

Exploration versus exploitation in space, mind, and society

Thomas T. Hills¹, Peter M. Todd², David Lazer^{3,4,5}, A. David Redish⁶, Iain D. Couzin^{7,8}, and the Cognitive Search Research Group⁹

¹ Department of Psychology, University of Warwick, Coventry, UK

² Cognitive Science Program, Indiana University, Bloomington, IN, USA

³ Department of Political Science, Northeastern University, Boston, MA, USA

⁴ College of Computer and Information Science, Northeastern University, Boston, MA, USA

⁵ Harvard Kennedy School, Harvard University, Cambridge, MA, USA

⁶ Department of Neuroscience, University of Minnesota, Minneapolis, MN, USA

⁷ Department of Ecology and Evolutionary Biology, Princeton University, Princeton, NJ, USA

⁸ Department of Collective Behaviour, Max Planck Institute of Ornithology, Konstanz, Germany

⁹ Melissa Bateson (Centre for Behavior and Evolution/Institute of Neuroscience, Newcastle University), Roshan Cools (Radboud University Medical Center, Department of Psychiatry; Donders Institute for Brain, Cognition and Behaviour, Centre for Cognitive Neuroimaging), Reuven Dukas (Department of Psychology, Neuroscience, and Behavior, McMaster University), Luc-Alain Giraldeau (Faculté de Sciences, Université du Québec à Montréal), Michael W. Macy (Department of Sociology, Cornell University), Scott E. Page (Center for the Study of Complex Systems, Departments of Political Science and Economics, University of Michigan), Richard M. Shiffrin (Department of Psychological and Brain Sciences, Indiana University), David W. Stephens (Ecology, Evolution, and Behavior, University of Minnesota), and Jeremy W. Wolfe (Brigham and Women's Hospital and Harvard Medical School)

Search is a ubiquitous property of life. Although diverse domains have worked on search problems largely in isolation, recent trends across disciplines indicate that the formal properties of these problems share similar structures and, often, similar solutions. Moreover, internal search (e.g., memory search) shows similar characteristics to external search (e.g., spatial foraging), including shared neural mechanisms consistent with a common evolutionary origin across species. Search problems and their solutions also scale from individuals to societies, underlying and constraining problem solving, memory, information search, and scientific and cultural innovation. In summary, search represents a core feature of cognition, with a vast influence on its evolution and processes across contexts and requiring input from multiple domains to understand its implications and scope.

Exploration versus exploitation

Search, or seeking a goal under uncertainty, is a ubiquitous requirement of life. Animals forage for food, territory, and mates. Humans engage in a wide variety of search behaviors, from looking for lost keys, to finding financial opportunities, to seeking the meaning of existence. Moreover, they search in a wide range of spaces, including visual scenes, memory, social networks, information databases, and more abstract environments such as problem spaces associated with cures for disease, product design, and governmental policy. Search in each of these domains involves trade-offs between exploiting known opportunities and

exploring for better opportunities elsewhere. Here we describe how the exploration–exploitation trade-off is fundamental to understanding cognitive behavior at various levels, from its evolutionary origins to the function of cognitive control across domains. How cognitive systems handle this trade-off also has broad implications, from harmful mental disorders to the progress of social innovation. The ubiquity of this trade-off and cognition's solutions to it implicate search as a common framework for understanding many aspects of cognition.

Fields ranging from animal foraging to computer science use different terminology to describe what are fundamentally similar strategies for handling the exploration–exploitation trade-off. In decision-making research, it is referred to as exploitation versus exploration [1,2]; in visual attention, the contrast is on focused versus diffuse search [3,4], in foraging, intensive versus extensive [5], in memory, local versus global [6], and in artificial intelligence, depth-first versus breadth-first search [7].

The diversity of terms used to describe this trade-off reflects a common problem across many cognitive domains (Table 1). In response, numerous solutions have arisen. In computer science, rules that achieve appropriate reinforcement learning have elements of exploitation intermixed with exploration [8]. Ecological search strategies often involve intensive local foraging mixed with occasional exploration phases that move animals from one cluster or region of resources to another [9,10]. Similarly, social insect swarms, such as those of many ants and honeybees, exhibit phases of diffusive and focused search in response to changing resource distributions, as perceived at the level of the colony [11–13].

Traces of the historical significance of this trade-off are found in the evolution of cognition. Recent comparative

Corresponding author: Hills, T.T. (t.t.hills@warwick.ac.uk).

1364-6613/

© 2014 Elsevier Ltd. All rights reserved. <http://dx.doi.org/10.1016/j.tics.2014.10.004>

Table 1. Examples of trade-offs between exploitation and exploration across cognitive domains

Animal foraging	Exploiting a known berry bush versus exploring for new bushes
Visual search	Analyzing one spot on a chest radiograph versus looking for the next spot to check
Information search	Searching within a document versus searching for new documents; deciding when to accept an item on a menu versus continuing to look for new items
Search in memory	Trying to remember more African animals versus switching to Australian animals
Search in problem solving	Focusing on solutions that have worked in the past versus seeking new solutions
Social (group) learning	Learning or copying existing knowledge versus using innovation to seek new knowledge

neurobiological findings of dopaminergic and related mechanisms support a putative common ancestral precursor for foraging behaviors across many animal species. Our developing understanding of the shared control structures among neural correlates of response selection (e.g., prefrontal cortex and basal ganglia), exploration and memory (e.g., hippocampus), and visual attention (e.g., parietal cortex) has produced the realization that numerous goal-directed processes central to human cognition rely on the integration of search-related architectures. Findings such as these lead to the compelling conclusion that the same cognitive and neural processes underlie much of human behavior involving cognitive search – in both external and internal environments [9,14].

Beyond the individual, our understanding of the exploitation–exploration trade-off extends to the collective behavior of social insects, formal institutions, and populations organized around political and social identity. In all cases, performance depends on striking an optimal balance between imitating best practices and experimenting with innovations. Here we follow the path from individual to collective behavior, offering a tour of the unifying themes on search that encompass its evolutionary, ecological, neural, algorithmic, and social bases.

