

4

Searching for Fundamentals and Commonalities of Search

John M. C. Hutchinson, David W. Stephens,
Melissa Bateson, Iain Couzin, Reuven Dukas,
Luc-Alain Giraldeau, Thomas T. Hills,
Frederic Méry, and Bruce Winterhalder

Abstract

This chapter reports the discussion of a group of mostly behavioral biologists, who attempt to put research on search from their own discipline into a framework that might help identify parallels with cognitive search. Essential components of search are a functional goal, uncertainty about goal location, the adaptive varying of position, and often a stopping rule. The chapter considers a diversity of cases where search is in domains other than spatial and lists other important dimensions in which search problems differ. One dimension examined in detail is social interactions between searchers and searchers, targets and targets, and targets and searchers. The producer-scrounger game is presented as an example; despite the extensive empirical and theoretical work on the equilibrium between the strategies, it is largely an open problem how animals decide when to adopt each strategy, and thus how real equilibria are attained. Another dimension that explains some of the diversity of search behavior is the modality of the information utilized (e.g., visual, auditory, olfactory). The chapter concludes by highlighting further parallels between search in the external environment and cognitive search. These suggest some novel avenues of research.

Evolutionary Biology of Search

To begin, it may be useful to say something about the perspective we bring to the study of search. Our group is predominantly whole-organism biologists who investigate the mechanisms and adaptive significance of behavior. In doing this, behavioral ecologists commonly appeal to optimality or game-theoretical models, and these models, along with knowledge about animal genetics,

physiology, neurobiology, phylogeny and development, have guided our thinking about search.

For example, a classic optimality model considers when a foraging animal should stop feeding in a patch being depleted of prey and switch to a new patch, despite the cost of moving (Charnov 1976). The prediction most often tested is that increasing the travel time between patches should increase the time spent in each patch. This prediction has generally been confirmed, but less successful have been predictions about the absolute time spent in a patch (Nonacs 2001) and what cues to attend to so as to decide when to leave a patch (e.g., Roche et al. 1998; Hutchinson et al. 2008). Failures like this lead biologists to change or elaborate the basic model, for instance by incorporating additional aspects of the environment or by invoking some informational or cognitive constraint (e.g., Nonacs 2001; Hills and Adler 2002). Ideally, predictions are tested by manipulating the environment of an individual in the hope of a real-time response, but alternatives are to utilize variation among species or natural variation among individuals of a single species.

If what follows manages to say anything novel of interest to workers on cognitive search, we suspect that it will be because of, not despite, this perspective of the adaptation of behavior. Our biological perspective also brings to the table a greater diversity of search problems faced by different animals, and plants too (de Kroon and Mommer 2006), than by humans and our machines.

The Essence of Search

How would you define search? It is all too easy for a definition to use a near synonym like “locate,” which does not gain us much, or unintentionally to exclude phenomena such as searching internally for a solution to an anagram. Seeking a definition moved us beyond sterile questions of semantics, because it enabled us to recognize the essence of the search process that makes it distinct. We agreed that it would not be useful to define the term so broadly that it covered all adaptive processes.

Luc-Alain Giraldeau provided the initial insight. He proposed that for something to qualify as search there must first be a defined goal, such as food, mates, or particular information. The search itself then consists of acting to vary position according to some scheme that facilitates finding the goal. We definitely do not mean to restrict “vary position” to moving in space; rather, we include movement in other dimensions, such as sampling at different times of day or shifting attention somehow in one’s brain. It seems an important component of the definition that the varying of the position is adapted toward efficient location of the goal, hence the importance of defining the goal first. Thus we would not consider as search the process by which sand grains get deposited by the wind on the lee side of a dune. Nor is it search if animals explore and learn about the environment incidentally, ahead of starting to seek a

goal (latent learning: Thistlethwaite 1951). Our opinion is that search does not start until that goal seeking starts.

The goal that we invoke here is the function of the behavior. Without getting into philosophical debates about teleology, biologists are happy to say that a character has a particular ultimate function if design considerations suggest that natural selection has adapted it for that purpose. We are not talking about the proximate goal that one must identify to understand the *mechanism* of a control problem such as search. In this perhaps we differ from some other groups in this volume. To bring out the distinction, consider the princess and monster game, a classic example from the theory of search games (Isaacs 1965). A princess and a monster are free to move around in a darkened room or other space. The monster's goal, in the sense we intend, is to catch the princess; but, since neither can detect the other until they collide, its proximate goal cannot be capture but merely to move in particular prespecified directions.

Formally there may be an additional part of the search process: the application of a stopping rule to decide when the goal has been attained. Some valid sorts of search may lack a stopping rule. For instance, one can imagine a chemotactic bacterium following a gradient to the source; when it reaches the source it need not apply a stopping rule but oscillate around the source, its goal seeking continuing. If there is a stopping rule, its application is itself part of the search process. Note that a stopping rule may test the environment repeatedly during a search even though it triggers stopping of the task only once.

