Modeling human memory retrieval and computer information retrieval: What can the two fields learn from each other?

Peter W. Foltz[1]

Department of Psychology
University of Colorado
Boulder, CO, 80309
pfoltz@boulder.colorado.edu

Abstract

Models of human memory and computer information retrieval have many similarities in the methods they use for representing and accessing information. This article examines the methods and representations used in both human memory modeling and computer information retrieval and discusses similarities and differences From these similarities and differences, the features that lead to successful retrieval in both human memory and computer information retrieval domains can be determined. An analysis of these features can then help in the future design of both human and computer retrieval models.

Introduction

A recent estimate puts the amount of information accumulated by an adult over a lifetime to be on the order of 10^9 bits (Landauer, 1986). Modern computers are also able to store about this same amount of information. While this great capacity permits humans and computers to store much information, one of the keys to the flexibility of both the human brain and the computer is the ability to retrieve the relevant information quickly and accurately. Without fast access to information, both humans and computers could not make the almost instantaneous actions for which they are both known. Thus, a primary issue in both computer and human information retrieval is how to access the correct information in an efficient manner.

Human and computer memories differ in many ways. While human memory is neurally based and has evolved over millions of years, computer memory is silicon based and has existed for only half a century. The approaches to the study of the two also differ. In human memory retrieval, the primary approach has been to predict memory retrieval phenomena through development of psychological models of memory and to test these models via laboratory experiments. However, in computer retrieval, rather than concerns of trying to explain how the retrieval works, the primary issues involve designing the most efficient manner to store and retrieve the information.

While both fields have different goals and work on different types of storage, they have similarities in the methods used for representing and accessing information. This article examines methods and representations used in both human memory modeling and computer information retrieval From their similarities and differences, the features that lead to successful retrieval in both domains can be determined. An analysis of these features can aid in the design of both human and computer retrieval models. This article is a summary of findings from Foltz (1991).

Stages of Retrieval

The retrieval of information from both human memory and computers occurs in three stages: generating the retrieval cues, using the cues to retrieve information, and verifying that the retrieved information is what is desired. These stages differ in the type of cognitive activity used. To retrieve information, cues first need to be generated. This process is strategic, involving controlled processing. A person develops cues that describe the information desired based on the current context. The generated cues are then used for the retrieval. In the case of information retrieval, the cues are transmitted to the computer, while in human retrieval, the cues are used automatically. In both cases, the actual retrieval process is automatic, i.e. not under a person's strategic control. Once information has been returned from the retrieval process, it is again under the control of a strategic process which evaluates the information to determine if it is what was desired. The retrieval of information may not be just a single cycle of these three stages, but may involve several iterations. Information retrieved in a previous iteration may be added as additional cues for the next retrieval. In many human and computer retrieval models, the automatic and the strategic components are often treated separately. In this article, the models are organized based on whether an automatic or strategic component of retrieval is being examined.

Psychological approaches to retrieval

Human memory models are developed with the primary goal of trying to explain how our memory works. The approach tries to answer questions about how we store, retrieve and represent information, and why we may fail at retrieval. While there have been many models of memory retrieval, this section focuses on three different classes of models of automatic retrieval from long term memory.

---

[1]This research was supported by the Army Research Institute, project MDA 903-86-CO143. The author thanks Walter Kintsch, Anders Ericsson, Reid Hastie, Gerhard Fischer, Peter Polson, Adrienne Lee, Susan Dumais, and Thomas Landauer for helpful discussions and comments on versions of this paper.

Automatic memory models

Compound cue. The compound cue model of memory (Raaijmakers & Shiffrin, 1981; Ratcliff & McKoon, 1988), assumes associative connections between retrieval cues and items in memory. The models serve as a general theory of retrieval from long-term memory and has been used to predict such phenomena as the recall of word lists and the range and decay of priming. In the model, each item in memory has a probability of being sampled that is dependent on the associative strength of that item to a set of probe cues in comparison to the strengths of all other items to the same set of probe cues. While cues are connected to memory items, there are no connections between the items in memory. Thus, connections between memory items are only mediated through connections from two memory items to the same cue. Retrieval of information is treated as an iterative process, adding retrieved memory items as new retrieval cues until a certain number of failures to retrieve any new items occurs. Thus, retrieval in the compound cue model is entirely dependent on the cues used.

