Published in final edited form as:

Psychol Rev. 2015 July ; 122(3): 570–574. doi:10.1037/a0039248.

Hidden Processes in Structural Representations: A reply to Abbott, Austerweil, and Griffiths

Michael N. Jones¹, Thomas T. Hills², and Peter M. Todd¹

¹Indiana University, Bloomington

²University of Warwick

Abstract

In recent work exploring the semantic fluency task (Hills, Jones, & Todd, 2012) we found evidence indicative of optimal foraging policies in memory search that mirror search in physical environments. We determined that a two-stage cue-switching model applied to a memory representation from a semantic space model best explained the human data. Abbott, Austerweil, and Griffiths (in press) demonstrate how these patterns could also emerge from a random walk applied to a network representation of memory based on human free association norms. However, a major representational issue limits any conclusions that can be drawn about the process model comparison: Our process model operated on a memory space constructed from a learning model, while their model used human behavioral data from a task that is quite similar to the behavior they attempt to explain. Predicting semantic fluency (e.g., how likely it is to say *cat* after *dog* in a sequence of animals) from free association (how likely it is to say *cat* when given *dog* as a cue) should be possible with a relatively simple retrieval mechanism. The two tasks both tap memory, but they also share a common process of retrieval. Assuming that semantic memory is a network from free association behavior embeds variance due to the shared retrieval process directly into the representation. A simple process mechanism is then sufficient to simulate semantic fluency because much of the requisite process complexity may already be hidden in the representation.

The semantic fluency task (SFT; e.g., “name all the animals you can in a minute”) is widely used¹ to study memory retrieval in both experimental and clinical settings. A key finding in SFT is that participants tend to produce temporal bursts of semantically related items with longer lags between bursts. This “patchy” response pattern has led to the proposal that memory retrieval in the task is the product of two distinct processes: a local search process that generates a series of related items based on inter-item similarity, and a global process that moves between local regions of the search space when these regions become depleted (e.g., Troyer, Moscovitch, & Winocur, 1997; Gronlund & Shiffrin, 1986).

In previous work (Hills, Jones, & Todd, 2012), we compared the response patterns in SFT to patterns seen when animals are foraging for food, and found evidence that humans searching memory in SFT produced statistical signatures that are characteristic of optimal foraging

Correspondence Dr. Michael Jones, Dept. of Psychological and Brain Sciences, 1101 E. 10th St., Indiana University, Bloomington, IN, 47404, jonesmn@indiana.edu, Phone: (812) 856-1490, Fax: (812) 855-4691.

¹A PubMed search for semantic/category/verbal fluency returned 5,194 papers that have used the task.

policies observed in animals searching for resources in physical space (Charnov, 1976). This correspondence suggests that the mechanisms for searching memory may have been exapted from mechanisms that evolved to forage for resources in the physical world (Hills, 2006; Hills et al., 2015). However, there are alternate explanations that may also account for the behavioral patterns: It may simply be the case that memory is organized in such a patchy fashion that a model randomly producing items in relation to their similarity to previously produced items would naturally generate data that appear indicative of optimal foraging.

We tested a variety of models of search, based on classic cue-combination memory search models (Anderson, 1990; Raaijmakers & Shiffrin, 1981) applied to a spatial representation of semantic memory constructed by a corpus-based distributional mechanism (BEAGLE; Jones & Mewhort, 2007). The model that best explained the human data was one which was able to switch between two cues: Local similarity was used to generate items until no other item was close enough to pass a threshold, and then the model switched to a global frequency cue to select the next item (and search reverted again to local similarity).

Abbott, Austerweil, and Griffiths (in press; hereafter AAG) point out historical issues of model “mimicry” that may well be at play in our analysis. Our findings may be dependent on the memory representation we used.² Furthermore, our rejection of a one-stage random selection model, in favor of a two-stage local vs. global switching model, is potentially at odds with the success of random walk models in other areas of psychology (e.g., Griffiths, Steyvers, & Firl, 2007). To illustrate the issue of model mimicry, AAG first apply a random walk model to a network representation constructed from our BEAGLE space, and replicate our finding that such a model is insufficient to simulate the human data. They then apply a random walk to a network representation constructed from free association norms (Nelson, McEvoy, & Schreiber, 2004) and find that this model reproduces the behavioral markers indicative of optimal foraging that our local/global cue-switching model is able to. They propose that a random walk over a free-association network is an alternative account of SFT that is equally plausible to our cue-switching model over a semantic space.

