UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Route choice in individuals—semantic network navigation

Permalink

Journal

Authors

Beckage, Nicole
Steyvers, Mark
Butts, Carter

Publication Date

Peer reviewed

Route choice in individuals—semantic network navigation

Nicole M. Beckage (nbeckage@uci.edu), Mark Steyvers (mark.steyvers@uci.edu)
Department of Cognitive Sciences, Social and Behavioral Sciences Gateway Building
Irvine, CA 92697 USA

Carter T. Butts (buttsc@uci.edu)
Institute of Mathematical and Behavioral Sciences & Department of Sociology
Donald Bren Hall Irvine, CA 92697 USA

Abstract

In a novel experimental task, individuals are asked to navigate from a start word to a goal word through a semantic network. In this forced-choice task, individuals perform with a high success rate (73%) and frequently navigate to the target in the minimal number of required steps (22%). We utilize these experimental results to explore different search and decision strategies. Our descriptive modeling results suggest individuals are not guessing at random (or utilizing only local information) and that knowledge of the global structure is necessary for individuals to succeed. We further show that a latent semantic space model, such as word association space, can capture much of the global semantic knowledge necessary to explain participant decisions. We suggest that performance in this task might capture some of the underlying structure of semantic memory and, importantly, search within memory.

Keywords: Semantic network, navigation, semantic memory, network navigation, search in memory

Introduction

Much work within computer science and informatics has looked at humans as information foragers (Fu & Pirolli, 2007). In many analyses, foragers rely on the structure of the environment for information foraging cues. While we know humans are able to search and gather information from a variety of environments (Fu & Pirolli 2007; MacGregor et al, 1986), we ask what happens when individuals are themselves responsible for both the structure in the environment as well as for searching on and within the structured environment. We use semantic navigation as a task in order to examine this aspect of search.

Semantic navigation includes any and all orientation and search within semantic knowledge. This could be due to communication between individuals, comprehension of auditory and visual language or encoding and retrieval of vocabulary within memory. Semantic space is unique in that it has been shaped not only by individual experience but also through cultural and historical contexts. Unlike searching on the web (Fu & Pirolli, 2007) database menus (MacGregor et al, 1986) or Wikipedia (West et al, 2009), semantic search requires searching on a naturally evolved but explicitly learned representation.

This added level of implicit knowledge of the structure of the environment might allow for foragers within semantic space to utilize the environment more effectively and quickly. We know from past work, that humans are already very good at navigation even in foreign environments such as the web (Fu & Pirolli, 2007). We set out here to see if individuals can explicitly navigate semantic space and how much individuals rely on local information, available in many types of foraging tasks, as well as global information, which is available through previous linguistic experience. We give individuals a start location in the semantic network and ask them to build chains of associates to get to a target location. This data provides us with the decisions of individuals which can in turn tell us about the underlying navigation process and semantic environment.

Results from computer science have also suggested what types of networks are most easily navigable and Kleinberg (2000) has done simulation work considering what properties networks must have for humans to successfully navigate through them. This work operates within a *message passing* paradigm, in which navigation occurs via a series of independent, uncoordinated routing decisions in which each node selects a neighbor to serve as the next decision maker in the passing chain. An important complement to this family of problems is a *route choice* paradigm, in which a single decision maker identifies the entire path to be followed.

More recently, work exploring human navigation of semantic network paths (from start to goal) has been conducted within the route choice paradigm. Specifically, in relation to this work, this has been studied in Wikipedia where individuals are given a start Wikipedia page and asked to navigate to a goal page (West et al, 2009). Whereas this work does rely on human cognition, the results of this work focus mostly on computer science implications. We hope to expand this work by exploring the cognitive implications of an individual's decisions.

To achieve this goal, we consider how humans navigate through realistic semantic network representations and, therefore, consider a novel task in which individuals are asked to navigate in a predefined semantic space. Because so little work has been done on semantic network navigation by a single individual, we set out in this paper to answer some fundamental questions. The most obvious being whether individuals *can* navigate a semantic network without being given explicit global information regarding network structure. To foreshadow our results, they can and will do so quite well in specific situations. This leads to other questions, such as what type of information might individuals be using in making routing decisions; how much local information is utilized in our specific network representation versus how much knowledge comes from global language knowledge. More generally, what can this tell us about human cognition, navigation and search?