One important aspect of search is that there is some uncertainty in the location of the goal. If you can see the target and then walk straight toward it, that does not seem like search, although others coined the phrase "nonexploratory search" to cover situations where there is no uncertainty. Compare leafing through a book to find a particular passage with using the subject index: the former represents search with uncertainty, whereas an index is like a lookup table in computer programming, which is constructed to avoid repeated search or calculation. What about a blind organism that can apply a deterministic algorithm to locate a target reliably, say by chemotaxis? If it is absolutely always able to find the target, this behavior seems analogous to walking straight toward a visible target. Now consider the ability of some ants to return straight to their nests using solely path integration of their wiggly outward route (Müller and Wehner 1988). That does not initially sound like search, but actually their method of path integration is a clever approximation rather than exact (Müller and Wehner 1988), and they routinely must apply backup search mechanisms (Wehner 2003; Merkle and Wehner 2010). So, if we define search as involving uncertainty, recognizing a phenomenon as search may require us to know about the proximate mechanism and its performance.

We wondered whether a characteristic of search is that uncertainty tends to be reduced, or at least not to increase, at each step. One exception is the case when the search is for a mobile target known to be initially within some distance but able to move away (Foreman 1977), although perhaps search still

tends to delay the increase of uncertainty compared to random movement by the searcher. Real searches for a particular mobile prey can often fail, but this should not stop us considering the strategy that maximizes the probability of capture as a search.

Another aspect of most search is that it is sequential. By this we do not mean to exclude cases of multiple agents working in parallel and maybe sharing information; still each agent individually is searching sequentially. By “sequential” we intend to capture the idea that several steps must be taken to reach the target; a single-step process of selection between options is not search. The options change at each step and information gained from earlier phases should inform the choices made at later steps. A revealing example in this context is the secretary problem (Freeman 1983), the archetypal case of sequential search, which has been applied to model mate choice. Candidates of different qualities appear in random order one at a time; the object is to select a candidate of good quality, and each of a sequence of decisions is whether to accept the current candidate or continue inspecting further candidates. In this case, the only scope for varying “position” is the gain in information from inspecting the next candidate, but the crucial aspect is that information on the qualities of candidates inspected at earlier steps should determine whether search is terminated at later steps.

We tried, but failed, to agree on a single-sentence definition of search, preferring instead to list the key components: a functional goal, uncertainty about goal location, the adaptive varying of position, and often a stopping rule.

Nonspatial Search

The term search is most directly associated with seeking items in space, for instance, searching for your keys. But the term is also used in nonspatial contexts, and we thought it worth constructing a list of these in the hope of recognizing novel analogies between different domains.

Information

Many of the examples below fit into the larger category of information search. Searching for information is an implicit component of most search models, because finding the right target requires first acquiring relevant information (Vergassola et al. 2007). But sometimes we might consider that finding particular information is a goal in itself (Inglis et al. 2001). For instance, an explanation for why animals will work to explore suboptimal food sources (contrafree-loading) is that they gain the knowledge to utilize these sources if the currently better source disappears (e.g., Bean et al. 1999). Models from foraging for food have been reapplied to searching for information both on the Web and in our brains (Pirolli 2007; Wilke et al. 2009).

Quality

Selecting between mates is an example of a search over items of different qualities. In real life, candidates may well be distributed in space, but idealized models such as the secretary problem ignore this spatial component: the only decision is whether to continue search, not where to move, and it depends only on the qualities and number of earlier items, not on their positions. An interval between the inspection of successive items may represent a travel cost of moving between them, but the spatial aspect only makes a qualitative difference if there were some correlation of quality with position or if, for instance, checking one shop rather than another is more attractive because a third shop is closer to the first. Speed dating and comparison shopping on the Internet provide examples where a spatial component seems largely lacking. Nevertheless the money that advertisers pay to appear at the top of a Google search suggests that even slight spatial differences may be prominent to us.

Time

Animals may have to learn when during the day events are likely to occur (e.g., Biebach et al. 1994) or how long a resource takes to renew after the last visit (e.g., Henderson et al. 2006). Sampling over time so as to predict when an event will reoccur in the future is a search process. For instance, many of us will have learned from trial and error what times of day we will be best able to find a parking space near work. When bees start to learn when a food source is available, their sampling is biased earlier in the day than when they had first experienced the reward on previous days, which is adaptive in searching for the “opening time” of the source (Moore and Doherty 2009). Resampling a patch more intensely immediately after it depletes may also represent an adaptive search strategy in time (by analogy with area-restricted search: Gibson et al. 2006).

Correlation Structure and Learning

Just as there may be an association between a time of day and the occurrence of an event, other events or conditions might be associated with each other. Many forms of learning have been designed by natural selection for detecting this correlational structure in the world and responding to it adaptively. Habituation, which occurs following repeated exposure to the same meaningless stimulus, enables organisms to identify and ignore irrelevant events that do not predict meaningful events. Similarly, classical conditioning, which occurs when a previously neutral stimulus is temporally correlated with a meaningful event, enables organisms to identify and respond appropriately to events that predict meaningful events.