From spatial foraging to foraging in mind

Individual organisms must strike the proper balance between global exploration and local exploitation to survive – exploring sufficiently to find resources and exploiting sufficiently to harvest them. This optimal control problem is well-studied in animal foraging, where numerous species focus on finding resources in a cluster (or patch) until their acquisition rate falls below some threshold, as described by the marginal value theorem [15]. At that point, organisms typically leave to seek a new patch [10,16]. Related spatial search strategies apply when there are not clear boundaries around resources. In such cases, predictions can be made about how particular strategies are appropriate to different resource ecologies and cognitive capacities of organisms (Box 1).

Visual search in humans involves at least two nested senses of exploration and exploitation. In a typical visual search task, observers would look for targets among distractor items in a series of scenes or displays. Observers might be said to exploit each scene until the time comes to move on and explore the next scene [17]. On a finer scale, observers explore from item to item in a scene, spending differing amounts of time exploiting each attended item. Figure 1 illustrates some of the factors governing this

within-scene exploration. Under reduced-stimulus conditions, observers come close to optimally balancing within-scene exploration and exploitation. Similar issues arise when deciding to move between visual scenes. For example, radiologists may have many images to scan and must balance local search time in each against the need to

Box 1. Animal foraging

How long should a foraging animal stay within a given resource patch? In Charnov's (1976) marginal value theorem, a forager with perfect knowledge of the environment spends t time units exploiting and depleting a resource patch and takes τ time units to travel from one patch to another. Within each patch, the forager gains an amount of food specified by the function $g(t)$, with the rate of gain decelerating the more the patch has been exploited. The mean intake rate for an animal exploiting a sequence of these patches is:

$$\frac{g(t)}{(\tau + t)} \quad \text{[I]}$$

The value of t with the highest possible intake rate satisfies the 'marginal value' condition:

$$g'(t) = \frac{g(t)}{(\tau + t)} \quad \text{[II]}$$

This means that the forager should leave a patch when its instantaneous rate of intake in the patch $g'(t)$ falls to the mean intake rate for the environment as a whole. Hence, foragers should leave patches relatively early in rich environments but leave later in poorer environments. Qualitatively, this prediction is well supported [10,88].

Strategies change with increasing uncertainty about the quality of the current resource or the likelihood of finding something better. With non-depleting patches, a forager should 'stay forever' in a patch of the best type once it finds it; for example, trap-building spiders occupy permanent feeding sites after a period of sampling. This is a type of optimal stopping problem, comparing the value of the current resource to the forager's expectations about future rewards (cf. [15] versus [89]).

In situations with poorly defined patch boundaries (e.g., a bumblebee exploiting clumps of flowers), foragers often rely on area-restricted search, intensifying their search in a local area after encounters with food but exploring more as food encounters decrease [9,90]. This can be modeled as satiety S controlling frequency of turning r with an increasing function $r = R(S)$. Now, turning will be highest in areas of highest prey density. Such a process can generate 'prey taxis', in which foragers climb gradients of prey abundance [90].

At larger spatial scales, resource gradients can be more complex; gradient-following rules may then fail by becoming trapped on inferior local maxima far from global peaks of resource abundance. Making occasional large movements independent of recent experience may then be beneficial. One strategy that accomplishes this is a Lévy flight, which is operationally defined as any heavy-tailed search strategy such that path lengths, l , are drawn from a power-law distribution, with $P(l_j) = l_j^{-\mu}$ and $1 < \mu < 3$ [91]. Under what conditions Lévy flights are common or optimal, however, is hotly debated [5,92,93].

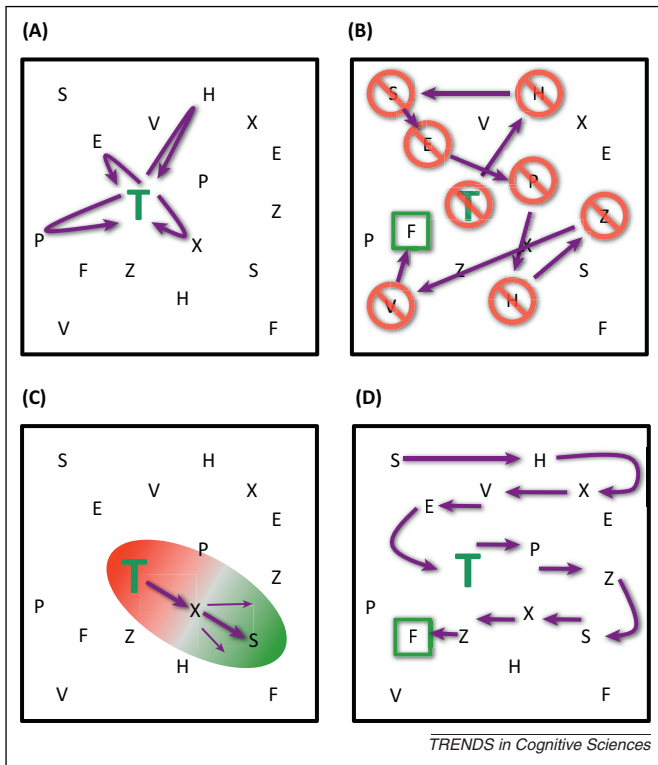


Figure 1. Local exploration in human visual search. In visual search, people look for specific letter targets; for example, Fs in an array of letter distractors. The local exploitation step is the act of recognizing a letter and determining whether it is your target. The local exploration step is the act of selecting the next letter (accomplished at a rate of about 20–40 letters/s). All else being equal, visual attention is drawn to salient items in the field (A). How then are we to avoid perseverating on one incorrect but vivid letter? One answer (B) is to rely on the phenomenon of ‘inhibition of return’ (IOR) [94]. If one attends or fixates on an item and then deploys one’s gaze or attention away from that item, it becomes harder to bring the gaze or attention back to the original item than to move it elsewhere. It was originally believed that IOR would permit attention to sample the display without replacement. Unfortunately, further research found that visual search was not markedly impaired when IOR was blocked and observers had to sample with replacement. A more moderate view might hold that IOR serves to bias exploration toward new items (C) even if it does not absolutely prevent return to a rejected item. However, given enough time, observers can adopt strategies that allow them to prevent perseveration [95]. Thus, for example, one might ‘read’ a display from side to side and top to bottom. This more controlled, prospective strategy (D) would avoid sampling with replacement but would slow the rate with which items can be processed. In some cases, a more chaotic strategy will get you to the target more quickly [96].

examine the next radiograph in the stack. The decision to move to the next case is driven by the visual properties of the stimulus [3], the probability of finding a target [18], the reward structure [19], the number of objects to scan in the stimulus [20], and the history of errors [17]. All of these are also factors common to search in other domains.

