



Published in final edited form as:

J Exp Psychol Gen. 2010 November ; 139(4): 590–609. doi:10.1037/a0020666.

The Central Executive as a Search Process: Priming Exploration and Exploitation across Domains

Thomas T. Hills¹, Peter M. Todd^{2,3,4}, and Robert L. Goldstone^{2,3}

¹Cognitive and Decision Sciences, University of Basel

² Program in Cognitive Science, Indiana University

³ Department of Psychological and Brain Sciences, Indiana University

⁴ School of Informatics and Computing, Indiana University

Abstract

The trade-off between exploration and exploitation is common to a wide variety of problems involving search in space and mind. The prevalence of this trade-off and its neurological underpinnings led us to propose domain-general cognitive search processes (Hills, Todd, & Goldstone, 2008). Here, we propose further that these are consistent with the idea of a central executive search process that combines goal-handling across subgoal hierarchies. The present study investigates three aspects of this proposal. First, the existence of a unitary central executive search process should allow priming from one search task to another, and at multiple hierarchical levels. We confirm this by showing cross-domain priming from a spatial search task to two different cognitive levels within a lexical search task. Second, given the neural basis of the proposed generalized cognitive search process and the evidence that the central executive is primarily engaged during complex tasks, we hypothesize that priming should require ‘search’ in the sense of a self-regulated making and testing of sequential predictions about the world. This was confirmed by showing that when participants were allowed to collect spatial resources without searching for them, no priming occurred. Finally, we provide a mechanism for the underlying search process and investigate three alternative hypotheses for subgoal hierarchies using the Central Executive as a Search Process model (CESP). CESP envisions the central executive as having both emergent and unitary processes, with one of its roles being a generalized cognitive search process that navigates goal hierarchies by mediating persistence on and switching between subgoals.

Keywords

working memory; central executive; foraging; search; priming; subgoal hierarchy

Corresponding author: Thomas Hills, Institut für Psychologie, Missionstrasse 64A, 4056 Basel, Switzerland, Phone: +41 61 267 35 21, Fax: +41 61 / 267 04 41, thomas.hills@unibas.ch.

Publisher's Disclaimer: The following manuscript is the final accepted manuscript. It has not been subjected to the final copyediting, fact-checking, and proofreading required for formal publication. It is not the definitive, publisher-authenticated version. The American Psychological Association and its Council of Editors disclaim any responsibility or liabilities for errors or omissions of this manuscript version, any version derived from this manuscript by NIH, or other third parties. The published version is available at www.apa.org/pubs/journals/xge

³We also verified the results presented here using string-edit (Levenshtein) distance (Sankoff & Kruskal, 1983). String-edit distance measures the minimum number of character insertions, deletions, or substitutions required to transform one string into another. The results using this distance measure did not differ in statistical significance from the results presented here.

Problem solving is often characterized as a search process, involving a trade-off between exploiting old solutions and exploring new ones. Such trade-offs influence a wide range of everyday decisions. Problems like buying and selling stocks, staying in or abandoning relationships, gambling further or cutting your losses, and having your regular lunch or trying “the special” are all examples of a decision between persisting with past actions or trying something different. Problems such as these are similar to animal foraging problems, in that they require a search process that can resolve the exploration-exploitation trade-off—choosing between staying in, and hence exploiting, locations where resources have been found in the past or switching to other locations, where the quality and quantity of resources may be less predictable. They also have much in common with characterizations of central executive processing, whose goal-handling and attention-focusing properties have been a keystone in the working memory literature (Baddeley, 1996; Miyake & Shah, 1999; Hazy, Frank, & O'Reilly, 2006). Tasks that tap into executive processing, like the Wisconsin Card Sorting Task and the Verbal Fluency Task, do so in part by inviting participants to mediate this trade-off between persisting with one line of thinking or switching to another (e.g., Berg, 1948; Troyer, Moscovitch, Winocur, Alexander, & Stuss, 1998). As with so many problems in our everyday experience, finding a good solution involves being able to do both.

The widespread discussion of the exploration-exploitation trade-off in numerous fields reflects an interesting opposition in the cognition literature. On the one hand, specialized mechanisms have been hypothesized for search, problem solving, and decision-making in a variety of domains, following arguments regarding the advantages of specifically tuned over general purpose designs. This has further led to suggestions that the mind incorporates numerous autonomous and domain-specific neural modules (Cosmides & Tooby, 1994; Barrett & Kurzban, 2006; also see Samuels, 1998; Fodor, 1983), or tools in an “adaptive toolbox” (Gigerenzer, Todd, & the ABC Research Group, 1999), each of which is designed to handle a specific class of problems. An extended version of this argument is that the mind lacks a central executive, and that cognitive control is principally an emergent process of competing subsystems, with conflict resolution in the absence of a metacognitive control system, as in, for example, Barnard's (1985) model of Interacting Cognitive Subsystems and Anderson's (1993) ACT-R framework.

On the other hand, the underlying structure of many problems, regardless of domain, can hinge on a decision of when to maintain or shift strategies and attentional resources (e.g., Robbins, 1952). This suggests the usefulness of a domain-general persist-or-switch process, which mediates the needed trade-off between exploitation and exploration. Evidence is accumulating that a central executive in working memory may do exactly this, mediating the focus of attention between competing goals in complex tasks (Hasher & Zacks, 1988; Baddeley, 1996; Engle & Kane, 2004). As we describe below, there is also evidence that some neural capacities are devoted to this generic goal maintenance and attention shifting across domains, and may in fact represent a general purpose search architecture. Though domain-specific mechanisms certainly exist, solving search-related problems—where finding solutions may require aborting non-self-terminating subgoals—may often be a matter of supplementing a domain-general search process with domain-specific information.

In previous work, we found evidence supporting the existence of one such generalized cognitive search process. We did this by showing that experience with a visuo-spatial search task could alter search behavior in a lexical search task, possibly by priming a domain general search process (Hills, Todd, & Goldstone, 2008). Participants were first asked to find as many resource tokens as possible by actively searching in a virtual world. The tokens were distributed either diffusely or in clumps, with the consequence that participants gave up quickly on a given area and moved on, or perseverated on local resource patches, respectively. Following this spatial search, participants were asked to form a total of 30 words over a series of letter sets,

using one letter set at a time. When they felt that they had sufficiently exhausted the word possibilities of one letter set, they could move to another letter set. Participants who experienced clustered spatial resources stayed searching longer in each letter set than participants who experienced diffuse spatial resources. The results indicated that tendencies acquired in one task to explore versus exploit an environment spontaneously transferred to a superficially dissimilar task.

In this paper, we provide a mechanism accounting for the generality of the cognitive search process described above, and we do so by combining evidence for a central executive process for goal persistence with a model of problem solving using subgoal hierarchies. As we describe in more detail below, a search architecture based on these two elements combines the goal-handling abilities of the central executive with our understanding of the shared ecological and neural basis of a generalized cognitive search process. By this account, the central executive is responsible for altering subgoals in response to either changing task demands or diminishing returns on subgoal performance. Thus, our notion of the central executive attempts to extend our understanding of working memory by pointing out that many of the everyday problems we contend with share a common abstract structure: one that involves knowing when to abandon one action sequence and initiate another.

In what follows, we first provide the theoretical basis for conceiving of the central executive as a search process. We do this by combining two lines of evidence. The first is the biological evidence for a common underlying search process in the brain, which in turn supports a unitary, domain-general central executive and also provides a putative mechanism for its role as a process of subgoal persistence in cognitive control. The second line of evidence, following from the first, indicates how the central executive mediates goal persistence via inhibition and adaptive transitions of attention. Following this, we present two experiments that test predictions regarding priming at multiple hierarchical levels within a single task, and the necessity of “search” as opposed to just “movement” for priming behavior across domains. Search is defined here specifically as the *self-regulation* of transitions between exploratory and exploitative actions to achieve higher-level goals. Finally, we formalize a model of the central executive as a search process (CESP), and test three alternative hypotheses regarding how the navigation of subgoal hierarchies can produce the observed human behavior patterns.

The Biological Basis of a Domain-general Search Process

“Regulating the balance between exploitation and exploration is a fundamental need for adaptive behavior in a complex and changing world” (Cohen, McClure, & Yu, 2007, p. 934). This statement summarizes what we feel is one of the most important selective forces operating in the evolution of cognition. It encompasses both the ability to detect—via exploration—the resource contingencies available in different environments, and then to pursue—via exploitation—effective behavioral strategies that take advantage of those contingencies. This regulatory process is so central to adaptive success that it is found at both evolutionary and ecological (i.e., cognitive) time scales (Badyaev, 2005; Bedau & Packard, 2003; Baldwin, 1896; Stephens, 1991).

Taking a comparative approach to investigating the biological origins of the exploitation/exploration trade-off in cognitive control, Hills (2006) presented evidence for an early evolutionary origin of the neuromolecular architectures involved in this control. He also showed that, in humans, this modulation appears to have taken on a domain-general form that is common across many tasks, controlling both external and internal search.

The evidence for an early evolutionary origin stems from the fact that the relevant neural architectures controlling food- and foraging-related behaviors almost universally involve dopaminergic modulations of glutamatergic synaptic pathways, which are shared across a wide

variety of species, ranging from nematodes to mammals (e.g., Hills, Brockie, & Maricq, 2004; Due, Jing, & Klaudiusz, 2004; Baumann, Dames, Kühnel, & Walz, 2002; Anderson, Braestrup, & Randrup, 1975; Szczypka et al., 1999). The domain generality of dopaminergic mechanisms in executive processing in humans is supported by their involvement in both spatial and more abstract goal-directed cognition (Salas, Broglio, & Rodriguez, 2003; Houk, Davis, & Beiser, 1995; Reiner, Medina, & Veenman, 1998), numerous pathologies of goal-directed behavior (see Nieoullon, 2002; Hills, 2006), and strong evidence that deficits in dopaminergic processing account for general age-related deficits in executive function (e.g., Volkow, Gur, Wang, Fowler, Moberg, et al., 1998; Bäckman & Farde, 2005).

The fact that dopamine has often been called a “reward” or “novelty” detector is an indication of the field's appreciation of this generality. Novelty and rewards are domain-general classifiers. They are so general that the dopaminergic signals that indicate them can be reassigned from unconditioned to conditioned stimuli (Ljungberg, Apicella, & Schultz, 1992)—neutral stimuli can come to activate reward detectors by being associated with rewards. Thus, detecting reward contingencies can mean detecting different kinds of opportunities for exploiting available resources.

Support for the domain generality of dopaminergic processing also stems from the antagonistic role of tonic D1 and phasic D2 dopamine (DA) receptor-mediated effects (Grace, 1991; Grace, 2000; Cohen, Braver, & Brown, 2002; Goto, Otani, & Grace, 2007). The current hypothesis is that while phasic bursts of DA update localized contents of working memory by gating new information to the prefrontal cortex, tonic dopaminergic activity is a more global inhibitor of phasic activity, down-regulating sensitivity to competing stimuli. Changes in tonic activity also occur over a slower-time scale (Schultz, 2007). This has the consequence that tonic dopamine levels offer a general mechanism for mediating the tendency to exploit the existing foci of attention, or transfer attention to other potential targets, no matter what the domain.

Equally compelling biological evidence for domain-general regulation of exploration and exploitation comes from the widely distributed cortical projections of the locus coeruleus-norepinephrine (LC-NE) system (Aston-Jones & Cohen, 2005). Like the phasic-tonic account of dopaminergic processing, the LC-NE system operates along a similar principle of local versus global mediation of attentional focus. The activation of the LC is considered to follow the course of positive and negative reinforcements from the environment—and to respectively induce exploitative and exploratory behaviors in response (e.g., Cohen et al., 2007; Aston-Jones & Cohen, 2005; McClure, Gilzenrat, & Cohen, 2006).

In summary, the biological arguments for domain-general cognitive search processes stem from at least two separate, but potentially integrated, neuromodulatory systems. Both appear to control processing between exploitative and exploratory behavior, and they may do so by helping to maintain or shift attention between competing task demands. Though there has been limited investigation of the evolutionary origins of the LC-NE system, in the case of dopamine the evidence points predominantly to spatial search as one of its earliest domains of application. Thus, mechanisms devoted to the mediation of exploration and exploitation of goals may have originated to facilitate spatial search behavior and possibly other search behaviors that have similar cognitive requirements. This suggests three possibilities. Either this search control mechanism is now duplicated (exapted) for multiple competing and independent subsystems, or it is instantiated as a single domain-general search process, or both. If either of the latter two is true, then we should be able to prime search strategies across domains, for example from a visuo-spatial domain to a lexical domain, as if we were priming a unitary central executive process.