We debated extensively whether these types of learning should qualify as search. Although they share the goal of reducing uncertainty about the consequences of events, it is not clear that they involve any adaptive alteration of “position,” as required by our definition of search. Indeed, they appear to be passive processes that occur all the time with no clearly defined start or end, similar to latent learning.

Operant conditioning or trial-and-error learning, in which an animal learns the association between its actions and the occurrence of meaningful events, enables animals both to predict and to control these events. In addition to having the goal of reducing uncertainty about the consequences of actions, this form of learning additionally has the feature that an animal can actively explore the correlational structure of the world during acquisition by varying the circumstances in which it tries out actions. Therefore, we conclude that this form of learning has all the features that we have defined as characteristic of search. It is unclear whether operant conditioning always has a stopping rule.

Memory

Many kinds of memory retrieval are also search processes, involving cued activation of knowledge representations acquired from prior experience (Pachur et al., this volume). Importantly, memory retrieval also shares a parallel with spatial search in that similar items are retrieved near one another in time. Thus, in a free recall task where a person is asked to name as many different animals as possible, items remembered successively tend to lie in similar subcategories; for instance, first we might list pets, then birds, then animals from the Antarctic (Bousfield 1953).

Puzzle Solutions

Other sorts of search are solutions to puzzles, such as algebra or chess. We would be interested to learn how our minds organize the set of possible solutions, how we search through this landscape, and whether one could identify naturally occurring analogues to these sorts of problems for nonhuman animals.

Morphology and Physiology

All organisms are themselves the product of natural selection. We hesitate to call genetic evolution search because it involves neither a searcher nor a well-defined goal. However, analogous genetic algorithms have been constructed by humans to optimize the design of complex machinery such as turbine blades (Gen and Cheng 1997). This is a search process: the program is written so as to converge toward a specified goal. Similarly, Sherlock Holme’s search method for the truth, by eliminating all alternatives, has echoes in how our

immune system selects out all immune cells sensitive to self, thus enabling it subsequently to recognize non-self (the clonal selection theory). Animals may use other search heuristics to improve the design of aspects of their external phenotype-like burrows and tools. Thus spiders adjust the spacing between the lines of their web in response to the size of prey caught (Schneider and Vollrath 1998). Going beyond morphology, any homeostatic mechanism has the property of directing the state toward the neutral or set point. When there is imprecision, lags, or overshoot in the process, this seems like search, but something like a mechanical thermostat may lack the aspect of uncertainty required to fit our definition.

A Taxonomy of Search

Already we have mentioned a diversity of search problems. To recognize structural similarities between search in different domains, it helps to consider in what fundamental ways the problem of search can vary. This might also facilitate understanding why different methods of search are used in different search problems.

A distinction is often made between searching for one particular item (e.g., the dropped key to your house or a missing offspring) and searching for a class of items. Contrast the birdwatcher who goes out to a sewage farm on the off chance that something interesting will be there with the serious twitcher who flies out to Fair Isle specifically to see the rare American vagrant that was reported on Birdline. In practice, it is usually possible to recognize a continuum between these extremes: the twitcher would be satisfied by an even rarer species that turned up while he was there. Models of optimal search when the target is a specific individual known to lie within a specified area predict rather different behaviors (e.g., systematic searching, randomized strategies; Alpern and Gal 2003) than when any individual in a population will suffice (Hutchinson and Waser 2007).

Some searches, archetypally for a male mate or for a nest site, are one-shot processes: once you make your choice, you stop searching. In contrast, once a bird finds one worm, it immediately starts searching for another, so the problem is iterated. The iteration seems important mainly in affecting the opportunity costs; one reason that the bird is less fussy about the quality of a worm than of its mate is because spending more time searching for one food item detracts from time searching for the next food item. In this respect there is no fundamental difference from the effect of other costs of search, such as mortality risk and locomotion costs. There may also be external time constraints, such as the ending of the breeding season (e.g., Backwell and Passmore 1996).

Another aspect is the dimensionality and topology of the problem. Contrast one-dimensional searching for flotsam along a river bank (or between a succession of secretaries knocking at your door) with the extra freedom of movement

in search of two- or three-dimensional space: somewhere in between are ants exploring a tree or foragers relying on tracks through thick scrub; they face a network of restricted moves that creates a topologically very different search space than the almost unrestricted search of a shark in the ocean. Just as important as the topology are the movement rules allowed in this landscape (e.g., in the secretary problem whether recall of candidates inspected earlier is allowed).

In some searches the animal can be guided only by its past experience in the patch, as in area-restricted search for buried prey (e.g., Nolet and Mooij 2002). In other cases there are external cues, such as a pheromone plume, that assist in locating the target and perhaps in indicating target density (e.g., Waage 1978). Mueller and Fagan (2008) make a similar distinction. The experimental and theoretical analysis of how animals utilize cues such as gradients and landmarks is well developed (e.g., Fraenkel and Gunn 1961; Schöne 1984; Dusenbery 2001).

