

Worm holes in memory: Is memory one representation or many?

Dirk U. Wulff (Wulff@mpib-berlin.mpg.de)

Max-Planck-Institute for Human Development,
Lentzeallee 94, 14195 Berlin, Germany

Thomas T. Hills (T.T.Hills@warwick.co.uk)

Department of Psychology, University of Warwick
Coventry CV4 7AL, UK

Ralph Hertwig (Hertwig@mpib-berlin.mpg.de)

Max-Planck-Institute for Human Development,
Lentzeallee 94, 14195 Berlin, Germany

Abstract

The analogy of space to human cognition has a long-standing tradition. Our study aims to elaborate on the validity of this analogy for search in memory. Using the search of associative memory framework (SAM) we show that people are able to dynamically recruit independent memory representations in the recall of country names. By instructing participants to use specific recall cues we also show that despite a strong effect on the retrieval sequence, total recall from memory remains unaffected. Whereas these findings strongly support a higher dimensionality to memory than often assumed, the simultaneous finding of severe retrieval time costs for non-default representations suggests that the use of particular retrieval structures may be adaptive. In sum, our results support local-to-global memory search strategies similar to foraging strategies in space, but further suggest that memory is not constrained to one local representation, but may indeed support many.

Keywords: Free recall; verbal fluency task; memory representation; Search of Associative Memory (SAM).

Introduction

Memory has long been considered to represent a high-dimensional landscape over which we search for information. The recent proliferation of semantic space models, which acquire semantic similarity of words based on statistical processing of text corpora implicitly characterizes memory as embodying such a representation. This idea is not new. In his “Principles of Psychology” William James wrote “We make search in memory ... just as we rummage our house for a lost object” (1890, p. 654), suggesting that search in memory is comparable to search in space. But how comparable is it? Research on spatial imagery and cognitive maps suggests that mental operations share much in common with the way we move around the physical world (e.g. Kosslyn, Ball, and Reiser, 1978). Almost all models of long-term memory incorporate a dimension of similarity (inverse of distance) in order to explain priming and serial position effects (e.g. Anderson, & Pirolli, 1984; Brown, Neath, & Chater, 2007). Shepherd’s account of

distance in mental representations (Shepard, & Metzler, 1971), as well as models of categorization (e.g. Nosofsky, 1988), suggests a similar conceptual landscape, in which similar items reside near one another and less similar items reside further apart.

Underlying these approaches is an implicit assumption, one that is highlighted by James. If searching memory is like rummaging our house for a lost object, is there just one house (i.e., representation) or are there many? Could an item be in more than one representation, and if so, does one representation facilitate memory search better than another? If memories reside in multiple representations, this presents a problem for many existing models of knowledge representation—that is, especially those that produce but one representation. Moreover, multiple representations would exemplify a feature of memory that clearly separates memory from space: outside of wormholes, the only way to get from one point in space to another is to travel the distance between them. Memory, on the other hand, may have no such constraints.

Before we describe how we investigated the possibility of multiple memory representations, we first describe some of the previous research that has highlighted the relationship between memory and space.

Memory and Space

Following this analogy between space and memory, Hills and colleagues recently put forth a model wherein free recall from memory produces patterns that can be predicted by a classical theorem of optimal foraging theory, the marginal value theorem (Hills, Jones, Todd, 2012). This theory describes optimal switching between explorative and exploitative search in response to a patchy resource environments. Further supporting the notion that internal search may be similar to external search in space, Hills, Todd, and Goldstone (2008) found that priming search in space primes search in a lexical search task, suggesting that a shared cognitive process may search in both domains.

Drawing from the search of associative memory framework (SAM), Hills and colleagues' model expresses search as an alternation between local and global memory search. In local search, sequential items are retrieved based on similarity to the last recalled item together with a position invariant context cue. In global search, this context cue is used exclusively. Interpreting the network of similarities as a landscape of distances, local search is spatially confined with nearer items in memory being more likely to be retrieved. Global search, on the other hand, is independent of this landscape allowing for jumps across the landscape that may utilize an alternative representation—much as wormholes do in science fiction. These aspects of local and global search capture the different search modes typical of exploration/exploitation trade-offs found in patch foraging models. However, they appear to break with the spatial analogy to the degree that the search process can escape the confines of one representation by switching to another.

