending on whether they were asked for a maximum buying price or a minimum selling price.

Open Questions

Most search situations that are investigated in empirical studies are relatively artificial. It is not clear whether tasks studied in the laboratory make search more or less difficult compared to more natural search situations. On one hand, experimental search contexts usually do not offer as much information to help navigate the search process as more natural search environments. On the other, search environments outside the laboratory are also more complex, for instance because the target object is less well defined (e.g., find an appropriate partner to start a family), or because the search process is more difficult to control. Consequently, researchers need to study search and stopping rules also in real-world domains.

Individual Differences in Cognitive Search

The efficiency of an individual's adaptive control, in general, and of searching for objects, memory traces, and problem solutions, in particular, is known to vary with intelligence, operation span, and age. For instance, the development

of people's performance in visual search tasks across the life span shows an inverted U-shaped trajectory (Hommel et al. 2004). Older adults have particular problems with excluding irrelevant distractors; they seem to recheck more often when a target is absent. Interestingly, the opposite tendency is observed in decision making: older adults seem to search for less information than younger adults when making a decision (Mata and Nunes 2010). Nevertheless, older adults still show a general ability to adapt their information search to the structure of the environment. When a more extended search affords better decisions, older adults acquire more information than when extended search pays only little (Mata et al. 2007). This suggests that elderly people actively employ context-specific search control strategies, presumably to compensate for (real or assumed) effects of age-related cognitive decline.

Although investigations of search behavior in decision making have found evidence for the use of a considerable variety of strategies, variation in strategy use seems to be due primarily to external factors, such as time pressure and the statistical structure of the task, with individual differences playing only a minor role (Bröder 2011). There are, however, some exceptions. In addition to the age differences described above, search strategies in decision making have been shown to differ reliably as a function of expertise. For instance, Garcia-Retamero and Dhimi (2009) found that crime experts (burglars, police officers) tend to follow a strategy with simple search and stopping rules (take-the-best) to judge the security of a property, whereas novices (students) tended to follow a strategy involving more extensive search (see also Shanteau 1992).

To the extent that search is associated with a person's willingness to take risks, there is some evidence for gender differences in search. For instance, in a task where extended search increased gains but also the risk of a large loss, young male participants were more willing to take risks and to search longer than female participants (Slovic 1966). In addition, individual differences in motivation or persistence may lead people to stay engaged versus disengage from search (Dougherty and Harbison 2007).

Associations between control processes and working memory (as measured by operation-span performance) suggest that individual differences in search might also be related to individual differences in working memory. The influence of working memory on search seems to be due, in particular, to operational capabilities (i.e., manipulating material stored in working memory) rather than to storage capacity (i.e., the number of items that can be stored). Standard measures of working memory, such as the operation-span task, reading span, listening span, and symmetry span, rely on a process-versus-maintenance distinction: participants are asked to maintain a growing list of to-be-remembered items simultaneously while engaging in a processing task. In the operation-span task, for example, participants are presented with a list of letters serially (one at a time), with a simple mathematics problem interleaved between each successive letter presentation. Performance on the operation-span task is given by the number of letters correctly retrieved across multiple sequences of

to-be-remembered items. Individuals high in operation span show advantages in rejecting irrelevant information in memory and search-like tasks (e.g., Vogel et al. 2005) and in the attentional blink (Colzato et al. 2007). This might be due to the fact that both operation span and age have an impact on a person's ability to distinguish between signal (i.e., targets) and noise (i.e., distractors). Indeed, older adults seem to spend more time sampling sensory evidence to achieve a reliable signal-to-noise ratio than young adults (e.g., by engaging in more rechecking operations to make sure that the signal actually belongs to the searched object; Hommel et al. 2004). However, it is currently unclear which search behavior fosters a high signal-to-noise ratio. On one hand, having a reliable signal-to-noise ratio might reduce cognitive conflict (because relevant items and distractors can be better distinguished) and therefore foster exploitation. On the other, a higher signal-to-noise ratio will also increase the sensitivity to detect conflict signals, which should foster exploration. Overall, individual differences in the adaptivity of switching between exploration and exploitation can be due to both perceptual abilities to detect (external or internal) signals to switch and the ability to perform the switch (cf. Mayr 2001).

Engle and colleagues have found that an individual's working memory span predicts aspects of their performance on longer-term memory search tasks, such as free recall. Unsworth and Engle (2007) suggest that these differences are related to the difficulty of individuals with low working memory span to use cues effectively to constrain memory search (resulting in more intrusions from prior lists and fewer correct retrievals). Recent work suggests that more strategic aspects of search show reliable individual differences: people who tend to organize recalled items with a temporal strategy reliably recall more items than individuals whose recalls are temporally disorganized (Sederberg et al. 2010).

There is also some evidence for individual differences in visual search resulting from cultural influences. Specifically, Nisbett and colleagues (e.g., Nisbett and Miyamoto 2005) found that Asians, who show a more collectivist orientation, are more sensitive to context information, and thus seem to have a more divergent search behavior, than individually oriented Westerners. Similar attentional biases have also been found as a function of religious orientation (Colzato et al. 2010b), suggesting that cultural practices might shape the way individuals configure their cognitive system for search operations.

Future Directions

In our discussion of the cognitive mechanisms underlying search, we distinguished three different aspects of search: the initiation of search, the maintenance of search, and the termination of search. We discussed commonalities and differences between different types of cognitive search tasks (e.g., visual, memory, spatial, action search), potential proximate mechanisms, as well as

individual differences in search. Although several experimental paradigms have been developed to investigate the cognitive processes during the different stages in search, little is known about how these processes are implemented and biologically mediated. There is evidence that some search behavior is linked to dopamine (e.g., for an overview, see Hills 2006; Hills, this volume), yet to what degree does this link hold across different domains (such as search in memory, visual search, spatial search, and search for actions)? Moreover, the mechanisms underlying an individual's adaptive use of different search strategies (e.g., to compensate for age-related decline) are not well understood, nor are the mechanisms that drive cultural influences on the control of cognitive search.

An important application of research on cognitive search may lie in developing methods to train individuals to change their search behavior (e.g., brain training), for instance, to be more flexible in switching (exploration or innovation) or more persistent in concentrating (exploitation or focus). Important questions here include how long the training effects last (long-term or short-term), and whether they transfer from one domain to another (for evidence, see, e.g., Karbach and Kray 2009; Hills et al. 2010b).

Sherlock Holmes' adventures offer some inspiration for hoping that people can be trained to adopt different search methods in some instances. Holmes often attempted to instruct Watson about his investigative methods in searching for clues and solutions to puzzling mysteries. Even though this was not always a success, in *The Disappearance of Lady Frances Carfax*, Holmes and Watson's search ends successfully as they manage to find the Lady, at the mercy of Shlessinger, just in time before he could bury her alive.

