

Convergent and Divergent Operations in Cognitive Search

Bernhard Hommel

Abstract

Human goal-directed action emerges from the interaction between stimulus-driven sensorimotor online systems and slower-working control systems that relate highly processed perceptual information to the construction of goal-related action plans. This distribution of labor requires the acquisition of enduring action representations; that is, of memory traces which capture the main characteristics of successful actions and their consequences. It is argued here that these traces provide the building blocks for off-line prospective action planning, which renders the search through stored action representations an essential part of action control. Hence, action planning requires cognitive search (through possible options) and might have led to the evolution of cognitive search routines that humans have learned to employ for other purposes as well, such as searching for perceptual events and through memory. Thus, what is commonly considered to represent different types of search operations may all have evolved from action planning and share the same characteristics. Evidence is discussed which suggests that all types of cognitive search—be it in searching for perceptual events, for suitable actions, or through memory—share the characteristic of following a fixed sequence of cognitive operations: divergent search followed by convergent search.

Introduction

The phylogenetic development of humans and many other species is characterized by a transition from reactivity to proactivity. In contrast to the dominant experimental paradigm in behavioral and neuroscientific research—where the presentation of experimenter-controlled stimuli marks the starting point of theoretical and empirical analysis—humans rarely await environmental triggers to get going but are instead driven by internal needs, goals, and passions. Proactive behavior guided by internal anticipations and predictions requires executive functions that operate off-line rather than in real time. Indeed, beginning to plan an action only after its execution conditions have been

encountered often makes little sense, which is why, for instance, goalkeepers prepare themselves for jumping, catching, and pushing the ball long before they see it coming.

Planning an action in the absence of the object as well as the situational cues to which it relates requires the means to represent and simulate them in advance (i.e., off-line). This calls for cognitive abilities that go beyond what our basic sensorimotor online systems have to offer (i.e., the systems that translate stimulus input into motor output more or less in real time and that we share to some degree with almost all living organisms). In humans and perhaps in other primates, this reliance of planning ahead on longer-term internal representation has promoted the development of a dual-system architecture. In humans, for example, manual actions emerge from the interaction between a stimulus-driven sensorimotor online system (mediated by the dorsal pathway; e.g., Milner and Goodale 1995) and a slower-working (ventral) control system that relates highly processed perceptual information to the construction of goal-related action plans (Glover 2004; Hommel et al. 2001a, b). The emergence of off-line systems did not lead to the replacement of online systems; having both systems provides the opportunity to restrict off-line action planning to the specification of the goal-relevant features of an action but to leave the filling-in of the less important motoric details to the sensorimotor online system (Turvey 1977).

This distribution of labor requires the acquisition of enduring action representations; that is, of memory traces that capture the main characteristics of successful actions and their consequences. I will argue that these traces provide the building blocks for off-line action planning, which renders the search through stored action representations an essential part of action control. In other words, action planning requires cognitive search (through possible options) and might have led to the evolution of cognitive search routines that we now can also employ for other purposes, such as searching for perceptual events and through memory. Thus, what are commonly considered to be different types of search operations may all have evolved from action planning, suggesting that they might share the same characteristics.¹ Indeed, I will argue that all types of cognitive search (be it in searching for perceptual events, for suitable actions, or through memory) share the characteristic of following a fixed sequence of cognitive operations.

¹ From the perspective of a cognitive system, there is no logical difference between searching for a visual target; searching through (i.e., retrieving) one's memory of one's last birthday present; searching for (i.e., selecting) an action alternative suitable to produce a particular song by means of a keyboard; searching for a particular metaphor; or searching for the possible uses of a pen. In all these cases, the searcher consults an internal representation of past and/or present events and matches their content against some sort of search template, which represents the target or intended outcome, until some reasonable match is obtained. In that sense, there is no logical reason to believe that the cognitive operations underlying visual search, memory search, action selection, or the production of options in verbal planning, creativity, or problem-solving tasks differ in principle, and I know of no empirical evidence that enforces this assumption.

