Using Controlled Knowledge Search to Retrieve Cross-Contextual Information

Michael Wolverton
Barbara Hayes-Roth
Knowledge Systems Laboratory
Stanford University
701 Welch Rd, Bldg. C
Palo Alto, CA 94304

Abstract

Indexing cases on features has been shown to be a successful method of retrieving experiential knowledge for many single-task systems. However, discovering a general set of indices for retrieving cross-contextual information in a multi-use knowledge base is impractical. We present a method for retrieving cross-contextual information based on *knowledge-directed spreading activation* in a semantic network. This method uses task-specific knowledge to guide a spreading activation search to a case or concept in memory which meets a desired similarity condition. We describe a specific instantiation of this method that retrieves analogies for the task of creative design, and also how the behavior of this method will differ from indexing schemes under varying knowledge base conditions.

1 Introduction

Researchers in artificial intelligence (AI) have long searched for ways to improve existing techniques for information retrieval (IR). Statistical approaches to the analysis of text and the retrieval of documents have so far been the most successful methods applied to IR. However, their absolute performance in retrieving (only) relevant documents to a user's query leaves a great deal of room for improvement. One possible explanation for the seemingly low rate of recall with these techniques is the limited amount of information about a document that can be represented merely by counting words in the text, and also the limited amount of information a user can specify in his query using only lexical information. For this reason, IR and AI researchers have been trying to find more robust ways of representing information contained in documents and also ways of finding relevant items using these new representations.

One retrieval technique which has been studied in this context is *spreading activation* within a semantic network. In this approach, all information in the system is represented in a global semantic network (a large labeled graph): nodes in the network represent primitive concepts, and links between nodes represent relations between those concepts. More complex concepts – including, e.g., summary information of documents known to the system – are represented as subgraphs within the larger semantic net. Retrieval is a process in which activation is spread outward from the nodes in some retrieval cue subgraph (representing either a user request or some target concept to be matched) to neighboring nodes, and repeating this process until some other subgraph in the semantic network accumulates enough activation to be considered retrieved. Both Rau [Rau, 1987] and Cohen and Kjeldsen [Cohen and Kjeldsen, 1987] have applied spreading activation to general IR problems. This general class of approaches has the following advantages over traditional word-based IR approaches:

(1) Information can be stored without explicit consideration of how it might be used in the future. I.e., information need not be stored with any special-purpose indices (except the links to the rest of the knowledge base). And the way the information is represented makes it amenable to a number of other possible AI reasoning techniques to which simple text is not amenable.

(2) The user can be given a much more descriptive language in which to express his query. In Rau's system SCISOR, for example, the user makes his request for information using a restricted-vocabulary English.

(3) Spreading activation can be triggered by any concept activation in the semantic net, so information retrieval can be a by-product of many system activities other than explicit user requests.

Spreading activation has been applied to IR problems with success in small narrowly-defined domains. However, since spreading activation is essentially a blind breadth-first search mechanism, it has increasingly greater difficulty retrieving information the larger and more unconstrained the knowledge base becomes.

This is especially true when one considers the problem of retrieving subgraphs which are *semantically distant* from the retrieval cue – those for which most corresponding nodes between the cue and the retrieved graph have a long path of nodes and links between them. These types of retrieval problems arise especially when the information required is to be utilized for some sort of *creative* problem solving. For example, researchers will often analyze citation indices, going through chains of citations and ending up with diverse documents from a variety of domains which they see as relevant to their research. Another example of this type of retrieval, and one which seems to occur fairly frequently in some types of human reasoning, is the process of retrieving "far-flung" analogies – cross-domain connections between two concepts not normally thought of as connected. This type of analogy is important to a number of creative tasks in humans, including creative design [Koestler, 1965, Hughes, 1971, Kock, 1978].

We have developed an extension of the spreading activation models used by Rau and Cohen and Kjeldsen called *knowledge-directed spreading activation* (KDSA), for retrieving far-flung analogous concepts from a large diverse knowledge base. This method is based on controlled search in a general semantic network, rather than indexing in a small or single-domain case library. It uses task-specific knowledge to guide a series of spreading activation searches from the target concept to a semantically distant base concept. This knowledge is applied in the evaluation of intermediate concepts retrieved by a standard spread of activation, and by the modification of weights controlling the spread of activation based on those evaluations. The method is implemented in a system which retrieves cross-domain analogies for use in a creative design process.

Section 2 describes this method in detail. Section 3 steps through an example of a retrieval of a useful analogy for creative design using this method. Section 4 presents two hypotheses about KDSA which we plan to test using our implementation. And section 5 discusses the connections between KDSA and other research.