AAG analyzed the representations used by each model (BEAGLE space vs. free association network) and determined that the behavioral bursts in SFT were better represented as clusters in the association network than they were in the BEAGLE space, and this embedded structure is what allowed their random walk model to mimic the optimal foraging predictions of our cue-switching model. They conclude, “a random walk on a semantic network produces optimal foraging behavior, but a random walk on a corresponding spatial representation does not. Consequently, it seems that there is something special about the semantic network representation that allows the simple random walk to appear similar to optimal foraging.”

²We have also noted this in prior work, as individual differences and personal experience can additionally influence the structure of semantic representation (Hills, Mata, Wilke, Samanez-Larkin, 2013; Hills & Pachur, 2012). Indeed, in those two studies and in Hills et al. (2012), multiple representations were compared against one another, including frequency, hypergraphs (categorical structures), and proximity measures (such as semantic space and social proximity). In all cases, results favored a two-stage model over a one-stage model, with search in a proximity-based local representation interspersed with occasional long-distance transitions using an alternate representation (e.g., frequency).

We certainly concur with AAG that this is a research area ripe for model comparison, and that constructing competing models of SFT is important to help design new experiments that constrain theoretical accounts and home in on the true generating process. However, the competing accounts being compared must be equivalent or no firm conclusions can be drawn and we risk rejecting what may well be an excellent model of the processes and representations that humans actually use. For example, the casual reader may conclude from the AAG article that a network representation of memory is more like what humans have in their heads than a spatial model, and Occam's razor then favors the more parsimonious random walk as the process that humans use when retrieving items in SFT over our local/global cue-switching model. The free-association network is a model of memory structure that better represents the behavioral clusters seen in SFT, and this allows a simple random walk to successfully produce patterns consistent with optimal foraging.

But there is a deep and broadly important issue concerning the structural representations of memory, and the "something special" that AAG suggest about the semantic network used in their model compared to the spatial representation used in ours. Simply put, the spatial representation was based on a statistical model of learning, whereas the network representation was quite literally built from human behavioral data on a task similar to SFT. The process used in SFT is already partially embedded in AAG's network representation.

Hidden Process Complexity in the Structural Representation

A model of a cognitive phenomenon requires an account of both structure and process, and how the two interact (Estes, 1975). Too often in cognitive modeling, the structural representation is ignored in favor of the process account (Johns & Jones, 2010). AAG are explicit in their model that memory is an associative network over which a random walk operates. Similarly, we were explicit that memory in our model is a spatial representation over which a local/global search process operates. But these two proposals differ in a crucial way.

BEAGLE (Jones, Kintsch, & Mewhort, 2006; Jones & Mewhort, 2007) is a mechanistic model of how humans construct a semantic space from statistical redundancies in the environment (trained on a natural language corpus). In contrast, free association is a behavioral task. The free association norms (Nelson et al., 2004), used by AAG to construct their network, are aggregated human responses to cue words—behavioral data. BEAGLE explains, for instance, how the pair *dog-cat* become similar to one another in memory by applying a learning mechanism to environmental regularities that humans learn from. Indeed, in early language development the structure of actual language is a better predictor of word learning than are free association norms (Hills, Maouene, Riordan, & Smith, 2010; Hills, 2012).

Free association norms are simply a tabulation of how often humans give, for instance, *cat* as a first response to *dog* as a cue. Free association is not a pure readout of the *structure* of semantic memory—it also incorporates the *process* of memory retrieval. A concern then is that AAG's "representation" of memory is based on human responses from an association task that is already very similar to the SFT that the model is meant to explain. Hence, much

of the requisite variance needed to account for the process of SFT may already be embedded in the representation. As we describe in the next section, the problem here—proposing that memory is represented as what people do in a task very similar to SFT, and then concluding that a simple process model is sufficient to simulate the SFT behavioral data—is an instance of what we refer to as a *Turk problem* in cognitive modeling. Turk problems are so-named in reference to the 18th century chess-playing automaton, and we use it to express the “man in the machine” issue of using human behavioral data as a cognitive model’s mental representation. Rather than explaining anything, this approach primarily predicts one kind of behavior (here SFT) from another very similar behavior (free association).