Searching for Suitable Actions

Evolution does not care so much about the deep insights and interesting memories an organism may or may not have, but selects instead for appropriate actions. So how do we identify and choose appropriate actions, and how do we tailor them flexibly to the current situational requirements? In view of the dominant stimulus-response paradigm in the experimental analysis of action control and other cognitive processes, it is not surprising that many theoretical approaches to action selection attribute the greatest responsibility in the selection process to the stimulus. Preparing for an action is viewed as the activation of task-relevant stimulus-response rules or associations, which then make sure that processing a task-relevant stimulus leads to the spreading of activation to the associated response code so that this code is likely to win the internal competition for controlling the output (e.g., Anderson 1993; Cohen et al. 1990; Kornblum et al. 1990). Obviously, this approach presupposes either extensive experience of the agent with the task at hand or some sort of instruction describing which rules are acceptable in a given situation—much like in the standard experimental setup. How people choose actions under less constrained conditions and how they can ever act in the absence of stimuli remains unclear (Hommel et al. 2001b).

Carrying out an action presupposes the existence of a goal,² the intention to create a particular outcome. This requires some sort of anticipation regarding the action's outcome, some expectation that the action will produce particular effects, and some motivation to produce them. The question of how these anticipations are created and how they guide the eventual selection of one concrete action has been addressed by two different approaches: the *ideomotor approach*, which focuses on the perceptual aspects of action outcomes, and the *motivational approach*, which emphasizes their affective implications.

The ideomotor approach to goal-directed action (James 1890; for an overview, see Shin et al. 2010) assumes that agents automatically register the perceptual consequences of their movements and integrate the motor patterns underlying the movements with the representations of the consequences they produce (see Figure 14.1, left panel). Hence, moving in one's environment is assumed to lead to the acquisition of bidirectional associations between movement patterns and codes of their perceptual outcomes. Given this bidirectionality, agents can then use the associations in either direction and thus intentionally reactivate a particular motor pattern by anticipating ("thinking of") its

² In this chapter, I make the assumption that all sorts of cognitive search, including perceptual search and action selection, are under the control of goals. How, according to which principles, and by what kind of mechanism goals are searched for and identified will not be discussed. However, it may well be that what I consider a goal is no more than the next level of what is actually a multilayered search-driven decision-making hierarchy. Accordingly, goals might be selected according to the same principles, and by means of the same mechanisms, that are involved in selecting a visual target or a manual action.

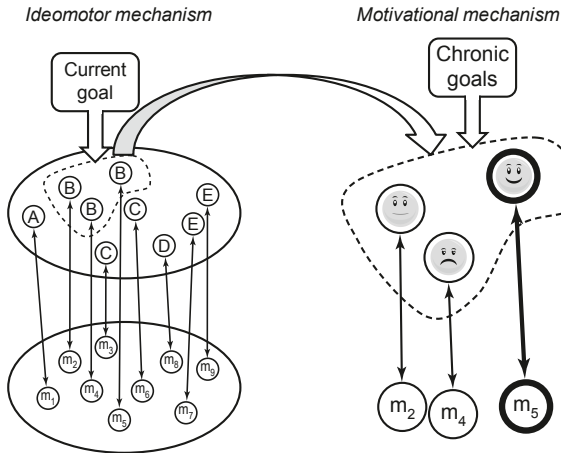


Figure 14.1 Distribution of labor between (A) the preselection of actions (from motor patterns, $m_1 - m_9$) associated with goal-related action effects (B in this example) and (B) the weighting of the preselected actions according to the state they are expected to produce. Chronic goals bias this process toward alternatives which satisfy them; in this example, m_5 is the most likely candidate for execution.

sensory consequences. The ideomotor approach has received ample empirical support. Novel action-produced perceptual effects are indeed spontaneously acquired and integrated with the corresponding action in adults, children, and infants, so that effect-related stimuli become effective primes of that action (for an overview, see Hommel 2009). Brain-imaging studies suggest that the hippocampus provides the bidirectional link between action plans stored in and/or generated by the supplementary motor area and the perceptual representations of action effects in the respective sensory cortices (Elsner et al. 2002; Melcher et al. 2008).

According to the ideomotor approach, translating an intended goal into actual action requires the cognitive representation of the desired sensory consequences or, more precisely, of the sensory implications of the desired effect. Once this representation has been formed or activated, the first step of action selection can be considered a kind of feature match: the desired outcome's sensory consequences (i.e., the description of the action goal) can be matched against the sensory consequences of all the actions in the agent's repertoire (see Figure 14.1, left panel). The result of this matching operation is the activation of all candidate actions that would be suited to create the intended effect in principle, or at least an effect that is perceptually similar.