The Central Executive as a Search Process

Although a diversity of ideas have proliferated around Miller, Galanter, and Pribram's (1960) initial use of the concept of *working memory*, it is generally accepted that working memory comprises that part of human cognition that handles activation of information and execution of plans related to the achievement of current goals (see Miyake & Shah, 1999). In seminal work, Baddeley and Hitch (1974) argued that this conceptualization of working memory could be divided into two subcomponents: a set of slave systems, which handle the storage of domain-specific information (e.g., the visuo-spatial sketchpad and the phonological loop); and a central executive. Baddely (1986) later proposed that the central executive might be consistent with the functioning of Norman and Shallice's (1986) *supervisory attentional system*, guiding and updating behavior adaptively, appropriate to the current context. More recently, the central executive has become associated with the cognitive control of attention and with the ability to maintain goals and processing in working memory in the face of interference while handling complex tasks (e.g., Cowan, 2005; Baddeley, 2007; Engle, Kane, & Tuholski, 1999; Oberauer, Süß, Wilhelm, & Sander, 2008).

Our definition of a domain-general cognitive search process is a potential component of the more expansive central executive processes proposed above. The biological evidence suggests a central executive that is involved with domain-general high-level cognitive processing, devoted to mediating attentional processes to achieve goals. We further conceive of this central executive process as *searching* through high-dimensional problem spaces, because reducing the difference between one's present state and one's goal state often involves a process of dynamically switching between subgoals—involving the exploration and exploitation of alternatives. Furthermore, on complex problems where an executive control of attention is required, the central executive can work to achieve goals by perseverating on subgoals associated with particular actions, and choosing new subgoals when prior subgoals do not lead to adequate progress. We recognize that the central executive as proposed by some may be broader or narrower in purpose than we claim here. Our claim is that the central executive as a general search process, one that searches by navigating subgoal hierarchies, is a minimal functional requirement of the central executive.

Search through subgoal hierarchies—If the central executive instantiates a search process, then it needs an environment over which to forage. One proposed representation of this search environment is in terms of subgoal hierarchies (e.g., Botvinick, 2008; Braver & Bongiolatti, 2002). Subgoal (or action) hierarchies and their corresponding reward structures are appropriate natural environments for searching by means of a cognitive process that mediates between persistence and switching. Miller et al.'s (1960) initial description of a “working memory” was made with respect to precisely these kinds of hierarchically structured plans of action. Moreover, hierarchically structured actions are employed by numerous cognitive architecture models, such as the General Problem Solver (Newell & Simon, 1972), ACT-R (Anderson, 1993), and SOAR (Laird, Newell, & Rosenbloom, 1987). Thus, the central executive in working memory may achieve goals not by searching through literal space, but by searching spaces of subgoals and their associated actions, which in turn facilitate movement in either external or internal domains.

Figure 1 presents one possible subgoal hierarchy for the lexical search task we investigate in Experiments 1 and 2. The aim is to discover English words that can be made from a series of scrambled letter sets, with the primary goal being to find a total of 30 such words across multiple letter sets. The subgoals represent possible methods for achieving the primary goal—including searching within a letter set and switching between letter sets. As with many real world tasks, there are multiple possible trajectories through a tree of subgoals. The role of the search process

is to choose subgoals such that an appropriate trajectory can be found to achieve the top-level goal.

Subgoal hierarchies like Figure 1 represent hypotheses about goal structures, through which possible courses of action can be found. Like production rules in ACT-R and SOAR, alternative hierarchies can be tested experimentally to investigate the actual goal structure of the task environment. As we demonstrate in more detail with the CESP model simulation, formalizing and testing alternative hierarchies is a way to uncover the goal spaces utilized to solve a given problem.

Two open questions still remain about the role of the central executive as a search process. First, what is the evidence that it plays a role in mediating switching between subgoals, and second, is the central executive really a unitary domain-general construct, or is it an emergent property of competing subsystems?

The central executive, inhibition, and task switching—“The ability to switch flexibly between tasks is the pinnacle of human cognition and the hallmark of executive control” (Logan, 2004, p. 3). This quote is fundamentally similar to that by Cohen et al. (2007) above, both referring to a cognitive ability to transfer attention flexibly and adaptively in response to task demands. The flip side of this perspective is to say that the pinnacle of human cognition is the ability to maintain attention in the face of distraction—that is, appropriate switching also implies focusing and *not* switching, at least some of the time. What evidence is there that a central executive mediates the control of attention between competing goals?

Some of the most compelling evidence comes from studies demonstrating that working memory capacity is strongly correlated with the ability to maintain focused attention. In an elegant study of the *cocktail party phenomenon*, Conway, Cowan, and Bunting (2001) first divided participants into two groups based on their operation span—a test of memory span while performing a distracting task. Participants were then presented with a dichotic listening task in which they were asked to shadow (repeat) words heard in the right ear, while an irrelevant message was played in the left ear. Unbeknownst to participants, the message in the left ear also contained the participant's name. When later asked if they had heard their name, participants with lower operation span were significantly *more* likely to have heard their own name—being less able to control their attention to avoid distraction—than those with high span.

In a study of behavioral inhibition, Unsworth, Heitz, and Engle (2004) found a similar result when participants were asked to inhibit a visual looking response in an anti-saccade task. Participants with low spans were less able to inhibit looking at the target than participants with high spans. In both the Conway et al. (2001) and the Unsworth et al. (2004) studies, a course of action associated with a specific goal is pitted against actions associated with attending to distracting stimuli—distracting stimuli representing competing subgoals. Maintenance of one subgoal requires not switching to the other. Having a higher working memory span could therefore be associated with greater ability to maintain attention to the items being kept in memory. Numerous studies have directly investigated this sustained-attention/switching hypothesis of the central executive, and have found consistent evidence in support of it (Hasher & Zacks, 1988; Kane & Engle, 2000; Rosen & Engle, 1998; Kane, Bleckley, Conway, & Engle, 2001; Kane & Engle, 2003; Hasher, Lustig, & Zacks, 2008; also see Braver, Gray, & Burgess, 2008 for a similar account based on proactive and reactive control).

In many problems, the ability to regulate the maintenance and transfer of attention between subgoals is critical to successful outcomes. The results just reviewed support the idea that this self-regulated switching process is a key component of the central executive.

Central executive or executive committee—The second problem to be addressed is whether or not the central executive is an emergent property of dynamic, interactive, competing task demands—what Baddeley (1996) called an executive committee—or is it a unitary, domain-general system that serves to modulate attention, regardless of the competing tasks? In other words, is the control of attention the outcome of competition decided only by the opponents (the committee), or decided in part by a referee (the executive)? Both perspectives are well represented in the literature.

Numerous cognitive architecture models produce central executive functions by processing interactions among competing subsystems. For example, the model of Interactive Cognitive Subsystems (Barnard, 1985), ACT-R (Lovett, Reder, & Lebiere, 1999), and the Executive-Process/Interactive-Control model (Kieras, Meyer, Mueller, & Seymour, 1999) all use competition between domain-specific subsystems to determine the outcome of cognitive control—there is no true “central executive.” On the other hand, systems such as the Embedded Processes Model (Cowan, 2005) and Executive Attention Theory (Engle & Kane, 2004; Kane, Conway, Hambrick, & Engle, 2008) explicitly posit a central executive that is shared across tasks to schedule subsystem activation.

Still other proposals, such as O'Reilly and Frank's (2006) Pre-frontal Cortex and Basal Ganglia Model and similar neural based connectionist models (e.g., Braver & Cohen, 1999; Durstewitz, Kelc, & Güntürkün, 1999), sit on the fence with respect to the unitary nature of the central executive. For example, they often assume a shared set of neural subregions (e.g., the thalamus, hippocampus, striatum, and prefrontal cortex) or neuromodulatory control mechanisms (e.g., dopamine) across tasks, but the goals and their associated behaviors compete for activation in the prefrontal cortex without a higher-level arbiter. They also typically remain agnostic as to whether there is a global mechanism for mediating exploration and exploitation across tasks (e.g., see Munakata, Morton, & O'Reilly, 2008; Braver et al., 2008; cf. Rougier, Noelle, Braver, Cohen, & O'Reilly, 2005; also see Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006).

One possible resolution is provided by phasic-tonic accounts of neuromodulation. That is, phasic-tonic dopamine or norepinephrine models of cognitive control are potentially consistent with a shared, domain-general mechanism for mediating exploratory/exploiting behavioral patterns, in a way that integrates the executive and committee views of search control (see Cohen et al., 2002; Aston-Jones & Cohen, 2005). The phasic-tonic accounts provide a means for a self-organized, committee system of goal-goal competition (via local phasic activity) mediated by a domain-general mechanism for increasing or reducing goal perseveration (via tonic activity). To the extent that processes like the DA and NE mechanisms are true global modulators of attentional control, we should be able to observe an effect of exploration/exploitation priming across domain-specific tasks, and at multiple hierarchical levels within those tasks. Experiment 1 examines this possibility.

Experiment 1

The first experiment was designed to investigate priming at multiple levels in a hierarchical problem solving task. Each participant worked through a series of three phases. The first phase was an initial training session with the lexical search task. In the second phase, participants were randomly assigned to either a clustered or diffuse treatment condition of a spatial search task. In the third phase, participants returned to the lexical search task.

In our earlier cross-domain priming study (Hills et al., 2008), the letter sets were presented in a fixed order, which added order effect confounds to the investigation of any priming within letter sets. To overcome this problem—and to establish the robustness of the earlier priming results—each participant in the present experiment was assigned an independently randomized sequence of letter sets.

The abstract lexical search task presented participants with two simultaneous search decisions: how to search for the next word within a letter set, and when to leave one letter set for the next letter set. A wait time of 15 seconds was imposed between the participant's expressed desire (via an on-screen button click) to move to the next letter set and the next letter set's presentation. This prevented participants from simply moving rapidly over multiple sets. By imposing this transition cost or “travel time,” the abstract lexical search problem becomes analogous to the typical patch-search paradigm in the animal foraging literature, in which animals must decide between foraging in one patch or taking the time to travel to the next patch (e.g., Charnov, 1976; Hills & Adler, 2001). In both cases, “travel” (between patches or between letter sets) is costly in the sense that resources cannot be acquired while traveling.

More broadly, the commonality in both our spatial and lexical search tasks is that participants must choose between continuing to try to harvest resources (particular pixels on the screen or words in the letter set) in their current “location” (screen region or letter set) or “moving” to and then checking a new location. See Table 1 for a fuller description of the analogy between the search tasks. The basis for this analogy may seem highly abstract, but if people have a general cognitive search process dedicated to making decisions about whether to exploit a current resource patch or explore for new patches, then they are implicitly making a common decision under both tasks about when to switch the “location” that is currently being searched. In terms of subgoal hierarchies, the common element is the decision to stick with a specific subgoal or to switch between subgoals.

We can obtain evidence for a common underlying process governing search across domains by examining the nature of cross-task priming. If participants in the diffuse treatment of the spatial search task are generally primed to explore (i.e., switch between subgoals more often) more than exploit (i.e., switch less often), then we predict that when they are transferred to the lexical search task, they will shift letter sets more frequently (exploring more). Conversely, if participants searching through clustered resource patches in the foraging task are generally primed to exploit resource patches, then we predict that they will shift letter sets relatively infrequently (exploiting each letter set for a longer period of time).

A further prediction of a common search process is that *within* letter set word transitions will also be affected by the priming. We formalize these predictions explicitly when we present the CESP model (see below), but here we lay out their general predictions. First, it may be that participants only ever sample from the letter set, without regard to their previous solutions. Here the terminal subgoals are to sample from the letter set, compose a test word from the sampled letters, and then try to match the test word to the known lexicon. In this case, similarity between adjacent solutions should *not differ* between the conditions, because the test word production process always starts from the same location (with the letter set itself). Note that we define the test word as a product of working memory manipulations; it may or may not be a valid solution. In the same way that a chess player explores different possible solutions to a chess problem before actually selecting one, a test word represents a possible solution in the current letter set.

A second hypothesis is that because clustered participants exploit resource patches for longer, they will also exploit previous solutions for longer. If this is true, the terminal subgoals are to *modify the previous solution* to compose a test word, and then check it against the lexicon.