We consider semantic space because individuals receive a great deal of linguistic input through a variety of different media. While we do not believe that every individual has the same semantic representations or knowledge, an important goal of language is communication, which facilitates the need for convergence to a similar, if not identical, representation. The fact, however, that imposing a pre-determined structure does not disallow success within this task suggests that even an impoverished representation of semantic knowledge still contains enough information for participant success.

We begin by describing the semantic network and the experimental task. Then we discuss the performance results and examine them in light of descriptive and cognitive interpretations. We consider descriptive statistics and qualitative models to help build the foundation for future modeling work. Our results importantly suggest that individuals have a specific route choice strategy, and that this strategy is greatly impacted by the similarity of an option to the end word. That is to say, individuals have some idea of distance from their current location to the goal and are often able to use this global information to navigate to the goal. With these main results, we then discuss the future for navigation models and their impact on our understanding of human cognition, navigation and search.

Methods

Semantic Network

Our task is rooted in the idea that individuals use both global and local information from the network. However, it is difficult to measure a semantic network for each individual, and moreover, sampling an individual's semantic network may bias the network and participant responses. To get away from these issues, we assume that individuals have similar semantic representations and that these represent-ations can be approximated by a network based on the Florida Association Norms (Nelson et al., 2004). While it seems unlikely that each individual has precisely the same network, an important goal of language is communication with others, suggesting that convergence on the same underlying network would be highly beneficial. Such a network could be recovered through an aggregation process, such as the Florida Association Norms (2004).

The Florida Association Norms (2004) were generated by asking participants to indicate the first semantically related word that came to mind when given a cue word. Because this was asked of many participants, we have many different associations as well as a population level proportion of responses to each cue. For example the word DOG might often elicit CAT but a measurable proportion of participants may respond with BONE. We consider a directed link between words to exist from cue to response if the cue word reliably generated the response word. Each association also has a weight equal to the probability of its elicitation. For example we consider both CAT and BONE to have a link from DOG but the link to CAT receives a higher weight since more individuals responded with CAT. This network is not symmetric. For example there exists a link from CAT to DOG but not from BONE to DOG.

Altogether 5008 words are included in the association norms with most responses being asked about as cues. However, in our experiment, we use a subset of the network. We trim the 5008-word network by including only words that had more than three words leading in as well as three leading out; we further removed the weakest connections when there were more than 12 associates. In cases where there were multiple associates with minimum strength, all were removed even when the resulting set was less than 12. We trimmed the network so that individuals would have fewer choices to sift through, were less likely to end up selecting an option that led to limited choices and to prevent trials with only a few successful paths. Limiting in-associations resulted in removing words like MOO since it is only generated in response to COW and thus all successful paths require going through COW. Further, LEFT elicits only the response RIGHT so we exclude LEFT and other similar words since it results in loops, or in the more general case, very limited options. This trimming resulted in a smaller network consisting of 2392 words. This network maintains the small world structure of the full network with a short average path length (4.19), a small overall diameter (8) and is a fully connected graph.

Word Navigation Task

In this task, individuals were given a start word and asked to navigate to a goal word. They were presented with between 3 and 12 associations of the start word and asked to pick the option that they believed would get them closer to the goal word. The selection was then centered in the screen and the next available options were generated from this word. See Figure 1 for a screen shot from the actual experiment. Each subject repeated this process until he or she reached the goal word or made a total of 25 choices (steps). Individuals could also select an undo button, which took them back to their previous decision. This incremented the 25-step count and could be repeated until the start word was reached.