Identifying action opportunities is commonly not part of experimental analyses, where the options are almost always specified by the task and/or arbitrarily defined by the experimenter. Accordingly, it is not surprising that this aspect of action selection is not very well understood while much more is known about selection of actions from prespecified response sets. Apart from

the above mentioned stimulus-centered accounts, which assume some sort of translation of stimulus information into response activation, research on this topic has emphasized two (related) motivational criteria that underlie action selection: reward and efficiency. Generations of learning theorists have pointed out that carrying out some actions provides more reward than carrying out others, and that this is likely to affect the probability with which an action is selected. Recent neuroscientific findings have provided strong support for the idea that action selection is systematically biased by the anticipation of reward or punishment (Schultz 2006) and/or the related affective states (Damasio 1996). Another line of research that has focused on the impact of efficiency on action selection showed that agents prefer action variants that imply less cognitive effort (e.g., Kool et al. 2010) and metabolic cost (e.g., Chapman et al. 2010). If one considers that both reward and efficiency correspond to something like chronic goals and that they are likely to be correlated with specific affective states, these findings seem to fit with the assumption that the anticipation of reward and/or positive affect biases decision making toward the associated action (see Figure 14.1, right panel).

Ideomotor and motivational approaches capture important aspects of the internal search for the action that is best suited to reach an intended goal. Interestingly, the purposes that ideomotor and motivational processes seem to serve are complementary (de Wit and Dickinson 2009): defining which actions would be suited to reach a particular goal (the purpose of ideomotor mechanisms) does not itself provide sufficient criteria for making the eventual selection, whereas comparing candidate goals with respect to the reward they may provide or the effort they require (the purpose of motivational mechanisms) presupposes some rather limited set of action alternatives that are all suitable in principle. This suggests that ideomotor and motivational mechanisms operate in a sequence, as indicated in Figure 14.1, with motivational mechanisms selecting from the set provided by ideomotor mechanisms.

In the present context, it is important to note that this suggested sequence of operations implies a succession of two rather different search modes. Ideomotor mechanisms start with one representation, the description of the goal, and try to diverge and activate as many perceptually related representations as possible. In contrast, motivational mechanisms start with a limited number of representations and then try to converge onto one optimal solution. In the following sections, I will discuss evidence which suggests that

1. convergent and divergent search operations can be found and distinguished in various types of cognitive search, including the search for perceptual targets and the search through memory for problem solutions;
2. all these types of search are likely to consist of a fixed sequence of divergent search operations followed by convergent search;

3. convergent and divergent search are likely to require different configurations of cognitive control.

To substantiate these claims, I will proceed by discussing evidence for convergent and divergent search operations in the context of searching for to-be-perceived objects (e.g., as in visual search), and in the context of searching through memorized objects and events (e.g., as with problem solving). I will conclude by suggesting a rudimentary control architecture that may underlie convergent and divergent search and present some evidence supporting this suggestion.

Searching for Perceived Targets

People tend to spend a great deal of their time searching for objects and other people—just think of parents looking for their kids, or scientists looking for a particular paper that they could swear was on their desk a few minutes ago. Searching for external events has been mostly studied in the visual modality, and there is consensus that at least two different types of visual search exist: feature search (e.g., looking for a red target among green distractors) and conjunction search (e.g., looking for a green X—i.e., the conjunction of the color green and the shape X—among red Xs and green Os) (Wolfe 1994; Wolfe, this volume). These two types of search differ in ease and efficiency: searching for a feature goes fast and is not much affected by the number of distractors, whereas searching for a feature conjunction is slow and highly sensitive to the number of distractors. These different characteristics have motivated the assumption that feature search can proceed in parallel and in a more or less bottom-up fashion, whereas conjunction search requires serial operations that are controlled top-down.

Perceptual search processes are commonly studied and theoretically addressed under complete neglect of action-related processes. The underlying idea is that the control of perceptual search is primarily input control whereas action-related processes deal with output control: two types of control that most researchers consider independent and unrelated (Johnston et al. 1995). Recent observations, however, tend to undermine this implicit conviction. As summarized elsewhere (Hommel 2010), a number of findings suggest that the efficiency of searching for a particular feature depends on the action carried out to signal the presence of the target or of actions that are being planned in the context of the search operation. For instance, searching for shape-defined targets is more efficient after preparing a grasping action, whereas searching for location- or intensity-defined targets is more efficient after preparing a pointing action (Fagioli et al. 2007; Wykowska et al. 2009). Hence, visual search is modulated by, and thus cannot be independent of, action planning.