Participants in the clustered treatment should then generate solutions that are *more similar* to their previous solution than participants in the diffuse treatment.

A third hypothesis is that clustered participants persevere longer on the process of iteratively manipulating test words. Here the terminal subgoals are to *modify the most recent test word* and check it against the lexicon. Now clustered participants may produce words that are *less similar* to their previous solution, because they make more manipulations to the previous solution before switching to another subgoal.

We argue that the outcome of the priming between solutions is dependent on the identity of the terminal subgoals in the subgoal hierarchy—and more generally, on the structure of the solution space and the strategies associated with navigating it. In the present case, the three hypotheses represent three different sources of subgoal perseveration: Hypothesis 1) always sample from the letter set, Hypothesis 2) modify the previous solution from the current letter set, when available or Hypothesis 3) modify the previous test word from this letter set, when available. Following the experiments, we formalize and simulate these hypotheses in the CESP model.

Methods

Participants—75 students from Indiana University participated in the experiment. Participants received course credit for participation. We established a completion criterion of at least 20 words found in the third word puzzle phase within the allotted time of approximately 1 hour (minus the time for the lexical search training, maze training, and spatial search task). Participants were informed that their goal was to find at least 30 words. Six participants failed to meet the above criterion.¹ Two participants also stayed longer than 5 standard deviations beyond the mean staying time in one or more letter sets, more than 200 seconds longer than the next closest participant, and these participants were also removed from the analysis. Participants were randomly assigned to the two separate treatment groups by the computer. After removing the outliers, this yielded 35 participants in the clustered treatment and 32 participants in the diffuse treatment.

Apparatus and procedure—Participants were seated in front of a computer and asked to follow written instructions that appeared on the screen. Instructions guided participants through four activities, beginning with a training session in the abstract lexical search task, followed by a maze training, a spatial search task, and then a final session in the abstract lexical search task. The initial training session in the lexical task and maze training were to familiarize the participants with the controls for the separate tasks.

Lexical search task: Participants were asked to find one or more words made up of at least four letters from each of a sequence of letter sets containing 6 letters, four consonants and two vowels (e.g., the letter set SULMPA can be used to form, among other words, SLAP and PLUM). Plurals and proper names were disallowed. Once a letter set was displayed, participants could type in as many words as they wanted, or click on a button to move to the next set. Letter sets were constructed randomly using only the twenty most common letters in English (i.e., excluding K, V, X, Z, J, and Q). Previous work has shown that participants in this task are sensitive to letter frequency, which is associated with number of possible solutions (Wilke, 2006), and we did not want there to be obvious cues to the ‘patch size’ for each letter set (i.e. the number of possible words present in a letter set). 18 letter sets were constructed. All participants saw the same four letter sets during the training phase. The remaining 14 letter

¹Due to technical limitations, participants only had one hour to complete the task. Using the data from only those subjects who found 30 words in the final lexical search phase did not alter the pattern or significance of the results presented.

sets were randomized across subjects. Immediately after entering a word, participants were given on-screen feedback as to whether it was correct (i.e., an English word, not plural, and formed from the appropriate letters) or incorrect. There were on average 14.7 ($SD = 5.5$) valid solution words per letter set, judged according to the wordsmith.org anagram dictionary. Participants could leave a letter set at any time but had to wait fifteen seconds before the next letter set was shown. After leaving a letter set, participants could not visit it again. All participants used in this study found the required number of words before seeing all the letter sets. Participants received instructions and training on one letter set before moving to the pretest session. In the pretest, participants were asked to find words in three letter sets, with no directions regarding how many words to find before moving on to the next letter set. The pretest session ended when participants left the last of these three letter sets. In the post-test lexical search session (following the spatial search treatment), participants were told that they needed to find a total of 30 correct words across any number of letter sets to finish the experiment. They were also told that they could spend as much time as they liked on any given letter set, but that they should allocate their time appropriately so as not to stay too long or too short in any given letter set. Analyses were performed on correct solutions only.

Spatial search task: In the spatial search task, participants controlled the movement of a foraging icon using the ‘J’ and ‘L’ keys representing ‘turn counterclockwise’ and ‘turn clockwise’, respectively.² The keys initiated turns of 35 degrees per step, and forward speed was approximately 20 pixels (steps) per second, which was constant through the trial. To familiarize participants with the controls, participants first had to complete a maze training task, involving the navigation of a two-dimensional maze. Following completion of the maze, participants entered the foraging treatment. Here participants saw an initially blank 200×200 pixel field (which wrapped around at the edges) with their icon in the center. They were told to move the icon to find as many hidden “resource” pixels as they could in the allotted time. Time was indicated by a sweeping clock-hand in the upper-right corner of the screen (clock units indicated remaining time steps, 20 per second)—this was meant to provide a sense that time in the trial was passing without providing explicit information about when the trial would end. Participants were randomly assigned to one of two resource distributions, ‘clustered’ or ‘diffuse,’ consisting of 3124 resource pixels within 4 diamond-shaped patches of 781 contiguous pixels each or 700 patches of 5 pixels each, respectively. Patches were assigned random locations. Figure 2 shows example resource distributions for the two spatial search treatments (above) and typical participant foraging paths for each treatment (below). Resource pixels were not visible to participants until they were encountered, nor were the participants' paths visible at any time. Once a resource pixel had been found by a participant (by moving over it) it was shown in green and remained visible on the screen for the remaining duration of the trial. A participant searched through five spatial search trials for 2 minutes each, with a different random arrangement of patch locations in each trial.

Results and Discussion

Overall, the results replicate the pattern found in Hills et al. (2008). The average turning angle for the clustered treatment group was 20.1 degrees and for the diffuse treatment group 11.5 degrees per 0.2 seconds ($t(65) = 4.27, p < 0.001$). Participants who were in the clustered spatial search treatment stayed longer in each letter set in the abstract lexical search task than did participants exposed to the diffuse spatial search treatment (clustered mean = 85.6 seconds, $SD = 28.32$, diffuse mean = 70.0 seconds, $SD = 25.8, t(65) = 2.39, p < 0.05$). Figure 3A presents this pattern in terms of time spent in the first seven letter sets in the posttest (after spatial search) minus the mean time spent across the three letter sets in the pretest, (more than 80% of the participants reached the 7th letter set).

²In Hills et al., 2008, ‘forward’ and ‘stop’ keys were also available. They were removed in this study.

Our dependent measure is the difference in time spent in each sequence between the posttest and the pretest. To control for the possibility that individuals who stay varying amounts of time in a letter set may be more or less susceptible to the spatial foraging manipulation, we used the pretest time, the position in the sequence of letter sets (i.e., first letter set, second, third, etc.), and the spatial distribution treatment as independent variables in an unbalanced equivalent of the repeated measures analyses of variance (i.e., the linear mixed effects model, Laird & Ware, 1982). The unbalanced test is needed because participants submitted differing numbers of solutions for each letter set. The main effects of pretest time and sequence position were both significant ($F(1, 64) = 38.2, p < 0.001$; $F(1, 512) = 24.51, p < 0.001$, respectively).⁴ The treatment effect was also significant ($F(1, 64) = 6.68, p = 0.01$), with participants from the clustered spatial condition staying significantly longer in the letter sets than participants in the diffuse spatial condition. For completeness, we find the same results for treatment even if the dependent measure is absolute posttest time with treatment as the only predictor ($F(1, 65) = 6.05, p = 0.02$).

Including the amount of resources found in the spatial search task had no effect on this result and was not itself a significant predictor of time spent in the letter set sequences ($p > 0.1$). The two groups did not differ in the number of sequences visited (mean for both groups = 9.22 sequences, $SD = 2.0, p > 0.05$), but did differ in the time taken between word submissions (clustered mean = 17.31 sec, $SD = 4.56$, diffuse mean = 15.00 sec, $SD = 5.03, t(65) = 1.96, p = 0.05$), and the total time taken to complete the task (clustered mean = 821 sec, $SD = 254$, diffuse mean = 670 sec, $SD = 180, t(65) = -2.8, p < 0.01$).

In Figure 3A, the fact that the posttest-pretest differences are greater for the participants in the clustered than the diffuse conditions indicates that clustered participants stayed longer in letter sets (in the posttest) following the spatial search treatment. Diffuse treatment participants, on the other hand, may have stayed a shorter period of time (at least for the first sequence, though this effect gets weaker for later letter sets). The gradual decline in time spent in the letter sets is picked up in the sequence position effect, with participants eventually staying for shorter durations in each letter set as the number of letter sets they visit increases. Principally, the treatment effect demonstrates that the distribution of resources in a spatial search task can influence the search for resources in a subsequent task, even if the environment in which the subsequent search is performed is radically different from the spatial search task and only involves space in a metaphorical manner. Moreover, the transfer of search tendencies across two domains, typically taken to represent independent subdomains in working memory (Baddeley, 1996; Logie, Zucco, & Baddeley, 1990), suggests that the transferred process is associated with a central executive process and this processing is domain general, i.e., shared across tasks.

A further test of this domain generality is whether or not the observed search priming effect at the level of the letter sets can also be witnessed in the search *within* a letter set. A participant's solution sequence within a letter set can be described as a path through a similarity space, where distances between words correspond to the similarity measure. A similarity distance measure has been developed in earlier research with anagrams, which found that the anagram solutions of non-experts are generated using serial-hypothesis testing in which letters are sampled from the letter set and manipulated to produce successive hypotheses (Mayzner, Tresselt, & Helbock, 1964; Mendelsohn, 1976; Novick & Sherman, 2008). In this case, the distance from the letter set to the solution is assumed to be dependent on the number of manipulations required

⁴Further analysis found that the initial pretest score did influence the magnitude of the resulting change. A regression predicting the absolute difference between posttest-pretest using pretest as the independent variable, found a significant effect of pretest score, with larger pretest scores indicating larger subsequent change in posttest-pretest ($B = 0.22, t(65) = 3.12, p < 0.01$). Thus, including the pretest as a predictor is warranted (also see Bonate, 2000).

to reach the solution from the letter set. The manipulation distance is inversely related to the bigram similarity—that is, the number of bigrams shared between two words or sets of letters. The assumption is that two words that are more similar to one another, in terms of sharing more bigrams, require fewer manipulations in working memory to move from one to the next. For example, BOAT and BOLT share 3 bigrams (start-B, BO, and T-end), for a bigram similarity of 3. This similarity measure is supported by the observation that greater bigram similarity between the letter set and the solution decreases the time to find that solution (Gavurin & Zangrillo, 1975).

Based on these results, we used bigram similarity to construct a search representation for anagram solutions. Because we use anagrams with multiple solutions of varying size, our measure of bigram similarity is the ratio of shared bigrams to the number of total bigrams in the larger word (or letter set). Our assumption is that participants are searching locally in word solution space and that this will be detectable in the bigram similarity between successive solutions. Just as foragers often search near where they have found resources in the past, so participants trying to come up with words from a letter set may search for words close to where they have found other solutions in the recent past. We argue that this nearness can be detected in the participant's solutions, and that it can provide information about where participants are perseverating in their goal structure. It may represent perseveration on a prior solution, the original letter set, or the process of iteratively altering a test word. To test among these alternative hypotheses, we computed the shared bigram similarity between successive submitted words, and the shared bigram similarity between each submitted word and the letter set itself.

Figure 4A provides a visual representation of all participants', in both conditions, sequential solutions in the letter set NSBDOE. Nodes represent solutions and arrows between nodes represent the production of one solution following another. Darker lines indicate that more participants made a particular solution-solution sequence. Larger node sizes indicate that more participants produced that solution before moving on to the next letter set. From this figure, one can see that words that follow one another tend to share similar letter arrangements. For example, BONE following DONE was a common transition, as was DOSE following NOSE (both pairs share three bigrams). However, transitions between less similar words are more rare (e.g., pairs like DENS- BOND and SNOB-BEND share no bigrams and zero links in Figure 4A).

Figure 4B presents the bigram similarity between the present solution (N) and previous solution (N-1), penultimate solution (N-2), and the current letter set, averaged first within participants, and then across participants, for all solutions in both treatment conditions. A paired t-test reveals that the previous solution (N-1) is more similar to the present solution than either the penultimate solution ($t(66) = 5.42, p < 0.001$) or the letter set ($t(66) = 9.44, p < 0.001$). Because the letter set is usually larger than the words produced from it, we can statistically control for this bias by computing the bigram similarity ratio as shared bigrams over the number of bigrams *in the smaller word*. Doing this, however, does not alter the above results; the previous solution still shares the highest bigram similarity with the current solution (as above, $t(66) = 5.152, p < 0.001$, and $t(66) = 12.97, p < 0.001$, respectively).

