

Original citation:

Hills, Thomas Trenholm, Todd, Peter M. and Jones, Michael N.. (2015) Foraging in semantic fields : how we search through memory. *Topics in Cognitive Science*, 7 (3). pp. 513-534.

Permanent WRAP url:

<http://wrap.warwick.ac.uk/71020>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work of researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher statement:

"This is the peer reviewed version of the following article Hills, Thomas Trenholm, Todd, Peter M. and Jones, Michael N.. (2015) Foraging in semantic fields : how we search through memory. *Topics in Cognitive Science*, 7 (3). pp. 513-534. Doi:10.1111/tops.12151 which has been published in final form at <http://dx.doi.org/10.1111/tops.12151> . This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Self-Archiving."

A note on versions:

The version presented here may differ from the published version or, version of record, if you wish to cite this item you are advised to consult the publisher's version. Please see the 'permanent WRAP url' above for details on accessing the published version and note that access may require a subscription.

For more information, please contact the WRAP Team at: publications@warwick.ac.uk

warwick**publications**wrap

highlight your research

<http://wrap.warwick.ac.uk/>

Running head: HOW WE SEARCH THROUGH MEMORY

Foraging in semantic fields: How we search through memory

Thomas T. Hills¹, Peter M. Todd², Michael N. Jones²

¹Department of Psychology, University of Warwick

² Cognitive Science Program and Department of Psychological and Brain Sciences,
Indiana University Bloomington

Words: Abstract (249) + main text and figures (8069) + references (1348) + endnotes
(154) = 9666

Contact information:

Thomas Hills

University of Warwick
Department of Psychology
Gibbet Hill Road
CV4 7AL, Coventry UK
Tel: +44 (0) 2765 75527
Fax: +44 (0) 2765 75527
E-mail: t.t.hills@warwick.ac.uk

Abstract

When searching for concepts in memory—as in the verbal fluency task of naming all the animals one can think of—people appear to explore internal mental representations in much the same way that animals forage in physical space: searching locally within patches of information before transitioning globally between patches. However, the definition of the patches being searched in mental space is not well specified. Do we search by activating explicit predefined categories (e.g., pets) and recall items from within that category (categorical search), or do we activate and recall a connected sequence of individual items without using categorical information, with each item recalled leading to the retrieval of an associated item in a stream (associative search), or both? Using semantic representations in a Search of Associative Memory (SAM) framework and data from the animal fluency task, we tested competing hypotheses based on associative and categorical search models. Associative, but not categorical, patch transitions took longer to make than position-matched productions, suggesting that categorical transitions were not true transitions. There was also clear evidence of associative search even within categorical patch boundaries. Furthermore, most individuals' behavior was best explained by an associative search model without the addition of categorical information. Thus our results support a search process that does not use categorical information, but for which patch boundaries shift with each recall and local search is well-described by a random walk in semantic space, with switches to new regions of the semantic space when the current region is depleted.

Keywords: optimal foraging, verbal fluency, semantic space, natural categories

Introduction

Producing sets of things from memory is both common and critically related to higher cognitive function. We construct and recall grocery lists, things needed for camping trips, the best set of people to invite to a coming wedding, and potential names for a new pet. Previous research has shown that when we search for such items in memory, we produce clusters of related items and transition between clusters in much the same way that an animal forages within and between patches of food (Bousfield & Sedgewick, 1944; Raaijmakers & Shiffrin, 1981; Troyer, Moscovitch, & Winocur, 1997; Hills, Jones, & Todd, 2012). People's performance in terms of how many items they find in memory over time is highly influenced by when and how quickly they transition between these mental patches or clusters (e.g., when recalling animal names, Hills et al., 2012). The mean number of related items produced in each cluster is also related to other aspects of cognition, such as executive capacity (Rosen & Engle, 1997; Troyer, Moscovitch, Winocur, Alexander, & Stuss, 1998; Hills & Pachur, 2012), age (Hills, Mata, Wilke, & Samanez-Larkin, 2013), and clinical conditions including Alzheimer's disease, depression, and Parkinson's disease (Fossati, Le Bastard, Eris, & Allilaire, 2003; Haugrud, Crossley, & Vrbancic, 2011; Murphy, Rich, & Troyer, 2006; Raoux et al., 2007; Taler & Phillips, 2008).

Knowing how clusters of related items are retrieved from memory is thus central to our understanding of how we use memory for recalling concepts more generally. However, despite recent advances in understanding long-term memory recall as a dynamic "foraging-like" process, where the searcher makes global transitions between

local “patches” in the memory space (e.g., Hills et al., 2012; Hills & Pachur, 2012; Hills et al, 2013; Todd, Hills, & Robbins, 2012), what exactly the patch is in memory space is still a matter of some debate. A patch comprises the items that are found between the global long-distance jumps that indicate a patch-leaving event; it is the local region of the search space through which memory moves. We are trying to understand how that space is defined in memory search. When we search for multiple concepts in memory, what constitutes the patch (i.e., the local region of the memory space) from which we are retrieving? Two major possibilities are the center of ongoing study: Do we search by activating entire predefined categories as patches, and recall items from within those category patches (categorical search)? Or do we activate and recall a connected sequence of individual items, with each new item recalled leading to a moving window of new retrieval candidates (associative search), consistent with a random walk between items? These hypothesized processes are not mutually exclusive—different individuals could do one or the other, and a given individual might combine the two at the same time. In both cases, a searcher may leave a patch when the local resources remaining in the local region are sufficiently low: With categorical search the individual may leave a patch when a category has few or no remaining items, and with associative search they may leave when the latest item found has only weak associations with other remaining items in memory.

In this paper, we compare these two hypotheses with competing models applied to data from a recall study using the verbal fluency task, in which participants say all the types of animals they can remember in three minutes. We do this by identifying likely patch switches in participants’ sequences of retrieved items and then testing whether response times, word similarity measures, and patch-leaving rules better fit categorical or

associative search, and by testing the fit of both types of models to the actual sequences of words that individual participants produce. Before we confront these two hypotheses with our data, we first outline the evidence for patch-based memory foraging for categorical and associative definitions of mental patches.

Foraging in memory patches

The patchiness of memory recall is evident in studies of free recall from natural categories, with clustered recall of related items noted in the earliest such studies (Bousfield & Sedgewick, 1944; Johnson, Johnson, & Mark, 1951). Because these clusters often appear to represent semantic associations, Gruenewald and Lockhead (1980) called these clusters “semantic fields” after Ipsen’s (1924) *Bedeutungsfeld*. More recent work on free recall from natural categories and list learning consistently finds that groups of semantically similar words are produced together (Bousfield, 1953; Howard, Addis, Jing, & Kahana, 2007; Gruenewald & Lockhead, 1980; Romney, Brewer, & Batchelder, 1993).

Groupings of semantically similar words during recall are consistent with a cognitive search process that modulates between local and global memory cues. Indeed, this two-stage search policy is common to several different models of long-term memory retrieval (Raaijmakers & Shiffrin, 1981; Metcalfe & Murdock, 1981). One of the most successful models of memory retrieval from natural categories is Search of Associative Memory (SAM; Raaijmakers & Shiffrin, 1981; Gronlund & Shiffrin, 1986; Walker & Kintsch, 1984). In SAM, the memory probe used to activate and retrieve memory items alternates between local information (such as item-by-item similarity or a subordinate

category) and global information (such as item frequency). During epochs of probing memory locally, sequences of similar items are retrieved. When local resources are sufficiently depleted, the memory probe temporarily transitions to global information, and the first item of the next cluster is recalled from some new location in the memory space. (For a discussion of cues of depletion in memory search see Davelaar & Raaijmakers, 2012; Dougherty & Harbison, 2007; Harbison, Dougherty, Davelaar, & Fayyad, 2009; Hills et al., 2012.) Thus, memory foraging is modeled in SAM as being adaptively modulated between local exploitation of patches and global exploration to find new patches when previous ones are depleted.