The reason why action planning and search are related may have to do with the distribution of labor between the off-line perception-action system and the online sensorimotor system discussed above (Hommel 2010). While this distribution makes action planning and execution more flexible, it also raises a number of serious control problems. For instance, how does the sensorimotor system know which information is relevant for steering the motor activity selected by the perception-action system? As Wykowska et al. (2009) and Hommel (2010) have suggested, this problem might be solved by the perception-action system through increasing the gain of feature information coming from action-relevant feature dimensions (see Figure 14.2). For example, when preparing for a grasp, the perception-action system might increase the weight given to feature values coming from feature maps coding for shape and orientation (in addition to some weighting of location codes to end up at the right place), whereas preparing for a pointing action might lead to comparatively stronger weighting of feature values coded on location maps. Interestingly, the stronger weighting of stimulus attributes coded on feature maps that provide task-related information has been assumed to be part of the mechanism underlying the attentional control of visual search processes (e.g., Found and Müller 1996; Wolfe et al. 2003). If the functionality attributed to this mechanism is what action control provides, it makes sense to assume that what we call visual

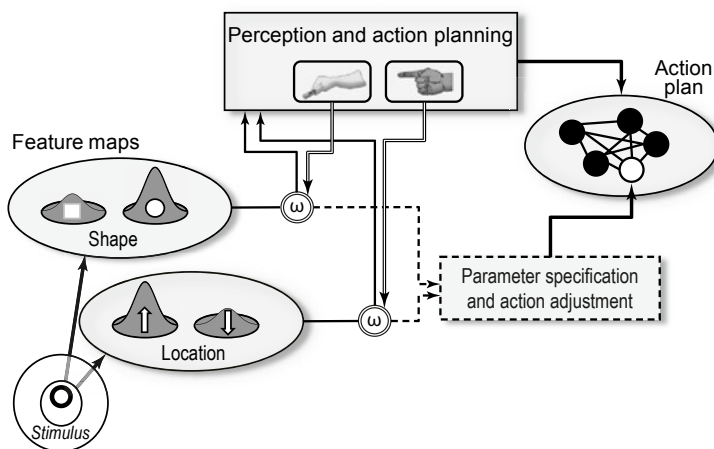


Figure 14.2 A process model of action-induced attention (after from Hommel 2010). Feature maps provide information both for off-line perception and action planning and for online specification of current action parameters. Perception provides contextual information, and action planning prepares an action plan with some parameters specified in advance (forward modeling; see black nodes) and others left for online specification (white nodes). To make certain that online specification uses appropriate information, the perception-action system modulates the output gain ω from the feature maps, so that information from goal-relevant feature maps has more impact on sensorimotor processing.

attention is a phylogenetic derivative of action control (Hommel 2010), which again would render observations of interactions between action control and attention less surprising than it might seem.

If visual attention is really an evolutionary by-product of improving action control mechanisms, one would expect perceptual search processes to show the same characteristics as action control. Thus, if the search for the right action proceeds through a sequence of divergent and convergent search operations, one would expect the same sequence for perceptual search. That seems to be far-fetched at first sight, especially if we consider the classical paradigms employed to study feature and conjunction search. Take, for instance, a display in a typical feature-search task: In what sense would searching for a red circle surrounded by twenty green circles require any sequence of divergent and convergent processes?