The finding that solutions are most similar to the solutions that precede them is consistent with the hypothesis that participants are using the previous correct submission as a starting point for the next exploratory search (Hypothesis 2 and 3), and inconsistent with the hypothesis that they are only using the letter set (Hypothesis 1). If they were predominantly using the letter set, then the average similarity between words should be approximately equal without regard to its order in production. Instead, the evidence suggests that participants use prior solutions as waypoints for further search into the solution space.

To evaluate the effect of the two spatial environment conditions on the between solution similarity, we calculated the mean similarity between successive solutions generated by each participant, averaged separately over participants in the two treatment groups. Figure 3B shows that the mean bigram similarity for successive solutions in the clustered group ($M = 0.35$, $SD = 0.07$) was significantly lower than that for the diffuse group ($M = 0.39$, $SD = 0.06$; $t(65) = 2.41$, $p = 0.02$). Using a linear mixed effects model, we also found a significant effect of treatment condition on bigram similarity ($t(65) = 2.28$, $p = 0.03$), and this holds even after controlling for the number of words submitted per letter set ($t(65) = 2.42$, $p = 0.02$). This suggests that experience searching in a particular spatial distribution also affects within-letter set word search in terms of solution similarity. However, the direction of the result supports Hypothesis 3: The diffuse spatial resource condition led participants to produce words that were *more* similar to the preceding solution than did the clustered treatment condition (we found evidence for the same pattern in the data from Hills et al., 2008, data not shown).

Participants who experienced clustered resources do *not* appear to perseverate more on their previous solution—which would have increased their between-word similarity—but may perseverate more on the subgoal of iteratively manipulating the test word. When we present the CESP model, we demonstrate how search in subgoal hierarchies can generate the above patterns.

Experiment 2

Experiment 1 established that the spatial search task was sufficient to induce priming at multiple levels in the lexical search task. However, in experiment 1, three alternative hypotheses for the priming are possible: 1) the priming may be elicited by *motor* perseveration, because the clustered condition required individuals to initiate multiple tight turns in order to stay in one location, or 2) the priming may be a consequence of a high level *schema* associated with clustered or diffuse resource distributions irrespective of search strategies, or 3) the priming may be mediated by an implicit search mechanism, that calibrates a tendency for exploratory or exploitative behavior, and thus mediates the maintenance and transfer of attention between subgoals.

For the motor hypothesis, perseveration at a motor level may initiate a similar kind of perseveration on cognitive processes, through the continuation of behavioral tendencies. For the schema hypothesis, simply harvesting from a resource distribution may elicit a particular pattern of expectations (potentially metaphorical) about how resources should be distributed in future environments—and this carries over from one task to the next. Only the implicit search mechanism requires an overt self-regulating search process—the other two hypotheses only require that individuals harvest from resource distributions.

This distinction between the need for effortful self-regulation or simply following overt cues provided by the environment is also important with respect to *task switching*. In numerous studies, task switching costs associated with a potential central executive have been found under effortful cognitive control conditions involving self-regulation in the absence of overt cues (Logie et al., 1990; Baddeley et al., 1998; Logan, 2004). However, others have found less convincing evidence for a relation between task switching and potential executive processes, such as working memory span, and have called for a more detailed analysis of how different tasks may interfere with one another (Oberauer et al., 2008; Hale, Myerson, Emery, Lawrence, & Dufault, 2008). One compelling possibility is that central executive and associated working memory processes are only required when self-regulation is required to manage task switching, i.e., in the absence of overt cues.

In the second experiment, we wanted to determine whether or not an implicit search process is required in the spatial search task in order to elicit priming in the lexical search task. To this

end, we replicated the experiment from Hills et al. (2008) except that the full resource distribution was made visible to participants at the beginning of the task. Consequently, if priming is caused simply by the behavior of harvesting resources by either a scattered or concentrated method (moving to them and passing over them), then we should find priming in this visible-resources setting as well. However, if the priming requires modulating an implicit executive search process, then priming should not occur when all resources are visible and no exploration or location prediction is required.

Methods

Participants—45 participants from Indiana University participated in the experiment in exchange for course credit. The completion criterion for being included in the analysis was achieving at least 20 submitted words in the final lexical search task within the allotted time of one hour (participants were informed that their goal was to find at least 30 words)—4 participants failed to meet this criterion. Participants were randomly assigned to the two separate treatment groups by the computer, with 20 participants in the clustered treatment and 21 in the diffuse treatment.

Apparatus and procedure—This task is identical to that described for Experiment 1.

Lexical search task: This is the same as in Experiment 1, except that the letter sets were not in random order. We let all participants experience the same order as presented in Hills et al. (2008); we did this to reduce the noise associated with different presentation orders for the letter sets, and thus to increase our sensitivity to the priming effects associated with the resource distributions.

Spatial search task: The main difference between this task and that described for Experiment 1 is that now participants were told to move their icon to collect as many colored pixels as they could in the allotted time. Unlike Experiment 1, where participants needed to hunt for *hidden* pixels, in this task pixel resources were *visible* to the participants and only needed to be harvested. All other features of the spatial search task were as in Experiment 1.

Results and Discussion

When searching for hidden resources, as reported in Hills et al. (2008) and replicated in Experiment 1, the turning angles for foraging paths by participants in the clustered resource group were significantly larger than they were for the diffuse group. In the present experiment, where participants could see the resources and only needed to harvest them, the difference between the two groups was even larger (clustered group mean turning angle = 43 degrees per step, $SD = 19.2$, diffuse group = 18.5 degrees per 0.2 seconds, $SD = 9.8$, $p < 0.001$). This increased difference in turning angle is presumably because the differences between the visible resource distributions are highly conspicuous. Thus, if the priming effect is not found when participants can see the resources, it is not simply because participants are not behaviorally invoking different patterns of action in the different spatial resource distributions.

Figure 5A presents the mean differences in posttest minus mean pretest switching times for the first 7 letter sets (visited by over 80% of participants) for the two treatment groups. While the general pattern is similar to that found in Figure 3A—with both groups spending less time in successive letter sets—the two groups do not significantly differ from one another in their overall mean staying times (using a linear mixed effects model as described for Experiment 1, the results for treatment group were $F(1, 38) = 2.10$, $p = 0.15$, and for sequence position were $F(1, 264) = 15.28$, $p < 0.001$). For completeness, we also present results for bigram similarity between successive solutions in Figure 5B, where the two treatment groups again did not significantly differ, whether or not we controlled for number of words produced per letter set

($p > 0.05$). Note that the sample size used in Experiment 2 ($N=41$) is the same as that used in Hills et al. (2008), where we did find an effect when resources were hidden. These results indicate that priming, if it does occur with visible resources, is much weaker than that found when resources must be actively searched for using a self-regulating process that must mediate between persistence and switching.

Summary of Experimental Results

The findings of Experiments 1 and 2 support our main thesis that the same underlying control process is used across and within domains to mediate search-related problem solving. Thus, we find that participants primed with clustered resource distributions in space act as if resources are more likely to be co-located (or clustered) in a non-spatial, verbal cognitive environment, and vice versa for diffusely distributed resources. We found this priming at two levels of analysis. The first effect was at the between-letter set level of departure times, with participants in the clustered spatial treatment leaving letter set patches more slowly than participants in the diffuse spatial treatment. We also found an effect of the priming within the letter set.

We take these findings to be a strong indication that a common underlying search mechanism is being used at two different hierarchical levels within this task—determining both when to leave a patch to seek the next one, and where to search next *within* a patch. Both levels are clearly important in search problems. The first involves making decisions, typically on the timescale of minutes, based on the number of solutions remaining in a letter set. The second involves making decisions, typically on the timescale of seconds, about how to approach the production of the next test word. Yet, as we will argue below through our model of hierarchical search, search at the two levels can be explained by the same underlying mechanism, and is in fact predicted by a domain general central executive search process that perseverates on terminal subgoals.

Finally, we also tested whether search (rather than just harvesting) was necessary to induce the multi-level priming effects. We found that it was: there was not significant evidence of any priming effect at either hierarchical level when participants were shown the location of resources and merely had to harvest them. This is consistent with the notion that self-regulation in the absence of overt cues is a central part of priming the central executive search process. We now turn to a model that builds on these accounts of a cognitive search process to explain the above results.

A Model of Subgoal Persistence as a Central Executive Search Process

Our model of the central executive search process (CESP) treats problem solving as a process of selecting subgoals in a hierarchical tree. Shifting between subgoals is mediated by a single parameter, which controls the relative stability of the current subgoal. This is consistent with the phasic-tonic accounts of exploratory/exploitative modulation, which have been demonstrated to enhance the robustness of goal-related firing using biophysically detailed models of the prefrontal cortex (Durstewitz et al., 1999; Durstewitz, Seamans, & Sejnowski, 2000). It is also consistent with interference and inhibitory accounts of working memory capacity (e.g., Hasher et al., 2008; Kane & Engle, 2003) and the difference between proactive and reactive control in working memory variation (Braver et al., 2008).

Our primary goal is to investigate whether or not differences in persistence among subgoals in a subgoal hierarchy is sufficient to explain priming effects at both levels of our lexical search task. We also ask exactly what hierarchical subgoal structure is most consistent with the observed data—i.e., where is the source of subgoal persistence. The basic effects the model attempts to explain are 1) how subgoal switching can generate earlier or later departure times

in each letter set, and 2) how subgoal switching can influence successive word similarity within a letter set.

Like many other interpretations of working memory processing (e.g., Garavan, 1998; Sternberg, 1966; Schneider & Shiffrin, 1977; O'Reilly & Frank, 2005), we take the attentional component of working memory to be able to hold but one goal representation in working memory at a time. In the present context, this means that participants can only compose and test against their lexical memory store one test word (i.e., one subset of letters that might be a viable solution) at a time. A critical consequence of this constraint is that participants must choose how to deploy attention between possible subgoals—e.g., the terminal branches of the tree in Figure 1—to produce potentially viable test words. It is this selective attention process that we propose guides the search for words across letter sets and may account for the multi-level priming effects. Though various implementations are possible for subgoal persistence, the central mediator of subgoal persistence in the CESP model is the subgoal updating parameter, which stochastically determines how long attention will persist at a given terminal subgoal. We use the term updating because the parameter provides a measure of how often the system changes (or updates) its goal representation.

Using the CESP model, we examined three alternative hypotheses in the lexical search task associated with each of three different terminal subgoals (as described in Experiment 1): Hypothesis 1) always sample from the letter set, Hypothesis 2) modify the previous solution from the current letter set, when available, and Hypothesis 3) modify the previous test word from the current letter set, when available. Note that for each hypothesis, increasing subgoal stability allows participants to persist with the given terminal subgoal for a longer period of time before switching to a new subgoal. Figure 6 presents the subgoal hierarchies associated with each of the above hypotheses. Processing starts with the main goal of finding 30 words, and then moves down to a terminal subgoal. Moving upwards in the hierarchy happens either when subgoals are completed, or as a consequence of a goal switching process—this being mediated directly by the subgoal updating parameter.

Figure 7 presents the relationship between the CESP model and each of the three subgoal hierarchies presented in Figure 6. In each case, the model begins with a new letter set presentation, and samples from that letter set to produce a test word. The models then differ in how the subgoals produce each subsequent test word. The bold lines in Figure 7 represent switching between terminal subgoals in the lexical search task—pointing out the primary difference between the three hypotheses.

For all three hypotheses, the test words are repeatedly produced and then tested against the lexicon until updating occurs, with probability α . Upon updating, there is an additional probability, $1/\beta$, that the process will move still further upwards in the goal hierarchy to switch letter sets (the letter set switch update). If letter set switching does not occur, then there is still a chance that the source for the next test word will be re-determined (the source update). With probability γ the source is the letter set, from which a new test word is constructed by sampling. Otherwise, if a previous solution exists (and the hypothesis allows for it), the new test word is composed as a modification of the previous solution. γ is a free parameter in the model, which we vary over its range, from 0 to 1, in the simulations that follow.