Recently, Abbott et al. (2012) proposed a slightly different model to explain the findings of Hills and colleagues with a stronger focus on the underlying representation. Their theoretical approach expresses the patterns of free recall as a random walk through a single partly connected graph (see Steyvers, & Tenenbaum, 2005). Their full random walk model also incorporates local-to-global transitions. However, compared to Hills et al. the switches between subsequent cues a) are random and thus independent of local retrieval success, except in the case of allowing more time for a global jump to occur and, more importantly b) do not imply any change of representation but rather a reset to the start point of search (similar to the executive search process model used to describe search for anagrams in Hills, Todd, & Goldstone, 2010). Additionally, the simulation results of Abbotts et al. also indicated that apparent optimal retrieval patterns were possible without local-to-global transitions. Together, this work provides only weak support for multiple representations and the presence of context-based local-to-global transitions in memory.

Present Study

Overall, the spatial analogy for memory and memory search has been successful, but existing theoretical approaches offer alternative hypotheses. Moreover, all of the previously described approaches have assumed that there is only one representation of memory that allows for local search. To investigate the potential for multiple representations in memory, we had people search their memory for country names. Friedman and Dewinstanley (2007) showed descriptively that at least three independent factors predict country retrieval: geographic distance to the previously recalled country, phonetic similarity to previously recalled countries and characteristics of the particular country itself (e.g. their frequency in the news). The availability of these multiple

objectively determinable cues within a single recall category suggest the potential for multiple representations, and allow us to address the following questions within a computational framework.

First, we want to assess if these three cues are used and how and in what form they are integrated in the retrieval of countries. Is the default local search representation (similarity-based) best characterized by space (e.g., Euclidean distance) or some other representation? Further, are the local search cues integrated dynamically with the global cue (item-based). To address this, we collected uninstructed recall data where people were simply asked to name all the countries they can think of. Our second question addresses more specifically the question of multiple representations: Specifically, how does a voluntary change of retrieval cues influence recall? Assuming a unitary underlying representation, changing the retrieval cues should harm retrieval with respect to response times, number of retrieved items, or both. Provided representations are independently accessible, retrieval sequences should reflect changes in the representation, and may further reflect differences in the accessibility of information with a representation. To this end we ran two instructed conditions in which participants were asked to base their recall on the letters of the alphabet or on geographical neighbors.

Method

Participants We collected data from 71 students at the University of Basel. The sample had an average age of 24.7 and 71% of the participants were female. Participation in the study was rewarded either by course credit or a fixed payment of 7 Swiss francs. Additionally the participants received 0.25 Swiss francs for every recalled country.

Procedure Participants were seated in front of a computer. First, each participant's typing speed was assessed. Next, they received the instructions to the country fluency task. In the control condition participants were asked to type all the countries they can think of, but were not given direction with regard to how to retrieve countries. In the alphabet condition participants were instructed to proceed by the letters of the alphabet. In the neighbor condition participants were asked to always first attempt to recall a neighboring country (with a shared border) before recalling a country from elsewhere. In both the alphabet and the neighbor condition participants were also instructed that whenever there were unable to recall a country by the first letter or a neighbor they could recall any other country.

Scoring All country entries were checked for spelling and validity. Only the 193 current members of the United Nations as well as Kosovo, Taiwan, Vatican City and Palestine were accepted as valid countries. Some synonyms were accepted, for example "Holland" for Netherlands. As approximations for the spatial

representations (i.e., retrieval structures) we implemented three measures: *Distance*, calculated as shortest Euclidian distance between the borders; *Neighbor*, indicating if two countries shared a border; and *Geodesic* distance, determined by shortest number of border crossing required to move from one country to the other. All information used for the spatial representations were based on the CIA World DataBank II. To measure phonetic representation, three further retrieval structures were implemented: *Levenshtein* distance, which indicates the number of orthographic edits; *Initial* letter, indicating if consecutive items share the same first letter; and *Phonetic* similarity following Friedmann and deWinstanley (2007), with phonetic similarity indicating the same letter in the first or the last three positions of two countries. To estimate the frequency we took two measures: *Google*, indicating the log number of hits for a country generated by a Google search, and *News*, indicating the log number of mentions in the weekly newspaper Die Zeit, which is widely read in all German speaking countries.

Modeling The model framework we used to simulate the search process is based on SAM (Raaijmakers & Shiffrin, 1981). The foundational assumption of the model is that recall is achieved by probing retrieval structures in memory with a specific cue set, that is, the memory probe. With I representing a possible target item for recovery in the search space, the probability of retrieving I is computed as the product of the individual retrieval strengths for I across a probe set of M cues, with $S(Q, I)$ representing the semantic similarity between cue Q and item I . This is incorporated into an overall probability of retrieval for item I via the ratio rule:

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

where N represents the total number of items available in the category for retrieval and β represents the saliency (or attention weight) assigned to a given cue.