Single-cell recordings in monkeys provide considerable evidence for such a sequence (Lamme and Roelfsema 2000). Facing a number of stimuli is assumed to trigger a nonselective (i.e., not yet attentionally modulated) spread of neural activation throughout the visual cortex all the way up to frontal areas—the so-called “fast feedforward sweep.” It is so fast that after about 100 ms, even the highest levels of visual coding (i.e., brain systems coding for complex stimulus characteristics and stimulus identities) have responded to a presented stimulus. Neuroscientific methods allowed for following the spread of stimulus-induced activation throughout the entire brain and revealed that the speed of spreading is mainly determined by the brain’s hierarchical structural and functional architecture—with each layer adding about 10 ms (Lamme and Roelfsema 2000; Tovee 1994). Neurally speaking, the fast feedforward sweep can be considered decidedly divergent, as it activates as many stimulus-related representations as possible, presumably including various alternative interpretations of a given stimulus (Marcel 1983): it also activates both representations of currently relevant, attended stimuli and stimulus features as well as representations of irrelevant stimuli and features to the same extent. Even so, this nonselective spread of information might well be sufficient for performing a number of tasks, such as the detection of the presence of a particular feature (Treisman and Gelade 1980).

The visual fast feedforward sweep is reliably followed by a second phase of neural activation with entirely different characteristics. This so-called “recurrent” processing wave works its way back to early visual areas and differentiates relevant and irrelevant (attended and unattended) information by selectively enhancing that part of the sweep-induced activation that relates to the relevant stimulus (features) (e.g., Chelazzi et al. 1993; Lamme and Spekreijse 1999). This recurrent wave is apparently necessary for the emergence of conscious representations (Lamme 2003) as well as for the segregation and integration of stimulus features (Lamme 2003; Lamme and Roelfsema 2000). This implies that the fast feedforward sweep may often be sufficient to detect particular features but that searching for feature conjunctions requires recurrent

processing. If we consider that the latter serves to integrate stimulus features, this scenario fits perfectly with the feature integration theory (FIT) suggested by Treisman and Gelade (1980). As FIT implies, searching for a feature may be mastered by monitoring the activation level of dedicated feature maps. If the monitored level increases during the feedforward sweep, the participant does not need to await the recurrent processing wave to give a response. This can explain why the search for simple features is often fast and insensitive to the number of distractors (for additional views, see Wolfe, this volume). When searching for conjunctions of features, however, detecting the presence of a particular feature is insufficient. Rather, the features making up the conjunction would need to be integrated, which, according to FIT, is a serial process; thus, search time increases with the number of visible objects being considered. If we consider that the recurrent processing wave is selective and converging onto one given object, conjunction search may indeed require a whole sequence of convergence operations (i.e., a sequence of recurrent waves targeting alternative objects).

As can be seen, conjunction search can be characterized as a sequence of divergent processing (the fast feedforward sweep) followed by a convergent processing (the recurrent wave). But what about feature search? It is interesting to note that this kind of search does not really capture the ecological essence of everyday search performance. People are commonly looking for objects (or people) that in some cases may have features with a particular pop-out quality but are not selectively defined by them. Hence, we rarely search for single features. If we have to serially process a visual scene to locate a conjunctively defined target, we do not scan the scene randomly but are instead guided by features that are part of the conjunction (Wolfe 1994). This suggests that the main function of the divergent feedforward sweep is to determine the feature database which the following convergent operations can use to home in on possible targets. We can thus conclude that at least the bulk of everyday visual search can be aptly characterized as a sequence comprising a divergent, stimulus-driven spread of activation—an operation that seems to serve the purpose of identifying as many candidate targets as possible—followed by a convergent, goal-driven selection of one specific event representation. Exactly the same sequence is seen in the case of searching for appropriate actions.

Searching for Solutions

Even though the ultimate purpose of selecting to-be-perceived targets and to-be-produced actions relates to external, environmental states of affairs, the search operations involved are without exception targeting internal representations. As we have seen, some characteristics seem to be shared by both search operations aiming at representations of currently perceived events, as in visual search, and search operations aiming at representations of future events, as in

action planning. Given that the representations themselves are likely to differ in various ways, this commonality in search is remarkable and suggests that the characteristics we can identify are not restricted to direct interactions with our environment.

Indeed, divergent and convergent operations can also be found in problem solving and similar entirely internal search processes. They are particularly obvious in the domain of creative thinking. Even though the importance of human creativity cannot be overestimated, the processes underlying it are understudied and poorly understood (Sternberg et al. 2002). In part, this is due to strong traditions in this field that either focus on creativity as a personal trait—hence, as a characteristic that a given person does or does not have—or emphasize the products, rather than the functional characteristics, of the creative process (see Brown 1989; Runco 2007). Only more recently have researchers begun to agree that truly creative acts do not reflect the operation of just one process, brain area, or intellectual faculty but rather the interplay of multiple cognitive processes and neural networks (e.g., Dietrich 2004; Eysenck 1993; Heilman 2005). Still, there is no agreement as to what these processes and networks might be and how they are to be identified.