2 Knowledge-directed Spreading Activation

2.1 General Method

The general execution of KDSA is shown in Figure 1. This method is designed to be a component of a larger intelligent agent architecture[1]; here the search for new analogies goes on in parallel with, and interacts with, the intelligent agent's other activities (problem solving, learning, handling sensory input, etc.).

All world knowledge is represented in a single semantic network. Within that semantic network, small subgraphs of nodes and links which represent aggregate concepts are explicitly grouped together as *conceptual graphs*[2]. Individual conceptual graphs are treated the same as primitive nodes – i.e., they can be associated to other nodes via links, and they can themselves be parts of larger conceptual graphs. In the discussion below, conceptual graphs will be referred to merely as "concepts".

The basic search of memory is conducted by a spreading activation mechanism. Spreading activation is a general model of concept association in semantic networks that has developed in psychology and AI over the years; our model is based loosely on the version developed by John Anderson in his ACT* cognitive architecture [Anderson, 1983]. The spreading activation process starts when a group of nodes are assigned an initial level of activation (a numeric value). The mechanism then goes through a series of cycles. At each cycle, each activated node transmits some portion of its activation to each of its neighbor nodes. The amount of activation transmitted is dependent on the amount of activation associated with the original node and a numeric strength attached to the link connecting the two nodes. At each cycle, each activated node also loses some of its activation to a decay process. Whenever a node or group of nodes accumulates enough activation to pass a retrieval condition (e.g., the activation becomes greater than a global threshold), it is considered retrieved and is passed on to the encompassing architecture for further processing.

Note that the basic spreading activation model we have described here is an exhaustive knowledge search mechanism. Some method of controlling the search is necessary for the system to retrieve the types of far-flung base concepts described in section 1. Anderson and others (while not concentrating on retrieving far-flung analogies) have used priming methods, where the strengths on links are increased each time they are used, to cause the mechanism to prefer some paths over others. KDSA, by contrast, uses feedback from the analogues retrieved so far to focus the search.

The agent architecture encompassing KDSA begins the retrieval process when some executing task requests an analogy and designates a target concept. This initial request causes some nodes in the semantic network – those representing the target concept plus possibly others representing desired features of the solution, etc. – to be assigned activation, and this assignment begins the spread of activation in memory. When a concept is retrieved by the spread of activation, the *matching component* computes a mapping between it and the target concept, and then forms an evaluation of the mapping

[1] For our particular implementation, we are using a new experimental version of the BB1 agent architecture [Hayes-Roth, 1990].

[2] Our present implementation and discussion uses on this conceptual graph representation of concepts because of its representational power. However, the general model presented here is equally applicable to a basic semantic network representation without the conceptual graphs.

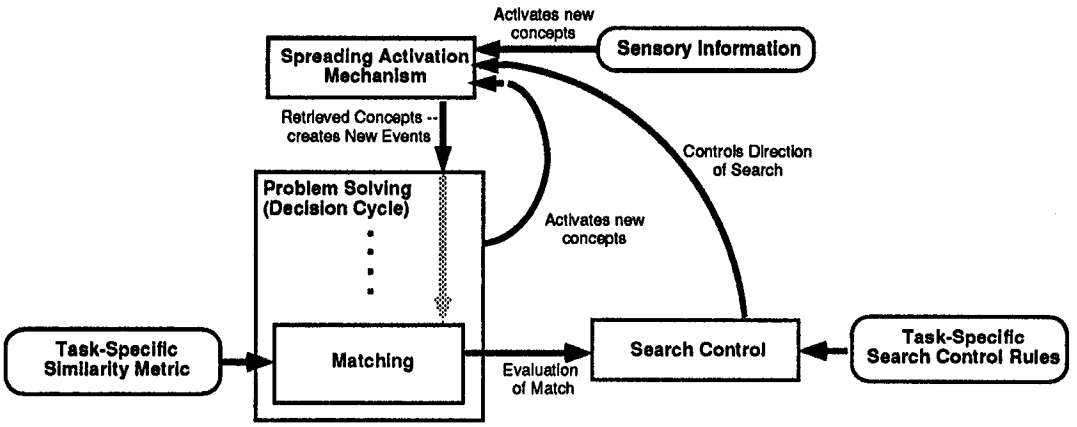

Figure 1: Knowledge-Directed Spreading Activation.

based on the type of analogy demanded by the task (the task of invention, for example, demands a far-flung analogy with a strong match between the functional descriptions of the two concepts). This evaluation is passed on to the *search control component*, which uses its task-specific heuristics to focus the spreading activation search in directions that are more likely to lead to highly-evaluated analogies for the current task. At any time during this cycle, other concepts may be activated by the agent's other activities, such as ordinary problem solving or processing sensory input. In this way KDSA simulates the individual possibly "stumbling across a solution". The process ends when the matching component finds a retrieved concept whose evaluation is high enough, or when all the activation has decayed with no successful analogy found.