Free association is a dependent variable—it is something that we need to explain, not a pure measure of how memory is organized. Nelson and McEvoy have noted this in virtually all of their publications on free association, but the practice of using their association norms as a proxy for human memory is still widespread. For example, Nelson et al., (2004) originally emphasized the potential of using free association norms “for evaluating models of semantic representation” (p. 403), and noted that free association is a good predictor of many memory tasks because it uses some of the same retrieval processes used in those tasks, as well as tapping some commonalities of semantic memory. Although free association was intended as a dependent variable, Nelson, McEvoy, and Dennis (2000) note that its primary use is as a yardstick for stimulus control—but very little work has been done to investigate the yardstick itself: “Unlike printed frequency norms in which words are counted across samples of text, words are produced in free association, and we know little about the representations and processes involved” (p. 888).

The tasks of free association and SFT are extremely similar, not only because they both tap semantic memory representations, but also because they draw on similar retrieval processes. The similarity is clear in AAG’s analysis showing that the clusters in the free association network are already very similar to the SFT response clusters. SFT can be thought of as a continuous version of the free association task constrained to a given category (animals). In fact, Bousfield and Sedgewick (1944) originally referred to the task as “sequences of restricted associative responses.” In introducing the task, they note that it “is essentially an extension of the classic word association” task (p. 149).

Consider the argument *ad absurdum*: An even simpler zero-parameter model that perfectly accounts for the SFT data can be built based on the assumption that human memory is represented as an associative chain using data from SFT. This is obviously absurd because it suggests that no process is required, and just simulates SFT data from SFT data. While it is a very parsimonious and predictive model, it has no explanatory power. Predicting behavior from similar behavior without generating new understanding is the problem underlying the idea that SFT can be simulated from free association data.

Hence, we do not believe that the critical difference between our model and the AAG model is a space versus a network or a random walk versus a cue-switching model. From a graph theoretic perspective, spaces are networks—fundamentally, both are matrices. Moreover, as noted above, both models use random walks over local proximity-based representations. The real issue is, in AAG’s words: “...a random walk on a semantic network derived from free

association produces phenomena suggestive of optimal foraging, while a random walk on a spatial representation generated by BEAGLE does not. This raises the natural question: Why? What is the critical difference between these two representations?" Our response is simple: much of the variance required to simulate the SFT *process* is already partially embedded in the free association representation.

BEAGLE may very well be wrong as a model of human semantic representation. But a much more informative model comparison would be to pit BEAGLE (a spatial model) against a network created with a Topic model (Griffiths, Steyvers, & Tenenbaum, 2007)—Topic models have been used to generate networks, and are a natural contender. Then, both learning models could be applied to the same experience (a text corpus), and the resulting network and space could be fully crossed with a random walk or cue-switching model.

Notably, the problem of inferring process without attending to the influence of the search environment has also led to controversy in the ecological search literature which inspired this cognitive work. Lévy flights have been proposed to parsimoniously model animal foraging because, not surprisingly, sampling from a power-law distribution of path lengths produces the power-law distribution of path lengths that animals often produce when foraging for food. However, closer inspection of animal behavior and the resource environments in which they forage has found that animals change their behavior to increase time in resource-rich environments while minimizing time in resource-poor environments. This behavior (often called area-restricted search) also produces power-law distributed path lengths (Hills, Kalf, & Wiener, 2013; Plank & James, 2008; Benhamou, 2007) and is more optimal than Lévy flights (Plank & James, 2008; Ferreira, Raposo, Viswanathan, & Luz, 2012). Hence, assumptions that overly simplify the search environment are likely to get the search process wrong.

The Representational Turk Problem in Cognitive Modeling

As indicated above, the practice of using free association norms as a proxy for human semantic memory has a considerable history in cognitive modeling (e.g., De Deyne, Navarro, & Storms, 2013; Steyvers, Shiffrin, & Nelson, 2004). But it is important to remember that any representation based on behavioral data also contains variance from the process used to produce the response. Oddly, the opposite problem is seen in the semantic representation literature: Complex learning models are used to generate estimates of memory structure, but then a remarkably simple metric (e.g., vector cosine) is used to predict human data (see Jones, Willits, & Dennis, 2015 for a review). But a cosine is certainly not a process model. Hence, the semantic representation literature needs to likewise consider process accounts (retrieval, similarity judgments, etc.) when fitting representational models to human data.