Guilford (1967) was one of the first to distinguish between two basic types of thinking that might underlie creative acts: (a) divergent thinking serves the purpose of producing as many possible solutions for a given problem as possible and (b) convergent thinking serves to find the one best solution to problems that require the satisfaction of multiple constraints. Two classical tasks provide good examples: The *alternate uses task* (Guilford 1967) requires participants to name as many appropriate uses of a simple object, such as a pen, as possible, which calls for a literal “brainstorm” through memory. In contrast, the *remote associations task* (Mednick 1962) presents participants with three concepts (e.g., time, hair, and stretch) per trial, who are then asked to identify the one concept that is best related to all three (long).

Unfortunately, the distinction between divergent and convergent thinking is seldom heeded in creativity studies. Instead of studying both types of process together, they often employ only divergent tasks (for overviews and discussion, see Baas et al. 2008; Davis 2009) or convergent tasks (e.g., Isen et al. 1987) or ad-hoc developed and difficult-to-categorize tasks to study “the creativity” (for an overview, see Plucker and Makel 2010). This seems particularly problematic as divergent and convergent thinking not only differ with respect to their computational goals but also seem to rely on different functional and neural mechanisms (cf. Dietrich 2004). A first hint is provided by the observation that individual convergent-thinking performance is not correlated with divergent-thinking performance (Akbari Chermahini and Hommel 2010) and that performing convergent- and divergent-thinking tasks induce opposite mood states (Akbari Chermahini and Hommel 2012). Moreover, there is evidence that divergent-thinking performance relates to the individual dopamine level of participants in the form of an inverted U-shape, with medium levels allowing

for the best performance, whereas convergent-thinking performance shows a linear, negative relationship with dopamine levels (Akbari Chermahini and Hommel 2010).

Even though many researchers have focused on one type of creative thinking or the other, entire creative acts (e.g., the invention of a new tool or the conception of a new painting) are likely to require both: a first phase of brainstorming that identifies as many options as possible and a subsequent phase of zooming in on one option and thinking it through. Indeed, Wallas (1926)—as various authors since—suggested that creative acts run through four stages:

1. Preparation, where the problem is investigated
2. Incubation, where the problem is thought about unconsciously
3. Illumination, where ideas come together to form a possible solution
4. Verification, the stage in which the chosen option is evaluated and confirmed

Even if more processes are likely to contribute to a creative act, it makes sense to characterize the first two stages as emphasizing divergent processes and the final two stages as emphasizing convergent processes.

Control States Underlying Convergent and Divergent Search

Summarizing the discussion so far, there is increasing evidence that convergent and divergent search operations can be observed in a broad range of cognitive activities, including the search for perceptual events, the search through problem solutions, and the search through representations of possible actions. Two types of operations seem to come as a fixed sequence, with divergent search followed by convergent search—a sequence that may sometimes be cycled through repeatedly, such as when the sought-for target, action alternative, or memory is not found and/or when goals are changing. Moreover, there are reasons to assume that the two types of search operations are controlled by different cognitive control states. Let us now consider how these states may be characterized.

Figure 14.3 sketches the basic idea underlying many biological models of decision making (for a review, see Bogacz 2007). Most models assume that the representations of multiple alternatives, such as A and B in the figure, compete against each other. As alternative decisions are commonly mutually exclusive, collecting more evidence for (or increasing the tendency toward) one alternative increases the activation of the corresponding representation (e.g., A), which leads to the suppression of other alternatives (such as B). If the evidence is clear-cut, decision making might proceed automatically: at some point, sufficient evidence is collected for one of the alternatives and/or the competing alternatives have received sufficient suppression so that the winner can be determined. However, biological systems are noisy and evidence is not always as




Figure 14.3 Possible mechanisms involved in decision making. The goal-relevant alternative A is supported by the goal representation (1) but competes with choice alternative B through mutual inhibition (2). Thus, in addition to the competition, bias is provided by the goal.

clear-cut as one might wish, so many researchers have assumed contributions from top-down processes that bias decision making toward goal-consistent solutions (e.g., Duncan et al. 1997).