However, SAM's definition of the local patch search process is left open to different possibilities. For example, when retrieving different types of animals, the local search could make transitions based on item-wise similarity (e.g., going from “dog” to “cat” because of the semantic similarity between them) or based on subordinate category membership (e.g., “dog” and “cat” are produced in succession because they are both in the subcategory “pet”), or both.

Some of the earliest investigations of associative versus categorical retrieval processes had participants memorize lists of words from several categories and found that retrieval was faster within than between categories (e.g., Pollio, Richards, & Lucas, 1969). This was interpreted as support for a categorical retrieval process. More recently, the categorical hypothesis has been revived as the *cluster-switching hypothesis* (Troyer et al., 1997). This involves “clustering”—the activation and production of words within a semantic subcategory—and “switching”—making the transition from one subcategory to

another (Gruenewald & Lockhead, 1980; Troyer et al., 1997; Troyer, Moscovitch, Winocur, Leach, Freedman, 1998; Robert et al., 1998).

Under the cluster-switching hypothesis, people search initially for a cluster (or identify an initial item within a cluster) of categorically-related words (the “patch”) and then search within that cluster, moving between clusters when the current cluster becomes depleted. For example, when recovering items from the category of “animals” one may first find the cluster “pets” and then recover individual pets, such as “dog”, “cat”, and “gerbil”, before proceeding to the next cluster, for example “farm animals”. This approach assumes that people have well-defined subcategory boundaries for semantic patches; testing such models requires specifying those subcategories, either by hand (as in Troyer et al., 1997) or automatically (Goñi et al., 2011).

In earlier work, we used a two-stage, local-to-global associative search hypothesis (i.e., a random walk over a semantic representation with global switches when the local region is depleted) to model semantic fluency data for the “animals” category (Hills et al., 2012) and found that it explained much of participants’ performance. Other approaches to memory retrieval from natural categories have similarly assumed a random-walk associative retrieval process (Abbott, Austerweil, & Griffiths, in press; Thompson & Kello, 2014). In a related social fluency task (“say all the people you know”), a direct comparison of associative and categorical models found support for an associative retrieval process (Hills & Pachur, 2012). That is, the pattern of recall was better predicted by whether the participant thought that two people knew one another than by whether the two people were in the same social category (e.g., family). However, that study lacked an exhaustive list of social categories, and therefore might have

misrepresented the categories used by participants. Thus, the question of whether and when people use associative or categorical search in free recall remains open. In this paper, we explicitly compare the associative process with the competing categorical search hypothesis.

What is the patch in memory foraging?

The metaphor of foraging in memory fits well with local-to-global search models because they capture key features of patch-foraging behavior in animals. This process of memory retrieval is similar to patterns of animal foraging called area-restricted search, in which animals stay near patches where they have found resources in the past but move away from locations where resources have not been found (Karieva & Odell, 1987; Grünbaum, 1998; Hills, Kalff, & Wiener, 2013). Moreover, both animals seeking food and individuals searching patchy memory must decide when to leave one region of the memory space and move to another, and both appear to involve similar control decisions as studied in optimal foraging theory (Charnov, 1976; Pirolli, 2007; Hills et al., 2012; Wilke, Todd, & Hutchinson, 2009). Finally, implementation of cognitive control in humans shares biological bases with mechanisms for controlling foraging in non-human species (e.g., using dopaminergic neuromodulation) and may therefore share a common evolutionary history (Hills, 2011; Todd, Hills, & Robbins, 2012; Hills et al., 2015). But to get further mileage out of the metaphor of patch-based memory foraging and generate new questions and insights, we need to specify its central component: What is the patch?

To an animal such as a ladybird beetle eating aphids, a patch might comprise an individual aphid, a cluster of aphids, a branch of a plant, a whole plant, or even a field of

plants (Karieva & Odell, 1987). In this sense, a patch is a useful conceptual construct to define a local region of the search space. However, it does not need an explicit boundary except in relation to the organism's behavior. Determining the boundaries of patches as perceived and searched by an organism is not trivial; numerous ecological approaches have been proposed (e.g., Fauchald, 1999; Grünbaum, 1998). Uncovering the nature of patches or local search regions in memory presents a potentially more challenging problem, because we must also make assumptions about each participant's internal memory structure—in particular, over what kinds of representations do people search when recalling information from memory. Critically, *a patch is a consequence of both the structure of the representation and the mechanism used to recall information from that structure*. Our approach is to compare predictions of the hypothesized associative and categorical patch search models, which can be seen as widely-separated points on a continuum between using previous items versus over-arching subcategories to activate and find new items in memory.

To provide some intuition for the distinction between associative and categorical patches, an example using Troyer et al.'s (1997) category codings is helpful. With these codings, one can define transitions across patch boundaries in two distinct ways (see Figure 1).¹ For the categorical search model, a switch between patches occurs whenever a person produces an item that is not in the same subcategory as all the items produced since the last patch switch (e.g., in Figure 1, leaving the “PETS” subcategory when moving from “dog” to “wolf” and leaving the “CANINE” subcategory when moving from “wolf” to “cow”). Therefore, to be defined as a categorical patch, all successively-produced items in a patch must share at least one common subcategory membership. The

associative patch encompasses a succession of items in which each neighboring pair is linked by at least one common subcategory membership; therefore, a patch switch only occurs when two consecutive items do not share *any* subcategory (e.g., in Figure 1, from “wolf” to “cow”). Thus, associative switches are always determined relative to the previous item alone.

Note that, by definition, associative switches are also categorical switches, but the reverse need not be true. Thus, the simple existence of one patch type or another is not diagnostic of the process used to produce patches: A person using an associative search model could occasionally produce category switches within associative patches (e.g., the “dog”-“wolf” categorical switch that is not an associative switch in Figure 1). As well, a person using a categorical search model could by chance produce items that are associates across the categorical boundary, and thus also produce categorical but not associative switches (again, “dog”-“wolf”). The problem of identifying the process that produces a patch type is therefore challenging, and relies on comparing patch transitions as identified by the two models with other features of the production data, such as reaction times and between-item similarity.

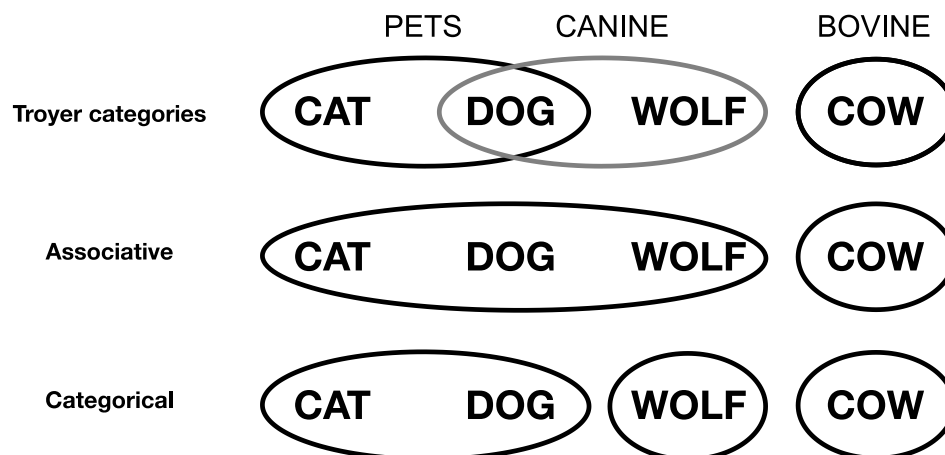


Figure 1: Illustration of switches in two patch types. The items retrieved sequentially from memory in response to the global cue “animals” are “cat”, “dog”, “wolf”, and “cow”. The top panel shows the subcategory membership as assigned by Troyer et al., PETS, CANINE, and BOVINE, which we use as proxy for defining switches. The middle panel indicates how the associative model would assign patch boundaries, with “cat”, “dog”, and “wolf” all produced in a sequentially connected patch of subcategories. The lower panel indicates how the categorical model would assign patch boundaries, with the transition from “dog” to “wolf” considered a switch because “wolf” is not in the category intersection of “dog” and “cat” (i.e., not in the subcategory PETS). The transition from “dog” to “wolf” is a categorical-only switch, while “wolf” to “cow” is a switch for both models.

