



Cognitive Science (2015) 1–15

Copyright © 2015 Cognitive Science Society, Inc. All rights reserved.

ISSN: 0364-0213 print / 1551-6709 online

DOI: 10.1111/cogs.12249

The Role of Semantic Clustering in Optimal Memory Foraging

Priscilla Montez, Graham Thompson, Christopher T. Kello

Cognitive and Information Sciences, University of California, Merced

Received 25 October 2013; received in revised form 10 October 2014; accepted 11 November 2014

Abstract

Recent studies of semantic memory have investigated two theories of optimal search adopted from the animal foraging literature: Lévy flights and marginal value theorem. Each theory makes different simplifying assumptions and addresses different findings in search behaviors. In this study, an experiment is conducted to test whether clustering in semantic memory may play a role in evidence for both theories. Labeled magnets and a whiteboard were used to elicit spatial representations of semantic knowledge about animals. Category recall sequences from a separate experiment were used to trace search paths over the spatial representations of animal knowledge. Results showed that spatial distances between animal names arranged on the whiteboard were correlated with inter-response intervals (IRIs) during category recall, and distributions of both dependent measures approximated inverse power laws associated with Lévy flights. In addition, IRIs were relatively shorter when paths first entered animal clusters, and longer when they exited clusters, which is consistent with marginal value theorem. In conclusion, area-restricted searches over clustered semantic spaces may account for two different patterns of results interpreted as supporting two different theories of optimal memory foraging.

Keywords: Lévy flights; Marginal value theorem; Category recall; Foraging

1. Introduction

Semantic knowledge is composed of conceptual categories and their inter-relations. There are categories for natural kinds like animals and plants, for artifacts like tools and clothes, and social divisions like friends and enemies, just to name a few. Conceptual categories can have different kinds of relations to each other, such as hierarchical relations when one category is subdivided into smaller ones and contained within larger ones

Correspondence should be sent to Christopher T. Kello, University of California, Merced, 5200 North Lake Road, Merced, CA 95343. E-mail: ckello@ucmerced.edu

(Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). Perhaps the most basic type of relation is one of similarity—dogs generally are more semantically similar to wolves than they are to wrenches, for instance, although similarity is always dependent on context (Goldstone, Medin, & Halberstadt, 1997; Rips, 1975; Tversky, 1977).

Similarity relations and hierarchical relations, among others, have been theorized in terms of a *semantic space* populated with conceptual categories and their members (Lund & Burgess, 1996). This space is searched to find conceptual information needed during language comprehension and production, for instance, along with many other cognitive activities (Collins & Loftus, 1975). From this theoretical perspective, one is led to ask how semantic spaces are organized and how this organization affects search processes and behaviors (Collins & Quillian, 1969). One can assume that semantic memory has evolved and developed to minimize time and access useful semantic information.

The same basic question about search processes has been posed in the animal foraging literature (Pyke, 1984), where semantic spaces are replaced by habitats, and semantic information is replaced by food and other physical resources that animals forage for: How are animal search processes adapted to their environments? Foraging researchers assume that animals have evolved to use efficient search strategies, where efficiency is defined as maximizing resources found per unit time. A globally optimal foraging process is difficult to define, but one can instead formulate optimal strategies for specific situations, given simplifying assumptions.

Two such classes of optimal foraging strategies have been debated in recent years. One class is based on *marginal value theorem* (Charnov, 1976), which assumes that resources are distributed in clusters, also referred to as “patches.” This assumption is consistent with the basic fact that semantic knowledge is organized into conceptual categories (Hills, Jones, & Todd, 2012). Clustering is hypothesized to be adaptive, partly because semantically similar categories tend to be accessed at similar points in time as thoughts and actions unfold (Lucas, 2000). This common property of semantic processing means that clustering should facilitate recall of multiple items by way of inherent priming or cueing as individual items are accessed in succession. Indeed, studies have found faster successive recall times for semantically similar versus dissimilar items (Gruenewald & Lockhead, 1980; Howard & Kahana, 2002).

Marginal value theorem focuses on the potentially adaptive nature of clustering by simplifying the forager as an agent continuously faced with just one decision—either to continue foraging the current patch, or move to another patch that is hopefully more bountiful. Focusing on this decision affords a simple, optimal strategy: Stay if the current rate of resource acquisition is at or above the mean expected rate over all patches, and go if the current rate drops below the mean expectation. Studies of animal foraging have found evidence consistent with marginal value theorem (Schoener, 1987), although there have also been a number of inconsistencies in the literature (Nonacs, 2001).