Every search cue generates a retrieval strength $S(Q, I)$ for each item based on the items similarity to the last item, e.g. in terms of Euclidian distance, or the item's own qualities, e.g. frequency in the newspapers or the Eigenvector of similarity-based cues. Using a maximum likelihood method, we fit β s to each participant's data, using the participant's individually generated sequence of items. This produced a log-likelihood fit, which was penalized based on the number of free parameters via the Bayesian information criterion. Results are presented as the median improvement in the Bayesian information

criterion relative to a random model specifying that all remaining items in the search space are equally likely to be retrieved, with greater values of BIC indicating a better fit.

We examined various static and dynamic models, using spatial, phonetic and frequency cues. In our terminology, static models rely on the same set of cues over the entire retrieval interval. Dynamic models on the other hand allow for a switching between cues. In SAM similarity-based cues are dropped when a threshold of retrieval failures is reached. This we modeled using an additional threshold parameter on the retrieval strengths of similarity-based local cues.

Results

Which cues are used?

As a first step in the analyses the predictive power of individual cues¹ was tested in a single cue version of our retrieval model. Figure 1 shows the BIC advantage over the random model for the unconstrained as well as the Alphabet and Neighbor condition. The results indicate that not all retrieval structures are equally predictive in the uninstructed (control) condition. Spatial representations seem clearly to be the dominating cue in this condition. Next to spatial information, only frequency of mentions in the news had predictive power. In contrast to the results of Friedmann and deWinstanley




Figure 1: Median of differences in BIC between the random model and single cue retrieval models.

¹ The eigenvectors of the similarity-based cue representations were dropped from the analyses as none were predictive.

(2007) when our participants were not instructed to use any particular cue they do not seem to rely on any phonetic cue. The overall pattern changes substantially in the instructed conditions. When instructed to recall by the letters of the alphabet, the Phonetic and Initial cue models fit very well, whereas the individual spatial models do not exceed chance level. On the other hand, the fit in the Neighbor condition is best fit by the spatial models. Thus, people appear to have changed their retrieval behavior in both conditions.

Are cues integrated dynamically or statically?

Contrary to the expectation that all cue classes – spatial, phonetic and frequency – contribute to the fits, the single cue data indicates that only two of three classes of cues are used in the individual conditions. We further asked how the cues are integrated and if the type of integration holds over the conditions and the particular cues used in these conditions. We compared two models: *static*, with both cue classes are used over the entire retrieval interval, and *dynamic*², using the similarity-based cue together with the context cue or, when the similarity-based cue falls below a fitted threshold, the context cue alone.

Figure 2 illustrates the results for these models. In line with the single cue models the models combining News with spatial cues fit the data of uninstructed and Neighbor condition best and provide a poor fit for the Alphabet condition. The Alphabet condition was best fit by phonetic information combined with frequency in the news. In regard to the state of integration the models also show a clear pattern. Irrespective of the condition, the dynamic models provide a better fit to the data than the static integration models.

Does cue use affect performance?

The analyses thus far show a) that cues can be voluntarily changed and b) that this however has no effect on the dynamic integration of local information with global frequency information. But did the controlled choice of a particular retrieval impact memory accessibility? Figure 3 shows the results for number of countries retrieved. An analysis of variance reveals that the slight advantage in the Neighbor condition is not greater than we would expect by chance ($F_{2,68}=.73, p=.48$). Thus, the overall accessibility in terms of number of countries was not dependent on using a particular cue. Item response times on the other hand reveal a substantial detriment when countries were retrieved by the letters of the alphabet. The median retrieval time in the alphabet condition (mdn

= 7.7s) was about five times higher than the retrieval times in the uninstructed (mdn = 1.4s) and neighbor condition (mdn = 1.5s). Consequently, participants in the Alphabet condition also keep on searching for a much longer period than in the other two conditions. This is likely due to participants feeling that they could not go back to previous letters. Clearly, however, not all alphabet-based responses are slow. About 21% of the response times in the alphabet condition fall below the medians of the other two conditions. Thus in a number of cases Alphabetic retrieval was faster than in the uninstructed and Neighbor condition.