We can define these patch types more rigorously by specifying the algorithms that would produce them. In the case of the categorical search model, item recovery would proceed as follows:

1. Start with the global cue (e.g., animals);
2. Pick a (new) subcategory label (which could be done by picking an item in some way, e.g. by frequency, and identifying its most prevalent subcategory);
3. Retrieve the first/next item in that subcategory in some way (e.g., sample randomly within the subcategory or based on similarity to the previous item retrieved—we discuss this further below);
4. Repeat step 3 until the conditions of a local stopping rule are met (e.g., no items are left in the subcategory, or too much time has elapsed without finding an item);
5. Return to step 2 and repeat until the conditions of a global stopping rule are met (e.g., the experimenter indicates that time is up).

For the associative model, retrieval would proceed as follows:

1. Start with the global cue (e.g., animals);
2. Retrieve the first item in a new patch with probability proportional to its frequency;
3. Retrieve the next item with probability proportional to its semantic similarity with the previous item retrieved;
4. Repeat step 3 until the conditions of a local stopping rule are met (e.g., there are no items found with sufficiently high semantic similarity to the previous item);
5. Return to step 2 and repeat until the conditions of a global stopping rule are met (e.g., the experimenter indicates that time is up).

(Note that in this model, the searcher does not need to have a subcategory in mind; clustering and switching occurs instead through semantic similarity between successive items.)

As mentioned above, people may use either of the above algorithms, or a combination of both in the same or different searches. For example, as we note above, a categorical search process could use associative search within a category. In the present study we aim to rule out some of these possibilities and to identify the predominant search process used by each individual, and hence the nature of the memory patches that they produce.

The Present Study

To test whether people forage in memory with search strategies that produce patches that are predominantly of one type or the other, we used the semantic fluency task for the category of “animals” to examine the sequence of items recovered from memory.

Categorical and associative patches can be defined in terms of a predetermined set of subcategories, as just explained; for this purpose, we used an augmented version of the hand-coded subcategorization of animal terms derived by Troyer et al. (1997) described further below. We compare evidence for the two patch search types first in terms of their predictions regarding the patterns of response times and semantic similarity between successive pairs of recalled items. Specifically, both types of search predict that response times between items should be highest when switching between patches, but they may differ in terms of the predicted similarity between any two items in a patch. An

associative patch search predicts that similarity between two items within a patch is lower when they are separated by more intervening items. Categorical patch models may differ in this regard depending on how search is assumed to occur within patches (see below). To determine item similarity, we use the BEAGLE lexical semantic memory model of Jones and Mewhort (2007). Following this analysis, we test whether an associative or categorical patch search model better fits each participant's individual sequence of retrieved items.

Methods: Modeling search in semantic memory

To model search in semantic memory, a structural representation of the search space is required, along with a model of the search process. We represent the structure of semantic memory using both hand-coded (Troyer) and statistically derived (BEAGLE) schemes (described next). These are used within a process model of semantic search based on local and global transitions with respect to item patches.

The actual search data analyzed in this study was taken from the semantic fluency memory recall experiment in Hills et al. (2012). Participants (141 undergraduates at Indiana University participating for course credit) at computers in a lab were asked to enter via the computer keyboard as many types of animals as they could in three minutes, with animal name and time of entry recorded for each item. Each animal entry was only visible until the "return" key was pressed, so that participants could not see previous entries. Data were hand-corrected for spelling and non-animal names were removed.

Representing the Structure of Semantic Memory

The Troyer Categorization Scheme. The original Troyer et al. (1997; see also Troyer, 2000) categorization scheme contains 155 animals put into 22 non-exclusive categories, for example, “African animals,” “water animals,” “beasts of burden” etc. This scheme was constructed from the actual sequences of words produced by their participants, and we extended it with the 214 additional animal names generated by our participants, placing them into the same 22 Troyer subcategories based on Wikipedia descriptions of those animals (see Hills et al., 2012). Support for the Troyer et al. categorization comes via its utility in detecting specific clinical backgrounds (e.g., Raoux et al., 2007; Murphy, Rich, & Troyer, 2006; Fossati, Le Bastard, Ergis, & Allilaire, 2003; Troyer et al., 1998) and its previous validation in memory tasks (Hills et al., 2012).

BEAGLE Semantic Representations. The statistically derived semantic representations we use here (taken from Hills et al., 2012) were computed using the BEAGLE model (Jones & Mewhort, 2007; Jones, Kintsch, & Mewhort, 2006), trained exclusively on contextual information, which produces similarity structure quite comparable to other high-dimensional semantic space models (e.g., Landauer & Dumais, 1997; Lund & Burgess, 1996). BEAGLE was trained on a 400 million-word Wikipedia corpus (Willits, D’Mello, Duran, & Olney, 2007), and then its memory representations were used to compute the pairwise cosine similarity matrix for a list of 765 animals. The corpus was preprocessed to fuse any multiword animal exemplars into a single lemma prior to training (e.g., *killerwhale*, *kingcobra*). BEAGLE provides measures of pairwise similarity between words based on their co-occurrence, and the co-occurrence of related words, in a natural language corpus; these similarities have been shown to successfully predict

lexical decision times (Jones & Mewhort, 2007) as well as semantic priming data (Jones et al., 2006). Pairwise similarities computed by BEAGLE for a large class of animals offer a structural representation of the semantic search space, which we can then use to predict the retrieval of association-based or category-specific instances from memory.

Modeling the Semantic Memory Search Process

After testing general predictions of the two types of patches using the representations just described, we model individual-level semantic fluency retrieval sequence data. The model framework we use to simulate the process of search in semantic memory is common to both the SAM and ACT-R memory model architectures (described in Raaijmakers & Shiffrin, 1981; Anderson, 1993). The foundational assumption of our model framework (Hills et al., 2012) is that recall is achieved by probing retrieval structures in memory with a specific set of cues, called the memory probe Q . For any possible target item I that could be found in the semantic space for a particular category, the probability of retrieving I is based on the product of the individual activation strengths, A , for I across all M cues in the memory probe Q , relative to the sum of all such products for all other items. The overall probability of retrieval for item I based on memory probe Q using the ratio rule is then:

$$P(I_i | Q_1, Q_2, \dots, Q_M) = \frac{\prod_{j=1}^M A(Q_j, I_i)^{\beta_j}}{\sum_{k=1}^N \prod_{j=1}^M A(Q_j, I_k)^{\beta_j}}, \quad 1$$

where N represents the total number of items I available for retrieval, and β_j represents the saliency (or attention weight) assigned to a given cue Q_j .

We employ here a dynamic, two-stage model that shifts between using a global cue, word frequency, and one or two local cues, the previous item or two items recalled. Frequency represents a *global* search cue, which generates a retrieval strength $A(Q, I) = \log(\text{freq}(I))$ for each item I based on the log of that item's frequency of occurrence in the Wikipedia corpus, with Q here simply being the category "animals". (Other global cues, such as eigenvalue and typicality of an item measured as summed cosine similarity to all other items in the same category, were found to be less effective predictors of free recall behavior.) The previous-item cues are *local* search cues, which generate a retrieval strength for a new item I based on its semantic similarity with previously-generated items—here, when the probe cue is the single previously-generated item, $Q_i = I_{t-1}$, then $A(Q, I) = S(I_{t-1}, I)$, where $S(I_{t-1}, I)$ is equal to the cosine similarity between the two items as computed by BEAGLE.