Because the options individuals received were based on the association norms, we had participants complete a quick version of the association norms task. We selected 50 words that were included in the original norming study but were excluded from our experiment based on our above network trimming. After participant completion of the association task, they received verbal instructions for the word navigation task—they were told that the choices they would be offered were generated in the same manner as the task that they just completed. This was done to aid in task understanding and minimize frustration during the task. Participants then began the computer task in which they again received written instructions, an example trial and 3 practice trials before a final set of written instructions. The example trial explained the layout of the experiment whereas the practice trials allowed them to try simplified variants of the word navigation task. On the practice trials,

participants received feedback as to how many steps it took them as well as what an optimal (fewest number of intermediate words) solution would have been. After the three practice trials, they were given a few lines reminding them of the goal of the task and the opportunity to ask the experimenter should they have any questions.
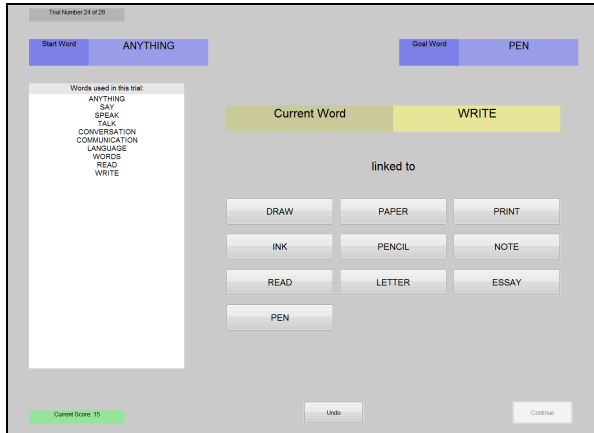

Figure 1: Screen shot of the experiment.

The test trials consisted of 28 trials divided into 4 blocks. Each block contained 7 trials, 3 requiring minimally 3 decisions, and 2 requiring each of 4 or 5 decisions. Trials were prescreened in a pilot study. Problems that were successfully completed in 15 steps or less by at least 1 out of 3 participants were selected. The block order was randomized and trials within each block were also randomized. When the 25-step limit was reached a screen popped up that said "Thank you for trying. You were # steps away," where # was the number of words between the last word clicked and the goal. They could always see previously selected words as well as the start word, goal word and current word. At the end of the first block, participants received feedback on the overall number of steps taken in that block. In each subsequent block, they could see their current block score at the bottom left of the screen. After completion of each block, a screen reported their overall performance on the completed block as well as the minimum score on any block thus far.

Overall, 53 undergraduates at University of California, Irvine were run in an experiment that lasted maximally 1 hour. Two participants did not complete the task in the allotted time and their data were excluded before analysis. All participants received course credit for completion of the experiment. To prevent meaningless clicking, an experimenter was within earshot for the length of the experiment and participants were warned that they would not receive credit unless they completed at least one trial. Every participant satisfied this requirement.

Results

Experimental Results

The first important result of this work is that individuals can reliably navigate semantic networks, moving from start to goal words in a relatively small number of steps. Every trial was solved by at least 15.1% of individuals with the average trial being solved 73.3% of the time and maximal success rate at 92.9%. The information in the semantic network is sufficient for individuals to navigate effectively. Individual performance over all trials varied from 28.6% correct to 92.9%. Further 22.2% of trials were solved in the shortest number of intermediate steps.


Figure 2: Subject performance on trials of varying min. length. Dark gray bars: proportion of trials correctly solved, light gray: proportion of trials solved in minimal number of steps. The solid line indicates unweighted random walk performance and the dashed, weighted random walk.

We expected to see a difference in success rate based on the number of minimum decisions required and figure 2 shows a general population level trend that trials requiring minimally 3 decisions have the highest success rates; however, the variation across trials suggests that more is going on. Figure 2 shows the results for each problem organized by minimum number of required steps. The first frame contains trials with minimally 3 decisions, the second 4 decisions, etc. The dark gray bar indicates the proportion of individuals who correctly solved that trial in 25 steps or less. The light gray bar indicates the proportion who solved that specific trial in the minimum number of decisions. The trials are rank ordered from least to most solved within each set. We see a general trend here that problems requiring fewer minimum steps are solved more often. This is an interesting finding since participants are not told the minimum length of a trial. We are also not considering the strength of the connections, the number of options or how quickly these problems were solved—instead the fact that the minimum number of steps can be used to help explain performance suggests that the information individuals are using during this task is sensitive to distance in the network. Salient information about distance seems to be present locally given the correlation between distance and

performance. Further, we find a similar relationship between optimal performance and minimum distance with shorter trials finished more optimally.