Marginal value theorem also can be applied to semantic memory search, specifically for the category recall task in which participants are asked to recall as many members of a category as they can from long-term memory. A well-studied example is to name as

many animals as possible for some number of minutes. One can imagine semantic memory as a space with animals arranged in various clusters like mammals, herbivores, farm animals, marine animals, and so on. The category recall task effectively asks participants to forage through semantic memory space for animal targets. If one assumes an attentional focus on one region of space at a given time, then one can simplify the memory search process as a repeated stay-or-go question. Animal names should be recalled repeatedly from an attended region until recall rate drops below the mean expected rate. Then the focus of attention should move to another region, presumably incurring a time cost to search and refocus, but gaining the benefit of a hopefully higher rate of recall in an unforaged region.

Hills et al., (2012) found evidence in semantic memory search data that was consistent with marginal value theorem (but see Austerweil, Abbott, & Griffiths, 2012). The authors analyzed time intervals from an animal category recall experiment. The intervals occurred either within or between clusters of animals as they unfolded during a recall session. Those occurring at the start of clusters were below the mean on average, and intervals tended to increase as items in a cluster were recalled, until intervals went above the mean at which point participants were likely to transition between clusters. This pattern of results is consistent with the marginal value theorem, but the theorem does not address how clusters are organized in relation to each other in semantic memory, or how each given cluster is chosen for foraging.

We can go back to the analogy with animal foraging to further consider the aspects of search left unaddressed by marginal value theorem. In particular, clusters of resources in habitats have spatial relations with each other, and animals choose clusters by moving to them. Animal foraging is inherently about movement through space, which is abstracted away by marginal value theorem for the sake of framing optimality. This point raises the question of whether optimality can be framed explicitly in terms of movement through space. Viswanathan et al. (1996) proposed *Lévy foraging* as an optimal movement-based strategy that accounts for a ubiquitous pattern of behavior observed throughout foraging studies.

Lévy foraging is based on analyzing animal foraging movements as a series of segmented paths of lengths L . From this perspective, one can ask if there is an optimal distribution of foraging path lengths, independent of movement directions, given certain simplifying assumptions. To address this question, let us first consider whether certain path length distributions are commonly observed in nature. It turns out that many studies of animal foraging for many species in many habitats have found path lengths to be distributed like an inverse power law function, $P(L) = 1/L^\alpha$ with $\alpha \sim 2$ (e.g., see Humphries, Weimerskirch, Queiroz, Southall, & Sims, 2012; Kello et al., 2010). The same is true for category recall, as first shown by Rhodes and Turvey (2007) and then replicated by others (Radicchi & Baronchelli, 2012; Thompson & Kello, 2014). To draw an analogy with memory foraging, Rhodes and Turvey reasoned that times between recall events should roughly correspond with foraging path lengths, in which case inter-response intervals (IRIs) should exhibit the same power law observed in animal foraging. Results bore out this expectation.

It is hotly debated exactly how well these observed distributions, be they spatial lengths or temporal intervals, are described by a power law function (e.g., see James, Plank, & Edwards, 2011). Limitations on measurement and on environmental scales will place limits on observable scales of variation, which means that only truncated, imperfect power laws can be observed. And indeed, there is consensus that findings are often characterized by “power law-like” path length distributions, with varying degrees and kinds of truncation depending on the particular study and method of analysis. The debate concerns whether other functions may better characterize the data, particularly those that can model deviations from a power law at the smallest or largest scales of measurement. Herein, we will focus on the general pattern of results showing power law-like path length distributions, and set aside the more contentious issue of the best fitting statistical function.

Findings from foraging studies suggest that there is something beneficial about power law-like path lengths. To formalize what this might be, Viswanathan et al. (1999) simplified foraging trajectories by treating them as random walks, that is, series of movement steps with each direction chosen uniformly at random, and each step length sampled from some probability distribution. Random walks with power law distributed path lengths are known alternately as *Lévy walks* or *Lévy flights*, and Viswanathan et al. showed that Lévy walks with $\alpha = 2$ maximizes the rate of finding items when they are sparsely and randomly distributed in an unbounded two-dimensional space. This maximum is relative to random walks with other path length distributions.

The analysis by Viswanathan et al. (1999) suggests an optimal basis for observed power law-like path length distributions, but only for classes of memoryless search processes. Simple search agents like bacteria may be effectively memoryless, but it is questionable whether human search through semantic memory is enough like a random walk for Lévy flights to be useful models. For instance, participants very rarely repeat an item already recalled, but a random walk has some probability of retracing its steps (albeit Lévy walks reduce this probability compared with other kinds of random walks). Also, a random walk is blind to clustering and other patterns in the spatial distribution of items being searched for.