Allowing that memory search involves local-to-global transitions (see Hills et al., 2012), our model exploits the patchy structure of the memory environment, switching from patch to patch by changing the contents of the memory probe where local-to-global transitions occur. Specifically, when leaving a patch, the model switches from the use of the previous-item cue(s) (similarity-based local search) to the frequency cue (context-based global search) to find a new appropriate patch, and then back again to the previous-item cue(s) as the new patch is entered. For example, a sequence of DOG-CAT-HAMSTER-HORSE may transition from the local previous-item cue to the global frequency cue following HAMSTER, and thereby retrieve the high frequency item HORSE that is not semantically similar to HAMSTER. (For examples of this modeling approach applied to other types of search, see Hills & Pachur, 2012). Where exactly

these patch switches are defined to occur depends on whether we are using an associative patch definition or a categorical patch definition; these two different definitions are incorporated into the two versions of the model we compare below, along with one or two previous-item cues in each case. We find the best-fitting model of each type for each participant by using maximum likelihood estimation to fit the β parameters to each participant's individually generated sequence of items (for more details, see Hills et al., 2012; Hills, Mata, Wilke, Samanez-Larkin, 2013)².

Results

Participants produced a combined total of 5187 valid animal entries. As reported in Hills et al. (2012), the mean number of animal words per participant was 36.8 (SD = 8.5); the maximum number produced by any participant was 61 items and the minimum was 16. The mean recovery time per word—averaging inter-item retrieval times first within and then between participants—was 5.4 seconds (SD = 1.5).

Aggregated analysis of categorical and associative search models

We first investigated the statistical properties of recalled items with respect to the categorical and associative patch definitions given earlier. Using the hand-coded Troyer subcategories with the categorical patch definition, there were a mean of 18.1 patch switches per participant (SD = 4.7; representing 50% of productions), and the associative definition yielded a mean of 17.2 switches (SD = 4.5; representing 47% of productions). Mean items retrieved per patch for the categorical and associative definitions were 2.0 (SD = 0.5) and 2.1 (SD = 0.5), respectively. These statistics are highly similar when

compared between the two different patch types ($p > .1$), because most patch switches identified by one model are also identified by the other model. There were 2425 switches that were identified by both the categorical and associative patch definitions, and 134 more patch switches that were categorical-only (i.e., categorical-but-not-associative switches). These categorical-only switches were spread over 86 (61%) of participants.

The relatively few categorical-only switches clearly indicates the challenge of identifying the patches that people search. For example, consider the set of first categorical-only patch switches from the first seven participants in our data set: aardvark → elephant; gazelle → deer; hamster → fish; moose → muledeer; buffalo → bison; rabbit → bunny; lion → lioness. These are clearly within-subcategory associative transitions, but because they follow a stream of productions from another subcategory, they represent a (categorical) switch between subcategories. For example, while aardvark and elephant are both in the African animal subcategory, aardvark follows bat in this participant's sequence of retrievals, and because both of those are in the insectivore subcategory, when this subcategory is left and elephant is retrieved, this counts as a categorical switch. Thus, these categorical-but-not-associative transitions mark instances where individuals appear to make a categorical switch between subcategories while maintaining an association between the successive items, blurring the distinction between the category types when considered just at the level of between-item similarity; hence we turn to further measures.

Analysis of the IRTs and semantic similarity of patch transitions

Next we looked for temporal evidence that people are searching through patches in either an associative or categorical manner. We examined how long it takes participants to retrieve items following different transition types defined in terms of the Troyer subcategories: associative switches (which are also always categorical switches), categorical-only switches (which are not associative switches), and non-switches. To determine which type of proposed switches best reflects actual switches between patches made by participants, we compare the inter-item retrieval times (IRTs) produced by participants at those proposed switch locations. True patch switches should show longer IRTs than within-patch retrievals (Gruenewald & Lockhead, 1980). This is because patch switches involve a giving-up process to decide to leave the old patch, an exploration process to find a new patch, and the first item retrieval from the new patch. By comparing the IRTs of the different switch types, we can determine which kinds of patch transitions potentially represent true patch leaving.

As shown in Figure 2, the results of this test indicate that the different switch types take different lengths of time. The mean IRT of all associative switches—representing both associative and categorical switches—was 6.4s ($SD = 6.3$). The 134 categorical-but-not-associative switches had a mean IRT of 4.8s ($SD = 4.6$). Non-switches took 3.8s ($SD = 3.8$). Because switches take place at different ordinal positions (earlier or later in the series), in the different patch types, we also computed the weighted average of the IRTs for non-switches matched to the proportion of switches taking place at different ordinal positions in the different patch types, ranked according to their position in the patch. As the variances were different, we used a Wilcoxon rank sum test. Both categorical-only and associative switches took significantly longer than non-

switches (categorical only: $W = 205575$, $p < .001$; associative switches: $W = 2269717$, $p < .001$). Associative switches also took significantly longer than the categorical-only switches ($W = 132655$, $p < .001$) and significantly longer than the non-switch IRTs matched to the same ordinal position ($W = 2287289$, $p < .001$). However, categorical-only switches did not differ significantly from the non-switch IRTs matched to the same ordinal position ($W = 4746$, $p = .62$).

We also compared the cosine similarity for the word pairs that straddle different transition types, with the expectation that true switches would show lower similarity between the two words than non-switches. Figure 3 shows that both associative (and categorical) switches and categorical-only switches had lower mean cosine similarities than did non-switches. A Wilcoxon test indicates a statistically significant difference between non-switches and associative/categorical switches ($W=4417020$, $p < .001$), and between non-switches and categorical-only switches ($W=211037$, $p < .001$). Associative/categorical switches and categorical-only switches were also significantly different, with categorical-only switches producing words of higher similarity across the switch ($W=139726$, $p < .047$). Furthermore, both associative/categorical and categorical-only switches had smaller cosine similarities than their corresponding ordinal position-matched cosine similarities for non-switches (associative/categorical: $W = 3141269$, $p < .001$; categorical-only: $W = 2554$, $p < .001$).

An additional difference between the switch types is found in relation to the frequency with which they are produced over the 3-minute retrieval interval. Figure 4 shows the different switch types as a proportion of all retrievals over different quartiles of production numbers (e.g., first 25% of items, second 25%, etc.). A mixed-effects logistic

regression—with random intercepts for participants and slopes for retrieval intervals—finds a significant increase in associative/categorical switches over successive intervals ($\chi^2(1)=25.95, p<.001$).³ However, a similar test for categorical-only switches finds no difference as a function of retrieval interval ($\chi^2(1)=0.46, p=.49$).⁴

In summary, both switch types take longer to produce than non-switches, but switches defined as associative (and categorical) take longer to produce than categorical-only switches. Categorical-only switch types do *not* take longer than ordinal position-matched non-switches. Comparisons of semantic similarity reveal similar differences between switch types, with associative switches being less similar than categorical-only, and both groups less similar than the ordinal position-matched non-switches. Finally, only associative/categorical switches increase as a proportion of retrievals, while categorical-only switches do not. This further suggests that associative/categorical switches are a function of using an associative search process, because the few instances of categorical-only switches do not share the same properties.

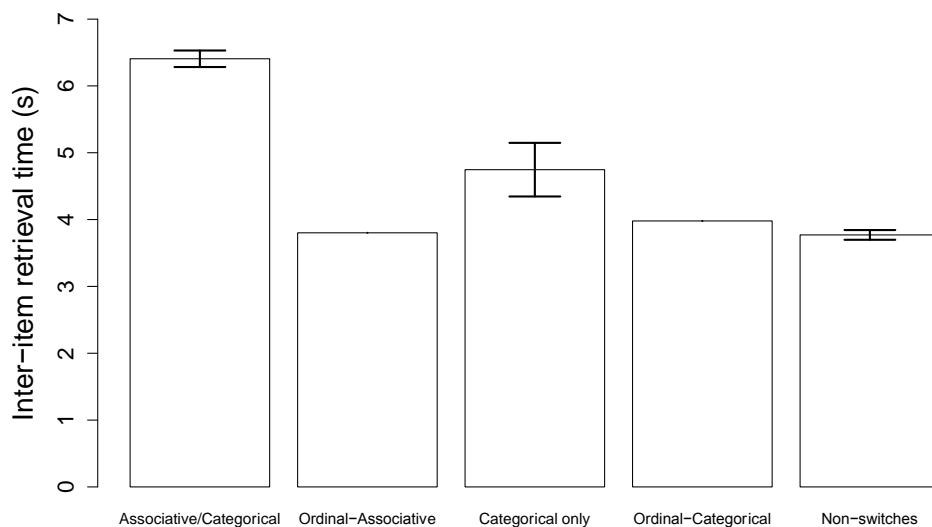


Figure 2. Comparison of mean inter-item retrieval times (IRTs) for different types of transitions between items. From left to right, associative (and categorical) switches, non-switches matched to the ordinal position of associative switches, categorical-only switches, non-switches matched to the ordinal position of categorical switches, and all non-switches. Error bars are SEM.