The source update occurs by replacing all letters currently being manipulated in working memory with a new subset from the chosen source. When the source is the letter set (or immediately following a letter set switch update), we use a sampling process consistent with earlier findings regarding bigram similarity (e.g., Gavurin and Zangrillo, 1975): Samples begin with chunks of 1 or more letters (up to four), with each letter after the first added with probability m . The maximum length is capped at 6—the length of the letter set. This initial

chunk can be read from the letter set in either direction (i.e., OB in the letter set could be read into working memory as OB or BO, with probability 0.5). After this, additional letters sampled individually from the letter set (without replacement) are added to the right hand side of the growing subset in working memory until a list of four or more letters is acquired. Each additional letter after the first 4 is added with probability m , up to 6 letters.

When the updated source is the previous correct submission, then this word is modified to produce a test word. Though participants may selectively swap in and out certain letters based on properties of specific letters or letter sequences, we make the simplifying assumption—which applies equally to all hypotheses and all treatment conditions—that a single non-used letter from the letter set is substituted in place of an existing letter in the previous solution. We also make the simplifying assumption that in cases where the previous solution is more than 4 letters long, one letter is removed prior to the initial substitution. Finally, each additional letter after 4 is added with probability m . Simply put, the ‘modify’ subgoal starts with a four-letter sample, substitutes one letter in the word with a non-used letter in the letter set, and then adds new letters with probability m .⁵

In this model, α defines the probability that a subgoal is abandoned. It is this α that we assume is being ‘tuned’ by searching in the spatial search task to either switch more or less often, depending on the experienced resource distribution. Upon deployment of the attentional system in a new situation or domain, α carries over a bias from the previous task. As with the phasic-tonic account of modulatory processing, α mediates a global bias towards exploitation or exploration.

Finally, in the data, the latencies to switch between letter sets (i.e., the time between the last word submission and the decision to switch to the next letter set) decreased with the number of solutions submitted for a given letter set. As shown in Figure 8, participants left the current letter set approximately 15 seconds on average after submission of the 7th word, but took on average 34 seconds to leave the letter set after submission of the 1st word. This pattern of decreasing latency with more submissions has been observed in other tasks involving serial search for items in memory (Dougherty & Harbison, 2007) and is consistent with the hypothesis that participants recognize they are in a depleting resource pool (see Iwasa, Higashi, & Yanamura, 1981). We capture this effect in CESP by making the probability of the letter set switch update be inversely proportional to a threshold parameter β that shrinks with each additional word submission according to the following equation:

$$\beta = \kappa \lambda^{\text{wordcount}} \quad [1]$$

To determine the appropriate values for the constants that specify β , we fit the above to the data for both conditions from Experiment 1 using the quasi-Newton method of Broyden, Fletcher, Goldfarb and Shanno (described in Nash, 1990), fitting both λ (0.86) and the scaling parameter κ (37.1).

Using the above implementation of CESP, we simulated 1000 participants at each of five α values, for each of the three hypotheses. The range of α values, between 0.05 and 0.25, were chosen such that the number of letter sets visited by the simulated participants is both more and fewer, respectively, than the mean number visited by our real participants (approximately

⁵Changes in the value of m had extremely modest effects on the overall pattern of results from the simulation. The main consequence of varying m was to increase the overall time to find a word (for large m). This is because most valid solutions are of the shortest length (i.e., four letters). This has the effect that more letter sets are visited, and between word similarity is lower. However, for all m that find 30 words in 14 or fewer letter sets, the pattern and general range for the number of iterations per letter set and the slope of between word bigram similarity is roughly identical to that shown for $m = 0.2$.

9 letter sets). Each simulated participant visited the same 14 letter sets presented to our real participants, in a randomized order. Also like our real participants, simulated participants completed the experiment when they found 30 words. Because the majority of the possible words for each letter set were common, high frequency words, we let the test of the lexicon be a simple test of whether or not the produced test word was among the valid answers for a given letter set.

Results of the Simulation

Figure 9 shows the results of the simulations, both with respect to total time within a letter set (measured in the number of words tested, i.e., terminal subgoal iterations) and bigram similarity between successive solutions. While all hypotheses show the pattern of reduced time in a letter set as the subgoal updating parameter increases, only Hypothesis 3 also shows the pattern of increasing between solution similarity. Hypothesis 1 shows a pattern of shorter letter set times for increasing α , but between solution similarity is unchanged for different values of the updating parameter. Because Hypothesis 1 never uses the previous solution as a source for producing new test words (γ is always 1), there is no bias to produce words more similar to the last solution.

Hypothesis 2 reveals a similar pattern of shorter letter set times with increasing α , but also shows nearly flat or a reduction of between-solution similarity as α increases. In general, the between-solution similarity is also much higher here than we find in experiment 1 or 2. Note that only for γ values near 1 do we find the pattern of reduced between-word similarity as updating increases. That is, solution perseveration only produces more similar solutions with less updating if participants update the source back to the letter set. When participants always source update to the previous solution ($\gamma = 0$), between-word similarity is unchanged as α varies, because within letter set updating is no different from simply perseverating on the existing subgoal (i.e., sampling from the previous solution).

Experiment 1's pattern of results is most similar to that produced by Hypothesis 3. Simulated participants both leave letter sets sooner and produce more similar solutions as subgoal updating increases. Following the logic of Hypothesis 3, reducing the updating rate allows participants to make more modifications to the previous solution before updating the source in working memory. They therefore find solutions that are more distant from the previous solution. Put another way, participants with lower α , primed by exposure to clustered resources, *stay in the current working memory patch longer* and thus find solutions that are “more well hidden” with respect to their starting point. Note, however, that as the source update becomes more biased towards returning to the letter set ($\gamma \rightarrow 1$), the pattern of less similar solutions with increasing α goes away. This is because source updates to the letter set produce solutions that are less similar to the last solution.

The effect of increasing γ is different from that on Hypothesis 2, because Hypothesis 2 continues to produce test words as modifications of the previous solution. Hypothesis 3 only uses the previous solution to generate the initial test word after each source update. If the probability of using the previous solution following a source update, γ , is zero, there is only one chance to use the previous solution for Hypothesis 3.

In Figure 10 we present the results of a simulation of Hypothesis 3, designed to closely mimic our test population in experiment 1. The simulation has 35 participants in each condition, using $\gamma = 0.5$ and values of α set to 0.1 (clustered) and 0.2 (diffuse). The pattern of results is similar to that found in Figure 3. A lower updating rate means more time is spent in a letter set, but solution similarity is lower between successive solutions.

The CESP model demonstrates how a search process operating over subgoal hierarchies is able to produce the behavioral results at two hierarchical levels that we found in Experiment 1. The model also demonstrates the plausibility of a minimal search architecture, controlled by a single subgoal updating parameter.

General Discussion

A central claim of the theory of a generalized cognitive search process is that search behaviors in different domains use a shared underlying neural architecture. We have argued that this shared architecture may represent evidence for a unitary, domain-general, central executive search process. As a consequence, the resource distribution encountered during search in one domain can alter subsequent search in another domain, because of cognitive parameters in the shared search architecture that have been tuned to perform appropriate switching between exploration and exploitation in the first environment. Once they have adapted to a particular environmental structure (e.g., clustered or diffuse resources), these parameters tend to exhibit some inertia such that their values will take time to adapt to the circumstances of any new search environment. If a second task is given relatively soon after the first task, the first task's search parameter values will thus automatically affect the second task. This can also influence search patterns at different levels in a hierarchical task, such as the lexical search task presented here.

For tasks like those described in the introduction, those that do not have overt stopping cues, a mechanism for eventual subgoal termination is critical. Looking for your car in a parking lot, trying to recall the author of a specific research article, and taking a bath to relax all require mechanisms for knowing when to stop. While many models have an updating process responsible for handling goals and the outcome of production rules (e.g., Braver & Cohen, 1999; Hazy et al., 2006; Laird et al., 1987), very few explicitly incorporate a “leaky capacitor” process like we use here for alternating between subgoals in the absence of explicit top-down or bottom-up control (but see Lovett et al., 1999). To our knowledge, none have proposed this as a unitary, domain general process for goal persistence. Along these lines, the present study makes several contributions to the theory of a generalized cognitive search process and its implications for a central executive.

First, our results provide supporting evidence for the domain generality of an underlying central executive search process operating over subgoal hierarchies. We found that when participants are trained to forage in either a clustered or distributed way in a spatial environment, this has the consequence of influencing their cognitive foraging strategy in a second environment—even though the two environments are drawn from domains often associated with independent slave systems in working memory, i.e., a visual-spatial search task and a verbal/lexical search task (Baddeley & Hitch, 1974). Consistent with the use of subgoal hierarchies, we also found priming at two levels in the lexical search task, with one level involving a decision about staying in or leaving a patch based on the number of remaining resources in a given letter set, and a second level involving a decision about where to search for the next word in the current patch based on the likelihood of finding a word similar to the previous solution.

It is certainly not a logical necessity that search between spatial and lexical tasks be related. As many proposed cognitive architectures clearly demonstrate, a domain general central executive could be an emergent outcome of competing subsystems (e.g., Barnard, 1999; Lovett et al., 1999). Our results speak against an explanation that is strictly determined by competing subsystems; there appears to be some global modulator that is carried over from one task to another. The evidence for this can be seen both via our priming evidence and in past research demonstrating interference in dual-task procedures utilizing “distinct” domains (e.g., Logie et al., 1990; D'Esposito, Detre, Alsop, Shin, Atlas, et al., 1995). It is also consistent with evidence

that the central executive is primarily a mediator of interference and inhibition in the control of attention (Hasher & Zacks, 1988; Engle & Kane, 2004).

Our results offer a putative mechanism for the tuning of this underlying search process. Our proposal of a tunable updating parameter operating over subgoal hierarchies is based on empirical evidence for a role of inhibition in working memory capacity and a biological basis for domain general mechanisms involved in the modulation of subgoal persistence. Though the CESP model proposes a mechanism based on an unconditional probability for subgoal abandonment, there are alternative processes that would share a similar function: for example, the probability to switch between subgoals may be conditional on the occurrence of recent success with a subgoal. Switching could also be based on tolerance for an absolute number of failures (similar to Harbison, Dougherty, Davelaar, & Fayyad, 2009) or on a certain rate of successes. More research will be required to discriminate between these alternative accounts, but all share the property of representing a general mechanism for persisting with subgoals.

The results of experiment 2 suggest that these mechanisms would be tuned only when tuning is consequential for behavior. When the environment provides overt cues, self-regulation is unnecessary. Similarly, lack of executive processing may also apply when learned action sequences are provided automatically, as when using heuristics (“find two words and leave”) or with expert performance. As a consequence of the subgoal updating parameter operating over subgoal hierarchies, it influences multiple levels within a given task. Giving up on one subgoal influences giving up on other subgoals in which that subgoal is nested. The results of Experiment 2 and the finding of multi-level priming suggest that higher-level schemas involving resource density, or metaphorical thinking in terms of space, are an unlikely explanation for our findings. Search, in terms of the self-regulated modulation of exploration and exploitation, appears to be required for the training and priming of the underlying process.

There are two potential reasons for this effect. One is that overt cues about how to distribute attention in a task do not engage central executive processes. That is, explicit directions about how to employ attention override the need to self-regulate, and thus preclude the need to train attentional sensitivity. This may explain some of the controversy over the extent of executive processing required for task switching (Logan, 2004; Oberauer et al., 2008), with the overt cues used in some tasks reducing the load for self-regulating executive processing. A second reason is suggested by the possible role of dopamine as cue for prediction error. Dopamine is hypothesized to facilitate reward detection and the ability to associate unconditioned stimuli with predictors of those stimuli. In the clustered spatial search context, resources are themselves the cues, and predict the co-locality of other resources. If generalized cognitive search processes are using a dopaminergic signal to tune goal updating, as proposed elsewhere (e.g., Hazy et al., 2006; O'Reilly, Braver, & Cohen, 1999), then this may account for the need for a search process: Dopaminergic signaling is not associated with the onset of predictable rewards, but only with the occurrence of unpredictable rewards.