Discussion

In this study we were interested in the utilization and integration of multiple cues in retrieval from memory. By having participants retrieve all the countries they know under three different instructions we were able to show that dynamic search models as proposed by Hills and colleagues (Hills, Jones, Todd, 2012; Hills, & Pachur, 2012) provide the best account for the data in all conditions. Further, the data clearly demonstrated that people are able to deliberately change the cues they are using (see Gronlund & Shiffrin, 1986). This however had no impact on how these cues were combined with a global representation of frequency. Finally, our data

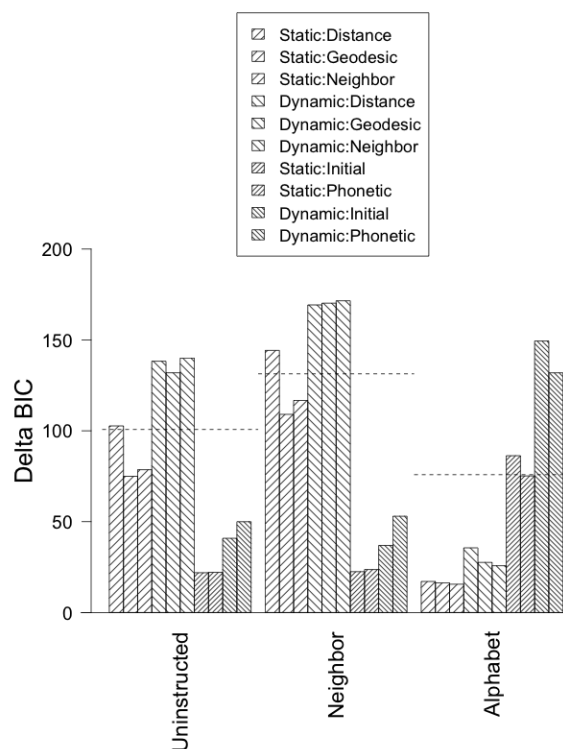


Figure 2: Median of differences in BIC between the random model and different cue integration models for the three conditions. All similarity-based cues are combined with the News cue. Dashed lines represent the best single cue model in the three datasets.

² The dynamic search model described in the text corresponds to the search models in Hills, Jones and Todd (2012). An alternative version of the model was also tested, that uses the similarity-based cue when above threshold and the context-based cue when below threshold. However, the results were indistinguishable.

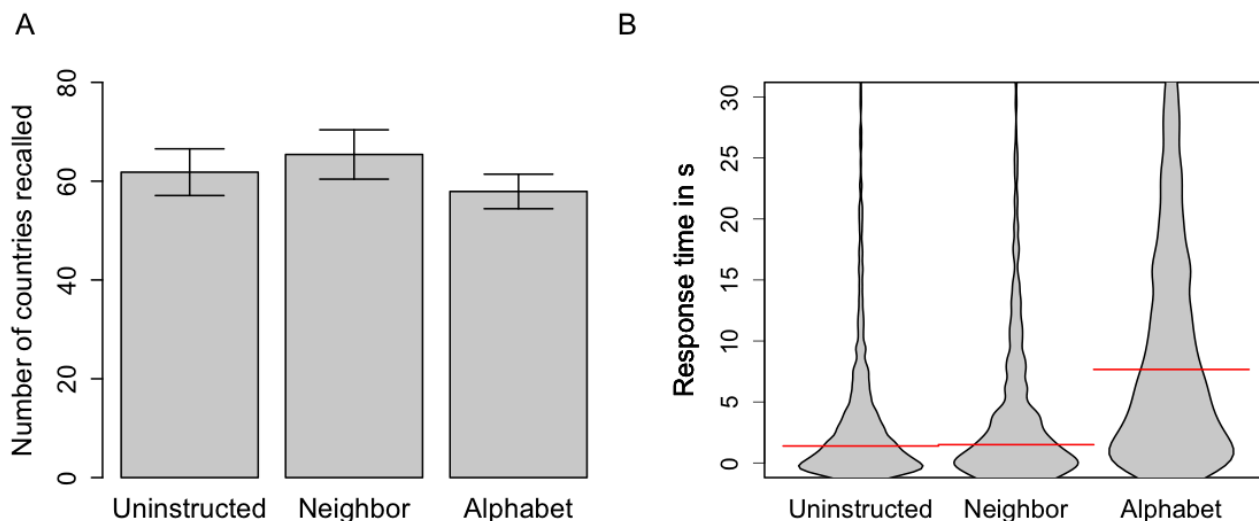


Figure 3: Number of items retrieved (A) and bean plots of item level response times in seconds corrected for typing speed (B) in the three conditions. Error bars represent standard error of the mean. Shapes in the right hand plot represent the density, the solid red line the median.

shows that changing the cue is not necessarily harmful to the recall performance – the same performance level was reached in terms of total items retrieved, despite dramatic costs in overall retrieval times.