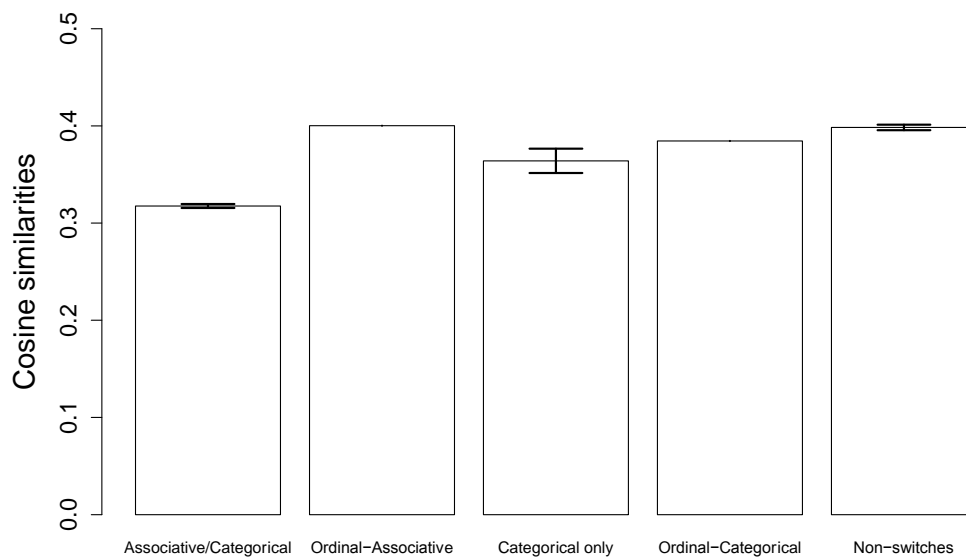


Figure 3. Comparison of mean cosine similarities between pairs of words at different types of transitions. From left to right, associative (and categorical) switches, non-switches matched to the ordinal position of associative switches, categorical-only switches, non-switches matched to the ordinal position of categorical switches, and all non-switches. Error bars are SEM.

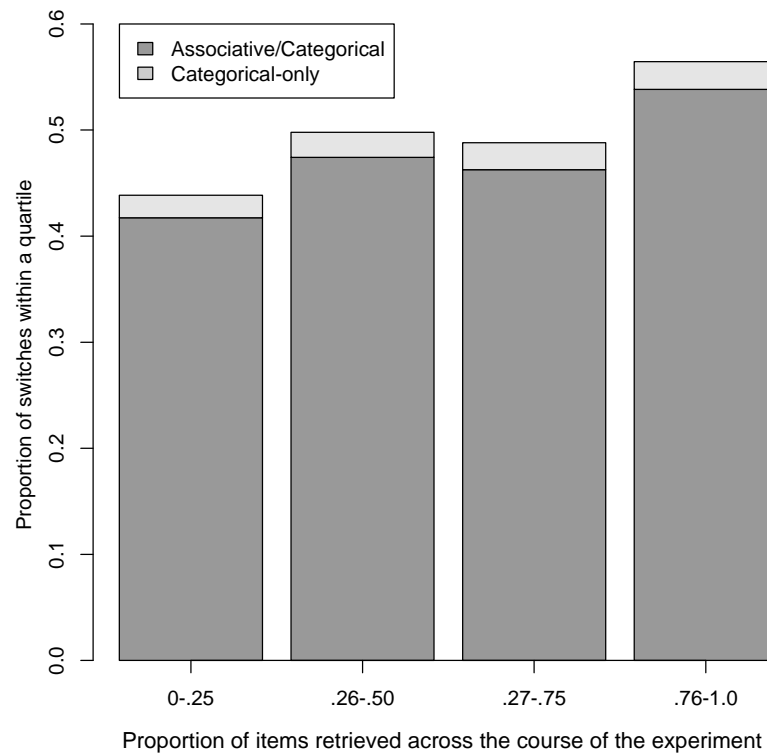


Figure 4. Proportion of different switch types out of all item retrievals, divided into the first, second, third, and fourth quartile of items retrieved across the course of the experiment.

Analysis of semantic similarity within a patch

When retrieving information from memory, if participants are using semantic associations between words to cue the next item retrieved, then words they produce consecutively should have heightened semantic similarity in relation to one another. This could be imagined as a wave of activation moving through the semantic space, with its peak located around the most recently recalled item. But if participants are instead using *only* categorical information to guide their recall (i.e., choosing the next item based on

subcategory membership and not based on similarity to the previous item), then we could expect that the semantic similarity between any pair of items within a categorically-defined patch should on average be the same for all pairs, so that two words produced near one another in a patch will not systematically be more similar to one another than more separated pairs within the same patch.

In previous work we computed the semantic similarity between each word retrieved and the last word in the same associative patch and found that similarity increased the closer the word was to the last one (Figure 1 in Hills et al., 2012).⁵ To make sure this proximity-similarity relationship was not just seen at the end of a patch (where words are lower frequency), in Figure 5 we show the semantic similarity between the last word in a patch and the words that come before it in the recall sequence (positions -5 to -1), and between the first word in a patch and the words that come after it (positions +1 to +5). Here we see that similarity does grow with proximity in both directions. However, the most dramatic change in similarity occurs for words just before and just after the last or first recalled word—suggesting local memory search reflects a Markov process, where only the previous item is relevant (consistent with Hills et al., 2012). This similarity pattern is fairly strong evidence for an associative search pattern, even if participants are also using categorical information in some way (e.g. for deciding

when to switch).

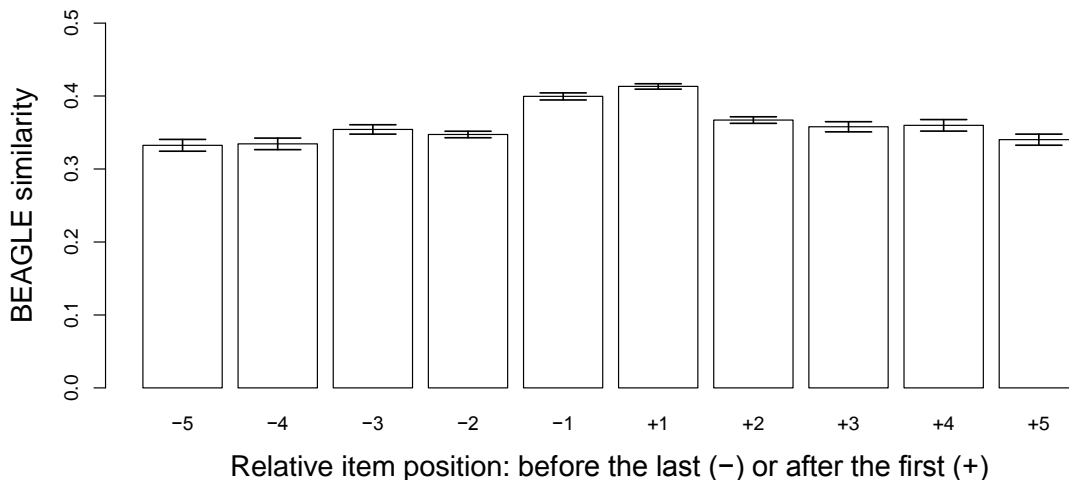


Figure 5: The semantic similarity between words at various positions upstream and downstream from a patch switch and the last (for positions -5 to -1) or first (+1 to +5) word within a patch. Error bars are SEM.