Because CESP provides a mechanism for navigating subgoal hierarchies, this also allows for the formal testing of possible underlying goal structures, as alternative hypotheses. Like a production rule in ACT-R or SOAR, a specific subgoal hierarchy represents a particular cognitive hypothesis. As with cognitive architectures and standard model testing, CESP both acts as a sufficiency proof of the psychological plausibility of the theory, and a mechanism for differentiating the necessary underlying architecture. Moreover, CESP could be implemented in many existing cognitive architectures. For example, biologically based models could incorporate a phasic-tonic activity that mediates competition between subgoals (e.g., see Durstewitz et al., 1999; Niv, Daw, Joel, & Dayan, 2006). Indeed, a global modulatory process is often implemented in neurocomputational models during their initial training (e.g., Braver & Cohen, 2000). Models like ACT-R and SOAR would need to incorporate a domain-general

modulator of subgoal persistence for non-terminating subgoals. This is unlikely to be complicated to implement, as proposals for domain-specific goal-maintenance mechanisms have already been developed for ACT-R (e.g., Lovett et al., 1999).

CESP is also a plausible explanation for many other kinds of working memory related processing besides search tasks. These include active maintenance, interference effects, and capacity effects. In each case, the updating parameter α tunes persistence on a specific subgoal, and thus modulates inhibition of other potential goals. In the case of the cocktail party phenomenon, revisited by Conway et al. (2001), ignoring irrelevant information requires inhibiting an action (i.e., switching to a new goal) to attend to the irrelevant information. Individuals with different baseline α parameters will attend to the inhibited information to different degrees, because they are more or less likely to update their current subgoal of attending to the relevant information. Similarly, working memory capacity is explained as a difference in the ability to inhibit actions (i.e., goals) that will distract attention away from the items that are to be remembered. As noted by others (Hasher & Zacks, 1988; Kane & Engle, 2000; Rosen & Engle, 1998), maintenance and inhibition may be two sides of the same process.

Finally, our results are also surprising with respect to the cognitive science literature on analogical transfer. Research has generally found that transferring schema from one problem to another is difficult for people (Gick & Holyoak, 1983). For example, solving one problem (e.g., in a military context) that involves the abstract schema “Defeating a centrally located negative presence by converging weak and separate forces onto it” does not spontaneously lead to better solutions of another example (e.g., in a medical context) of the same schema (Gick & Holyoak, 1983). Spontaneous transfer is particularly likely *not* to be found if only one example of the abstract schema is presented, and if the initial and transfer situations are highly dissimilar. Interestingly, both of these conditions apply to our experiment and yet we observe robust transfer of search strategy from spatial search to lexical search. Based on Experiment 2, our interpretation of this transfer is that it is not due to an abstract, explicit schema along the lines of Gick and Holyoak's convergence schema; instead, during the foraging task a central executive search process is “tuned up” to either shows a tendency to explore or to exploit its environment. Then, this tuning is simply left in place for the subsequent lexical search task. This kind of procedural priming may be a more powerful way to achieve transfer than relying on verbal abstractions or schemas (Goldstone, Landy, & Son, in press). Searching for other examples of cross-domain transfer like this could be an effective research tool for determining what kinds of cognitive processes have evolved to become general purpose rather than specialized to one domain.

The theory of generalized cognitive search processes and the CESP model make a number of testable predictions for targets of future study. First, if underlying search processes are genuinely domain general, then priming across domains should not be isolated to instances where the first task is literal search in 2D space. Such priming should be found between numerous tasks—to the extent that they involve a self-regulated giving up process—including, for example, semantic fluency tasks (Hills, Todd, & Jones, 2009), n-armed bandit tasks (Daw et al., 2006), and sampling of information prior to decision making (Hertwig, Barron, Weber, & Erev, 2004). Our results also suggest that priming will only be observed when the first task involves tuning attention in response to some level of unpredictability, i.e., involving self-regulation.

With regards to the biological evolution of the central executive, the results here speak mainly in support of 1) the preserved generality of the underlying biological control processes and 2) their relationship to search processing involving the regulation of exploratory and exploiting tendencies across tasks. That this can be modeled as search over goal hierarchies brings together the evidence for a shared neural ancestry involving spatial foraging and our current

understanding of a central executive process that handles goal-maintenance in tasks involving self-regulation.

While the argument for a shared neural ancestry is an evolutionary argument, we stress that it is one based on identifying shared biological mechanisms in existing populations. It is not an argument based on selection pressure or the adaptiveness of search in mental spaces. While selection has no doubt played a role in structuring human cognition, our theory of a generalized cognitive search process does not entertain any speculations about what those pressures may have been. It would seem to go without saying that an organism that can direct action in light of past experience benefits both from modulating action persistence and the retention of memories that might inform that modulation: this should be true both for 'simple' foraging and for higher level cognition. Given the substantial evidence that executive processing is associated with cognitive abilities (e.g., Kane et al., 2008), understanding the selection pressures that have shaped the exploration-exploitation trade-off in behavior, and the environments that support one versus the other, are clear targets for further research.

Further research is also needed to investigate the relationship between global neuromodulators, a central executive in working memory, and a domain general search process. For example, as enhanced dopaminergic processing in the striatum and frontal cortex have been associated with higher working memory capacity (Aalto, Bruck, Caine, Nagren, Rinne, 2005; Cools, Gibbs, Miyakawa, Jagust, & E'posito, 2008), manipulations that reduce phasic dopaminergic activity in these areas (e.g., haloperidol—see Meck, 1983; Maricq & Church, 1983), should reduce the time spent within each letter set and increase between-solution similarity. Moreover, time spent within letter sets and between solution dissimilarity should be correlated with working memory span. Similar manipulations for NE are also needed—both dopamine and the LC-NE systems may play distinct or integrated roles in global modulation of search. Likewise, individuals with larger working memory capacity—who are already thought to have less frequent goal updating—should spend more time in letter sets and produce between-solution similarities that are smaller (as opposed to participants with shorter memory spans). Answering these questions stands to offer much in terms of understanding the domain generality of a central executive and its underlying biological and cognitive architectures.

This leads to a final question: is the central executive nothing more than a search process? Positing the central executive *as* a search process leaves open the possibility that executive processing handles other components of working memory besides self-regulation of subgoal persistence. For reasons associated with both the comparative biology and the common structuring of the problem, we believe the central executive as a search process is well supported. However, it would not be inconsistent with our findings to suggest that the central executive also operates, for example, as a manipulator of particular task components, or as a monitor of task performance. As noted above, domain-general processes may be abundant once we go looking for them, and they may even operate independently of one another. If by a central executive one means domain-general cognitive processes associated with the control of attention to achieve specific goals, then it is not inconceivable that we will find a set of independent cognitive processes that achieve these ends. The real issue may be in what we define as domains. One could rightly argue that a *domain-general* central executive search process is really nothing more than a process *specific* to domains involving search (i.e., requiring mediation of the trade-off between exploration and exploitation). Such domains just happen to be particularly common in our present and past environments.

Acknowledgments

This work was supported by the National Institute of Health, T32 Grant # HD 07475 and the NSF REESE Grant 0910218.

References

- Aalto S, Brück A, Laine M, Nagren K, Rinne JO. Frontal and temporal dopamine release during working memory and attention tasks in healthy humans: A positron emission tomography study using high-affinity dopamine D2 receptor ligand [¹¹C]FLB 457. *The Journal of Neuroscience* 2005;25:2471–2477. [PubMed: 15758155]
- Anderson H, Braestrup C, Randrup A. Apomorphine-induced stereotyped biting in the tortoise in relation to dopaminergic mechanisms. *Brain Behavior and Evolution* 1975;11:365–373.
- Anderson, JR. *The architecture of cognition*. Harvard University Press; Cambridge, MA: 1983.
- Anderson, JR. *Rules of the mind*. Erlbaum; Hillsdale, NJ: 1993.
- Aston-Jones G, Cohen JD. An integrative theory of locus coeruleus-norepinephrine function: Adaptive gain and optimal performance. *Annual Review of Neuroscience* 2005;28:403–450.
- Bäckman, L.; Farde, L. The role of dopamine systems in cognitive aging. In: Cabeza, R.; Nyberg, L.; Park, D., editors. *Cognitive Neuroscience of Aging: Linking Cognitive and Cerebral Aging*. Oxford University Press; New York, NY: 2005. p. 58-84.
- Baddeley, AD. *Working memory*. Oxford University Press; Oxford: 1986.
- Baddeley AD. Exploring the central executive. *Quarterly Journal of Experimental Psychology* 1996;49A: 5–28.
- Baddeley, AD. *Working memory, thought, and action*. Oxford University Press; Oxford: 2007.
- Baddeley, AD.; Hitch, GJ. Working memory. In: Bower, GH., editor. *The psychology of learning and motivation: Advances in research and theory*. Vol. 8. Academic Press; New York: 1974. p. 47-89.
- Badyaev, AV. Stress-induced variation in evolution: From behavioral plasticity to genetic assimilation; *Proceedings of the Royal Society B*; 2005. p. 877-886.
- Baldwin JM. A new factor in evolution. *American Naturalist* 1896;30:441–451.
- Barnard, PJ. Interacting cognitive subsystems: A psycholinguistic approach to short term memory. In: Ellis, A., editor. *Progress in the psychology of language*. Vol. 2. Erlbaum; London: 1985. p. 197-258.
- Barnard, PJ. Interacting cognitive subsystems: Modeling working memory phenomena within a multiprocessor architecture. In: Miyake, A.; Shah, P., editors. *Models of working memory: Mechanisms of active maintenance and executive control*. Cambridge University Press; New York: 1999. p. 298-339.
- Barrett HC, Kurzban R. Modularity in cognition: Framing the debate. *Psychological Review* 2006;113:628–647. [PubMed: 16802884]
- Baumann O, Dames P, Kühnel D, Walz B. Distribution of serotonergic and dopaminergic nerve fibers in the salivary gland complex of the cockroach *Periplaneta americana*. *BMC Physiology* 2002;2:1–15. [PubMed: 11818034]
- Bechara A, Damasio A, Damasio H, Anderson S. Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition* 1994;50:7–15. [PubMed: 8039375]
- Bedau MA, Packard NH. Evolution and evolvability via adaptation of mutation rates. *Biosystems* 2003;69:143–162. [PubMed: 12689727]
- Behusi CV, Meck WH. What makes us tick? Functional and neural mechanisms of interval timing. *Nature Neuroscience* 2005;6:755–765.
- Berg AE. A simple objective technique for measuring flexibility in thinking. *Journal of General Psychology* 1948;39:15–22. [PubMed: 18889466]
- Bizo LA, White KG. Timing with controlled reinforcer density: Implications for models of timing. *Journal of Experimental Psychology: Animal Behavior Processes* 1997;23:44–55.
- Bonate, PL. *Analysis of pretest-posttest designs*. Chapman & Hall / CRC Press; Boca Raton, FL: 2000.
- Botvinick MM. Hierarchical models of behavior and prefrontal function. *Trends in Cognitive Science* 2008;12:201–208.
- Braver TS, Bongiolatti SR. The role of the frontopolar cortex in subgoal processing in working memory. *Neuroimage* 2002;15:523–536. [PubMed: 11848695]
- Braver, TS.; Cohen, JD. On the control of control: The role of dopamine in regulating prefrontal function and working memory. In: Monsell, S.; Driver, J., editors. *Attention and performance XVIII*. MIT Press; Cambridge, MA: 2000. p. 713-738.