1.1. Lévy flights and marginal value theorem in memory foraging

Lévy flights and marginal value theorem are based on different sets of assumptions and focus on different aspects of search behaviors, yet both have been applied to data from semantic category recall experiments. Lévy flights describe distributions of IRIs as foraging path lengths, whereas marginal value theorem describes mean times spent foraging clusters of items to be recalled. Each approach to memory search hypothesizes an optimal behavior, but they seem unrelated to each other, perhaps even at odds with each other. Marginal value theorem assumes clusters and maximizes the rates of items found within them. However, it does not specify their spatial/representational relations to each other, or how clusters are found and chosen from one to the next. Lévy foraging specifies spatial movements of the forager but canonically assumes random and sparse scattering

of items rather than clustering (but see Ferreira, Raposo, Viswanathan, & da Luz, 2012; Santos et al., 2007). Lévy flights are random, memoryless walks whose movements are uninfluenced by items found or other potential information gains as foraging goes on.

In this study, we investigated whether clustering in semantic space is a factor that underlies evidence for both Lévy flights and marginal value theorem in category recall experiments. Our experiment employed a version of the spatial arrangement method (Goldstone, 1994; Hout, Goldinger, & Ferguson, 2013) to elicit semantic spaces from individual participants (cf. Morais, Olsson, & Schooler, 2013). Each participant arranged items on a two-dimensional surface such that distances apart corresponded with item similarities. Manipulating objects on a 2D surface is a highly familiar task, and sensorimotor processes in general, including visuospatial processes, may provide an embodied basis for conceptual processing (Clark, 2008). Besides this theoretical rationale for using the spatial arrangement method, we also needed a 2D space to test for Lévy flights because they have been formalized and analyzed in two dimensions.

One might question whether semantic knowledge can be accurately represented in a 2D space. For instance, previous studies have used representations with greater dimensionality such as semantic networks (see Steyvers & Tenenbaum, 2005), high-dimensional vectors of semantic features (e.g. Cree, McRae, & McNorgan, 1999), textual co-occurrence vectors (Burgess, 1998; Jones & Mewhort, 2007; Landauer & Dumais, 1997), and learned distributed representations in connectionist models (Rogers & McClelland, 2004). 2D spaces may provide less fidelity compared with other representations, but for our purposes it is only necessary that participants can project their knowledge onto two-dimensions with enough fidelity to test Lévy flights and marginal value theorem.

Each participant was given a set of animal names printed on labels affixed to magnets. The sets of animal names were generated in a previous category recall experiment reported by Thompson and Kello (2013, 2014). In the classic category recall task (Bousfield & Sedgewick, 1944), participants recall as many items from long-term memory as they can within a minute or two. The number of responses serves as a standard measure of semantic fluency, and items tend to be recalled in semantic clusters (Patterson, Meltzer, & Mandler, 1971). Thompson and Kello gave participants 10 min to recall animal names, in order to generate enough responses to test for power law-like IRI distributions. Each participant recalled about 100 animals.

In the present experiment, participants did not know where the animal names came from—each participant simply placed each given set of names on a whiteboard, in whatever arrangement best reflected the similarities among animals according to their semantic knowledge. Similarity is a relative semantic judgment, and hence fundamentally influenced by context. For instance, semantic similarities to “chicken” are different for the category “animal” compared with “food.” Participants had as context the category label “animal,” and the set of animal names provided up front. Thus, each participant’s context was matched with a particular recall set, without providing information about specific clusters or recall order.

Even though we provided no information or instructions about clustering, one should expect participants in our experiment to cluster animals on the whiteboard into various

categories like land mammals, sea creatures, and so on. Each participant may arrange animals into different clusters, and clusters may align to differing degrees with those that emerged in the category recall series collected by Thompson and Kello (2013). Nonetheless, if there is common semantic structure to the animal category, then some correspondence should emerge between category recall series and subsequent arrangements of names on the whiteboard.

Each category recall series was matched with its corresponding spatial arrangements by expressing it as a trajectory over the whiteboard space. This was done after the spatial arrangement experiment was completed, using digital pictures taken of each arrangement. We used a whiteboard instead of a computer screen because the whiteboard afforded easy manipulation of individual labeled magnets or groups of them, thus minimizing any physical deterrent to arranging items. The whiteboard also provided a more visceral, embodied sense of the space over which animals were arranged.

Marginal value theorem predicts that, for search paths entering and leaving spatial clusters, IRIs should be shorter and longer, respectively, relative to their means. There is already some evidence for this prediction (Hills et al., 2012), but by projecting category recall series onto a physical space, we can measure movements both within and between categories. Measuring both kinds of movements allows us to test for the power law-like distributions of path lengths predicted by Lévy foraging, while simultaneously testing for evidence of marginal value theorem.