Fitting the associative and categorical models to individual behavior

Though the absolute number of categorical or associative switches differs by less than 10% over all item retrievals, we can test which version of the patch search model, categorical or associative, better fits the data at the individual level. To focus on the most important difference between the two models, we make them identical except with regard to when patch-switches occur: Because we found that the previous item I_{t-1} has the highest semantic similarity to the current item I_t in a patch (reported above), we use the previous item as the local cue for both categorical and associative model versions.⁶ Both versions of the model also use frequency as the global cue in the memory probe, because previous work found frequency to be an effective representation for item choice after a

patch transition (Hills et al., 2012). The associative and categorical versions of the patch search model we compare here thus differ just in where they make local-to-global transitions, step 4 in the algorithm for both types of models that we presented earlier. For the associative model, we determine the locations of the associative patch switches made by each individual by applying the associative patch-switch definition (using the hand-coded Troyer subcategories) to the individual's sequence of item data, and then we trigger the associative model to make its local-to-global transitions at those points; we do the same for the categorical model using the categorical patch-switch definition (which matches all of the associative switch locations, and additionally triggers a few more categorical-only switches).

We then fit these models to the sequence of items produced by each individual, using the log-likelihood fit penalized by the number of free parameters according to the Bayesian Information Criterion (BIC; for a general introduction to this approach, see Lewandowsky & Farrell, 2011).⁷ Results are presented as the median improvement in the BIC relative to a random model specifying that all remaining items in the “animal” search space are equally likely to be retrieved (i.e., using neither the global nor local cues). The top section of Table 1 shows that a comparison of these two model versions favors the associative model at the level of individual participants.⁸ While the two model versions produce similar median BIC scores overall, individually 42% of the participants are best fit by the associative search model (having the largest improvement in BIC), 20% are best fit by the categorical model, and the remaining 38% are fit equally well by both models.

We can also compare the two types of patch search by looking at the predictive power of items prior to the most recently produced item I_{t-1} . If search is purely associative, then the best predictor of the next item in a patch should always be the previous item retrieved in the patch (as shown in the previous section), meaning that all earlier-produced items would not add any predictive power to the model. If search is purely categorical, then an earlier item in the same patch should be as predictive as a later item, and together they may be still more predictive of future items: Either they strengthen the identification of the current category or, at least, they should on average be no further away from the last item produced. Therefore, adding an additional item from the category should increase the predictive power of the model. To test this, we fit models that added an additional local cue to the previous categorical and associative model versions, namely the item I_{t-2} , produced just before I_{t-1} (as long as it was also in the same categorical patch). As seen in the lower half of Table 1, the β values for I_{t-2} are roughly zero, indicating that earlier items do not add any predictive power. Moreover, these models show smaller improvements in BIC over the random model than models that do not use item I_{t-2} .

Taken together, these model comparisons and analyses above provide evidence that participants' behavior is generally better explained by an associative patch search model than by a categorical model. Most individuals do not appear to retrieve successive items based on subcategory membership, but instead appear to retrieve items that are semantically near to the most recently produced item. Moreover, there appears to be very little evidence for individuals using items prior to the most recently produced item to guide their search in patches, indicating that even if individuals are using categorical

information in some way, they are also strongly inclined to favor associative search within subcategories. Nonetheless, there are individual differences in search processes, with some people more guided by subcategory boundaries than others. Next we test whether differences in search process lead to differences in performance with respect to the number of items recalled.

Table 1: *Comparison of associative versus categorical search models: Switches are determined by the associative or categorical patch definitions using the hand-coded subcategories from Troyer et al. (1997), the global cue is word frequency, and the local cues are the previously-produced items I_{t-1} and I_{t-2} (1 or 2 items back). Median ΔBIC is the improvement over a random model that retrieves all items with equal probability.*

<i>Patch Model</i>	<i>Cue</i>	<i>Median β</i>	<i>Median ΔBIC</i>	<i>% best fit</i>
Associative	Global	7.22 (2.17)	100.12 (28.29)	42%
	Local item I_{t-1}	5.03 (1.67)		
Categorical	Global	7.30 (2.15)	100.12 (27.74)	20%
	Local item I_{t-1}	5.09 (1.67)		
<hr/>				
Associative	Global	7.28 (2.17)	98.05 (28.4)	
	Local item I_{t-1}	4.79 (1.70)		
	Local item I_{t-2}	-0.05 (4.42)		
Categorical	Global	7.28 (2.16)	98.05 (27.8)	
	Local item I_{t-1}	4.76 (1.76)		
	Local item I_{t-2}	1.27 (12.14)		

Does the search process correlate with performance?

We assigned each participant to the associative or categorical search model based on the quality of the model fits, as shown in the right hand side of the Table 1.

Comparing the performance in terms of number of items retrieved by those who were

best fit with the associative model versus those best fit with the categorical model yielded no statistical difference in the overall number of items produced (associative: $M = 38.2$, $SD = 7.97$ vs. categorical: $M = 39.39$, $SD = 7.97$, $t(53) = -0.67$, $p = .51$, $BF = 3.47$ in favor of the null). In comparison with individuals who were fit equally well by both models, those fit best by either the associative or categorical model did significantly better ($M = 38.56$, $SD = 7.32$, vs. 31.93 , $SD = 7.32$, $t(139) = 4.97$, $p < .001$, $BF = 7751.41$ in favor of the alternative). This latter result may however indicate that individuals who produce more items provide us with more data to better identify their true underlying search process—and therefore may simply represent an effect of statistical power.

Discussion

In this article we explored whether the patch boundaries produced in memory search are better explained by an associative or categorical search process, or possibly a combination of both. We first addressed the more general question of whether semantic memory search is patchy at all—that is, involving a two-stage search process between local and global representations. The evidence in support of patchy search came from the finding that associative/categorical between-patch transitions (identified using Troyer categories) take longer than ordinal position-matched non-transitions. This is consistent with our previous work fitting individual data (Hills et al., 2012), which pitted one-stage models against two-stage models using an associative patch model and found support for two-stage models.

Regarding patch boundaries, our results were predominantly consistent with the use of an associative search process. Associative search was supported both within patches for all participants and across patches for most participants. However, some

participants were nonetheless better explained by a categorical search process between patches, indicating they may be activating categories during the search process.

Nonetheless, categorical-only switches did not take longer than ordinal position-matched non-switches, suggesting they may simply be low-similarity transitions, which do not reflect a true ‘switch’ process. Via model comparison, we also found that the largest set of individuals were best fit by an associative search model, and some were fit equally well by both models. However, we did not find a performance benefit for using an associative search model: Our results indicated that among those individuals best fit by one or the other patch model, they recalled approximately equivalent numbers of items from memory. This is in contrast with previous work suggesting the benefits of categorical retrieval strategies in list recall (Pollio et al., 1969; see also Patterson, Meltzer, & Mandler, 1971), but the difference may arise because of a distinction between recall from natural categories and recall from learned lists.

Within patches, there is a clear increase in semantic similarity for items produced nearer to each other. This is consistent with our previous work (Hills et al., 2012) and the assumptions of more recent work (Abbot et al., in press; Thompson & Kello, 2014), which all share the common feature that memory search is associated with a random walk over some representation. Evidence for an associative process indicates that search cannot be a purely categorical process; semantic similarity matters even within patches. Additional evidence against purely categorical search was seen in the lack of improved performance when models used two previous items as local cues rather than just one. This would have added additional categorical information, but no additional associative information.