This scenario suggests that different control states might be created by modulating the strength of the top-down bias (control route 1) and/or local competition (control route 2; see Colzato et al. 2008). Strengthening top-down bias and/or increasing local competition would establish a relatively “convergent” control mode that goes for singular targets and “exclusive” decision making. In contrast, relaxing top-down control and/or decreasing local competition would establish a relatively “divergent,” integrative control mode that is able to tolerate the selection of multiple targets. Such focused and relaxed control modes may underlie convergent and divergent processing in perceptual search, creative thinking, and action selection, and thus represent general control states of the human cognitive system.

Interestingly, similar pairs of states have been claimed to exist in other cognitive domains as well. For instance, both functional (Dreisbach and Goschke 2004) and neuroscientific (Cools 2008; Cools and D’Esposito 2009) considerations suggest that executive control seeks a balance between two extreme control states: one mode guarantees the stability of goal representations in the face of obstacles and resistance, whereas the other mode allows for giving up and trading the present goal for a more reasonable or promising alternative. Cools and d’Esposito (2009; see also Cools, this volume) suggest that the stability part of this delicate balance might be mediated by the prefrontal dopaminergic pathway, whereas the flexibility part is mediated by the striatal dopaminergic pathway. It is interesting to note that the prefrontal mechanisms that Cools and d’Esposito consider relevant for maintaining stability have also been assumed to provide the top-down bias in competitive decision making (Desimone and Duncan 1995); this might suggest that there is a tight relationship between the control modes responsible for stability and for convergent

thinking. Indeed, the individual efficiency of both top-down control (Duncan et al. 1996) and convergent thinking (Akbari Chermahini and Hommel 2010) has been reported to correlate with intelligence. Conversely, a behavioral genetics study revealed that individuals with the DRD2 TAQ IA polymorphism, which results in a 30–40% reduction in dopamine D2 receptor density (the receptor type found primarily in the striatal dopaminergic pathway), show significantly better performance in divergent thinking (Reuter et al. 2006). This fits with the fact that antipsychotic D2-antagonistic drugs reduce the so-called “positive symptoms” of schizophrenia, which have been described as a kind of “widening of the associative horizon” (Eysenck 1993). It thus seems that the functional dialectic between convergent and divergent operations is mirrored to at least some degree in the relationship between stability and flexibility, and this seems to imply some overlap of the underlying neural substrate. Nevertheless, until now the logic of the stability-flexibility concept has only been applied to action goals, whereas the convergent-divergent concept can potentially be applied to any type of decision making—be it between to-be-attended targets, memory traces, representations of alternative actions, or goals. However, given that the search for a target, memory item, or action needs to be goal-directed, decisions between goals need to precede, and selected goals need to outlive, more specific decisions, which requires at least some sort of temporal hierarchy of decision making (cf. Hommel 2009).

A similar, possibly related pair of control states has been referred to as exploitation and exploration modes (Cohen et al. 2007; Daw et al. 2006; see also Daw, this volume, and Hills and Dukas, this volume). The concepts of exploitation and exploration are almost identical to what others have referred to as stability and flexibility, but exploitation-exploration approaches have focused more on the strategies driving control toward one or the other pole of this dimension and the information and neural signals informing such strategies. Moreover, although dopamine has been assumed to control the balance between stability and flexibility (Cools and D’Esposito 2009), the control of balance between exploitation and exploration has also been attributed to norepinephrine (Aston-Jones and Cohen 2005b) as well as to dopamine (Hills 2006; Hills and Dukas, this volume). Expectations and uncertainty are assumed to be important parameters, with moderate degrees of certainty and expected uncertainty promoting exploitation, and perfect certainty (producing boredom) and unexpected uncertainty (undermining confidence in one’s assumptions) promoting exploration (Cohen et al. 2007). Future research is needed to test the interesting hypothesis that the same information that promotes exploitation also induces a convergent operation style, while information that promotes exploration induces a divergent operation style.

If we assume that comparable convergent and divergent search modes exist in perceptual search, memory search, and action selection, and that in all these cases the search modes are controlled by the same cognitive control states, one would expect specific interactions between all sorts of tasks that are likely

to require the establishment of such control states. In particular, one would expect that interleaving or quickly switching between any two tasks would yield better performance if they call for the same (focused or relaxed) control state than if they imply different states. Two recent studies suggest that this is indeed the case.