2. Methods

We first summarize the methods and results of the category recall experiment reported by Thompson and Kello (2013), and then report the methods and results of the present spatial arrangement experiment. Nineteen participants were recruited for the prior category recall experiment, each of whom was given 20 min to type as many names of animals as they could recall from long-term memory. Responses were displayed on a monitor as they typed. IRIs were measured as the times from the last keystroke of one response, to the first keystroke of the next (IRIs did not include times taken to type responses). Animal names were corrected for spelling errors, and any names repeated by a given participant were removed (32 in all across 19 participants). The mean number of responses per participant was 117 ($SD = 38.6$), and IRI distributions were found to replicate the basic power law pattern reported by Rhodes and Turvey (2007). More specific results can be found in Thompson and Kello (2013).

For the spatial arrangement experiment, three sets of animal name responses were chosen from three participants in the previous category recall experiment. Sets were chosen from participants who recalled roughly equal numbers of responses close to the mean (103, 105, 107). Of all these responses, 27 animal names were recalled independently by all three participants, 80 were recalled by two participants, and the remaining 208 were recalled uniquely by one participant. Each of the three participant sets were used as stimuli for the present experiment.

2.1. Participants

Fifteen University of California, Merced undergraduates participated for course credit. Participants reported having normal or corrected vision and were physically able to stand, walk, and reach the majority of space on a whiteboard that served as the 2D semantic space.

2.2. Stimuli

Three sets of magnets were created by pasting a 2×2 inch square of poster board (3 mm thick) onto each magnet, and then pasting a label to each square with the name of an animal printed in 36 point Calibri font underneath a blue circle 1.5 inches in diameter. The blue circle was used to locate magnets in photos, explained later. In addition, 10 magnets with frequent plant names were created to use as a practice set.

2.3. Procedure

Each participant was randomly assigned to one of the three sets of animal magnets. First, the practice set of magnets was placed on a 3×3 foot square table, their positions scrambled randomly on the table for each participant. Next to the table was a 5 foot tall by 9 foot wide whiteboard attached to a wall, with its lower edge 40 inches off the ground. Participants were instructed to arrange the set of 10 plants on the whiteboard according to their similarities, such that more similar plants should be placed closer together. No other instructions were given about how to arrange the magnets or how to interpret similarity, except that all magnets needed to be placed somewhere on the whiteboard. Participants were allowed to prearrange magnets on the table before putting them up on the whiteboard and/or rearrange magnets on the whiteboard as they go up. Magnets were designed to afford ease of manipulation, individually or in groups, on the table or on the whiteboard.

After participants were done with the arranging of all the practice magnets onto the whiteboard, a tripod-mounted Nikon D5000 12.3 megapixel camera (Nikon, Melville, New York USA) was used to take a picture of the arrangement. The camera was placed 49 inches off the ground and 84 inches from the whiteboard, such that the whiteboard filled the camera's field of view. Upon completion of the practice trial, participants were given a whiteboard marker and instructed to circle any groupings of similar plants they had made. They were also instructed to label the groupings and draw lines connecting any two groupings if they perceived a broader similarity relationship between the groupings. A second photograph was taken after each participant was done with the marker.

After finishing with the practice set of plant names, the experimenter removed the practice magnets from the whiteboard and placed the assigned set of animal magnets on the table, scrambled randomly. Participants were given as much time as needed to arrange the animal names, and they took about 15–30 min to complete the task. A photograph of the whiteboard was taken every 4 min, plus a photo upon completion. Photos were also

taken before and after using the whiteboard marker to circle, label, and connect groupings. Three final example photographs are shown in Fig. 1.

3. Results

The final photographs were analyzed in terms of the positions of magnets on the whiteboard and in terms of groupings as indicated by circling. A simple computer vision algorithm was used to find the centroid of all blue circles within a given size range, and animal names were assigned to circle coordinates by hand. All 15 participants used groupings as their main organizational principle, although the shape of groupings varied from single rows or columns, to grids, to more haphazard configurations. Regardless of shape, the magnets within a given grouping were always spatially close to each other, on average, relative to magnets outside the grouping. The number of groupings ranged from 10 to 35 with a mean of 19, and the number of items per grouping ranged from 1 to 30 with a mean of 5.93. Groupings were labeled mostly according to animal taxonomies (e.g. “dogs”), but some labels were based on exemplars (e.g., “human-like”), contexts in which they are encountered (e.g., “farm” or “zoo”), physical similarity (e.g., “four-legged”), and other idiosyncratic categories (e.g. “unpleasant”). The labels indicate that spatial arrangements were based on a diversity of similarity relations, but the labels themselves are not used in subsequent analyses.