In the absence of external cues, the autocorrelation of items and their qualities in space and time provides the only information that the searcher uses. Autocorrelation in space is an integral part of models of area-restricted search (Benhamou 1992), but in other cases modelers have instead invoked discrete recognizable patches of items in a sea of absence (e.g., Charnov 1976). Which is more appropriate depends both on the actual distribution of targets and on the ability of the searcher to recognize the edge of the patch at a glance (Bond 1980). Autocorrelation in time involves the processes of depletion, disturbance, and renewal. Analyses of data derived from modern tracking technologies demonstrate the importance of considering autocorrelation in space and time at multiple scales simultaneously (e.g., Fauchald and Tveraa 2006; Amano and Katayama 2009).

An unduly neglected aspect of search is social interactions among searchers and targets. We devote the next two sections to considering how social interactions can transform the problem.

Social Interactions

Search is not always a single individual seeking an inanimate target that is indifferent to being located. There can be positive (mutualistic, +), negative (competitive, -), or indifferent (neutral, 0) relationships, to varying degrees, among social searchers, among social targets, or between searchers and targets, whether social or solitary (summarized in Figure 4.1). Interactions may be infrequent or nearly continuous. Once atune to these social possibilities, one can recognize a large set of possibilities that may shape the evolution of search behavior.

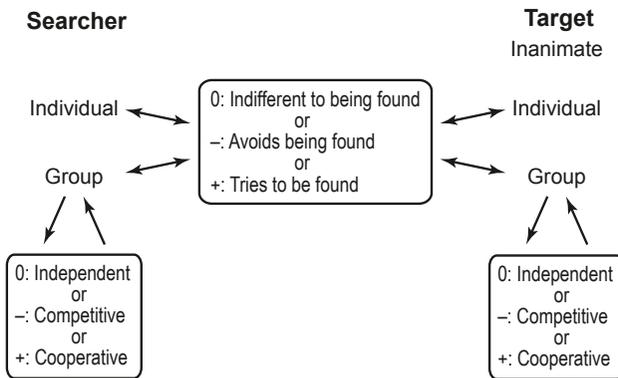


Figure 4.1 A schematic summary of the possible social interactions of searchers with other searchers, of searchers with their targets, and of targets with other targets.

Searchers

Normally we expect competition between foragers. Think of searching for a parking space as near as you can to a cinema: we probably suppose that the closest spots will tend to be occupied already, which makes us use different search strategies than if we expected randomly distributed spaces (Hutchinson et al. 2012). Because the strategies used by others determine the distribution of spaces, the situation is game theoretic. The converse case of cooperation or sharing of information among searchers also affects the effectiveness and appropriate choice of different search tactics. In some central-place foragers (social insects, camp-based hunter-gatherers), it is deliberate sharing of information that allows improvement in the locating of resources. But even if individuals do not deliberately signal to others, they may coordinate by copying processes, resulting in emergent search properties of collections of individuals in contact with one another (see Box 4.1).

Searcher-Target Interactions

Some sorts of targets want to be found (e.g., mates, +), some do their best to avoid being found (e.g., prey, -), and some are indifferent (e.g., a water source, 0). Mates may signal their presence to potential suitors; prey may adopt cryptic or evasive tactics. Each thereby may change what strategies are effective for the searcher. The brain presumably locates information for storage in a manner to facilitate its being found: memories in some sense want to be located. Similarly, we expect food-storing birds to hide their caches at a pattern of sites that facilitate rediscovery by themselves, but with the complication that the cues should not make it easy for competitors to pilfer (Cheng and Sherry 1992; Barnea and Nottebohm 1995; Briggs and Vander Wall 2004). Plants may distribute their flowers on an inflorescence so as to benefit from the search rules

Box 4.1 Collective Search

Generally speaking, adaptive search strategies act to match individual behavior to the relevant statistical properties of the environment. In terms of sensing, separating a relevant signal from environmental noise is often a challenge. At the level of individuals, sensory adaptation and simultaneous use of multiple modes of sensory information can allow individuals to respond dynamically to maximize the signal-to-noise ratio. When searching as a collective, however, strategies may be implemented both at the individual and group level. An illustrative example of collective search is to consider each individual as a sensing agent capable of detecting and responding to relevant environmental features, such as the estimated direction of a local resource gradient, but also to other individuals. If the environment has simple structure, such as a linear gradient and low noise, taxis up the gradient is relatively straightforward. In more complex environments, such as where local noise inhibits taxis or where simple gradient climbing can result in entrapment in local optima, collective strategies can facilitate much more effective search.