Modulation between exploration and exploitation is also fundamental to many non-spatial aspects of human behavior. Social search for potential mates (e.g., in marriage and divorce) may proceed from local search (e.g., people highly similar to the searcher) to global search as thresholds for mate quality are adjusted with age [21]. Humans searching for information in the external environment – information foraging (e.g., on the Web or through social networks) – also switch adaptively between local and global search: people leave a local patch of web pages when they perceive that its value has fallen below what can be found globally elsewhere [22]. In decision making, the search for cues on which to base choices often starts with those that have

been most useful in the past (local focus) and proceeds to others until enough cues have been found to select an option [23]. This pattern also arises in problem solving. For example, when engaged in an anagram task that involves making words from a set of random letters, people use the discovery of past solutions to determine how long to stay in a local ‘patch’ [24], as they also do when angling for fish in a sequence of virtual ponds on a computer screen [25]. Expecting resources near where (or when) other resources have been found may be a broadly adaptive cognitive bias – explaining, for example, the tendency for people and nonhuman primates to expect ‘streaky’ outcomes, sometimes called the hot-hand fallacy [26].

Exploration can be random, treating all courses of action more or less equally, or guided by beliefs about the structure of the environment [27–29]. The former, often called model-free, is driven by stimulus-response relationships and defines exploration as an increased probability of choosing an outcome associated with a lower expected value or of choosing any behavior with unknown consequences. Belief-guided search, often called model-based, involves a cognitive representation or map of the relational structure of the environment. Model-based decisions can inform both where and when to search by guiding search to regions that provide more information. An example of model-based ‘when’ search is exploring to see whether a cake is finished baking only after sufficient time has passed to allow it to do so. In the model-free case, there is no model of the temporal dynamics of baking, simply knowledge that baking cakes need checking. While model-free exploration is typically defined by choosing outcomes with lower expected rewards than exploitation, model-based exploration can lead to greater long-term rewards by using cognitive representations to decide both where and when exploration is most likely to pay dividends in the future [30,31]; from this perspective, learning can be conceived of as a foraging process [32].

Are internal and external search processes governed by common underlying mechanisms? Figure 2 illustrates a recent priming study that found that experience with more local exploitation in a spatial foraging task led to more focused production of word solutions in a subsequent anagram task [33]. The common mechanism is also supported by findings demonstrating that measures of executive capacity (e.g., working memory span) are related to search in various domains, both internal and external [34–37], including visual search [38,39]. Furthermore, dynamic switching between local and global search has been suggested to account for human decisions to move from one ‘region’ of memory to another – for example, when recalling types of animal in a verbal fluency task [6] – as well as decisions to leave visible patches (e.g., berry bushes in a simulated berry-picking task) [40].

The neurobiology of cognitive search

Cognitive search involves the capacity to maintain goals in stable neural representations and to relax these representations as needed to create opportunities for flexibility and exploration in either the internal or external world. These capacities allow us to search not only in the present, but also in the past and the future (Box 2).




Figure 2. Priming from external to internal search. **(A)** People initially searched for hidden targets (invisible to the participants, but shown in black) in environments with either clustered or diffuse resource distributions. The lower two panels show typical foraging patterns in the two environments. Following this period of foraging in external space, people's internal search behavior was quantified in a lexical search task. In the lexical search task, participants were asked to find 30 words across a series of letter sets, with multiple words possible per letter set. **(B)** A network representation of search in one letter set where participants found multiple words in the letter set NSBDOE. Nodes represent solutions and links between nodes represent transitions between solutions. Node size is proportional to the number of participants who provided that solution. Link thickness is proportional to the number of participants who made that transition. Bigram similarity between previous solutions and the original letter set shows a clear tendency to produce solutions that are more like the previous solution, $N - 1$, than like solutions two items back, $N - 2$, or the letter set. Priming is shown in panel **(C)**, with individuals in the clustered condition staying longer in letter sets (post-test–pretest time) than individuals in the diffuse condition. Reproduced from [33].

Research in neuroscience is uncovering core neural architectures for guiding search and mediating the exploration–exploitation trade-off – stability and flexibility – in relation to goal maintenance.

A host of neural circuits initiate and control search based on goals and motivations and the simultaneous identification of uncertainty with respect to their attainment. With the exploration–exploitation trade-off and its requisite goal maintenance and updating are the dorsolateral and medial prefrontal cortices [41,42]. These regions interact with numerous other brain structures associated with the control of attention, action evaluation, and outcome prediction, including the anterior cingulate cortex, the hippocampal formation, and the dorsal, ventral, medial and lateral aspects of the striatum. For example, the prefrontal areas interact with the hippocampal formation more predictably during flexible behaviors in rats [43–45] and humans [46,47] and this is reduced as habits form and dorsal and lateral aspects of the striatum learn action chains [48–50].

Although these interactions are governed by numerous neuromodulators, a key component is dopamine [51,52]. Dopamine is known to facilitate attentional control in response to unexpected stimuli [53], to modulate reward-seeking across species [9,54], and to manipulate goal maintenance and updating via phasic and tonic activation patterns [55,56]. How does dopamine mediate stability and flexibility in the brain? Dopamine receptor stimulation is proposed to increase the signal-to-noise ratio of neuronal firing [57]. This leads to increased robustness of goal

representations in the face of intervening distractors [58], maintaining the stability of goals and actions associated with their exploitation.

In addition to dopamine, the stabilization of search with respect to goal-relevant representations may also involve acetylcholine and noradrenaline (norepinephrine), which modulate levels of activation associated with expected and unexpected uncertainty [59], respectively – that is, whether the uncertainty originates from cues that are known to be unreliable or cues that are grossly inconsistent with the current context. In particular, noradrenaline has been proposed to mediate the shift from exploration to exploitation by changing the tonic background activity of neurons in the prefrontal cortex [60,61], thereby altering the threshold for new goal representations to compete for activation. When one course of action leads to less-than-expected rewards, neuromodulators in the brain lower the threshold for exploring alternatives [40,44]. Similarly, serotonin may also contribute to modulating search and, in particular, the propensity for model-based versus model-free search strategies [28,62].