Why might associative patch search strategies be widely used? One explanation is the likely high dimensionality of memory subspaces. Memory items can reside in multiple patches corresponding to multiple (sub)categories at the same time. Unlike leaving discrete patches of food in spatial foraging, leaving one category in semantic memory does not mean *all* current categories have been left as well—the searcher may still be retrieving a sequence of items that belong to a different category. That is, even when a series of retrieved items fall in the same superordinate category, our results suggest that each new item activates new traces in memory, making new routes through memory available. For the majority of participants, our results dispel the notion that categories are activated at all, and instead suggest that search follows a Markov process, where only the most recent retrieval is activated. Search in memory may be compared to a person walking through a forest at night with a flashlight; as each new location is reached, new areas become visible, while others recede into the darkness, and the whole path is linked as one “patch” even though the beginning and the end may have nothing in common.

The present model can be easily described as search in semantic networks: The nodes of the semantic network represent animals and edges are semantic relationships between words. Recent work has shown that the organization of semantic memory has a high clustering coefficient (Goñi et al. 2011; see also Steyvers & Tenenbaum, 2005)—it is a small world network. Such small-world representations offer the possibility that local clusters in the memory space are easily accessible via “local” routes. Thus, one need not make a jump between categories to go from the current “semantic field” to another, but may travel via semantic “brokers” between clusters. Along with the flashlight

metaphor, this view of memory invites the notion that the semantic fields simply shift along with the retrieval process, with the boundaries of more or less sparse patches always lying just over the horizon. To put this into numbers from our data, of the 347 animals produced by our participants, the number of items accessible by a series of intermediate words that never fall below the mean cosine similarity for switches (0.31 in Figure 3) is 321—with only 26 isolated animals, predominantly rare ones (e.g., pillbug, sea monkey, and yak). In other words, the majority of animal types in memory are connected in one large component, which nevertheless has structure, but also provides routes of high semantic similarity from one location to another. This small world structure may make memory more easily searched (e.g., Kleinberg, 2002), but research investigating the interaction of structure, process, and goals is needed to further develop this hypothesis.

An additional challenge still to be addressed in this research is our use of the Troyer et al. (1997) hand-coded subcategories to determine categorical and associative patch boundaries. At present this type of hand-coding is the best resource we have for identifying potential associative and categorical patch search strategies. However, as this article suggests, the cluster-switching model often used with the Troyer subcategories is unlikely to accurately describe the cognitive process used to search memory, to the extent that the cluster-switching model assumes that categories are activated. Furthermore, assuming that everyone uses these same categories does not capture the (possibly large) individual differences in category representations between people. This was one motivation for the Hills & Pachur (2012) study that had individuals specify their own social networks for comparison against their social free recall patterns. The Troyer

subcategories have also recently come into question regarding their predictive power in long-term memory search (e.g., Hills et al., 2013). Unfortunately, the patch-switching model that thus far performs better than the ones built on the Troyer subcategories—the associative similarity drop model of Hills et al. (2012)—do not specify underlying categories and so does not enable discrimination between different types of patch models, associative or categorical. Thus, like much of memory research, the present work makes assumptions about the representation and the processes doing the search over that representation. Future work could resolve some aspects of this problem by isolating categories from the BEAGLE representations using clustering algorithms, or using similar approaches applied to large collections of recall data (e.g., Goñi et al., 2011).

In sum, our results favoring associative search strategies suggest that local patches may be retrieved via similarity relations between items generated on the fly with respect to a given problem (e.g., Barsalou, 1983). As demonstrated by studies of neutral and adaptive evolution with multiple genetic loci (Gavrilets, 1997; Kauffmann & Levin, 1987), movement in high-dimensional spaces often violates our intuitive assumptions about nearness and categorizations in these spaces may be less meaningful than our intuition may at first suggest. In the present case, this means that following associative links through the high-dimensional semantic space can create a patch in memory retrieval that fluidly shifts from one subcategory to another. However, some domains or search tasks may be better adapted for categorical search (e.g., searching for a recipe for a particular type of dish), or influenced by the prior knowledge of the searcher (e.g. searching for a new car based on previous experience with different vehicle makes). In addition, it may be possible to use the above approach to study individual differences in

memory search among clinical populations, for example, by examining the influence of category priming on category-only switches, or examining switch-rates among individuals allowed to search across multiple categories at the same time (e.g., Maylor, Chater, & Jones, 2001). Hence, while our results provide support for an associative search process in a widely-studied domain, they do not eliminate the possibility that some individuals and circumstances may favor categorical strategies, or use still other representations besides those investigated here.

References

- Abbot, J. T., Austerweil, J. L., & Griffiths, T. L. (In press). Random walks on semantic networks can resemble optimal foraging. *Psychological Review*.
- Anderson, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Erlbaum.
- Barsalou, L.W. (1983). Ad hoc categories. *Memory & Cognition*, *11*, 211-227.
- Bousfield, W. A. (1953). The occurrence of clusters in the arrangement of randomly arranged associates. *Journal of Experimental Psychology*, *40*, 643-647.
- Bousfield, W.A., & Sedgewick, C.H. (1944). An analysis of sequences of restricted associative responses. *The Journal of General Psychology*, *30*, 149-165.
- Charnov, E.L. 1976. Optimal foraging: the marginal value theorem. *Theoretical Population Biology*, *9*, 129–136.
- Davelaar, E.J., & Raaijmakers, J.G.W. (2012). Human memory search. In P.M. Todd, T.T. Hills, and T.W. Robbins (Eds.), *Cognitive search: Evolution, algorithms, and the brain* (pp. 177-194). Strüngmann Forum Reports, vol. 9. Cambridge, MA: MIT Press.
- Dougherty, M. R., & Harbison, J. I. (2007). 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*, *33*, 1108-1117.
- Fauchald, P. (1999). Foraging in a hierarchical patch system. *The American Naturalist*, *153*, 603-615.
- Fossati, P., Le Bastard, G., Ergis, A.M., & Allilaire, J.F. (2003). Qualitative analysis of verbal fluency in depression. *Psychiatry Research*, *117*, 17–24.
- Gavrilets, S. (1997). Evolution and speciation on holey adaptive landscapes. *Trends in Ecology and Evolution*, *12*, 307-312.

- Goñi, J., Arrondo, G., Sepulcre, J., Martincorena, I., Vélez de Mendizábal, N., Corominas-Murtra, B., Bejarano, B., Ardanza-Trevijano, S., Peraita, H., Wall, D.P., & Villoslada, P. (2011). The semantic organization of the animal category: Evidence from semantic verbal fluency and network theory. *Cognitive Processes*, *12*, 183-196.
- Gronlund, S. D., & Shiffrin, R. M. (1986). Retrieval strategies in recall of natural categories of categorized lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*, 550-561.
- Gruenewald, P. J., & Lockhead, G. R. (1980). The free recall of category examples. *Journal of Experimental Psychology: Human Learning and Memory*, *6*, 225-239.
- Grünbaum, D. (1998). Using spatially explicit models to characterize foraging performance in heterogeneous landscapes. *The American Naturalist*, *151*, 97-115.
- Harbison, J. I., Dougherty, M. R., Davelaar, E. J., & Fayyad, B. (2009). On the lawfulness of the decision to terminate memory search. *Cognition*, *111*, 397-402.
- Haugrud, N., Crossley, M., & Vrbancic, M. (2011). Clustering and switching strategies during verbal fluency performance differentiate Alzheimer's disease and healthy aging. *Journal of the International Neuropsychological Society*, *17*, 1153-1157.
- Hills, T. T. (2011). The evolutionary origins of cognitive control. *Topics in Cognitive Science*, *3*, 231-237.
- Hills, T. T., Jones, M. & Todd, P. M. (2012). Optimal foraging in semantic memory. *Psychological Review*, *119*, 431-440.
- Hills, T. T., Kalff, C., & Wiener, J. M. (2013). Adaptive Lévy processes and area-restricted search in human foraging. *PLoS One*, *4*, 1-7.