Modeling such behavior, Torney et al. (2009) considered the case of locating the source of a chemoattractant within a stochastically fluctuating advective flow; think of how blood from an injured swimmer might drift offshore toward a patrolling shark. This is a ubiquitous behavior important to the lives of many aquatic animals and observed over a wide range of scales. In this situation, individual-level search is particularly ineffective since the filamentous and turbulent structure of the plume confuses local search strategies, resulting in individuals following local optima and seldom being able to find the source itself. Similar problems apply when considering any spatially heterogeneous gradient of resource (including gradients of discrete resources). A highly effective strategy under such circumstances is for multiple individuals to reconcile their goal-oriented taxis with social interactions (i.e., affiliating or aligning with others). The central principle is that if organisms dynamically adjust how much they are influenced by social interactions based on their confidence in their own environmental assessment, they can, as a collective, find global optima. Thus, in the model of Torney et al. (2009), if an individual perceives an increasing local concentration of odor, it decreases the weight it places on social interactions. When concentrations are unpredictable or declining, individuals may instead place more weight on social interactions. This strategy does not require organisms to know the informational state of others explicitly but nevertheless can spontaneously create a time-varying spatial leadership structure in which individuals with low confidence follow spontaneously those who are obtaining relevant information from the environment. Thus individuals continuously adapt to the changing physical and social structure of their environment, giving them the capacity to respond to structural information over length scales much larger than their own range of perception.

Evidence that animals can, and do, adjust their sensitivity to the behavior of others comes from studies of schooling fish. For example, the context dependence of interaction ranges can explain the group-size distribution of schooling killifish (Hoare et al. 2004), and stickleback fish have been shown to restrict their schooling tendency when they can gather direct reliable information from the environment, but increase their tendency to group with others when this information is perceived to be unreliable or scarce (van Bergen et al. 2004).

of pollinators, but not necessarily in the very best interest of the pollinators (Jordan and Harder 2006).

A special case is when both searcher and target are mobile and can search for each other (technically, models of rendezvous: Alpern and Gal 2003). For instance, both sexes of some butterfly species fly to the tops of hills (hill-topping) to facilitate encounter there (Alcock 1987). In contrast, had they not evolved to utilize such an asymmetry in the environment, the optimal policy to maximize encounter is for both sexes to move as fast as possible in straight lines (Hutchinson and Waser 2007). A similar situation of symmetrical roles for searcher and target is mutual mate choice, which has been modeled both with and without competition between searchers through depletion (e.g., Collins and McNamara 1993; Johnstone 1997).

Targets

Targets, if they are living, have their own internal relationships that affect search. Many animals create exclusive home ranges to avoid competition with neighbors, and their consequent overdispersion should affect the search rules used by predators (Iwasa et al. 1981). Conversely, in the selfish herd model individuals hide behind neighbors so as to minimize their own chance of being selected by a predator, incidentally creating herds (Hamilton 1971), which may facilitate the search of the predator (Treisman 1975).

Open Questions

Behavioral ecologists have models for many of the behaviors mentioned above. The application of game theory can predict rather different outcomes than models that ignore the social interactions we have considered here (e.g., Johnstone 1997; Hamblin et al. 2010). But do animals, including humans in everyday life, also know to shift their search methods in social situations? For instance, if a traplining hummingbird tries to adjust its revisit rate to a particular flower, it is crucial for it to judge whether nectar supply has declined because of competition (when it should revisit sooner) or because the flower is producing at a slower rate (when it should revisit later; Garrison and Gass 1999). A different sort of question is whether models of social foraging may be relevant to cognitive search processes that involve parallel processing. For some more open questions, we now address in more detail the game-theoretic analysis of the well-studied producer-scrounger paradigm.

The Peculiar Social Dynamics of Selfish Parallel Search

When several selfish animals search in parallel for some resource—be it mates, nesting material, food, or information itself—they have an option that is never

available in an individual search process: they can either search for the resource themselves or search for other individuals that have already uncovered a resource item. For instance, when a group of pigeons search the ground for hidden seeds, only some do the actual searching whereas almost all of the group will gather at any one individual's discovery. The decision of whether to invest in one or the other search mode is modeled as an economic decision using a game-theoretic approach: the producer-scrounger (PS) game, where producer is the strategy of searching directly for the resource and scrounger is the alternative of searching for individuals that have uncovered the resource (Barnard and Sibly 1981). Very similar scenarios have received other names, such as tolerated theft in anthropology or free-loading in economics. Essentially the dynamics are always the same. The PS game has been modeled in many different ways (see Giraldeau and Caraco 2000; Arbilly et al. 2010) and has given rise to an extensive experimental research program (Giraldeau and Dubois 2008; Katsnelson et al. 2011). In all cases, the question is directed at predicting what fraction of the selfish parallel searchers searches directly for items.

All PS models predict that the strong frequency dependence of the scrounger strategy's payoffs leads to an equilibrium frequency of producers and scroungers characterized by equal payoffs for each strategy. In behavioral ecology, the usual account of how this equilibrium is attained involves invoking a mutant scrounger strategist originating within a population of pure producers. The rare mutant outperforms the producers and so the strategy spreads in the population over generations. As the scrounger strategy becomes more common, its fitness declines and eventually reaches the fitness of the producer strategy. At that point, no further evolution occurs because both strategies have equal payoffs. This equilibrium point is referred to as an evolutionarily stable strategy (ESS) because no other combination of strategies can do better within this population (Maynard Smith 1982).