The role of these brain regions and neurotransmitters naturally leads to predictions about pathologies of goal-directed and, specifically, search-related behavior associated with corresponding lesions or dysfunctions of neuromodulation. Many symptoms of Parkinson's disease, addiction, obsessive–compulsive disorder, schizophrenia, attention deficit disorder, and depression can be interpreted as deficits in search-related neural architectures and fall at the extremes of the exploration–exploitation trade-off [9,62,63].

Box 2. From foraging to deliberation

Cognitive search in humans offers the capacity to search within one's knowledge of the external world via simulation of that external world. This creates the capacity for deliberation. When faced with a complex and difficult choice, humans will deliberate over their options by constructing imagined future outcomes and evaluating them serially [97,98]. This is a search through an internal model of the world [99–101].

In the 1930s, Tolman observed that rats would sometimes pause and reorient back and forth at choice points and suggested that this behavior entailed a form of 'vicarious trial and error' – an internal search process [102]. It is now known that, during these behaviors, hippocampal representations encode future options [103] and hippocampal–prefrontal interactions increase [45,104], a process called episodic future thinking. Humans with damage to the hippocampal–prefrontal interacting system do not deliberate over choices and do not show these internal search processes [46,105].

The same trade-offs that impact external search processes affect this type of internal search. For example, should one continue to take time to explore possibilities (exploration) or should one stop and take the best choice found so far (exploitation)? Because the internal search is a computation, it takes time, which will slow one's interaction with the external world. Interestingly, there is strong evidence that mammalian decision making includes internal (cognitive) search, external (interactive) search processes, and cached (exploitation) processes [63]. Which process drives decision making depends on the time available for introspection and the information known about the world [106–108]. The advantage of internal search is that possible outcomes can be considered without the time commitment and inherent danger of implementing them. However, taking too much time to deliberate can be disastrous. At some point, the search (exploration) has to be stopped and the action (exploitation) taken. When one is truly time limited (for example, in sports), this internal, deliberative search process is often purposely bypassed by a less flexible (cached) action-selection system that has been learned through extensive training, often referred to as overlearned or automatic processing [108,109].

Collective problem solving in social species

Social systems also face the exploration–exploitation trade-off, but they mediate this trade-off in two distinct ways. Some social organisms, such as honeybees, and many ant species, delegate exploration and exploitation across different individuals so that individuals do not have to navigate the trade-off. Beginning with Karl von Frisch [64], biologists recognized that social insect colonies are

often divided into scouts who explore independently for new food sources, and workers who are recruited to exploit what others find (e.g., by pheromones in ants or the waggle dance in honeybees). The exploratory behavior of scouts is also modulated by dopaminergic agonists, but it is unclear to what extent this is homologous with the mechanisms described above [65]. Individuals also exhibit the flexibility to switch roles, with those failing to find recruitment signals within the colony switching to scouting behavior [66]. Consequently, the proportion of a colony participating in exploration increases if there are not enough existing foraging locations for the workers to exploit. Across species, the capacity for individuals to exploit resources discovered by others creates an additional exploration–exploitation trade-off beyond that posed for individuals – one that is a strength and a weakness of social systems (Box 3).

A second way that social systems can negotiate this trade-off is by altering the underlying connectivity between individuals and thus manipulating the ability of individuals to socially influence one another. The optimal level of connectivity also depends uniquely on the structure of the search problem. For search over smooth environments with single peaks, experiments on humans and simulations have found that fully connected groups can outperform less well-connected groups; however, as environments became more multi-peaked and rugged, less well-connected communication networks lead to longer periods of exploration and better long-term solutions [67,68]. Such diffuse communication systems may also resemble the foraging tactics of groups who can track moving resource gradients that are detectable collectively but are undetectable to individual foragers (Figure 3).

Both social–cognitive approaches to the exploration–exploitation trade-off – mixed roles across individuals and controlled communication connectivity between individuals – are found in cultural systems of human innovation. Federal systems rely on separate states to explore new policies and exploit what other states find [69]. Franchises experiment with new products and advertising plans in distinct locations and then carefully control the

Box 3. Social search

Many organisms utilize social information when searching for resources (such as food or an appropriate habitat) and when making decisions. In this way individuals can integrate personally-acquired information about their environment with social information, the latter based on the positions and behavior of others [74,110]. If effort is required to find items, organisms may limit their costs by exploiting the search effort of others. Thus, individuals can 'scrounge' information from 'producers' [111]. If too many individuals adopt a scrounging strategy, however, it pays to invest in personal search. This leads to a frequency-dependent balance of strategies within populations, although this often does not maximize population-level resource acquisition due to the game-theoretic nature of producer–scrounger behavior [111–114].

For search scenarios, organisms must detect – and climb – noisy, long-range environmental (e.g., temperature, salinity, resource) gradients. Here, social interactions can provide substantial additional benefit by allowing individuals, simply through grouping, to average their imperfect estimates of temporal and spatial cues (the so-called 'wisdom-of-crowds' effect) [115]. Due to the investment necessary to obtain personal information, however, this again sets the scene for

producers (searchers) to be exploited by others [116]. Alternatively, a recent experimental study of gradient detection in schooling fish discovered that distributed sensing, correlated with group size, facilitates adaptive responses to complex environmental gradients [74]. Thus, groups can exhibit a size-dependent exploratory awareness of the environment that is not present at the individual level.

In other organisms, such as the eusocial insects (e.g., ants, bees), additional communication strategies enhance search. Many ants deposit chemical trails during search, allowing others to determine which areas have previously been searched and where food has been found, along with (approximately) when, as indicated by the degree of trail evaporation. Thus, different chemicals employed simultaneously can guide potential searchers to engage in rapid, directed motion through areas of low interest and intensive search of other areas [117]. Similarly, honeybees modulate the search efforts of other group members through the waggle dance, which, performed within the nest, indicates to others the approximate direction, quality, and distance to resources. As with other search processes, some degree of error has been found to be adaptive in the dance, allowing an effective balance between exploration and exploitation [118,119].