The trend suggesting that problems with fewer minimum steps are easier does not capture the full complexity of the task or responses chosen. With a closer examination of the results in figure 2, we see that there are many trials that violate this trend. For example, there are a few trials of medium length that are more often solved than shorter trials (and some that are more often optimally solved). It is also interesting to note that the trial that is solved most often has one of the lowest rates of "optimal" performance. This may suggest that our definition of optimal is not the correct baseline for human performance.

We are also interested in capturing the descriptive trends within individual trials. To understand what these trends might look like, we explore one problem in depth. Figure 3 shows the breakdown of a single trial. Only correct responses are included and the weight of the arrow indicates the proportion of individuals who chose that path. This trial has a high success rate and a high percentage of minimum distance paths. The minimum distance path runs along the left. This figure helps illuminate the cognitive process that may be underlying the strategy.


Figure 3: Network based on successful trials; participant responses from start word ANYTHING to goal word PEN.

For instance we can see that each option of ANYTHING led to at least one correct response—that is to say, individuals do not remove their chance of success by picking the wrong option at the first decision. Another important feature is that there are multiple successful paths. Looking closely we can see that individuals are utilizing the undo button to further explore the semantic space. For example an individual at the word PAINT selected BRUSH, but then decided to go back to PAINT. While it is difficult to say exactly why s/he made that choice, it would seem plausible that s/he went back to PAINT because s/he felt

that it was closer to the goal word of PEN than BRUSH was. Another interesting result that we can see from a detailed examination is that the word definition might change based on the goal word. This is most clear in the path that goes from ANYTHING→EVERYTHING→ WORLD→EVERYTHING(undo)→LIFE→MAGAZINE→ PAPER→PEN. Here we see that the undo button was used to back up after making a decision to go to WORLD. Further, this individual selected the word LIFE from a list of words associated to EVERYTHING. This suggests that the definition coming to mind was one of living, however, with the goal of PEN, s/he utilized LIFE to get to MAGAZINE, suggesting an interpretation of life magazine. We can also see that s/he does a similar thing in going from MAGAZINE to PAPER (likely newspaper) but then from PAPER (something to write on) to PEN.

Random Walk model

Though the data suggest that individuals are able to solve these semantic navigation problems (to varying degrees) it is possible that participants are guessing and that the structure of the network allows for high rates of success. To test this assumption we considered two types of random walk models. The first is a random walk model that simply checks if the goal word is present and if it isn't, randomly selects from the available options. The second random walk model picks the goal word if present and otherwise randomly selects from the available options with a probability distribution equal to the association norms data (e.g. if CAT was the response to DOG 80% of the time, this random walk would pick CAT 80% of the time as well). In figure 2, the two random walk models are indicated by a solid red (unweighted) and dashed blue (weighted) line. Both random walk models perform worse than our participants. The general trend does not follow that of the participants—problems frequently solved by random guessing are not those that participants most often solved. This confirms the hypothesis that individuals are utilizing global information present in semantic space in the task.

Descriptive Geodesic model

Now that we know individuals are not guessing at random, we combine the results suggested by the data to build a descriptive model of the decisions individuals make. To do this we consider the geodesic distance (number of steps between two nodes) of the current word to goal word to see if individuals are more likely to select words with lower geodesic. We know that individuals do not always pick options that decrease the geodesic because that would result in optimal performance. However, we can plot the distribution of subject choices and the distribution of all options to see if some of the success can be explained by sensitivity to geodesic distance. Figure 4, top graph, shows the distribution of geodesic from current word (WRITE in figure 1) to goal word (PEN in figure 1) along the x-axis. Options (grey buttons in figure 1) to goal word fall along the y-axis. Here we can see that proportion of choices made