Spreading Activation. The concept of spreading activation has been widely used as a search mechanism of semantic networks (e.g. Anderson, 1983). Spreading activation models of retrieval use a semantic network of interconnected memory items, with connections representing semantic relationships. As in the compound cue model, initial items are activated by the cues given. The activation of these items then spreads along the connections to activate other associated items. The amount of spread from one item to another is typically based on the amount of activation received by the first item, the number of connections out of the first item and the connection strength between the items.

In terms of retrieval, compound cue and spreading activation models differ primarily on the range of activated nodes. In compound cue retrieval, the only activated nodes would be those directly associated with the cues. In spreading activation, retrieval is not only cue-dependent, but also dependent on the connections between memory items.

Distributed models. Distributed models of memory (Murdock, 1983; Hintzman, 1986), differ from the previous models in terms of their representation of information in memory. In these models, information is represented as a vector of features rather than discrete nodes in which a concept is stored. In retrieval, a vector representing the cues and context is compared to the stored vectors. The result represents the best match between the cues and the stored information while also capturing some of the similarities between the information that has been previously stored. In this manner, given a set of features, the retrieval process returns items that best match the set of features, although the items are slightly abstracted from when they were encoded due to their associations to other memory items.

One primary difference from the previous models is that the distributed models match based on features of the item and on the associations made with other encoded items. While the compound cue and spreading activation models have associations between similar items, they do not directly match based on the encoded features. However, the models all encode information items using some form of associations between the items and other information encoded with them and therefore all exhibit context dependence. They also all display cue-dependence in that the retrieval of information is primarily dependent on the cues given, although it can be mediated by the patterns of associations of other stored information.

Strategic retrieval

While the models described above address retrieval issues once the cues have been provided, they do not address many of the issues of human strategic control of retrieval. When we need to retrieve information, we must decide what are the best cues to use for the retrieval and evaluate what is returned to determine if it is what we want. In a protocol study of people recalling the names of their high school classmates, Williams (1978) identified some of the strategies used in retrieval. He found that people initially identified a context in which to search, then searched their memory, and then verified what was retrieved. The process was applied recursively, with people using the whole cycle to retrieve contexts for another retrieval. The Williams study highlights the need for an appropriate context for retrieval. Context effects have been widely researched in connection with Tulving's encoding specificity principle (Tulving & Thompson, 1973). Thus, information can only be retrieved given the appropriate context with which it was stored.

In addition to context, encoding strategies also play a major role in retrieval. Chase and Ericsson's (1981) studies on expert memory show that retrieval is not only dependent on effective retrieval strategies, but also on effective encoding strategies. Thus, specific strategies are developed for both encoding and retrieving information, with these strategies being highly person and domain specific.

The strategic aspects of retrieval described in this section point to the fact that retrieval is not just the process of searching a series of memory items given certain cues to see which items best match the cues. Instead, a context for the retrieval must be established. Individual cues, such as the few words used in memory models for cued recall, are seldom the only items used for retrieval. A person not only uses the word cue that was provided, but also the additional context of the situation when the original word was encoded for the retrieval. As information is retrieved, it too serves as a context for retrieving additional information. Specific strategies are similarly used when initially encoding the information, thereby facilitating the use of context to help retrieve it.

Computer information retrieval

The goal of Information Retrieval (IR) is to retrieve what is wanted while minimizing the retrieval of undesired information. There are again, the two interacting sub-processes of retrieval. There is a controlled process in which humans must decide what cues to provide to the IR system and evaluate what is returned. Once they have decided on the cues to use, they provide the cues to the automatic process of the system to retrieve the information.

Automatic retrieval models

As in human memory retrieval, retrieval of an item is dependent on the similarity between that item and the cues given. The variety of information retrieval models represent different methods of calculating and representing these similarities in order to maximize the effectiveness of the retrieval given the tasks and environment. In the retrieval of textual information, each document is treated as a set of features, in which each feature corresponds to a term used in the document. A standard method of retrieval is to represent each term as a vector with each vector element representing whether a particular document contains the term. Given a query consisting of terms, the best matching documents can be retrieved through Boolean operations.