Examples of embedding human behavior in a model to then predict similar human behavior are easy to find across the field. For example, Glenberg and Robertson (2000) demonstrated the superiority of their indexical hypothesis over Latent Semantic Analysis (LSA; Landauer & Dumais, 1997) at explaining human comprehension of sentence affordances. They found that LSA could not distinguish between afforded sentences (using a newspaper to protect

one's face from the wind) and non-afforded sentences (using a matchbook to protect one's face from the wind), but this distinction was naturally predicted by their indexical hypothesis. Humans rated the afforded sentences as easier to "envision" than the non-afforded sentences. However, the affordance predictions of Glenberg and Robertson's account were generated from sensibility ratings (virtual nonsense to completely sensible) made by the same subjects who produced the envisioning ratings (impossible to imagine to easy to imagine) that were used as the dependent variable. Not only were the two ratings on judgments that are almost identical, they were from the same participants. Hence, it is both impossible for a learning model like LSA to outperform humans, and is questionable what explanatory power comes from such an exercise.³

Certainly, the use of norms as a proxy for human representation has a very fruitful history in cognitive modeling (e.g., McRae, de Sa, & Seidenberg, 1997; Smith, Shoben, & Rips, 1974). But the representational Turk problem is becoming more common in cognitive modeling as we strive for parsimony under threat of Occam's razor. We have formal methods for penalizing models based on parametric complexity (Lewandowsky & Farrell, 2010). However, we only penalize the process account; the field does not yet have a good idea how to "tax" a model based on its representational assumptions. The representation is becoming the Cayman Islands of modeling, where complexity may be hidden from such taxation to support a simpler process mechanism, often unintentionally. Hummel and Holyoak (2003) have gone so far as to suggest that representational complexity may be the single most serious problem facing cognitive modeling as a scientific enterprise. Particularly as we see more large-scale models making use of naturalistic 'big data' in their representations (Griffiths, 2015; Jones, 2015), it is important to be vigilant about process variance hiding in the representation.

Acknowledgments

This work was supported by NSF BCS-1056744 and NIH R01MH096906 to M.N.J., and Swiss National Science Foundation Grant 100014 130397/1 and British academy mid-career fellowship MD130030 to T.T.H.

References

- Abbott JT, Austerweil JL, Griffiths TL. Random walks on semantic networks can resemble optimal foraging. *Psychological Review*. in press.
- Anderson, JR. *The adaptive character of thought*. Hillsdale, NJ: Erlbaum; 1990.
- Benhamou S. How many animals really do the Lévy walk? *Ecology*. 2007; 88:1962–1969. [PubMed: 17824427]
- Bousfield WA, Sedgewick CHW. An analysis of sequences of restricted associative responses. *Journal of General Psychology*. 1944; 30:149–165.
- Burgess C. Theory and operational definitions in computational memory models: A response to Glenberg and Robertson. *Journal of Memory and Language*. 2000; 43:402–408.
- Charnov E. Optimal foraging, the marginal value theorem. *Theoretical Population Biology*. 1976; 9(2): 129–136. [PubMed: 1273796]

³A rebuttal by Burgess (2000) summarized the flaws that made the Glenberg and Robertson (2000) paper an unfair evaluation of LSA. Although Burgess' rebuttal appeared immediately in tandem in the same issue as the Glenberg and Robertson paper, the rebuttal has only been cited 19 times compared to 485 citations for the Glenberg and Robertson paper.