- Braver TS, Cohen JD. Dopamine, cognitive control, and schizophrenia: The gating model. *Progress in Brain Research* 1999;121:327–349. [PubMed: 10551035]
- Braver, TS.; Gray, JR.; Burgess, GC. Explaining the many varieties of working memory variation: Dual mechanisms of cognitive control. In: Conway, ARA.; Jarrold, C.; Kane, MJ.; Miyake, A.; Towse, JN., editors. *Variation in working memory*. New York: Oxford University Press; New York: 2008. p. 76-106.
- Charnov EL. Optimal foraging: The marginal value theorem. *Theoretical and Population Biology* 1976;9:129–136.
- Cohen JD, Braver TS, Brown JW. Computational perspectives on dopamine function in prefrontal cortex. *Current Opinion in Neurobiology* 2002;12:223–229. [PubMed: 12015241]
- Cohen JD, McClure SM, Yu AJ. Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society of London B: Biological Science* 2007;362:933–942.
- Conway ARA, Cowan N, Bunting MF. The cocktail party phenomenon revisited: The importance of working memory capacity. *Psychonomic Bulletin and Review* 2001;8:331–335. [PubMed: 11495122]
- Conway ARA, Engle RW. Working memory and retrieval: A resource-dependent inhibition model. *Journal of Experimental Psychology: General* 1994;123:354–373. [PubMed: 7996121]
- Cools R, Gibbs SE, Miyakawa A, Jagust W, D'Esposito M. Working memory capacity predicts dopamine synthesis capacity in the human striatum. *The Journal of Neuroscience* 2008;28:1208–1212. [PubMed: 18234898]
- Cosmides, L.; Tooby, J. Origins of domain specificity: The evolution of functional organization. In: Hirschfeld, L.; Gelman, S., editors. *Mapping the mind*. Cambridge University Press; Cambridge: 1994. p. 85-116.
- Cowan, N. *Working memory capacity*. New York: Psychology Press; New York: 2005.
- Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ. Cortical substrates for exploratory decisions in humans. *Nature* 2006;441:876–879. [PubMed: 16778890]
- D'Esposito M, Detre JA, Alsop DC, Shin RK, Atlas S, Grossman M. The neural basis of the central executive system of working memory. *Nature* 1995;378:279–281. [PubMed: 7477346]
- Dougherty MR, Harbison JI. Motivated to retrieve: How often are you willing to go back to the well when the well is dry? *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2007;33:1108–1117.
- Due MR, Jing J, Klaidiusz RW. Dopaminergic contributions to modulatory functions of a dual-transmitter interneuron in *Aplysia*. *Neuroscience Letters* 2004;358:53–57. [PubMed: 15016433]
- Durstewitz D, Kelc M, Güntürkün O. A neurocomputational theory of the dopaminergic modulation of working memory functions. *Journal of Neuroscience* 1999;19:2807–2822. [PubMed: 10087092]
- Durstewitz D, Seamans JK, Sejnowski TJ. Dopamine-mediated stabilization of delay-period activity in a network model of prefrontal cortex. *Journal of Neurophysiology* 2000;83:1733–1750. [PubMed: 10712493]
- Engle, RW.; Kane, MJ. Executive attention, working memory capacity, and a two-factor theory of cognitive control. In: Ross, B., editor. *The psychology of learning and motivation*. Vol. 44. Elsevier; NY: 2004. p. 145-199.
- Engle, RW.; Kane, MJ.; Tuholski, SW. Individual differences in working memory capacity and what they tell us about controlled attention, general fluid intelligence, and functions of the prefrontal cortex. In: Miyake, A.; Shah, P., editors. *Models of working memory: Mechanisms of active maintenance and executive control*. Cambridge University Press; New York: 1999. p. 102-134.
- Fodor, JA. *The modularity of mind*. MIT Press; Cambridge, MA: 1983.
- Garavan H. Serial attention within working memory. *Memory & Cognition* 1998;26:263–276.
- Gavurin EI, Zangrillo RL. Effect of bigram cues on the solution of monosyllabic and bisyllabic solution-word anagrams. *Journal of General Psychology* 1975;92:225–229.
- Gick ML, Holyoak KJ. Schema induction and analogical transfer. *Cognitive Psychology* 1983;14:1–38.
- Gigerenzer, G.; Todd, PM.; the ABC Research Group. *Simple heuristics that make us smart*. Oxford University Press; New York: 1999.

- Goldstone RL, Landy DH, Son JY. The education of perception. *Topics in Cognitive Science*. in press.
- Goto Y, Otani S, Grace AA. The yin and yang of dopamine release: A new perspective. *Neuropharmacology* 2007;53:583–587. [PubMed: 17709119]
- Grace AA. Phasic versus tonic dopamine release and the modulation of dopamine system responsivity: A hypothesis of the etiology of schizophrenia. *Neuroscience* 1991;41:1–24. [PubMed: 1676137]
- Grace AA. The tonic/phasic model of dopamine system regulation and its implications for understanding alcohol and psychostimulant craving. *Addiction* 2000;95:S119–S128. [PubMed: 11002907]
- Hale, S.; Myerson, J.; Emery, LJ.; Lawrence, BM.; Dufault, C. Variation in working memory across the lifespan. In: Conway, ARA.; Jarrold, C.; Kane, MJ.; Miyake, A.; Towse, JN., editors. *Variation in working memory*. New York; Oxford University Press; New York: 2008. p. 194-224.
- Harbison JI, Dougherty MR, Davelaar EJ, Fayyad B. On the lawfulness of the decision to terminate memory search. *Cognition* 2009;111:397–402. [PubMed: 19344894]
- Hasher, L.; Zacks, RT. Working memory, comprehension, and aging: A review and a new view. In: Bower, GH., editor. *The psychology of learning and motivation: Advances in research and theory*. Vol. 22. Academic Press; San Diego: 1988. p. 193-225.
- Hasher, L.; Lustig, C.; Zacks, R. Inhibitory mechanisms and the control of attention. In: Conway, ARA.; Jarrold, C.; Kane, MJ.; Miyake, A.; Towse, JN., editors. *Variation in working memory*. New York; Oxford University Press; New York: 2008. p. 227-249.
- Hazy TE, Frank MJ, O'Reilly RC. Banishing the homunculus: Making working memory work. *Neuroscience* 2006;139:105–118. [PubMed: 16343792]
- Hertwig R, Barron G, Weber EU, Erev I. Decisions from experience and the effect of rare events in risky choice. *Psychological Science* 2004;15:534–539. [PubMed: 15270998]
- Hills TT. Animal foraging and the evolution of goal-directed cognition. *Cognitive Science* 2006;30:3–41.
- Hills TT, Adler FR. Time's crooked arrow: Optimal foraging and rate-biased time perception. *Animal Behavior* 2001;63:589–597.
- Hills TT, Brockie PJ, Maricq AV. Dopamine and glutamate control area-restricted search behavior in *Caenorhabditis elegans*. *Journal of Neuroscience* 2004;24:1217–1225. [PubMed: 14762140]
- Hills TT, Todd PM, Goldstone RG. Search in external and internal spaces: Evidence for generalized cognitive search processes. *Psychological Science* 2008;19:676–682.
- Hills, TT.; Todd, PM.; Jones, M. Optimal foraging in semantic memory. In: Taatgen, NA.; van Rijn, H., editors. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*. Cognitive Science Society; 2009.
- Houk, JC.; Davis, JL.; Beiser, DG., editors. *Models of information processing in the basal ganglia*. MIT Press; Cambridge, MA: 1995.
- Iwasa Y, Higashi M, Yanamura N. Prey distribution as a factor determining the choice of optimal foraging strategy. *American Naturalist* 1981;117:710–723.
- Kane MJ, Bleckley MK, Conway ARA, Engle RW. A controlled-attention view of working-memory capacity. *Journal of Experimental Psychology: General* 2001;130:169–183. [PubMed: 11409097]
- Kane, MJ.; Conway, ARA.; Hambrick, DZ.; Engle, RW. Variation in working memory capacity as variation in executive attention and control. In: Conway, ARA.; Jarrold, C.; Kane, MJ.; Miyake, A.; Towse, JN., editors. *Variation in working memory*. New York; Oxford University Press; New York: 2008. p. 21-48.
- Kane MJ, Engle RW. Working memory capacity, proactive interference, and divided attention: limits on long-term memory retrieval. *Journal of Experimental Psychology: Learning, Memory, Cognition* 2000;26:333–358.
- Kane MJ, Engle RW. Working-memory capacity and the control of attention: The contributions of goal neglect, response competition, and task set to Stroop interference. *Journal of Experimental Psychology: General* 2003;132:47–70. [PubMed: 12656297]
- Kieras, DE.; Meyer, DE.; Mueller, S.; Seymour, T. Insights into working memory from the perspective of the EPIC architecture for modeling skilled perceptual-motor and cognitive human performance. In: Miyake, A.; Shah, P., editors. *Models of working memory: Mechanisms of active maintenance and executive control*. Cambridge University Press; New York: 1999. p. 183-223.

- Laird JE, Newell A, Rosenbloom PS. SOAR: An architecture for general intelligence. *Artificial Intelligence* 1987;33:1–64.
- Laird NM, Ware JH. Random-effects models for longitudinal data. *Biometrics* 1982;38:963–974. [PubMed: 7168798]
- Ljungberg T, Apicella P, Schultz W. Responses of monkey dopamine neurons during learning of behavioral reactions. *Journal of Neurophysiology* 1992;67:145–163. [PubMed: 1552316]
- Logan GD. Working memory, task switching, and executive control in the task span procedure. *Journal of Experimental Psychology: General* 2004;133:218–236. [PubMed: 15149251]
- Logie RH, Zucco GM, Baddeley AD. Interference with visual short-term memory. *Acta Psychologica* 1990;75:55–74. [PubMed: 2260493]
- Lovett, MC.; Reder, LM.; Lebere, C. Modeling working memory in a unified architecture: an ACT-R perspective. In: Miyake, A.; Shah, P., editors. *Models of working memory: Mechanisms of active maintenance and executive control*. Cambridge University Press; New York: 1999. p. 135-182.
- Mangels JA, Ivry RB, Shimizu N. Dissociable contributions of the prefrontal and neocerebellar cortex to time perception. *Cognitive Brain Research* 1998;7:15–39. [PubMed: 9714713]
- Maricq AV, Church RM. The differential effects of haloperidol and methamphetamine on time estimation in the rat. *Psychopharmacology (Berl.)* 1983;79:10–15. [PubMed: 6403957]
- Mayzner MS, Tresselt ME, Helbock H. An exploratory study of meditational responses in anagram problem solving. *Journal of Psychology* 1964;57:263–274. [PubMed: 14143053]
- McClure, S.; Gilzenrat, M.; Cohen, J. An exploration-exploitation model based on norepinephrine and dopamine activity. In: Weiss, Y.; Sholkopf, B.; Platt, J., editors. *Advances in neural information processing systems*. Vol. 18. MIT Press; Cambridge, MA: 2006. p. 867-874.
- Meck WH. Selective adjustment of the speed of internal clock and memory processes. *Journal of Experimental Psychology: Animal and Behavioral Processes* 1983;9:171–201.
- Mendelsohn GA. An hypothesis approach to the solution of anagrams. *Memory & Cognition* 1976;4:637–642.
- Miller EK, Cohen JD. An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience* 2001;24:167–202.
- Miller, GA.; Galanter, E.; Pribram, KH. *Plans and the structure of behavior*. Holt; New York: 1960.
- Miyake, A.; Shah, P. *Models of working memory: Mechanisms of active maintenance and executive control*. Cambridge University Press; New York: 1999.
- Munakata, Y.; Morton, JB.; O'Reilly, RC. Developmental and computational approaches to variation in working memory. In: Conway, ARA.; Jarrold, C.; Kane, MJ.; Miyake, A.; Towse, JN., editors. *Variation in working memory*. New York; Oxford University Press; New York: 2008. p. 162-193.
- Nash, JC. *Compact numerical methods for computers: Linear algebra and function minimisation*. Adam Hilger; Bristol: 1990.
- Newell, A.; Simon, H. *Human problem solving*. Prentice-Hall; Englewood Cliffs, NJ: 1972.
- Nieoullon A. Dopamine and the regulation of cognition and attention. *Progress in Neurobiology* 2002;67:53–83. [PubMed: 12126656]
- Niv Y, Daw ND, Joel D, Dayan P. Tonic dopamine: Opportunity costs and the control of response vigor. *Psychopharmacology* 2006;191:507–520. [PubMed: 17031711]
- Norman, DA.; Shallice, T. Attention to action: Willed and automatic control of behavior. In: Shallice, T.; Norman, DA., editors. *Consciousness and self-regulation*. Plenum Press; New York: 1986. p. 1-18.
- Novick LR, Sherman SJ. The effects of superficial and structural information on online problem solving for good versus poor anagram solvers. *The Quarterly Journal of Experimental Psychology* 2008;61:1098–1120. [PubMed: 18938285]
- Oberauer, K.; Süß, H.; Wilhelm, O.; Sander, N. Individual differences in working memory capacity and reasoning ability. In: Conway, ARA.; Jarrold, C.; Kane, MJ.; Miyake, A.; Towse, JN., editors. *Variation in working memory*. New York; Oxford University Press; New York: 2008. p. 49-75.
- O'Reilly, RC.; Braver, TS.; Cohen, JD. A biologically based computational model of working memory. In: Miyake, A.; Shah, P., editors. *Models of working memory: Mechanisms of active maintenance and executive control*. Cambridge University Press; New York: 1999. p. 375-411.