In the vector space model (Salton & McGill, 1983), each document is represented as a vector of terms in a high dimensional space based on the indexing terms. Any query can also be represented as a vector in the space. The similarity between a query vector and document vectors can then be measured by the distance between the two vectors. Like the distributed human memory models, this similarity is calculated using the sum of the products of the vectors.

Another retrieval method used is probabilistic retrieval (Bookstein & Swanson, 1975). Documents are ranked based on two probabilistic relations between terms and documents, the probability that a given term occurs in the relevant set of documents and the probability that a term occurs in the non-relevant set of documents. By estimating probabilities from a collection of documents, the probability of the relevance of each document given the terms can be calculated.

The fact that many IR methods require using the exact words from the document to retrieve it highlights one of the deficiencies in current techniques. People seldom know which words will describe a document and there is a great variability in the choice of words between people. Thus, keyword matching can fail due to *polysemy* (multiple meanings for a word) and *synonymy* (multiple ways of refering to one concept).

While the models described above are based on relations between individual documents and the cues given to retrieve them, some retrieval models are more *network* based, in that they emphasize the connections between the documents. In this way, the documents form a semantic network of information in which documents on similar topics tend to cluster. A variety of methods have been used for creating network representations using the terms to determine feature similarity such as clustering and factor analytic methods. Retrieval of information can then be done by returning information occuring in regions of the semantic space. The problems of exact word usage are avoided since semantically similar documents can be retrieved together.

Strategic retrieval

In retrieval it is often not clear to the user how or what can be retrieved. Users may not know which terms to use due to problems of synonymy and because they are not familiar with what type of information they can retrieve from the database. There are also problems with the actual interaction with the system; users may not know how to form a query or use the query language. Thus, a user interacting with a retrieval system may need to use some conscious strategies. To ease these strategic problems, IR uses methods for interpreting what a user wants and ways of letting the user browse through the information.

In relevance feedback, (Salton & Buckley, 1990) the IR system returns a list of documents after an initial query. A user can then rate the relevance of the retrieved documents. This information is then used to perform additional retrievals. Since the new query contains more refined information about what the user considers relevant, it tends to return more relevant items. Users thus, are no longer dependent on providing accurate terms for their queries, but can just indicate documents that appear relevant to their interests. This reduces the effort of having to know what is in the database and allows users to have a lot more control of the iterative search process.

Relevance feedback is also similar to information browsers and hyptertexts. Information browsers use a set of rich connections between documents to allow a user to navigate through the space of information. Through navigation, a user can see the relationships between items in the semantic network and move to areas where the most relevant items are located. Thus, as in human retrieval, strategies can be used for finding the desired information.

Hybrid models

With the many similarities between human retrieval and computer information retrieval, there are some models and systems that have incorporated features of both. These models could be considered *hybrids* in that they do information retrieval, but apply psychological models of retrieval. These models cross the boundaries of being just an information retrieval or a human memory model, but use features from both fields in order to be an effective retrieval system. They therefore illustrate the complementary nature of human and computer retrieval models.

One of the primary ideas from the human memory literature that has been used in information retrieval is the concept of differential associative connections between items of information. For this reason, there have been a variety of retrieval models using spreading activation or connectionist models (Jones, 1986; Rose & Belew, 1989). In these systems, a user can activate certain terms for their query. Activation then flows from these terms to the documents and other terms until the network settles. The highest activated documents would then be retrieved.

Some information retrieval systems incorporate psychological models of users' retrieval strategies. One such system was Williams' (1984) RABBIT system, based on his research on memory retrieval. strategies. With RABBIT, people used an iterative process of giving a partial description of what they wanted, retrieved a general context and then used the information to narrow the cues to get the information. Thus, the system allowed users to do computer retrieval using familiar memory retrieval strategies.

A very different hybrid approach is Anderson's (1990) rational model of human retrieval, which combines features of both information retrieval with the constraints of human memory. It serves as a model of human retrieval, although it is built using features of a probabilistic model of library borrowing. Incorporating history of usage and context factors, his model uses some of the environmental features that have typically not been considered in other psychological models. This differs from the other hybrid models, which applied psychological principles to design an information retrieval model because it applies some of the principles from information retrieval to help constrain the design of a psychological model of human retrieval.