- De Deyne S, Navarro DJ, Storms G. Better explanations of lexical and semantic cognition using networks derived from continued rather than single-word associations. *Behavior Research Methods*. 2013; 45:480–498. [PubMed: 23055165]
- Estes WK. Some targets for mathematical psychology. *Journal of Mathematical Psychology*. 1975; 12:263–282.
- Ferreira AS, Raposo EP, Viswanathan GM, da Luz MGE. The influence of the environment on Lévy random search efficiency: fractality and memory effects. *Physica A: Statistical Mechanics and its Applications*. 2012; 391(11):3234–3246.
- Glenberg AM, Robertson DA. Symbol grounding and meaning: A comparison of high-dimensional and embodied theories of meaning. *Journal of Memory and Language*. 2000; 43:379–401.
- Griffiths TL. Manifesto for a new (computational) cognitive revolution. *Cognition*. 2015
- Griffiths TL, Steyvers M, Firl A. Google and the mind. *Psychological Science*. 2007; 18(12):1069–1076. [PubMed: 18031414]
- Griffiths TL, Steyvers M, Tenenbaum JB. Topics in semantic representation. *Psychological Review*. 2007; 114:211–244. [PubMed: 17500626]
- Gronlund SD, Shiffrin RM. Retrieval strategies in recall of natural categories and categorized lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 1986; 12:550–561.10.1037/0278-7393.12.4.550
- Hills T. Animal foraging and the evolution of goal-directed cognition. *Cognitive Science*. 2006; 30:3–41. [PubMed: 21702807]
- Hills T. The company that words keep: Comparing the statistical structure of child versus adult-directed language. *Journal of Child Language*. 2012; 40:586–604. [PubMed: 22584041]
- Hills T, Kalff C, Wiener J. Adaptive Lévy processes and area-restricted search in human foraging. *PLoS One*. 2013; 8:e60488. [PubMed: 23577118]
- Hills TT, Jones MN, Todd PT. Optimal foraging in semantic memory. *Psychological Review*. 2012; 119:431–440. [PubMed: 22329683]
- Hills T, Mata R, Wilke A, Samanez-Larkin G. Mechanisms of age-related decline in memory search across the adult life span. *Developmental Psychology*. 2013; 49:2396–2404. [PubMed: 23586941]
- Hills T, Maouene J, Riordan B, Smith L. The associative structure of language and contextual diversity in early language acquisition. *Journal of Memory and Language*. 2010; 63:259–273. [PubMed: 20835374]
- Hills T, Pachur T. Dynamic search and working memory in social recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2012; 38:218–228.
- Hills T, Todd P, Lazer D, Redish A, Couzin I. and the Cognitive Search Research Group*. Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*. 2015
- Hummel JE, Holyoak KJ. A symbolic-connectionist theory of relational inference and generalization. *Psychological Review*. 2003; 110:220–264. [PubMed: 12747523]
- Johns BT, Jones MN. Evaluating the random representation assumption of lexical semantics in cognitive models. *Psychonomic Bulletin & Review*. 2010; 17:662–672. [PubMed: 21037164]
- Jones, MN., editor. *Big data in cognitive science: From methods to insights*. Psychology Press: Taylor & Francis; 2015.
- Jones MN, Kintsch W, Mewhort DJK. High-dimensional semantic space accounts of priming. *Journal of Memory and Language*. 2006; 55:534–552.
- Jones MN, Mewhort DJK. Representing word meaning and order information in a composite holographic lexicon. *Psychological Review*. 2007; 114:1–37. [PubMed: 17227180]
- Jones, MN.; Willits, J.; Dennis, S. Models of semantic memory. In: Busemeyer, JR.; Townsend, JT., editors. *Oxford Handbook of Mathematical and Computational Psychology*. 2015.
- Landauer T, Dumais S. A solution to Plato's problem: The Latent Semantic Analysis theory of the acquisition, induction, and representation of knowledge. *Psychological Review*. 1997; 104:211–240.
- Lewandowsky; Farrell. *Computational Modeling in Cognition: Principles and Practice*. SAGE; 2010.
- McRae K, de Sa VR, Seidenberg MS. On the nature and scope of featural representation of word meaning. *Journal of Experimental Psychology: General*. 1997; 126:99–130. [PubMed: 9163932]

- Nelson D, McEvoy C, Dennis S. What is free association and what does it measure? *Memory & Cognition*. 2000; 28:887–899. [PubMed: 11105515]
- Nelson D, McEvoy C, Schreiber T. The University of South Florida free association, rhyme, and word fragment norms. *Behavior Research Methods*. 2004; 36:402–407.
- Plank M, James A. Optimal foraging: Lévy pattern or process? *Journal of The Royal Society Interface*. 2008; 5:1077–1086.
- Raaijmakers J, Shiffrin R. Search of associative memory. *Psychological Review*. 1981; 88(2):93–134.
- Steyvers M, Shiffrin RM, Nelson D. Word association spaces for predicting semantic similarity effects in episodic memory. *Experimental psychology and its applications: Festschrift for Lyle Bourne, Walter Kintsch, & Thomas Landauer*. 2004:237–249.
- Smith EE, Shoben EJ, Rips LJ. Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review*. 1974; 81:214–241.
- Troyer AK, Moscovitch M, Winocur G. Clustering and switching as two components of verbal fluency: Evidence from younger and older healthy adults. *Neuropsychology*. 1997; 11(1):138–146. [PubMed: 9055277]