Figure 3. Collective search allows individuals to track environmental gradients that occur at long length scales relative to the size of individuals, as well as to track gradients that they cannot individually detect. (A) The positions of 30 fish are shown as they negotiate, collectively, a moving light field toward the preferred darker regions. The snapshots are 2 s apart, with time progressing downward. (B) Performance, as measured by the time-averaged darkness level at fish locations, increases with group size. Experiments (points connected by dotted line) and simulations (red line) show this to be a function of individual fish simply manipulating their speed (moving slower in darker regions) while maintaining connectivity via locally mediated social interactions. Thus, counter to the conventional view of the ‘wisdom of crowds’, in which individuals pool imperfect estimates [120], here detection of the local light gradient is absent at the individual level (individuals employ only a scalar measure of the light intensity at their present location), yet exploration and exploitation emerge as a dynamic property of the collective. Reproduced from [74].

diffusion of these innovations [70]. Also, firms directly control roles and imitation by attempting to balance the exploration of new product-market domains with an exploitative focus on increasing existing product-market efficiency; firms better able to achieve this see higher sales growth [71]. Roles and connectivity can also be extended over time, to adapt to changing environments. For example, search for innovations by groups (including industries and nations) can be improved through balancing small bursts of social learning during periods of environmental change against long stretches of exploiting known strategies during periods of environmental stability [72]. In environments where solutions are hard to find but nonetheless indicated by other solutions, imitation can act as a form of group memory and enhance long-term performance by facilitating search around good solutions that have been

found in the past [73]. This is not unlike the search behavior observed among schooling fish [74].

In self-organized social domains, sometimes called coordination problems, people can self-assign roles to achieve the level of diversity in exploration and exploitation needed to track past success while simultaneously seeking better solutions [75]. Beneficial diversity can also come from differences in individual ability, but its usefulness depends on how information is aggregated across individuals and the kinds of problems they are solving [13,76,77]. Furthermore, exploration can be facilitated by dissent and ignorance: dissent can delay collective decisions, increasing the number of options considered by keeping exploratory discussion open [78], while individuals who participate in collective decisions with unbiased or uninformed preferences may tend to inhibit strongly opinionated minorities and promote a fair representation of opinions within groups [79]. Finally, in philosophy and political science, group deliberation is often viewed as a process of social search through a problem space in which different institutional rules affect the balance between exploring that space and exploiting the existing positions of individuals [80]. Thus, the capacity for different individuals to have different amounts and qualities of information extends the adaptive achievement of exploration–exploitation trade-offs from social foragers to more cognitively abstract domains of culture and policy.

Concluding remarks and future directions

This review provides a framework for envisioning cognition as a search process characterized by the exploration–exploitation trade-off, which scales from individuals to society (Figure 4). Viewing cognition as a search process highlights the potential conserved properties and problems, shared neural and cognitive mechanisms, and homologous or convergent evolutionary origins of many components of cognition [81]. Moreover, some components of cognitive search appear to be domain general, possibly facilitating the goal maintenance necessary for model-based internal foraging over varieties of cognitive representations [33].

Socially interacting groups regularly engage in search, acting as distributed cognitive systems to the extent that they collectively develop a cognitive representation of the environment that no one individual fully contains and behave based on that representation [82]. Consequently, they clearly face the same exploration–exploitation trade-off, which they can solve in novel ways such as altering the exploratory biases of individuals or altering the communication structure between individuals. Cognitive systems at the individual level utilize similar mechanisms. For example, individuals can alternate between exploration and exploitation in space and brains can alternate between exploration and exploitation by modulating connectivity via changes in tonic firing. Considering links between the individual and group levels invites questions about the extent to which social adaptations for the exploration–exploitation trade-off are built on cognitive faculties for mediating the trade-off within individuals. If spatial foraging provides a precursor for internal foraging, it may also have provided pre-adaptations for aspects of socially coordinated exploration.




Figure 4. Examples of cognitive approaches to the exploration–exploitation trade-off, as discussed in the text.

The cognitive search perspective also requires us to consider how cognitive representations of different types of search space – whether internal or external – interact with different possible search mechanisms (Box 4). Do some cognitive representations obey the same rules as physical 2D or 3D space or can those representations increase the dimensionality of search problems, in much the same way as genomic complexity increases the dimensionality of evolutionary search [83]? Does this increased dimensionality require new search strategies?

Box 4. Outstanding questions

- What cognitive representations of search spaces are there, what are their features, how do we switch between them, and how might they change over time?
- How do attentional processes guide information search in external and internal environments and how might this be related to other cognitive domains, such as categorization, learning, or language production?
- How are search processes biologically and cognitively adapted to the structure of different environments?
- Does role switching in social systems operate in the same way as role-switching decisions in the context of individual foraging? In addition, does a similar kind of group-level working memory capacity govern the exploration–exploitation trade-off in social systems?
- Can differences in exploration policies between social insect groups, franchises, and federal systems be linked to aspects of their internal structure?

A more complete understanding of cognitive search requires studies of the mechanisms for controlling and optimizing search as well as of how they interact with the structure of various internal and external environments. This will include understanding how search processes respond to both dynamic and noisy environments [84,85]. We also need to develop methodologies for disentangling changes in cognitive representations from changes in the search processes that navigate them. For example, is age-related cognitive decline in memory a consequence of problems with the search processes [37] or with changes in the memory representation, which may reflect degradation [86] or enrichment [87], or both?

Search also offers a framework for understanding cognitive problems in relation to their molecular, genetic, and neural underpinnings. Knowing how these building blocks influence search at one level of description can offer insights into why we see changes at other levels. As noted above, maladaptive states of both individual and group search lie at the extremes of too much exploitation (compulsiveness, perseveration, and groupthink) or too much exploration (impulsiveness, inattentiveness, and failure to leverage social information). Understanding how these extremes arise in response to age, pathology, genetics, environment, and social structure represents significant challenges to a wide range of fields in cognitive science – challenges for which a shared conceptual framework offers both insights and a basis for communicating our ideas and questions.