However, in almost every situation in which the predictions of the PS game have been studied experimentally, the equilibrium is reached quickly within a generation, not over evolutionary time. Moreover, individuals rarely search only as scroungers or as producers. Instead usually individuals alternate, sometimes rather quickly, between the two strategies. The process through which the population of parallel searchers reaches the equilibrium therefore involves selfish agents using a decision rule adapted to maximize their individual benefits but leading the group to an equilibrium point that is not an evolutionary equilibrium in the sense above, but rather a Nash equilibrium that is behaviorally stable. The optimal decision rule specifies the probability of playing scrounger given how often it is played in the population. In the case of two-person games, we know that such a decision rule can yield different equilibria than the ESS described earlier (McNamara et al. 1999).

The form of the optimal decision rule that evolves depends on the payoffs of playing each strategy as encountered over evolutionary time. But the assumption of the model is that the players themselves simply apply the hardwired

rule that has evolved, and thus respond not to the payoffs but only to the proportions of the strategies. In real life, however, given the diversity of foods and environments encountered each day and the differing payoffs of the two strategies in each, what might evolve instead is a rule that does try to sample the current payoff of each strategy and shift the probabilities of playing each accordingly. A number of early studies proposed learning rules that allow individuals to adjust their search strategy based on their experienced payoffs (e.g., Harley 1981). These learning rules must contend with the nontrivial problem of estimating the value of searching as a producer or as a scrounger while these payoffs keep changing as a result of other players also switching policy to learn both payoffs. How animals might discover their best policy remains a gap in our knowledge about the parallel search of groups of selfish agents.

When trying to derive lessons from this collective search to problems of cognitive search, one must first determine whether cognitive search might be represented as a collective of selfish agents. If it can, then no doubt the dynamics of the PS game will emerge. However, even if the collective search involves cooperating rather than selfish cognitive search agents, scrounging will likely remain an option. In such cases, research has shown that cooperative solutions to the PS game dynamics can often lead to cooperative producers extending their assistance to all other searching agents, which means more scrounging in cooperative systems compared to selfish ones (Mathot and Giraldeau 2010). It would be important, therefore, to investigate the extent to which cognitive search can be represented as a group of agents searching in parallel (cf. Minsky 1986).

Multiple Modalities and Search Cues in the External Environment

Within the animal kingdom, a wide variety of senses are known. In addition to the obvious senses of vision, hearing, chemosensitivity (olfaction, taste) and touch (somatosensitivity), animals may be sensitive to magnetic and electric fields, gravity, acceleration, time of day, the configuration of their own bodies (proprioception), and to the sensation of internal states, such as the fullness of the gut (visceral senses). Potentially any of these could be used to guide search. Which senses provide the most suitable cues to guide particular search problems is partly a function of the laws of physics and chemistry. For instance, in dense forest sound carries better than light; in turbulent wind, airborne pheromones will not allow precise location of a target. Even within each modality, selection of the signal allows some tuning of the physical properties; thus frequency affects sound transmission, and pheromones differing in their half lives are used in an adaptive way by ants to mark their trails with different permanences (Dussoutour et al. 2009). Sensitivity to a particular modality depends not only on the physical properties of the cue but on the sense organs

and brain of the searcher, which are constrained by their costs of construction and maintenance.

The physical differences between modalities can explain some of the variation of search strategy used by different organisms or the same organism in different contexts. Thus a moth can use its eyes to fly straight to a bright flower or use the moon as a distant navigational beacon (keeping it at a constant angle). And when searching for a mate releasing a pheromone it follows another distinctive search strategy, flying crosswind when not sensing the odor, and flying upwind when within the odor plume (Sotthibandhu and Baker 1979; Kennedy 1983).

Physical aspects may also explain both what modalities a species has evolved to use for search and which of these modalities it uses in particular circumstances. For instance, a pigeon may use a sun compass in clear weather but switch to a magnetic compass when the sun is obscured (Walcott 2005). Shine et al. (2005) consider why male garter snakes at low densities rely on following olfactory trails to find females (olfaction is accurate in distinguishing sex), whereas at high densities they switch to visual tracking (vision is not disrupted by the trails of rivals and greater speed is valuable in the more competitive situation). A common pattern is that searching animals switch between modalities sequentially as they approach the target and each sense gets into range. For instance, the digger wasp *Philanthus triagulum* hunting prey is first attracted visually by a smallish and moving object, then approaches closer downwind to check its scent, jumps on it, and can then use tactile or taste cues (Tinbergen 1958). Analogously, female sage grouse first assess males gathered in a lek on the basis of their calls and then visit only those passing this test for a closer inspection of display rate (Gibson 1996). Similar winnowing of options by one cue at a time is mirrored in strategies humans used in Internet shopping (Fasolo et al. 2005).