Each set of animal names was recalled in a particular sequence by a participant in the prior experiment by Thompson and Kello (2013). The original recall sequence was traced over each arrangement of names on the whiteboard, and Euclidean distance was recorded for each path segment, in camera pixel coordinates (example paths shown in Fig. 1). The mean length of each path segment was 832.1 pixels ($SD = 704.6$), and the aggregate distribution of path lengths for each animal set is shown in Fig. 2, in logarithmic coordinates. These path length distributions can be compared with IRI distributions observed in the prior category recall experiment, as shown in Fig. 3.

Two main results are illustrated in Fig. 2. First, the distributions of path lengths and IRIs both resembled inverse power law functions, albeit with roll-off for the highest values in the distribution (i.e., some degree of truncation). The exponents of these power law-like patterns can be estimated by the slopes of regression lines fit to logarithmically binned histograms (Newman, 2005), which were all between -1.0 and -2.0 . These roughly estimated exponents are within the range for Lévy flights, although they have often been observed to have slopes more consistently closer to -2.0 . The regression line slope is only meant to provide a rough comparison with power law fits in other studies—we make no claims about the goodness of fit to a power law, given that the whiteboard substantially limited the range of spatial scales measurable. The point we wish to make is that the power law-like distribution of path lengths appears to come from traversal over a spatially clustered set of items.

The other result is evidence for a relationship between spatial configurations of whiteboard arrangements, and orders in which animals were recalled from memory. This

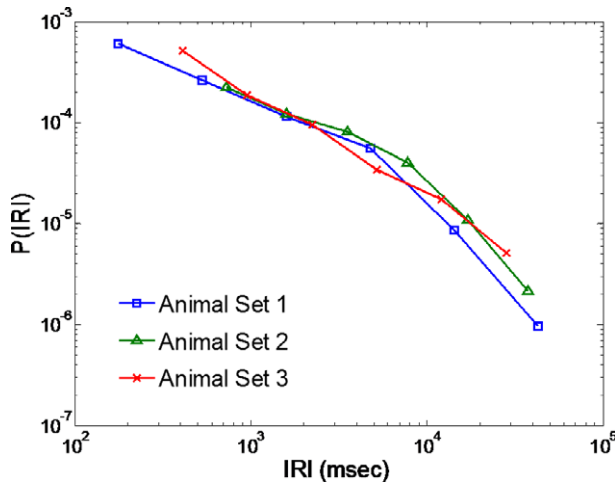


Fig. 2. Probability density function for whiteboard path lengths (left), aggregated for each of the three sets of animal names tested and plotted in logarithmic coordinates using logarithmically spaced bin sizes. The same is shown on the right for IRIs from Thompson and Kello (2013).

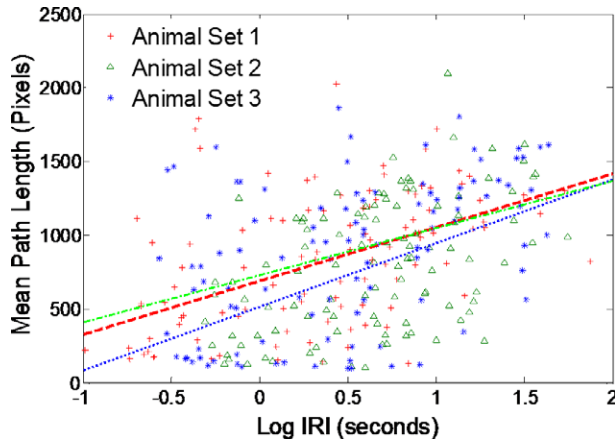


Fig. 3. Log IRI for category recall data from Thompson and Kello (2013) plotted against mean path lengths from the current experiment, averaged across each of the five participants who saw each of the three animal sets.

evidence is seen in comparing observed path lengths with those derived from randomized control sequences. For each whiteboard arrangement, a random path was traced through all the animal names, visiting each magnet exactly once. Random path lengths had a mean of 1126 ($SD = 638.9$), which is substantially longer than the mean of 832 for observed path lengths. As shown in Fig. 2, the biggest discrepancy is for the shortest path lengths: On average, 29% of original path segments traversed adjacent magnets (as determined by a distance threshold of 250 pixels), whereas for random path lengths it was less

than 6%. Moreover, randomized path lengths were much closer to the normally distributed group than the power law distributed, unlike observed path lengths. Thus, the observed inverse power law for path length distributions did not come from a random walk over clustered items.