The important components of the retrieval system are discussed in more detail below.

Matching Component Each time a concept is retrieved by the spreading activation search as a potential base concept, it is passed to the matching component. The matching component first forms the best possible partial mapping between the potential base and the target, and then it evaluates that partial mapping using its heuristics, specifically the subset of its heuristics which are active for the task requesting the analogy. These heuristics will base their evaluation primarily on the following features of the partial mapping: (1) how many nodes and links match between the target and potential base, i.e., the degree of isomorphism between the two graphs, (2) what *kinds* of nodes and links match, and (3) the level of abstraction of the mapping, i.e., the minimum path distance in the type hierarchy between corresponding nodes of the mapping. The evaluation consists of a numeric rating of the mapping, and a description of the shortcoming(s) of the mapping assigned by the heuristics. If the numeric rating is greater than a threshold value, the potential base is accepted as the final analogy, and the KDSA process halts. Otherwise, the evaluation is passed on to the search control component.

For example, if the task requesting the analogy demands a far-flung analogy, one good rule of thumb is that the structure of the target concept should not match the structure of the base concept (see section 2.2). The matching heuristic DISCOURAGE-STRUCTURE-MATCH does this by examining the degree and level of abstraction of the match between the structural components of the target and the base (we make the simplifying assumption that the knowledge base designer has specified the portion of concept representations which identify structure), and downgrading the rating of the mapping if the structural match is good. If the rating is downgraded, the heuristic will also attach a description of the shortcoming (e.g., :STRUCTURE HIGH) to the evaluation; this description may then be used by the search control component to decide to move the search away from structurally similar concepts.

Search Control The search control component uses evaluations from the mapping component and other information about the state of the search to influence the direction of the spread of activation. It uses heuristics to control the direction of the search in three ways: (1) it weakens the activation of concepts with poorly-evaluated matches with the target, (2) it strengthens the activation of concepts with highly-evaluated matches with the target, and (3) it modifies the strengths of links in order to search in more promising areas of the KB. The heuristics can modify link strengths either by link type (e.g., supertype links), or by the nodes the links are adjacent to (e.g., all links pointing to action nodes).

For an example of a heuristic that modifies link strengths, consider the rule of thumb given above: for far-flung analogies, it is a good idea to avoid structurally similar concepts. Thus, links which point to structural nodes – those nodes which are subtypes of the node STRUCTURE – should not spread much activation relative to other links. The search control heuristic

INHIBIT-STRUCTURE-NODES does this by weakening the activation-passing strength of all such links in the KB. This heuristic will normally be active at the early stages of the search for a far-flung analogy, though it may be overridden at later stages when the mechanism is searching in an area of the KB already structurally distant from the target concept.

The method of strengthening the activation of promising concepts allows us to use these concepts as "beacons" along the way to the final good analogy. This is very similar to the way that promising intermediate states are used in heuristic search techniques such as hill-climbing or best-first search [Pearl and Korf, 1987].

The use of the matching component of the mechanism to provide feedback to the spreading activation search provides a key distinguishing feature of our approach. Most previous approaches to analogy serialize the retrieval and mapping processes: first they retrieve a concept, then they try to map it, then if mapping fails they start at ground zero with retrieval again. By contrast, mapping in KDSA is an integral *part* of retrieval: mapping (the matching component) provides ongoing information to the retrieval mechanism (spreading activation and search control) throughout the duration of the retrieval process.

2.2 Specific Instantiation for Invention

The previous section described an architecture which uses evaluations of "promising" near-analogies to help guide a search through a large knowledge base. But how should the matching component determine whether a concept is a "promising" or "good" analogy for creative design? This section proposes evaluation criteria to be used by KDSA.

In general problem solving by analogy, we believe a given analogy will be likely to lead to a useful creative result if it possesses two characteristics:

(1) a high degree of match in features that are relevant to the goal at hand, and

(2) a low degree of match in all features irrelevant to the goal.

Point (1) above suggests that a solution to the goal based on the given analogy will be useful. Basing a solution to a problem on an analogous concept which served as a successful solution to a similar problem in the past makes it likely that the solution to the present problem will be analogously successful. For this reason, Carbonell [Carbonell, 1983] and others have based their retrieval of analogies in problem solving totally on the present goal. And Point (2) above suggests that the solution to the goal produced by the analogy will be novel. A base concept which is markedly different from the target concept is likely to be an analogy that others haven't thought of before (since analogy retrieving in humans seems to be based mostly on semantic similarity [Thagard *et al.*, 1990]), and therefore likely to lead to a solution that is markedly different from other solutions tried in the past.