by individuals, as indicated by the size of the box, look different than the full set of options. This is particularly pronounced at low geodesics (heavy weight along the near-diagonal indicates more optimal decisions). The difference between subject choices and options becomes almost unrecognizable as the distance between current and end word increases to a geodesic near 4 or greater. This suggests that individuals have knowledge of the general location of the goal word and that this becomes more accurate as they get closer to in minimal number of steps to the goal. It also suggests that individuals may be picking up on a gradient but that they might be guessing until they get close enough to the goal word to find the gradient.
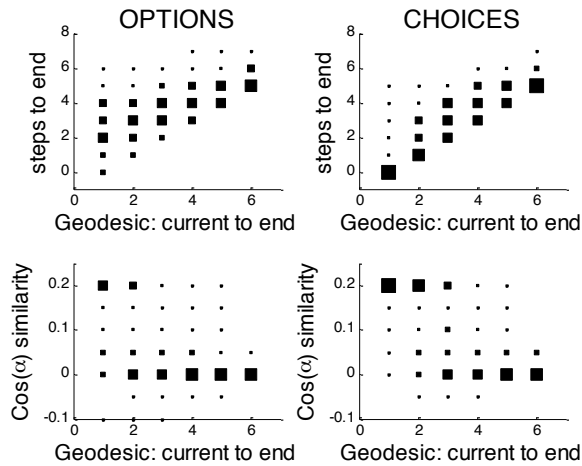

Figure 4: The top panels, based on geodesic distance, represent the distribution of options (left) and participant choices (right) with the size of the square indicating frequency of observation. The bottom panels capture local language-level information based on cosine similarity in latent space.

Latent Space model

While we consistently have been talking about navigation within a network, it is not necessary to assume individuals represent a complete network in memory. Instead it is likely that the representation is a reduction of this network—a summary of global information that allows individuals to locate their current position in the network as well as access local knowledge of nearby words. One reduction that has been widely studied in semantic memory is that of a reduction of the dimensionality of the semantic space by vector decomposition (latent space analysis, e.g. Landauer & Dumais, 1997; Burgess & Lund, 2000). While we considered a variety of latent spaces, we present data only on the symmetric WAS space from Steyvers et al. (2004). Our goal is to capture the same population level trends as we did previously with the descriptive geodesic model but with a semantic latent space that would be more cognitively plausible and require less information than the full network. To do this, we again consider figure 4 (lower panels) and

the geodesic between current and goal (x axis) but compare that to the latent space cosine similarity of options (panel 1, lower row) and subject choices (panel 2). Whereas, in the geodesic case, weight along the diagonal indicates more optimal choices, high cosine similarity suggests nearness to the target. In this graph, we expect higher weight on similarity judgments to capture more optimal choices. We again see a noticeable difference between selected choices in comparison to all options, moreover, the general trends of the latent space model follow a similar pattern to that of the geodesic in that individuals' choices are indistinguishable from guessing at higher geodesic. The latent space model offers an explanation—the cosine similarity is near zero for all options such that participants may not be able to use similarity and instead resort to guessing

Discussion and Implications

Our results suggest that individuals are succeeding by utilizing information present in the network in order to get closer to the goal word. We know that their decisions can be explained at least in part by the local information in the options, especially relative to the goal word. Their success is not, however, based on random guessing or strong associates. This is an interesting finding since it suggests that the information individuals are using is not captured by the environment of the free association task alone. Further, the paths individuals do end up utilizing appear to suggest that the semantic space may be changed and altered by the goal word implying that individuals have a direct influence on their environment. That is to say, the entire structure of the network may be influenced and changed based on the goal. Though we only gave one example in the text, it is not unique. Individuals often interpret words in light of the goal word as opposed to the current word. This adds a dynamic component to network structure that we know exists in memory and knowledge more generally. This task, further gives us a way to study the dynamic nature of semantic knowledge and the role of context in speech.

We also see that there is a large variance in potentially successful paths and that the shortest path is not always the most salient to individuals. While we have not specifically analyzed the difference between shortest paths and participant paths, West et al. (2009, 2012) have looked more closely into this question in Wikipedia navigation and suggest that shortest paths often require out-of-the-box thinking whereas paths that are a bit longer allow for a more obvious chain of associates. We hope to test this directly utilizing our data in the future.