If our spatialized recall paths represent memory foraging over the semantic space of animals, then path lengths should correlate with IRIs from the prior category recall experiment. This prediction is based on the premise that longer times between successive recalls reflect longer distances to travel in semantic space. Log IRI is plotted against path length in Fig. 3, where each point represents a single IRI from one of the three category recall participants, plotted against the mean of the five path lengths for the five participants who arranged the corresponding set of animals. Correlations were positive and moderately strong, with coefficients ranging from 0.45 to 0.48 across the three animal sets, all highly reliable at $p < .001$ ($df = 104-108$). IRIs were logarithmically transformed so that their dynamic range would be comparable to that of the mean path lengths. The latter span a limited range of scales because of averaging, and because of space limits imposed by the whiteboard. Correlations provide further evidence that paths traced over whiteboard arrangements provided a spatialized representation of memory foraging for the category of animals.

For our third and final analysis, we turn to the groupings circled by participants after arranging the magnets. As discussed earlier, semantic categories contain clusters at various scales and sizes, and clusters at one scale are represented by the groupings participants circled. The sizes of groupings were limited by the number of animals to arrange, the size of the board, and most likely an implicit expectation to produce a countable, manageable number of groupings (although participants were not explicitly instructed to do so). Nevertheless, these groupings afford a test of whether spatialized paths follow a pattern expected by optimal foraging theory: IRIs should be shorter than their mean when initially foraging a patch (grouping), and IRIs should be longer than expected when leaving a patch. This pattern is predicted because foragers should stay in a patch as long as the rate of recall is relatively fast and leave a patch when the rate of recall becomes relatively slow (Charnov, 1976).

To test for evidence of adaptive foraging, each path segment was labeled relative to transitions into and out of groupings. The first path segment into a newly visited grouping was labeled as 1, and the last path segment before leaving a given grouping was labeled as -1 . Subsequent and prior path segments were incrementally numbered higher and lower, respectively. Hills et al. (2012) reported this analysis for animal category recall data, using predefined animal patches, and their results are plotted alongside ours in Fig. 4. As predicted by adaptive foraging, IRIs were below their means when first recalling animals in a new patch, and above their means when leaving a patch. The one difference was that IRIs just prior to leaving a patch (-1) tended to be above the mean in the current experiment, but not in Hills et al. This difference may be due to differences as to how patches were defined, or differences in the lengths of recall sessions (~ 3 min in Hills et al., 10 min herein). But generally speaking, both patterns of results are consistent with adaptive foraging.

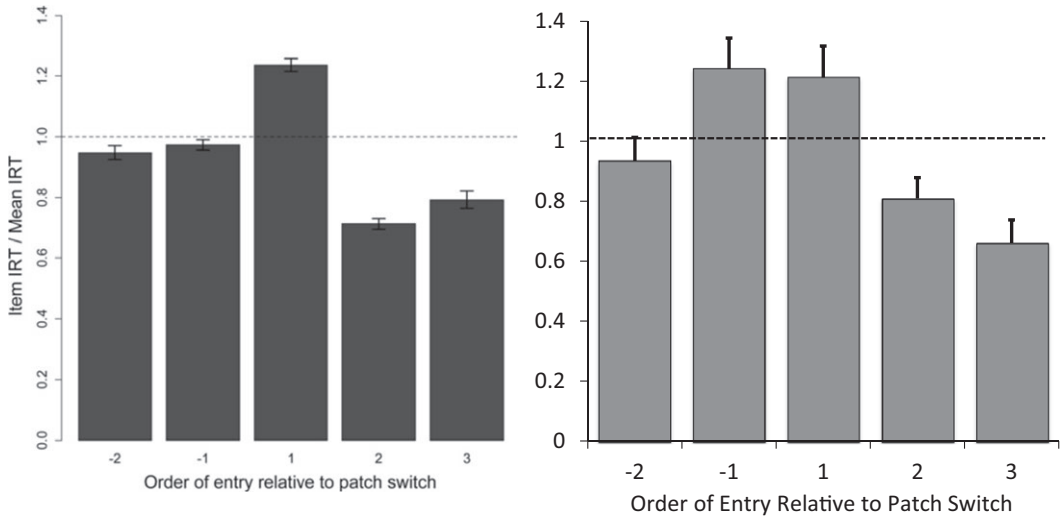


Fig. 4. Path lengths normalized by their means, and aggregated according to order relative to transitions between predefined patches (taken from Hills, Todd, & Jones, 2009) and for circled groupings in the current experiment (right).

4. Discussion

In the present experiment, searching through semantic memory was expressed as movement through a 2D semantic space. We collected individual datasets on semantic memory search in an animal recall task, and we compared these datasets against individual spatial arrangements of animal names on a whiteboard. Even though different participants performed the category recall task and the spatial arrangement task, strong correspondences were found between the corresponding temporal, and spatial behavioral measures. Thus, there appears to be a common structure to the semantic memory and knowledge that behaviors drew upon.