References

- 1 Cohen, J. *et al.* (2007) Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 362, 933
- 2 March, J.G. (1991) Exploration and exploitation in organizational learning. *Organ. Sci.* 2, 71–87
- 3 Treisman, A.M. and Gelade, G. (1980) A feature-integration theory of attention. *Cogn. Psychol.* 12, 97–136
- 4 Wolfe, J.M. *et al.* (1989) Guided search: an alternative to the feature integration model for visual search. *J. Exp. Psychol. Hum. Percept. Perform.* 15, 419–433
- 5 Benhamou, S. (2007) How many animals really do the Lévy walk? *Ecology* 88, 1962–1969
- 6 Hills, T.T. *et al.* (2012) Optimal foraging in semantic memory. *Psychol. Rev.* 119, 431–440
- 7 Korf, R.E. (1985) Depth-first iterative-deepening: an optimal admissible tree search. *Artif. Intell.* 27, 97–109
- 8 Sutton, R.S. and Barto, A.G. (1998) *Reinforcement Learning: An Introduction*, MIT Press
- 9 Hills, T.T. (2006) Animal foraging and the evolution of goal-directed cognition. *Cogn. Sci.* 30, 3–41
- 10 Stephens, D.W. *et al.* (2007) *Foraging: Behavior and Ecology*, University of Chicago Press
- 11 Bonabeau, E. *et al.* (1999) *Swarm Intelligence: From Natural to Artificial Systems*, Oxford University Press
- 12 Deneubourg, J.L. *et al.* (1986) Random behaviour, amplification processes and number of participants: how they contribute to the foraging properties of ants. *Physica D* 22, 176–186
- 13 King, A.J. *et al.* (2009) Is the true “wisdom of the crowd” to copy successful individuals? *Trends Cogn. Sci.* 13, 197–200
- 14 Hills, T.T. and Dukas, R. (2012) The evolution of cognitive search. In *Cognitive Search: Evolution, Algorithms, and the Brain* (Todd, P.M. *et al.*, eds), pp. 11–24, MIT Press
- 15 Charnov, E.L. (1976) Optimal foraging, the marginal value theorem. *Theor. Popul. Biol.* 9, 129–136
- 16 Stephens, D.W. and Krebs, J.R. (1987) *Foraging Theory*, Princeton University Press