Sequential application of each cue, one at a time, is one way in which cues of multiple modalities may be combined in search, but there are many other possibilities and many patterns have been observed (Candolin 2003; Hutchinson and Gigerenzer 2005). There is also a rich literature in human decision making on how we combine information from different cues when comparing two items (Payne et al. 1993; Gigerenzer et al. 1999; Bröder and Newell 2008). It seems that often we do apply one cue at a time, even for information already in memory, particularly if the information was not originally presented as a single image (Bröder and Schiffer 2003); a single image seems unlikely if the information comes from several modalities and appears at separate times.

We observe that one aspect of search shows a striking commonality across modalities. When searching for cryptic prey visually, humans and other animals tend to pick a characteristic feature and focus on that, filtering out other information (e.g., Dukas 2002). Such so-called search images improve performance at spotting the target prey and other objects sharing the feature, but decrease our ability to detect other dissimilar prey items (Dukas and Kamil

2001). Analogues of search images have been found in the auditory (Fritz et al. 2007) and olfactory domains (Cross and Jackson 2010). Even bacteria can tune their sensitivity to particular chemicals in their environment (Muller-Hill 1996), so it may be a rather general feature of search in the external world. The following might be an analogue in memory search: when people have to decide whether a sequence of sounds is a valid word, they recognize “robin” as a valid word more quickly if they have been warned that any valid word presented is likely to be the name of a bird (Neely 1977).

Humans form different neuroanatomical representations of memories depending on the sensory modalities they use to encode those memories (Markowitsch 2000). Do we also search for things in memory differently depending on the modality with which the memory was encoded? For instance, whereas it is straightforward to order colors and sounds along simple axes such as wavelength or loudness, with tastes there are no such obvious dimensions because of the physical basis of chemosensitivity. And even though most real visual stimuli are complex patterns which also cannot be readily ranked along a single dimension, the poverty of our language to describe tastes points to a difference in our ability to classify them. Does this mean that we store and access memories for tastes differently than we do for memories of visual objects? Is the process by which a wine expert deduces the origin and vintage of a wine from its taste different from how an expert attributes a painting?

Further Connections between Search in Behavioral Ecology and Cognitive Psychology

As our discussion above reveals, external search problems (often the domain of ecology) and internal or more abstract search problems (often the domain of psychology) are perhaps not as unrelated as they may at first appear. Here we explore some more of the potential connections between ecological and psychological perspectives by considering specific problems in cognitive psychology about which insights from ecological research offer new questions.

Interindividual Variation in External Information Search

For a grazing animal, exploration for new resources often goes on simultaneously with the exploitation of those resources. In other cases, for instance, when an animal or human is searching for a new home, an exploration phase precedes the exploitation. In the exploration phase, the search is only for information in the external environment.

One task that captures this distinction between exploration and exploitation is called the sampling paradigm (Hills and Hertwig 2010). In the sampling paradigm, a person is asked to make a decision between two options (Option A and Option B). The person is allowed to sample freely from these two options,

without receiving any direct reward, gaining only information that will later be useful. The options themselves are associated with specific payoff distributions (e.g., Option A pays \$3 with certainty, Option B pays \$32 10% of the time, and \$0 otherwise). So, for example, a person might sample from Option A several times and witness potential payoffs (e.g., \$3, \$3, \$3) and then sample from Option B (e.g., \$0, \$0, \$30). After some amount of sampling, the person makes a final consequential choice between the two options, and only then actually receives one payoff.

Studies of the sampling patterns in this information-search task reveal a bimodal distribution in the frequency with which individuals switch back and forth between Options A and B (Hills and Hertwig 2010). Some participants sample repeatedly from Option A, then they switch to sample repeatedly from Option B, and then they make a final decision. Others participants switch frequently between Option A and Option B. People who frequently switch tend to take fewer samples overall than those who switch less frequently. People who frequently switch are also more likely to make a decision consistent with a roundwise decision policy, one based on the number of times a sample from one option beats the preceding sample from the other option. People who switch infrequently are more like to choose the option associated with the higher expected value overall.

Individual differences in search behavior are not restricted to humans. In the fruit fly *Drosophila*, natural allelic variation in a protein kinase gene results in the “rover” and “sitter” dimorphism (Osborne et al. 1997). Rovers leave food patches more readily, visit more food patches, and revisit food patches less compared to sitters, which are more sedentary and aggregate within food patches (Nagle and Bell 1987; Pereira and Sokolowski 1993; Stamps et al. 2005). The same gene has been implicated in learning and memory traits in *Drosophila* larvae and adults (Méry et al. 2007; Reaume et al. 2011), and orthologues are involved in regulating food-related and social behaviors in a variety of other animals (Reaume and Sokolowski 2009). Do similar genetic differences underlie the variation in human search patterns? Moreover, might these differences in search behavior reflect differences in cognitive processing that influence a wide range of tasks involving cognitive search, including learning?