Conclusions

Models and representation of information

The retrieval models described in this article use a limited number of methods for creating the connections between cues and memory items, matching the information and representing the memory items. These features are outlined for some of the models in Table 1.

For memory items to be retrievable, there must be connections from the items to the all the cues that could be used to retrieve it. As seen in Table 1, the origin of these connections for all of the models derive from either similarity of the features in the cues and items or based on the co-occurrence of cues and items during encoding. In this manner, the connections in all models are based on temporal or featural co-occurence. Once these connections are created, the models use matching methods to retrieve items given cues. In some models, the matching methods use the connections directly as the retrieval method (e.g. compound cue model). However in most models, the retrieval is also based on what other information is encoded in memory. In spreading activation, retrieval is based both on the cue connections and the inter-item connections. Similarly in distributed models with a composite associative trace, retrieval is based not only on feature similarity between a cue and information items, but also on relationships between items added to the composite associative trace.

Based on the models discussed, there are two ways of representing the memory information, as a network, or as a set of features. In a network representation, the information is represented as discrete items. These items are connected to each other through representing activations between the items. In the feature representation, information is represented as a set of features rather than discrete elements. Rather than representing item A as strongly connected to B and weakly connected to C, it could be represented as sharing many features of B and few features of C. In this manner, connections between information in network representations are typically made through examining feature overlap. Therefore a network representation can be derived from a feature representation, as is typically done for spreading activation models.

Strategic retrieval

Humans are experts at using strategies to store and retrieve information from their own memory. They are familiar with the structure of the information stored and the retrieval cues that can be used since they did the initial encoding. This is not the case in computer retrieval. Users are seldom familiar with what information is available and how it is organized. This unfamiliarity hinders their ability to develop good retrieval cues to give to the system. Users are also not familiar with the ways of specifying the cues. Since most retrieval systems are term based, the exact terms must be specified to get the desired information.

Information retrieval systems are currently adding features to aid strategic retrieval. Iterative retrieval allows users to refine their query interactively as they become more familiar with the structure and semantics of the stored information. Relevance feedback similarly lets users narrow the context to where they want to search by providing additional context cues for the search. Browsing also permits users to navigate through a set of information, typically staying within the same context. Thus, a user is in control of the search process, making the decisions on which direction to go to find a relevant piece of information.

Implications for the two fields

The similarities between human and computer retrieval indicate the complementary nature of the two systems. Information retrieval systems function as extensions of our own memory. For this reason, there should be some similarities in how information retrieval systems are used and designed. One such feature is to allow users to employ familiar retrieval strategies when using the system. As described above, some systems have implemented these strategic aids. Nevertheless, computer retrieval systems do not make the large number of associations between information items as do humans. This is due to the fact that systems have minimal encoding skills and must rely on term based encoding as opposed to the human's semantic encoding. The systems also do not use much context in retrieval. Typically, retrieval is performed using just a set of words provided by a user, whereas human retrieval uses many additional contextual cues such as temporal and locational information. Thus, improvement in information retrieval models can be made through tailoring the systems to incorporate semantic relationships in encoding and to use contextual information for retrieval such as user profiles. Psychological models of memory and of retrieval strategies highlight the current abilities of the human retrieval system and can provide directions for information retrieval systems to augment the human's ability to find information.

Conversely, human memory retrieval can learn from the insights into computer information retrieval. Information retrieval researchers have developed efficent systems, taking into account such features as the frequency of occurrence of information, associative and feature based representations, and human retrieval strategies. These models have become fairly similar to those used in psychological memory modeling. This suggests that both fields are searching for analogous solutions. Developments of efficient information retrieval models can provide guidance for psychologists as to types of representations and algorithms to be used to model efficient memory retrieval. The improvements in retrieval when using relevance feedback suggests that the automatic process of retrieval need not be entirely efficient for human memory. With strategic intervention of relevance feedback, a query can be quckly narrowed down to the relevant information. This may be the case too in human memory. The actual automatic retrieval process need not be entirely efficient, since we have strategic skills for developing the context and revising the cues to narrow down on the desired information. There are few memory models that incorporate both the automatic retrieval component and the wealth of strategies used to develop and revise the cues.

In information retrieval, one of the key methods to developing a retrieval system is based on knowledge of task and environmental factors. As suggested by Anderson's (1990) work and from the ecological approaches to human memory, human memory models in the future may examine more of what types of tasks human memory is designed to do, and what are some of the environmental constraints in which the memory must operate.