- 17 Chun, M.M. and Wolfe, J.M. (1996) Just say no: how are visual searches terminated when there is no target present? *Cogn. Psychol.* 30, 39–78
- 18 Wolfe, J.M. *et al.* (2005) Rare items often missed in visual searches. *Nature* 435, 439–440
- 19 Kristjansson, A. *et al.* (2010) Fortune and reversals of fortune in visual search: reward contingencies for pop-out targets affect search efficiency and target repetition effects. *Atten. Percept. Psychophys.* 72, 1229–1236
- 20 Cousineau, D. and Shiffrin, R.M. (2004) Termination of a visual search with large display size effects. *Spat. Vis.* 17, 327–352
- 21 Hills, T. and Todd, P.M. (2008) Population heterogeneity and individual differences in an assortative agent-based marriage and divorce model (MADAM) using search with relaxing expectations. *J. Artif. Soc. Social Simul.* 11, 5
- 22 Pirolli, P. (2005) Rational analyses of information foraging on the web. *Cogn. Sci.* 29, 343–373
- 23 Gigerenzer, G. *et al.* (2012) Efficient cognition through limited search. In *Ecological Rationality: Intelligence in the World* (Todd, P.M. and Gigerenzer, G., eds), p. 608, Oxford University Press
- 24 Wilke, A. *et al.* (2009) Fishing for the right words: decision rules for human foraging behavior in internal search tasks. *Cogn. Sci.* 33, 497–529
- 25 Hutchinson, J.M.C. *et al.* (2008) Patch leaving in humans: can a generalist adapt its rules to dispersal of items across patches? *Anim. Behav.* 75, 1331–1349
- 26 Scheibehenne, B. *et al.* (2011) Expectations of clumpy resources influence predictions of sequential events. *Evol. Hum. Behav.* 32, 326–333
- 27 Dayan, P. and Berridge, K.C. (2014) Model-based and model-free Pavlovian reward learning: revaluation, revision, and revelation. *Cogn. Affect. Behav. Neurosci.* 14, 473–492
- 28 Blanco, N.J. *et al.* (2013) The influence of depression symptoms on exploratory decision-making. *Cognition* 129, 563–568
- 29 Otto, A.R. *et al.* (2014) Physiological and behavioral signatures of reflective exploratory choice. *Cogn. Affect. Behav. Neurosci.* 14, 1167–1183
- 30 Markant, D.B. and Gureckis, T.M. (2014) Is it better to select or to receive? Learning via active and passive hypothesis testing. *J. Exp. Psychol. Gen.* 143, 94–122
- 31 Sims, C.R. *et al.* (2013) Melioration as rational choice: sequential decision making in uncertain environments. *Psychol. Rev.* 120, 139
- 32 Metcalfe, J. and Jacobs, J. (2010) People's study time allocation and its relation to animal foraging. *Behav. Processes* 83, 213–221
- 33 Hills, T.T. *et al.* (2010) The central executive as a search process: priming exploration and exploitation across domains. *J. Exp. Psychol. Gen.* 139, 590–609
- 34 Hills, T.T. and Pachur, T. (2012) Dynamic search and working memory in social recall. *J. Exp. Psychol. Learn. Mem. Cogn.* 38, 218–228
- 35 Rakow, T. *et al.* (2008) Biased samples not mode of presentation: re-examining the apparent underweighting of rare events in experience-based choice. *Organ. Behav. Hum. Decis. Process.* 106, 168–179
- 36 Unsworth, N. (2007) Individual differences in working memory capacity and episodic retrieval: examining the dynamics of delayed and continuous distractor free recall. *J. Exp. Psychol. Learn. Mem. Cogn.* 33, 1020–1034
- 37 Hills, T. *et al.* (2013) Mechanisms of age-related decline in memory search across the adult life span. *Dev. Psychol.* 49, 2396–2404
- 38 Sobel, K. *et al.* (2007) Individual differences in working memory capacity and visual search: the roles of top-down and bottom-up processing. *Psychon. Bull. Rev.* 14, 840–845
- 39 Anderson, D.E. *et al.* (2013) A common discrete resource for visual working memory and visual search. *Psychol. Sci.* 24, 929–938
- 40 Wolfe, J.M. and Danielson, J. (2012) Visual foraging behavior: when are the berries riper on the other side of the screen? *J. Vis.* 12, 265
- 41 Fuster, J. (2008) *The Prefrontal Cortex*, Elsevier
- 42 Goldman-Rakic, P.S. (1995) Cellular basis of working memory. *Neuron* 14, 477–485
- 43 Jones, M.W. and Wilson, M.A. (2005) Theta rhythms coordinate hippocampal–prefrontal interactions in a spatial memory task. *PLoS Biol.* 3, e402
- 44 Hyman, J.M. *et al.* (2010) Working memory performance correlates with prefrontal–hippocampal theta interactions but not with prefrontal neuron firing rates. *Front. Integr. Neurosci.* 4, 1–13
- 45 Benchenane, K. *et al.* (2010) Coherent theta oscillations and reorganization of spike timing in the hippocampal–prefrontal network upon learning. *Neuron* 66, 921–936
- 46 Voss, J.L. *et al.* (2010) Hippocampal brain-network coordination during volitional exploratory behavior enhances learning. *Nat. Neurosci.* 14, 115–120
- 47 Schacter, D.L. and Addis, D.R. (2009) On the nature of medial temporal lobe contributions to the constructive simulation of future events. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 364, 1245–1253
- 48 Barnes, T.D. *et al.* (2005) Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. *Nature* 437, 1158–1161
- 49 Smith, K.S. and Graybiel, A.M. (2013) A dual operator view of habitual behavior reflecting cortical and striatal dynamics. *Neuron* 79, 361–374
- 50 Dezfouli, A. and Balleine, B.W. (2012) Habits, action sequences and reinforcement learning. *Eur. J. Neurosci.* 35, 1036–1051
- 51 Robbins, T. and Roberts, A. (2007) Differential regulation of fronto-executive function by the monoamines and acetylcholine. *Cereb. Cortex* 17, i151–i160
- 52 Cools, R. and D'Esposito, M. (2011) Inverted-U-shaped dopamine actions on human working memory and cognitive control. *Biol. Psychiatry* 69, e113–e125
- 53 Schultz, W. (1998) Predictive reward signal of dopamine neurons. *J. Neurophysiol.* 80, 1–27
- 54 Barron, A.B. *et al.* (2010) The roles of dopamine and related compounds in reward-seeking behavior across animal phyla. *Front. Behav. Neurosci.* 4, 1–9
- 55 Redish, A.D. *et al.* (2007) Reconciling reinforcement learning models with behavioral extinction and renewal: implications for addiction, relapse, and problem gambling. *Psychol. Rev.* 114, 784–805
- 56 Beeler, J.A. *et al.* (2010) Tonic dopamine modulates exploitation of reward learning. *Front. Behav. Neurosci.* 4, 1–14
- 57 Servan-Schreiber, D. *et al.* (1990) A network model of catecholamine effects: gain, signal-to-noise ratio, and behavior. *Science* 249, 892–895
- 58 Durstewitz, D. and Seamans, J.K. (2008) The dual-state theory of prefrontal cortex dopamine function with relevance to catechol-O-methyltransferase genotypes and schizophrenia. *Biol. Psychiatry* 64, 739–749
- 59 Yu, A.J. and Dayan, P. (2005) Uncertainty, neuromodulation, and attention. *Neuron* 46, 681–692
- 60 Aston-Jones, G. and Cohen, J.D. (2005) An integrative theory of locus coeruleus–norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.* 28, 403–450
- 61 Kehagia, A.A. *et al.* (2010) Learning and cognitive flexibility: frontostriatal function and monoaminergic modulation. *Curr. Opin. Neurobiol.* 20, 199–204
- 62 Huys, Q.J.M. *et al.* (2012) Bonsai trees in your head: how the Pavlovian system sculpts goal-directed choices by pruning decision trees. *PLoS Comput. Biol.* 8, e1002410
- 63 Winstanley, C.A. *et al.* (2012) Search, goals, and the brain. In *Cognitive Search: Evolution, Algorithms, and the Brain* (Todd, P.M. *et al.*, eds), pp. 125–156, MIT Press
- 64 Frisch von, K. (1967) *The Dance Language and Orientation of Bees*, Harvard University Press
- 65 Liang, Z.S. *et al.* (2012) Molecular determinants of scouting behavior in honey bees. *Science* 335, 1225–1228
- 66 Beekman, M. *et al.* (2007) What makes a honeybee scout? *Behav. Ecol. Sociobiol.* 61, 985–995
- 67 Mason, W.A. *et al.* (2008) Propagation of innovations in networked groups. *J. Exp. Psychol. Gen.* 137, 422–433
- 68 Lazer, D. and Friedman, A. (2007) The network structure of exploration and exploitation. *Admin. Sci. Quart.* 52, 667–694
- 69 Kollman, K. (2000) Decentralization and the search for policy solutions. *J. Law Econ. Organ.* 16, 102–128
- 70 Sorenson, O. and Sorenson, J.B. (2001) Finding the right mix: franchising, organizational learning, and chain performance. *Strat. Mgmt J.* 22, 713–724