Memory Search

Memory search can be characterized as search through information topologies stored in the brain (Davelaar and Raaijmakers, this volume; Hills and Dukas, this volume). What is the structure of these topologies? In the semantic fluency task, people are asked to recall as many items as they can from a specific category (e.g., “say all the animals you can think of”). In this task, humans often produce items as if they were retrieving them from memory clusters (Bousfield 1953). Some data suggests that semantic memory may reflect a clumpy or

patch-like structure (Steyvers and Tenenbaum 2005). This suggests that human memory search could follow similar foraging policies as described for animals foraging on spatial patches of prey (Hutchinson et al. 2008; Hills et al. 2009).

However, memory representations differ in potentially important ways from space. Understanding the nature of these potentially dynamic topologies may be critical to our understanding of how memory search works. For example, an item in memory can belong to different representations simultaneously: the word “cat” can belong to the category of “pets” as well as to the category of “predators.” The representation need not be based solely on semantic similarity but also, for instance, on phonological similarity (“cat” and “bat”). Thus words could potentially belong to more than one patch. Studies of memory search should ask what is the patch structure of memory and how are these patches used. For example, do the patch-like subcategories (e.g., pets) really reflect some special organizational linking of items in memory, or are items in memory evenly spread and the apparent patches simply the behavioral outcomes of individuals moving in memory from one item to a nearby item, what Pollio et al. (1968) called an “associative” search? Studies of memory search can potentially explore the cognitive mechanisms guiding search in similar ways to those used to study animal foraging. For example, increasing the costs necessary to switch between patches leads animals to stay in patches for longer periods of time (Nonacs 2001). The analogous manipulation in a memory-recall experiment could be accomplished by imposing external costs, for instance, by increasing the time it takes to be presented with a new category for free recall (if subjects are paid in terms of recalled items per unit time; cf. Wilke et al. 2009). One could also look at recall patterns from more or less sparse semantic domains: foods are highly semantically similar; occupations may be less so. Recent models of semantic space allow the objective computation of similarity between words based on large corpora of text, using word co-occurrence (Jones and Mewhort 2007). This offers innovative ways to represent the landscape over which memory searches.

Do nonhuman animals also search memory as if it had a patchy structure? One experiment that we considered involved training a pigeon to peck at several categories of images (e.g., cats, trees, and human faces). These could then be presented on a grid with numerous distractors (i.e., images not belonging to the target categories). After extensive training, the pigeons could then be asked to recall the locations of these targets on an unlabeled grid. Would they recall the items by category (e.g., first all the cats, then all the trees), as if the information were stored in semantic patches or in some other way such as spatial proximity?

Some animals possess cognitive maps which allow them to take novel routes between spatial locations. Might information in memory be stored as a cognitive map, allowing humans and nonhuman animals to link previously unlinked information adaptively?

Problem Solving

Another form of internal cognitive search involves manipulating the arrangement of information in working memory in such a way that it provides a solution to a problem. Chess players search for potential solutions to a chess problem, even if they have never before seen this particular arrangement of chess pieces (De Groot 1965). A similar kind of problem involving a search through arrangements is the Tower of London problem, involving the lawful rearrangement of colored balls on sticks to match a final target pattern (Shallice 1982). A novice player cannot solve such a problem by recalling the answer; it requires the active construction of a new solution, by cognitively simulating and searching through the possibilities.

Some nonhuman animals appear capable of this kind of problem solving. Jumping spiders plan paths before moving (Tarsitano and Andrew 1999). Some individual ravens faced with food suspended on a string discovered how to lift it up using beak and feet without trial-and-error learning of the process (Heinrich 1995). Emery and Clayton (2004) discuss other examples of such insight.

These kinds of planning associated with problem solving might be productively thought of as forms of route planning, similar to the way rats have been demonstrated to simulate exploration of space actively in so-called episodic future thinking (Redish, this volume). Are searches through configural solution spaces governed by similar kinds of strategies, as found in spatial search?

Language Acquisition

Is learning language also a kind of search process? Social animals may have as a developmental goal the acquisition of effective communication strategies. Human children learn language, learning both word meaning and grammar, and they do so in predictable ways. However, the process of language acquisition is still not well understood. Could it represent a search process? Goldstein et al. (2003) suggest this possibility by noting that human children share a phase of exploratory linguistic babbling similar to that found in birds; in both cases, the babbling appears to be “shaped” by interactions with adults. Analogously, male brown-headed cowbirds (West and King 1988) and satin bowerbirds (Patricelli et al. 2002) rely on feedback from females they court to refine their courtship behavior. Even *Drosophila* males show plasticity in the courtship dance, which is learned through successive interactions with females (Polejack and Tidon 2007). Though the question is rather broad, could goal-directed exploration characterize the learning of these various forms of communication?

Tailpiece

This chapter is not a comprehensive review of search, but rather reflects the esoteric choice of topics that matched our interests and expertise and that we had time to discuss. The topic seems endless because search is such a widespread phenomenon and research on it so multifarious. By thinking carefully about what is fundamental to search, and by recognizing some commonalities between research in different disciplines, we hope to have introduced a little more structure into the topic. There may never be a single overarching theory of search, but some imposed structure is helpful in recognizing how our own research relates to existing work, and in drawing attention to relevant gaps in our knowledge that require investigation.