What do our results mean with respect to our initial question whether memory is one or many representations? Clearly, people are able to change the cues they are using and our results further suggest that they may access alternative representations. However, these changes can come with costs. These costs can be interpreted as the result of different distances within a given representation. Thus, a wrong cue or retrieval structure might mean traveling greater or lesser distances in memory. Under this interpretation two speculations can be made. The comparable performance in terms of number of retrieved items would disappear under time pressure and, in the present case, the overall retrieval success is potentially a result of a relaxed retrieval failure threshold. It remains striking that the same overall performance was reached in our three conditions, as recognition data indicates that people have potentially about twice as many countries stored in their memory (Friedman & deWinstanley, 2007).

Moreover, in every case, models using frequency in the news alone for stretches of the retrieval interval performed much better despite being penalized by the extra threshold parameter. This, there appears to be at least one alternative representation allowing for non-spatial movement in memory space. Combined with evidence for dynamic switching, this breaks with the spatial analogy by allowing for travel through memory via multiple representations.

Both, the successful switch of representations under specific instructions and the independent use of frequency are difficult to explain within the model that is based on a unitary representation or space as proposed by Abbot et al. (2012) and others. Our results seem to be much better explained by frameworks allowing for the variable integration of multiple cues. The SAM-based memory search model developed by Hills and colleagues is but one model affording this possibility. Other models such as multi-trace memory models (e.g. MINERVA; Hintzman, 1986) or the recently proposed context maintenance and retrieval model (CMR; Polyn, Norman, & Kahana, 2008) are also in principle capable of utilizing multiple cues to varying degrees over time. However, in modeling and most experimental work the possibility of entirely switching between representations has been rather neglected. In our eyes this possibility should receive more attention in future research.

Assuming that our current findings are not constrained to the recruitment of a spatial versus a phonetic or alphabetic local search representation leads to the question of what is the right retrieval strategy to use. Clearly, our data shows that, without instructions, alphabetic and phonetic retrieval strategies receive little support. The data also suggests that this is done for a good reason, as response times tend to be on average larger when using phonetic cues. On the other hand, a substantial number of alphabet retrievals were at least as fast as retrievals based on spatial information. A savvy memory forager could potentially exploit this fact by adaptively switching between retrieval cues – that is, by taking dimensional short-cuts through memory space. In principle, this is no different from the short-cuts made

available by global transitions to frequency. But, cognitively, it represents the capacity to jump between local representations, or landscapes, in much the same way that children might enter an alternate universe by passing through a mirror.

three-dimensional objects. *Science*, 171, 701–703.
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.

References

- Anderson, J. R., & Pirolli, P. L. (1984). Spread of activation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 791–798.
- Abbott, J. T., Austerweil, J. L., & Griffiths, T. L. (2012). Human memory search as a random walk in a semantic network. *Advances in Neural Information Processing Systems*, 25.
- Brown, G. D.A., Neath, I., & Chater, N. (2007). A Temporal Ratio Model of Memory. *Psychological Review*, 114, 539–576.
- Friedman, W.J., & deWinstanley, P.A. (2006). The mental representation of countries. *Memory*, 14, 853–871.
- Gronlund, S. D., & Shiffrin, R. M. (1986). Retrieval strategies in recall of natural categories and categorized lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 12, 550–561.
- Hills, T. T., Todd, P. M., & Goldstone, R. G. (2008). Search in external and internal spaces: Evidence for generalized cognitive search processes. *Psychological Science*, 19, 676–682.
- Hills, T. T., & Pachur, T. (2012). Dynamic search and working memory in social recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38, 218–228.
- Hills, T. T., Jones, M.N., & Todd, P.M. (2012). Optimal foraging in semantic memory. *Psychological Review*, 119, 431–440.
- Hintzman, D. L. (1986). “Schema Abstraction” in a Multiple-Trace Memory Model. *Psychological Review*, 93, 4, 411–428.
- James, W. (1890). *The Principles of Psychology*. New York: Holt.
- Kosslyn, S. M., Ball, T. M., & Reiser, B. J. (1978). Visual images preserve metric spatial information: Evidence from studies of image scanning. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 46–60.
- Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 700–708.
- Polyn, S. M., Norman, K. N., & Kahana, M. J. (2008). A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review*, 116, 1, 129–156.
- Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review*, 88, 93–134.
- Shepard, R. N. & Metzler, J. (1971) Mental rotation of