The elicited semantic arrangements served to demonstrate evidence for two kinds of results supporting two different theories of optimal memory search. One kind of evidence is found in distributions of path lengths and search times that resemble inverse power laws. Some researchers explain these distributions in terms of Lévy flights as random, efficient search strategies. The other kind of evidence is found in path lengths and search times as function of entering versus leaving clusters of animal recalls—IRIs and whiteboard distances were shorter upon entering clusters and longer upon leaving them. Some researchers explain results like these in terms of marginal value theorem as a kind of normative search behavior.

Our whiteboard results indicate that *evidence* for both Lévy foraging and marginal value theorem can arise from search movements over a space of clustered items. The evidence goes beyond Lévy flights per se, because results indicate that search trajectories

were non-random and had memory. The evidence also goes beyond marginal value theorem because search performance reflected more than just the distinction of moving within versus between clusters—search also reflected spatial distances between individual items and categories. Also results did not perfectly fit either theory: Path length distributions deviated from strict power laws, and search trajectories remained within clusters longer than predicted on average.

We believe both of these theories have value, in that they point out important aspects of search that any complete theory must account for. We do not place stock in deviations from a power law because they may reflect, for example, particularities of the category recall experiments and spatial arrangements on whiteboards. We also do not place stock in deviations from the marginal value theorem because they may reflect how categories are defined, and non-stationarity in expected recall rate. Instead, we consider what theory might encompass findings previously used to support each theory.

A model that simulates our results lies beyond the scope of this paper, but one promising theoretical framework is *area-restricted search* (Benhamou, 2007; Grünbaum, 1998). The general idea is that the search process becomes biased toward shorter movements as items are found, and longer movements as items are not found. In other words, the search area becomes effectively restricted when items are found, which is based on the assumption that items are clustered in space. If so, the rate of finding items will increase by searching more intensively near items already found. Hills, Kalff, and Wiener (2013) recently applied area-restricted search to data from human foraging in a virtual environment, and other studies have shown that power law distributions in path lengths can emerge from area-restricted search processes (see Benhamou, 2007). It is not clear whether area-restricted search can be formulated to yield optimal memory search behaviors, but it has the potential for both patch-like search dynamics and power law-like path length distributions. It would be informative to further investigate whether area-restricted search can be usefully applied to memory foraging and other cognitive search tasks.

Acknowledgments

The authors thank members of the Cognitive Mechanics Lab for helpful comments and feedback. This research was funded in part by a grant from the National Science Foundation (BCS 1031903).

References

- Austerweil, J. L., Abbott, J. T., & Griffiths, T. L. (2012). *Human memory search as a random walk in a semantic network*. Paper presented at the Advances in neural information processing systems.
- Benhamou, S. (2007). How many animals really do the Levy walk? *Ecology*, 88(8), 1962–1969.
- Bousfield, W. A., & Sedgewick, C. H. W. (1944). An analysis of sequences of restricted associative responses. *The Journal of General Psychology*, 30, 149–165.