Another important finding is that individuals seem to be making more optimal decisions (ones that get them closer to the end word) when they are already close to the end word. This suggests that individuals can intuit how far away the goal word is without having exact knowledge of the space. This is a result that has been found in most navigation studies (e.g. West et al. 2009) but most studies suggest that the way individuals get closer to the goal is by navigating to hubs with many out-links and utilizing these hubs to get to

an area of the network. However, in our study, we thresholded much of the hub structure away by allowing maximally 12 options for any word. Participants could not simply navigate to a central hub and then jump towards the goal word. Instead, we believe that this ability to perform more optimally when closer comes from the fact that individuals have a semantic representation that allows them to compute distance between two words but that this semantic representation is limited to identification of a relative location. Further, the ability to identify a word as near requires a level of information about current location and goal location that is not always available, specifically when individuals are further away from the goal. Going back to Figure 4, we see that individual choices look very similar to options both in geodesic and WAS space when many words are needed to complete the trial. This suggests that, if individuals are far enough away, guessing might be their main strategy However, guessing may be the best thing for individuals to do since there is little information available to them (as captured by WAS) and, based on network structure alone, often places them in a better or equal position (in terms of geodesic) than before.

WAS space captures most of these global trends. Particularly, WAS space is often near zero unless there is a strong similarity between words—implying that they are close enough in the network that individuals can sense it. This space also captures the noisiness of relative distance. Since individuals only have a very general idea of goal location, any estimates of which choice is closer to the goal is less exact as the distance between choice and goal increases—which is captured by cosine similarity.

In the future we hope to extend our understanding of network navigation through a relational event model (Butts, 2008). That is to say, we can assume every decision is independent once we condition on the goal word and current word. With independence, we can apply a multinomial logistic regression on linguistic and network-based covariates. We hope to use this model to show how the current and goal word influence decisions as well as more specifically exploring the specifics of language in this task.

We also hope to experimentally test subjects on other types of networks. Since all subjects in this study are natural "experts" in language, there is still the question of whether a more limited level of prior knowledge of the underlying network is still adequate to allow successful navigation and search. Work on folk knowledge of networks suggests that individuals are not very good at reconstructing social networks (Freeman et al., 1987) but our results suggest that success on this task may not require an accurate or even complete network representation, since most individuals succeeded on a variety of problems even though our underlying network of the task is impoverished.

Not only do the applications extend beyond cognitive understanding, but the fact that individuals can navigate suggests that network representations are useful. While we consider semantic space here, many other types of knowledge can be represented as a network, such as social relationships or a schedule. We believe that the results in this paper speak much more broadly about navigation than they do about language navigation specifically. A model of network navigation may be useful in explaining search, decision-making and even memory. Network structure captures many naturalistic relationships. However, unless we understand ways in which individuals are able to navigate this type of structure, we cannot utilize this representation in cognitive architectures. With this paper, we've begun to address the first concerns of understanding how individuals navigate a network structure, providing us with a new direction for navigation within memory.

Acknowledgments

References

Burgess, C., & Lund, K. (2000). The dynamics of meaning in memory, *Cognitive dynamics: conceptual and representational change in humans and machines.*

Butts, C. (2008) A relational event framework for social action. *Social Methodology* 38,155-200.

Freeman, L.C., Romney, A.K., Freeman, S.C. (1987) Cognitive structure and informant accuracy. *American Anthropologist.*

Fu, W.T. & Pirolli, P. (2007). SNIF-ACT: A cognitive model of user navigation on the world wide web. *HCI*

Kleinberg, J. M. (2000) Navigation in a small world. *Nature*, 406(6798).

Landauer, T.K., & Dumais, S.T. (1997). A solution to Plato's problem: The Latent Semantic Analysis theory *Psychological Review*, 104, 211-240.

MacGregor J, Lee, E. & Lam, N (1986). Optimizing the Structure of Database Menu Indexes. *Human Factors*

Newman, M.E.J. (2003). The structure and function of complex networks. SIAM Review 45, 167–256

Nelson, D.L., McEvoy, C.L., & Schreiber, T.A., (2004). The University of South Florida word association, *Behavioral research methods.*

Steyvers, M., Shiffrin, R.M., & Nelson, D.L. (2004). Word association spaces for predicting semantic similarity effects in episodic memory. *Experimental cognitive psychology and its applications*

Watts, D.J. & Strogatz, S.H. (1998) Collective dynamics of "small-world" networks, *Nature*, 393, 440–442.

West R., Pineau, J. & Precup, D. (2009) Wikispeedia: An online game for inferring semantic distances between concepts. IJCAI '09