Hills et al. (2008) demonstrated that participants, who in a visual foraging task searched through clumpier distributions in space, spent more time on constructing possible words from a set of letters in a Scrabble-like task. One possible interpretation is that a clumpier environment is more likely to propagate a convergent control style than a more diffuse distribution of possible targets, and that a convergent style would lead to more endurance when working on a Scrabble problem. Along similar lines, Hommel et al. (submitted) had participants switch between blocks of convergent- and divergent-thinking tasks and other tasks that are commonly taken to tap into cognitive control processes. Tasks suspected to require rather strong top-down control—like Navon's (1977) global-local task, the Stroop task, and the Simon task—yielded better performance when mixed with a convergent-thinking rather than a divergent-thinking task. This fits with the prediction that types of tasks which rely on a rather focused control mode benefit more when mixed with each other than with a task that calls for a relaxed mode, such as the divergent-thinking task. Hommel et al. (submitted) also employed the attentional blink task, which has been suspected to benefit from lesser top-down control (Olivers and Nieuwenhuis 2006; Shapiro et al. 2006). As predicted, this task yielded better performance when mixed with a divergent-thinking task.

Conclusion

The evolutionary emergence of a cognitive off-line system that allows for both the anticipation and generation of external events has made perceivers/agents more or less independent of current situational circumstances and rendered them proactive rather than reactive. Proactive processes require choices, however, and choices imply the search for a suitable or, ideally, even the best option. Accordingly, humans have developed search strategies that serve two different goals. Divergent search operations identify useful and feasible options without necessarily comparing them, whereas convergent search operations try to pick the best (i.e., most rewarding and/or least demanding) option from this restricted set. There is evidence that these two types of operations can be found in perceptual and memory search, as well as in action selection—hence, in all sorts of searching through cognitive representations. Moreover, there is evidence that these two operations differ with respect to the neural underpinnings and that they are controlled by dissociable control states.

The observed similarities across various sorts of search processes suggest a common phylogenetic source, and I have speculated that the emergence of the

ability to plan actions prospectively (i.e., off-line, in the absence of response-triggering external cues) was the driving force. If planning involves decision making between alternative stored action representations, it can be considered a process of cognitive search, and it is possible that it represented the prototype for the development of other types of cognitive search (Hommel 2010). What I did not discuss was how the ability to plan ahead evolved. An interesting possibility is discussed in Hills and Dukas (this volume), who suggest that cognitive search—the internal checking of a number of representations for a match against some goal-relevant representation—might represent the internalization of the ability to search the environment overtly (Hills 2006; Hills et al. 2008). In other words, cognitive search through object and event representations might in some sense simulate overt, active search for external objects and events. This view is consistent with my suggestion that action control is the prime mover in the evolution of cognitive search (and other attentional operations; Hommel 2010), and it may help in extending the present discussion to the analysis of cognitive skills and processes in general.

Moreover, the perspective of Hills and Dukas (this volume) points to a possible origin of the divergent-convergent sequence in cognitive search operations that I have considered. Overt search, as in food foraging behavior, logically and empirically alternates between (overt) exploration (looking around for possible food) and exploitation (collecting and/or eating the food). The cognitive control of overt exploration behavior is likely to require a more divergent decision-making style, as discussed above, whereas the control of overt exploitation calls for a convergent style. This implies a systematic sequence of action-control styles over time, commonly beginning with the divergent control style, followed up by convergent control. If so, seeing the same sequence of control operations in various versions of cognitive search seems to be a logical consequence of the internalization of overt search behavior into a cognitive skill.

First column (top to bottom): Peter M. Todd, Eddy J. Davelaar, Bernhard Hommel, Sean M. Polyn, Eddy J. Davelaar in discussion, Sean M. Polyn, Peter M. Todd
Second column: Thorsten Pachur, Jeremy M. Wolfe, Jeroen G. W. Raaijmakers, Jeremy M. Wolfe, Michael R. Dougherty, Nathaniel D. Daw, Michael D. Lee
Third column: K. Richard Ridderinkhof, Jeroen G. W. Raaijmakers, Michael R. Dougherty, Nathaniel D. Daw, Michael D. Lee, Thorsten Pachur, K. Richard Ridderinkhof



From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.