- Burgess, C. (1998). From simple associations to the building blocks of language: Modeling meaning in memory with the HAL model. *Behavior Research Methods*, 30(2), 188–198.
- Charnov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theoretical Population Biology*, 9(2), 129–136.
- Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension: Embodiment, action, and cognitive extension*. Oxford, England: Oxford University Press.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82(6), 407–428.
- Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 8(2), 240–247.
- Cree, G. S., McRae, K., & McNorgan, C. (1999). An attractor model of lexical conceptual processing: Simulating semantic priming. *Cognitive Science*, 23(3), 371–414.
- Ferreira, A. S., Raposo, E. P., Viswanathan, G. M., & da Luz, M. G. E. (2012). The influence of the environment on Lévy random search efficiency: Fractality and memory effects. *Physica A: Statistical Mechanics and its Applications*, 391(11), 3234–3246.
- Goldstone, R. (1994). An efficient method for obtaining similarity data. *Behavior Research Methods, Instruments, & Computers*, 26(4), 381–386.
- Goldstone, R. L., Medin, D. L., & Halberstadt, J. (1997). Similarity in context. *Memory & Cognition*, 25(2), 237–255.
- Gruenewald, P. J., & Lockhead, G. R. (1980). The free recall of category examples. *Journal of Experimental Psychology: Human Learning and Memory*, 6(3), 225.
- Grünbaum, D. (1998). Using spatially explicit models to characterize foraging performance in heterogeneous landscapes. *The American Naturalist*, 151(2), 97–113.
- Hills, T. T., Jones, M. N., & Todd, P. M. (2012). Optimal foraging in semantic memory. *Psychological Review*, 119(2), 431–440.
- Hills, T. T., Kalf, C., & Wiener, J. M. (2013). Adaptive Lévy processes and area-restricted search in human foraging. *PLoS ONE*, 8(4), e60488.
- Hills, T. T., Todd, P. M., & Jones, M. N. (2009). Optimal foraging in semantic memory. In N. A. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st annual cognitive science society* (pp. 66–72). Austin TX: Cognitive Science Society.
- Hout, M. C., Goldinger, S. D., & Ferguson, R. W. (2013). The versatility of SpAM: A fast, efficient, spatial method of data collection for multidimensional scaling. *Journal of Experimental Psychology: General*, 142(1), 256–281.
- Howard, M. W., & Kahana, M. J. (2002). When does semantic similarity help episodic retrieval? *Journal of Memory and Language*, 46(1), 85–98.
- Humphries, N. E., Weimerskirch, H., Queiroz, N., Southall, E. J., & Sims, D. W. (2012). Foraging success of biological Lévy flights recorded in situ. *Proceedings of the National Academy of Sciences*, 109(19), 7169–7174.
- James, A., Plank, M. J., & Edwards, A. M. (2011). Assessing Lévy walks as models of animal foraging. *Journal of The Royal Society Interface*, 8(62), 1233–1247.
- Jones, M. N., & Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*, 114, 1–37.
- Kello, C. T., Brown, G. D. A., Ferrer-i-Cancho, R., Holden, J. G., Linkenkaer-Hansen, K., Rhodes, T., & Van Orden, G. C. (2010). Scaling laws in cognitive sciences. *Trends in Cognitive Sciences*, 14(5), 223–232.
- 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(2), 211–240.
- Lucas, M. (2000). Semantic priming without association: A meta-analytic review. *Psychonomic Bulletin & Review*, 7(4), 618–630.
- Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments, & Computers*, 28(2), 203–208.

- Morais, A. S., Olsson, H., & Schooler, L. J. (2013). Mapping the structure of semantic memory. *Cognitive Science*, 37(1), 125–145.
- Newman, M. E. J. (2005). Power laws, Pareto distributions and Zipf's law. *Contemporary Physics*, 46(5), 323–351.
- Nonacs, P. (2001). State dependent behavior and the marginal value theorem. *Behavioral Ecology*, 12(1), 71–83.
- 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.
- Pyke, G. H. (1984). Optimal foraging theory: A critical review. *Annual review of ecology and systematics*, 15, 523–575.
- Radicchi, F., & Baronchelli, A. (2012). Evolution of optimal Lévy-flight strategies in human mental searches. *Physical Review E*, 85(6), 061121.
- Rhodes, T., & Turvey, M. T. (2007). Human memory retrieval as Lévy foraging. *Physica A*, 385, 255–260.
- Rips, L. J. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning and Verbal Behavior*, 14(6), 665–681.
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: MIT Press.
- Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8(3), 382–439.
- Santos, M. C., Boyer, D., Miramontes, O., Viswanathan, G. M., Raposo, E. P., Mateos, J. L., & da Luz, M. G. E. (2007). Origin of power-law distributions in deterministic walks: The influence of landscape geometry. *Physical Review E, Statistical, Nonlinear, and Soft Matter Physics*, 75(6), 061114–061116.
- Schoener, T. (1987). A brief history of optimal foraging ecology. In A. Kamil, J. Krebs, & H. R. Pulliam (Eds.), *Foraging behavior* (pp. 5–67). New York: Springer.
- 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.
- Thompson, G., & Kello, C. T. (2013). Searching semantic memory as a scale-free network: Evidence from category recall and a wikipedia model of semantics. In M. Knauff, N. Sebanz, M. Pauen, & I. Wachsmuth (Eds.), *Proceedings of the 35th annual meeting of the cognitive science society*. Austin, TX: Cognitive Science Society.
- Thompson, G. W., & Kello, C. (2014). Walking across wikipedia: A scale-free network model of semantic memory retrieval. *Frontiers in Psychology*, 5, doi:10.3389/fpsyg.2014.00086.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84(4), 327.
- Viswanathan, G. M., Afanasyev, V., Buldyrev, S. V., Murphy, E. J., Prince, P. A., & Stanley, H. E. (1996). Levy flight search patterns of wandering albatrosses. *Nature*, 381(6581), 413–415.
- Viswanathan, G. M., Buldyrev, S. V., Havlin, S., da Luz, M. G. E., Raposo, E. P., & Stanley, H. E. (1999). Optimizing the success of random searches. *Nature*, 401(6756), 911–914.