The insights from the two fields can provide guidelines to aid development of better retrieval models. These models will more likely fall under the category of hybrids incorporating the best features and constraints of the two types of retrieval. In the long run, these insights can provide both a better understanding of how our own retrieval system works and how to develop better external retrieval systems.

References

Anderson, J. R. 1983. *The architecture of cognition.* Cambridge, MA: Harvard University Press.

Anderson, J. R. 1990. *The adaptive character of thought.* Hillsdale, NJ: Erlbaum.

Bookstein, A., & Swanson, D. R. 1975. Probabilistic models for automatic indexing. *Journal of the American Society for Information Science, 26,* 1 45-50.

Chase, W. G., & Ericsson, K. A. 1981. Skilled memory. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition.* Hillsdale, NJ: Erlbaum.

Foltz, P. W. 1991. Models of human memory and computer information retrieval: Similar approaches to similar problems. Tech Report 91-03, Institute of Cognitive Science, Univ. Of Colorado, Boulder. (Submitted for publication)

Hintzman, D. 1986. "Schema abstraction" in a multiple-trace memory model. *Psychological Review, 93,* 411-428.

Jones, W. P. 1986. On the applied use of computer models of memory: The Memory Extender personal filing system. *Intl. Journal of Man-Machine Studies,* 25, 191-228.

Landauer, T. K. 1986. How much do people remember? Some estimates of the quantity of learned information in long-term memory. *Cognitive Science, 10,* 477-493.

Murdock, B. B. 1983. A distributed memory model for serial-order information. *Psychological Review, 90,* 316-338.

Raaijmakers, J. G. W., & Shiffrin, R. M. 1981. Search of associative memory. *Psychological Review, 88,* 93-134.

Ratcliff, R., & McKoon, G. 1988. A retrieval theory of priming memory. *Psychological Review. 95,* 385-408.

Rose, D. E., & Belew, R. K. 1989. Legal information retrieval: A hybrid approach. *Proceedings of the International Conference on AI & Law,* Vancouver, BC.

Salton, G., & Buckley, C. 1990. Improving retrieval performance by relevance feedback. *Journal of the American Society for Information Science. 41,* 4, 288-297.

Salton, G., & McGill, M. J. 1983. *Introduction to Modern Information Retrieval.* McGraw-Hill: NY.

Tulving, E., & Thompson, D. M. 1973. Encoding specificity and retrieval processes in episodic memory, *Psychological Review, 80,* 352-373.

Williams, M. D. 1978. The process of retrieval from very long term memory. Tech report 7801. Center for Human Information Processing, Univ. California, San Diego.

Williams, M. D. 1984. What makes RABBIT run? *Intl. Journal of Man-Machine Studies, 21,* 333-352.

Retrieval Model	Example	Domain	Matching Method	Origin of connections	Representation
Compound Cue	SAM Gillund & Shiffrin (1984)	Human	Activation of items associated with cues	Cues to item strength based on STM co-occurrence	Network
Distributed	TODAM Murdock (1982)	Human	Correlation of item feature vectors to cue vector	Feature similarity	Feature
Spreading Activation	ACT* Anderson (1983)	Human	Activation of items spread from cues and other items	Based on degree of association between items	Network
Rational Memory model	Anderson (1990)	Human	Based on probabilities of history and context	Based on probabilities of history and context	Network
Inverted index	Salton & McGill (1983)	IR	Boolean operations on feature vectors	Feature similarity (term overlap)	Feature
Vector Space	SMART Salton & McGill (1983)	IR	Distance between item and cue vectors	Feature similarity (term overlap)	Feature
Probabilistic	Bookstein & Swanson (1975)	IR	Prob. a cue is in a relevant item vs. prob. in a non relevant item	Feature similarity (term overlap)	Feature
Spreading Activation	Memory Extender Jones (1986)	IR	Activation of cues spreads to cues and other items	Semantic network based on feature (term) similarity	Network/ Feature
Connectionist	Rose & Belew (1989)	IR	Activation of cues spreads to other cues and items	Feature (term) similarity and user feedback	Network/ Feature

Table 1. Representations and methods used for human and computer retrieval models.