- O'Reilly RC, Frank MJ. Making working memory work: A computational model of learning in the prefrontal cortex and basal ganglia. *Neural Computation* 2006;18:283–328. [PubMed: 16378516]
- O'Reilly, RC.; Munakata, Y. Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain. MIT Press; Cambridge, MA: 2000.
- Reiner A, Medina L, Veenman CL. Structural and functional evolution of the basal ganglia in vertebrates. *Brain Research Reviews* 1998;28:235–285. [PubMed: 9858740]
- Robbins H. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society* 1952;58:527–535.
- Rosen VM, Engle RW. Working memory capacity and suppression. *Journal of Memory and Language* 1998;39:418–436.
- Rougier, NP.; Noelle, DC.; Braver, TS.; Cohen, JD.; O'Reilly, RC. Prefrontal cortex and flexible cognitive control: rules without symbols; *Proceedings of the National Academy of Sciences*; 2005. p. 7338-7343.
- Salas C, Broglio C, Rodriguez F. Evolution of forebrain and spatial cognition in vertebrates: Conservation across diversity. *Brain, Behavior and Evolution* 2003;62:72–82.
- Samuels R. Evolutionary psychology and the massive modularity hypothesis. *The British Journal for the Philosophy of Science* 1998;49:575–602.
- Sankoff, D.; Kruskal, JB. Time warps, string edits, and macromolecules: The theory and practice of sequence comparison. Addison-Wesley; Reading, MA: 1983.
- Schneider W, Shiffrin RM. Controlled and automatic human information processing: I. Detection, search and attention. *Psychological Review* 1977;84:1–66.
- Schultz W. Multiple dopamine functions at different time scales. *Annual Review of Neuroscience* 2007;30:259–288.
- Sternberg S. High-speed scanning in human memory. *Science* 1966;153:652–654. [PubMed: 5939936]
- Stephens DW. Change, regularity, and value in the evolution of animal learning. *Behavioral Ecology* 1991;2:77–89.
- Szycypka, MS.; Rainey, MA.; Kim, DS.; Alaynick, WA.; Marck, BT.; Matsumoto, AM., et al. Feeding behavior in dopamine-deficient mice; *Proceedings of the National Academy of Sciences USA*; 1999. p. 12138-12143.
- Troyer AK, Moscovitch M, Winocur G, Alexander MP, Stuss D. Clustering and switching on verbal fluency: The effects of focal frontal- and temporal-lobe lesions. *Neuropsychologia* 1998;36:499–504. [PubMed: 9705059]
- Unsworth, N.; Heitz, RP.; Engle, RW. Working memory capacity and the control of attention. In: Engle, RW.; Sedek, G.; Hecker, U.; McIntosh, DN., editors. *Cognitive limitations in aging and psychopathology: Attention, working memory, and executive functions*. Oxford University Press; New York: 2004.
- Volkow ND, Gur RC, Wang G, Fowler JA, Moberg PJ, Ding Y, Hitzemann R, Smith W, Logan J. Association between decline in brain dopamine activity with age and cognitive and motor impairment in healthy individuals. *American Journal of Psychiatry* 1998;155:344–349. [PubMed: 9501743]
- Wilke, A. Evolved responses to an uncertain world. Free University of Berlin; 2006. Ph. D. thesis <http://www.diss.fu-berlin.de/2006/14/indexe.html>

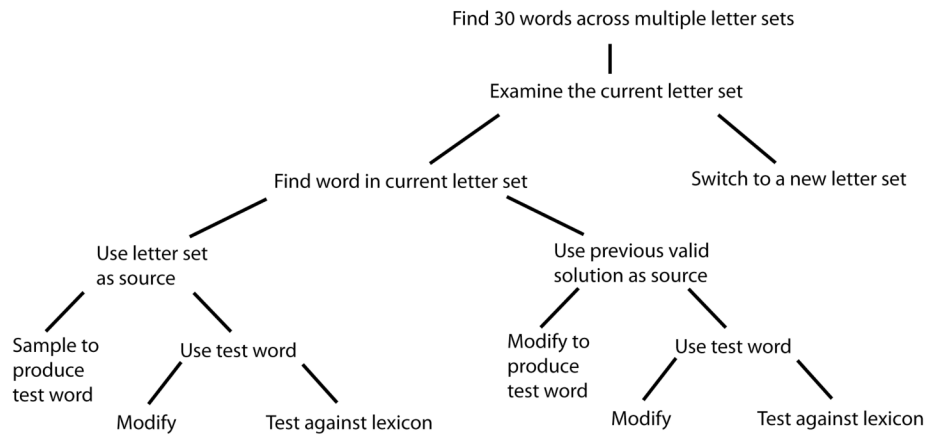


Figure 1.

One possible subgoal hierarchy for the lexical search task. The primary goal is to find 30 words across multiple letter sets. Branches of the tree represent possible subgoals for activation in working memory. Individuals start with the top level goal and make decisions between alternative subgoals to search the problem space.

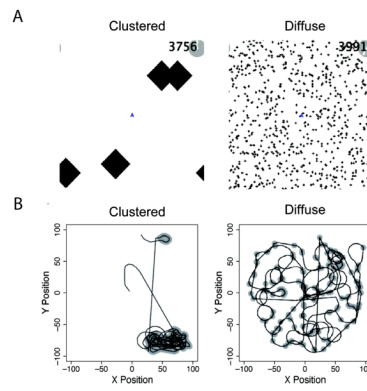


Figure 2.

A. Examples of clustered and diffuse resource distributions. Black pixels represent resources (not seen by participants until they pass over them in Experiment 1, or seen at the beginning of the task in Experiment 2). B. Example paths for single participants in clustered and diffuse treatments. Grey circles are positioned over the pixels where participants found a resource.




Figure 3. Priming results from Experiment 1. A. Time spent in each letter set, for each treatment group, across the first 7 letter sets. B. The bigram similarity between successive solutions within a letter set, for each treatment group. Error bars are standard error of the mean.

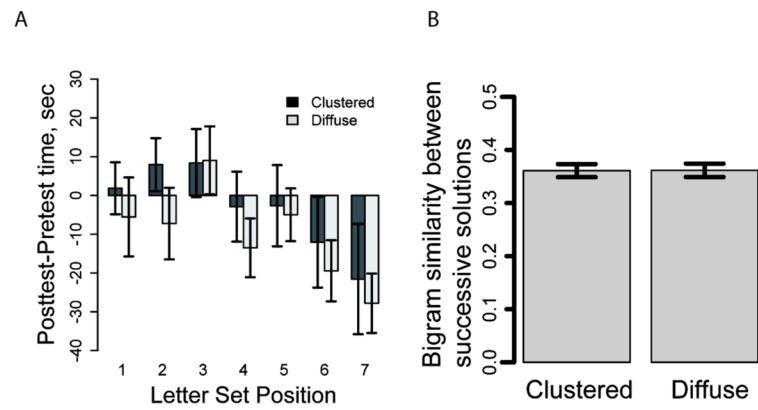


Figure 5. Priming results from Experiment 2. A. Time spent in each letter set, for each treatment group, across the first 7 letter sets. B. The bigram similarity between successive solutions within a letter set, for each treatment group. Error bars are standard error of the mean.




Figure 6.

The subgoal hierarchies for the three hypotheses. Moving from the top to the bottom hierarchy, bold text represents new subgoals added to the tree above. This has the consequence that Hypotheses 1 is a subset of Hypotheses 2, etc. Goal processing starts at the top and chooses between subgoals until it reaches a terminal branch. Subgoal processing is assumed to persist until either the subgoal is achieved and processing moves up to the parent goal, or the updating process allows switching to a goal at a higher level in the tree. Some subgoals have multiple, self-terminating components in their implementation, and these should be read from left to right.

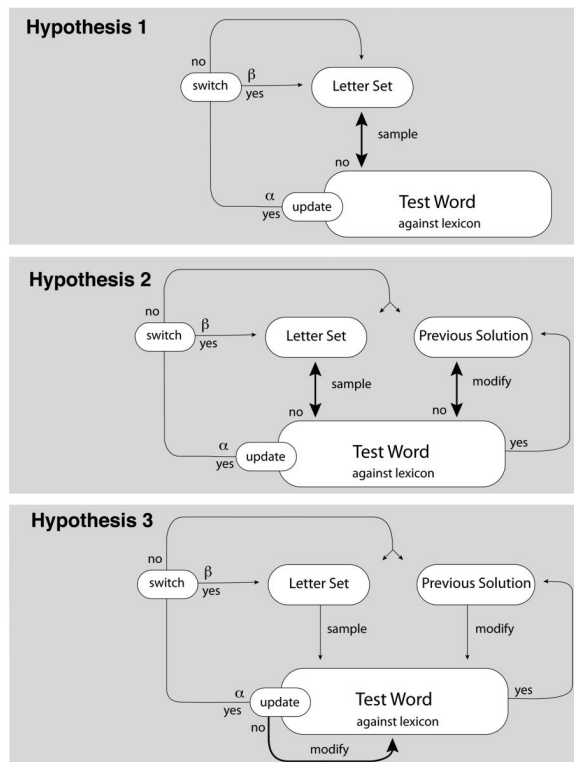


Figure 7. The CESP models for the three hypotheses. The models are derived from the subgoal hierarchies in Figure 6. Bold lines represent switching between the terminal word production and testing subgoals. Checking for updating occurs after each word is tested against the lexicon. For example, with Hypothesis 1, participants start with the letter set, and then sample from the letter set to produce a test word, which is tested against the lexicon. The sampling and testing is repeated until either a solution is found or, with probability α , the process updates to switch between subgoals at a higher level in the tree.




Figure 8.

The latency to leave the letter set—the time between the last submission and the decision to switch to a new letter set—as a function of number of words submitted in the current letter set. Participants take longer to decide to leave a letter set if they have found fewer words. Error bars are standard error of the mean. The line is the best fit line using Equation 1.

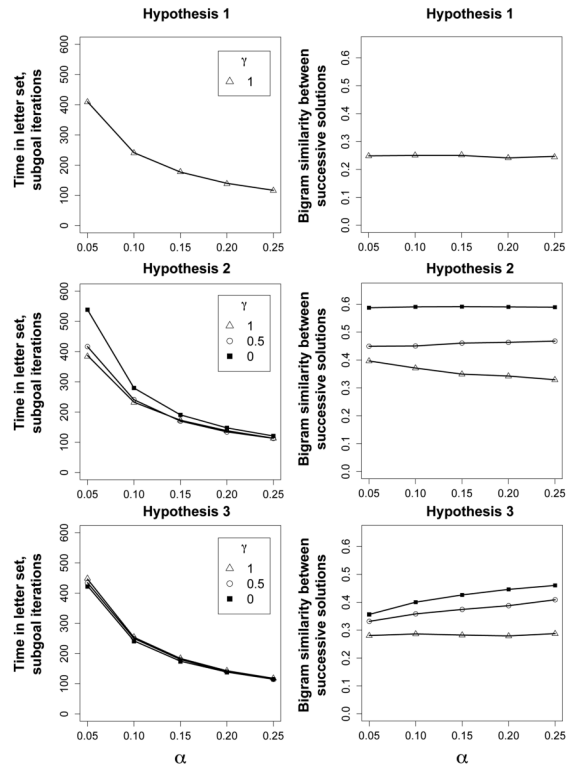


Figure 9. Results of the CESP model for the three hypotheses presented in Figure 7. 1000 participants were simulated for each value of α , for each of the three hypotheses. Iterations are equivalent to the number of times the algorithm tested the test word against the lexicon. γ is the probability of returning the test word source to the letter set (as opposed to the previous solution), following a source update. Only Hypothesis 3 shows the upward trend in between solution similarity observed in Figure 3.




Figure 10. Results of the CESP simulation of Hypothesis 3 in Figure 7, with 35 participants in each condition. A. The mean time spent in each letter set for the two different values of the subgoal updating parameter, α . Subgoal iterations represent the number of times a test word was produced and tested against the lexicon. B. The bigram similarity as a function of α .

Table 1

Analogies between spatial and lexical search tasks

Task	Spatial search	Lexical search
Resource	Pixels that change color	Words hidden in letter sets
Resource patches	Resource-filled regions of the screen	Level 1: Sets of letters Level 2: Sources of solutions within a letter set
Cost of transitioning between patches	Time needed to travel between patches	Level 1: 15 seconds between letter sets Level 2: Time to retrieve a letter sample from a new source
Exploration strategy	Move to unexplored screen regions and check for resources	Level 1: Give up on a current letter set and receive a new set Level 2: Give up on current solution source and transfer attention to another
Exploitation strategy	Remain in spatial region	Level 1: Keep trying to find words in current letter set Level 2: Keep using the same solution source