- Hills, T. T., Mata, R., Wilke, A., & Samanez-Larkin, G. (2013). Mechanisms of age-related decline in memory search across the adult life span. *Developmental Psychology, 49*, 2396-2404.
- Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., Couzin, I. D., & Cognitive Search Research Group. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences, 19*, 46-54.
- Howard, M. W., Addis, K. A., Jing, B., & Kahana, M. J. (2007). Semantic structure and episodic memory. In K. T. Landauer, D. McNamara, S. Dennis, & W. Kintsch (Eds.), *LSA: A road towards meaning*. Mahwah, NJ: Lawrence Erlbaum Associates (pp. 121- 141).
- Ipsen, Gunther, 1924. Der Alte Orient und die Indogermanen. In: Johannes Friedrich, Johannes Baptista Hofman, Wilhelm Horn (Eds.), *Stand und Aufgaben der Sprachwissenschaft: Festschrift fur Wilhelm Streitberg*. Heidelberg: Winter (pp. 200-237).
- Johnson, D. M., Johnson, R. C., & Mark, A. L. (1951). A mathematical analysis of verbal fluency. *The Journal of General Psychology, 44*, 121-128.
- Jones, M. N., Kintsch, W., & Mewhort, D. J. K. (2006). High-dimensional semantic space accounts of priming. *Journal of Memory and Language, 55*, 534–552.
- Jones, M.N., & Mewhort, D.J.K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review, 104*, 1-37.
- Karieva, P., & Odell, G. (1987). Swarms of predators exhibit “preytaxis” if individual predators use area-restricted search. *American Naturalist, 130*, 233-270.
- Kauffman, S., & Levin, S. (1987). Towards a general theory of adaptive walks on rugged

- landscapes. *Journal of Theoretical Biology*, 128, 11-45.
- Kleinberg, J. (2002). Small-world phenomena and the dynamics of information. *Advances in neural information processing systems, 1*, 431-438.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104, 211-240.
- Lewandowsky, S., & Farrell, S. (2011). *Computational Modeling in Cognition: Principles and Practice*. Los Angeles: Sage.
- Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavioral Research Methods, Instruments, & Computers*, 28, 203-208.
- Maylor, E. A., Chater, N., & Jones, G. V. (2001). Searching for two things at once: Evidence of exclusivity in semantic and autobiographical memory retrieval. *Memory & cognition*, 29(8), 1185-1195.
- Metcalfe, J. & Murdock, B. B. (1981). An encoding and retrieval model of single-trial free recall. *Journal of Verbal Learning and Verbal Behavior*, 20, 161-189.
- Murphy, K. J., Rich, J. B., & Troyer, A. K. (2006). Verbal fluency patterns in amnesic mild cognitive impairment are characteristic of Alzheimer's type dementia. *Journal of the International Neuropsychological Society*, 12, 570-574.
- Patterson, K. E., Meltzer, R. H., & Mandler, G. (1971). Inter-response times in categorized free recall. *Journal of Verbal Learning and Verbal Behavior*, 10(4), 417-426.
- Pirolli, P. (2007). *Information Foraging Theory: Adaptive Interaction with Information*.

Oxford University Press, New York, NY.

Pollio, H. R., Richards, S., & Lucas, R. (1969). Temporal properties of category recall.

Journal of Verbal Learning and Verbal Behavior, 8, 529-536.

Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory.

Psychological Review, 88, 93-134.

Raoux, N., Amieva, H., Le Goff, M., Auriacombe, S., Carcaillon, L., Letenneur, L., &

Dartigues, J. (2007). Clustering and switching processes in semantic verbal fluency

in the course of Alzheimer's disease participants: results from the PAQUID

longitudinal study. *Cortex*, 44, 1188-1196.

Robert, P. H., Lafont, V., Medicin, I., Berthet, L., Thauby, S., Baudu, C., & Darcourt, G.

(1998). Clustering and switching strategies in verbal fluency tasks: comparison

between schizophrenics and healthy adults. *Journal of International*

Neuropsychological Society, 4, 539-546.

Romney, A.K., Brewer, D.D., & Batchelder, W.H. (1993). Predicting clustering from

semantic structure. *Psychological Science*, 4, 28-34.

Rosen, V. M., & Engle, R. W. (1997). The role of working memory capacity in retrieval.

Journal of Experimental Psychology: General, 126, 211-227.

Steyvers, M., & Tenenbaum, J. B. (2005). The large-scale structure of semantic

networks: Statistical analyses and a model of semantic growth. *Cognitive Science*,

29, 41-78.

Taler, V., & Phillips, N. A. (2008). Language performance in Alzheimer's disease and

mild cognitive impairment: A comparative review. *Journal of Clinical and*

Experimental Neuropsychology, 30, 501-556.

- Thompson, G. W. & Kello, C. T. (2014). Walking across Wikipedia: A scale-free network of semantic memory retrieval. *Frontiers in Psychology, 5*, 86.
- Todd, P.M., Hills, T.T., and Robbins, T.W. (Eds.) (2012). *Cognitive search: Evolution, algorithms, and the brain*. Strüngmann Forum Reports, vol. 9. Cambridge, MA: MIT Press.
- Troyer, A. (2000). Normative data for clustering and switching on verbal fluency tasks. *Journal of Clinical and Experimental Neuropsychology, 22*, 370-378.
- Troyer, A. K., Moscovitch, M., & Winocur, G. (1997). Clustering and switching as two components of verbal fluency: Evidence from younger and older healthy adults. *Neuropsychology, 11*, 138-146.
- Troyer, A. K., Moscovitch, M., Winocur, G., Alexander, M. P., & Stuss, D. (1998). Clustering and switching on verbal fluency: The effects of focal frontal- and temporal-lobe lesions. *Neuropsychologia, 36*, 499-504.
- Troyer, A. K., Moscovitch, M., Winocur, G., Leach, L., & Freedman, M. (1998). Clustering and switching on verbal fluency tests in Alzheimer's and Parkinson's disease. *Journal of the International Neuropsychological Society, 4*, 137-143.
- Walker, W. H., & Kintsch, W. (1985). Automatic and strategic aspects of knowledge retrieval. *Cognitive Science, 9*, 261-283.
- Wilke, A., Todd, P.M., and Hutchinson, J.M.C. (2009). Fishing for the right words: Decision rules for human foraging behavior in external and internal search tasks. *Cognitive Science, 33*, 497-529.
- Willits, J. A., D'Mello, S. K., Duran, N. D., & Olney, S. (2007). Distributional statistics and thematic role relationships. In D. S. McNamara & J. G. Trafton (Eds.),

Proceedings of the 29th Annual Meeting of the Cognitive Science Society (pp. 707-712). Nashville, TN.

Endnotes

¹ Note that in this example we use the Troyer et al. categories as a proxy for places in memory where items are likely to be sufficiently far from one another to warrant the inference of a between-patch transition. In practice, our models will ask if category membership provides any additional information for predicting the search process beyond that provided by local semantic similarity (associative information) produced by BEAGLE.

² All repetitions were removed from the data (representing less than 1% of retrievals). Following Hills & Pachur (2012), we used sampling-without-replacement to compute probabilities of retrieval.

³ This pattern is also matched by a reduction in BEAGLE similarity over production quartiles, either measured by number of productions ($\chi^2(1)=181.34, p<.001$) or time (e.g., first 45 seconds, second 45 seconds, and so on, $\chi^2(1)=202.84, p<.001$).

⁴ A test dividing production into temporal quartiles leads to the same statistical pattern of increasingly frequent associative/categorical switches ($\chi^2(1)=28.42, p<.001$) but no change in categorical-only switches ($p=.39$).

⁵ In Hills et al. (2012) we used the word ‘categorical’ to define patches. However, to be clear, in that article we used the associative definition of patches described in this article.

⁶ Using the subcategory as a cue has a number of problems, the most important of which is that it is not straightforward how to penalize out-of-category retrievals. The performance of the models is highly dependent on this penalization.

⁷ Using AIC does not change the conclusions reported here.

⁸ The associative model in Figure 1 corresponds to the ‘combined cue dynamic model’ using Troyer et al. categories in Table 1 of Hills et al. (2012), which we here compare with a model that adds categorical-only switches. This model and the similarity drop model both performed better in model comparisons than ‘static’ models that assumed memory search was a one-stage random walk. (All other models in this paper are new.)