- 71 He, Z.-L. and Wong, P.-K. (2004) Exploration vs. exploitation: an empirical test of the ambidexterity hypothesis. *Organ. Sci.* 15, 481–494
- 72 Rendell, L. et al. (2010) Why copy others? Insights from the social learning strategies tournament. *Science* 328, 208–213
- 73 Wisdom, T.N. et al. (2013) Social learning strategies in networked groups. *Cogn. Sci.* 37, 1383–1425
- 74 Berdahl, A. et al. (2013) Emergent sensing of complex environments by mobile animal groups. *Science* 339, 574–576
- 75 Roberts, M.E. and Goldstone, R.L. (2011) Adaptive group coordination and role differentiation. *PLoS ONE* 6, e22377
- 76 Page, S.E. (2007) *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*, Princeton University Press
- 77 Krause, S. et al. (2011) Swarm intelligence in humans: diversity can trump ability. *Anim. Behav.* 81, 941–948
- 78 Nemeth, C.J. and Wachtler, J. (1983) Creative problem solving as a result of majority vs minority influence. *Eur. J. Soc. Psychol.* 13, 45–55
- 79 Couzin, I.D. et al. (2011) Uninformed individuals promote democratic consensus in animal groups. *Science* 334, 1578–1580
- 80 Habermas, J. and McCarthy, T. (1985) *A Theory of Communicative Action*, Beacon Press
- 81 Todd, P.M. et al., eds (2012) *Cognitive Search: Evolution, Algorithms, and the Brain (Strüngmann Forum Reports)*, MIT Press
- 82 Goldstone, R.L. and Gureckis, T.M. (2009) Collective behavior. *Top. Cogn. Sci.* 1, 412–438
- 83 Gavrillets, S. (1997) Evolution and speciation on holey adaptive landscapes. *Trends Ecol. Evol.* 12, 307–312
- 84 Gureckis, T.M. and Love, B.C. (2009) Learning in noise: dynamic decision-making in a variable environment. *J. Math. Psychol.* 53, 180–193
- 85 Otto, A. et al. (2009) Navigating through abstract decision spaces: evaluating the role of state generalization in a dynamic decision-making task. *Psychon. Bull. Rev.* 16, 957–963
- 86 Borge-Holthoefer, J. et al. (2011) Modeling abnormal priming in Alzheimer's patients with a free association network. *PLoS ONE* 6, e22651
- 87 Ramscar, M. et al. (2014) The myth of cognitive decline: non-linear dynamics of lifelong learning. *Top. Cogn. Sci.* 6, 5–42
- 88 Nonacs, P. (2001) State dependent behavior and the marginal value theorem. *Behav. Ecol.* 12, 71–83
- 89 Beachly, W.M. et al. (1995) On the economics of sit-and-wait foraging: site selection and assessment. *Behav. Ecol.* 6, 258–268
- 90 Kareiva, P. and Odell, G. (1987) Swarms of predators exhibit 'prey taxis' if individual predators use area-restricted search. *Am. Nat.* 130, 233–270
- 91 Viswanathan, G. et al. (1999) Optimizing the success of random searches. *Nature* 401, 911–914
- 92 Hills, T.T. et al. (2013) Adaptive Lévy processes and area-restricted search in human foraging. *PLoS ONE* 8, e60488
- 93 Ferreira, A.S. et al. (2012) The influence of the environment on Lévy random search efficiency: fractality and memory effects. *Physica A* 391, 3234–3246
- 94 Klein, R. (1988) Inhibitory tagging system facilitates visual search. *Nature* 334, 430–431
- 95 Peterson, M.S. et al. (2001) Visual search has memory. *Psychol. Sci.* 12, 287–292
- 96 Wolfe, J.M. et al. (2000) Attention is fast but volition is slow. *Nature* 406, 691
- 97 Buckner, R.L. and Carroll, D.C. (2007) Self-projection and the brain. *Trends Cogn. Sci.* 11, 49–57
- 98 Redish, A.D. (2013) *The Mind within the Brain: How We Make Decisions and How those Decisions Go Wrong*, Oxford University Press
- 99 Kurth-Nelson, Z. et al. (2012) A theoretical account of cognitive effects in delay discounting. *Eur. J. Neurosci.* 35, 1052–1064
- 100 Otto, A.R. et al. (2013) Working-memory capacity protects model-based learning from stress. *Proc. Natl. Acad. Sci. U.S.A.* 110, 20941–20946
- 101 Genovesio, A. et al. (2014) Prefrontal–parietal function: from foraging to foresight. *Trends Cogn. Sci.* 18, 72–81
- 102 Tolman, E.C. (1938) The determiners of behavior at a choice point. *Psychol. Rev.* 45, 1–41
- 103 Johnson, A. and Redish, A.D. (2007) Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J. Neurosci.* 27, 12176–12189
- 104 Schacter, D.L. and Addis, D.R. (2007) The cognitive neuroscience of constructive memory: remembering the past and imagining the future. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 362, 773–786
- 105 Hassabis, D. et al. (2007) Patients with hippocampal amnesia cannot imagine new experiences. *Proc. Natl. Acad. Sci. U.S.A.* 104, 1726–1731
- 106 Keramati, M. et al. (2011) Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLoS Comput. Biol.* 7, e1002055
- 107 Johnson, A. et al. (2012) The hippocampus and exploration: dynamically evolving behavior and neural representations. *Front. Hum. Neurosci.* 6, 216
- 108 van der Meer, M. et al. (2012) Information processing in decision-making systems. *Neuroscientist* 18, 342–359
- 109 Daw, N.D. et al. (2005) Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* 8, 1704–1711
- 110 van Bergen, Y. et al. (2004) Nine-spined sticklebacks exploit the most reliable source when public and private information conflict. *Proc. Biol. Sci.* 271, 957–962
- 111 Barnard, C.J. and Sibly, R.M. (1981) Producers and scroungers: a general model and its application to captive flocks of house sparrows. *Anim. Behav.* 29, 543–550
- 112 Vickery, W.L. et al. (1991) Producers, scroungers, and group foraging. *Am. Nat.* 137, 847–863
- 113 Giraldeau, L.-A. and Caraco, T. (2000) *Social Foraging Theory*, Princeton University Press
- 114 Giraldeau, L.A. et al. (2002) Potential disadvantages of using socially acquired information. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 357, 1559–1566
- 115 Simons, A.M. (2004) Many wrongs: the advantage of group navigation. *Trends Ecol. Evol.* 19, 453–455
- 116 Guttal, V. and Couzin, I.D. (2010) Social interactions, information use, and the evolution of collective migration. *Proc. Natl. Acad. Sci. U.S.A.* 107, 16172–16177
- 117 Dussutour, A. et al. (2009) The role of multiple pheromones in food recruitment by ants. *J. Exp. Biol.* 212, 2337–2348
- 118 Deneubourg, J.L. et al. (1983) Probabilistic behaviour in ants: a strategy of errors? *J. Theor. Biol.* 105, 259–271
- 119 Weidenmüller, A. and Seeley, T.D. (1999) Imprecision in waggle dances of the honeybee (*Apis mellifera*) for nearby food sources: error or adaptation? *Behav. Ecol. Sociobiol.* 46, 190–199
- 120 Galton, F. (1907) Vox populi. *Nature* 75, 450–451