For the particular problem of creative design, the goal of the problem is usually expressed as a functional specification, so functional properties of potential base concepts are the features we want to match in the analogy. In particular, we believe an analogous concept will be likely to be useful in a creative design process if:

(1) it has a very similar function to the function specified in the design goal (at least on an abstract level), and

(2) it has very different structural and behavioral characteristics from the existing device which is being improved, or structural and behavioral characteristics which are novel to the current design domain.

In the above points, "similarity" and "difference" are measured with a combination of isomorphism (a complete one-to-one mapping between the features involved), and semantic distance between corresponding features.

3 Example

This section contains an example demonstrating the working of KDSA. It shows one possible way that the retrieval process could proceed for an analogy between an irrigation system and the circulatory system. Here the agent architecture is given a goal to invent a new type of irrigation system – specifically, one which is less wasteful of water than existing sprinkler irrigation systems. Given this goal, KDSA retrieves CIRCULATORY-SYSTEM as an analogue to use to guide the design of the new system. This analogy could potentially lead an inventor to conceive of a design like drip irrigation – i.e., an irrigation system which, like the circulatory system, takes its delivery substance directly to each individual destination.

Figure 2 shows simple representations for four devices which may be found in a multi-domain KB. Each concept contains simple graphs representing the device's function, behavior, and structure. The representations used here are clearly too simplistic to appear in any real knowledge base; we are using them only for demonstration purposes.

After the nodes representing SPRINKLER-IRRIGATION-SYSTEM are activated and designated as the target, the retrieval process for the irrigation redesign example proceeds as follows:

(1) Activation spreads from these nodes to the neighboring nodes comprising several other concepts. The concepts PLUMBING-SYSTEM and FOUNTAIN accumulate enough aggregate activation to be passed to the matching component.

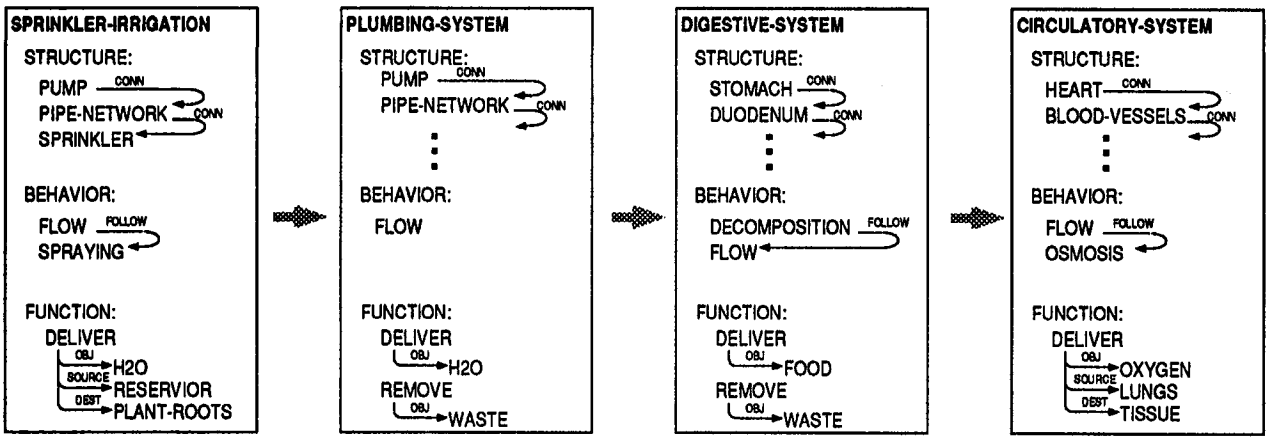

Figure 2: Example search sequence of KDSA.

(2) The matching component rejects FOUNTAIN as not promising – the function (beautification) does not match the function of SPRINKLER-IRRIGATION-SYSTEM's (delivery), and the match on structural and behavioral nodes is very high. The search control component uses this match evaluation to zero the activation of FOUNTAIN, thus removing that area of the search.

(3) PLUMBING-SYSTEM is rated as a promising near-miss. It performs one function which is the same (DELIVERY) as SPRINKLER-IRRIGATION-SYSTEM's, and its behavior is quite different. It is classified as a near-miss rather than a final analogy because it shares too many structural aspects with the target. This evaluation is passed on to the matching component, where the heuristic INHIBIT-STRUCTURE-NODES weakens the link strength on links leading to structural components. This causes PLUMBING-SYSTEM's behavioral and functional descriptions to play a larger role in the next phase of spreading activation.

(4) Spreading activation retrieves DIGESTIVE-SYSTEM next. This concept is rated by the intelligent matcher as a promising near-miss – it performs much the same function (WASTE-REMOVAL is a type of delivery), and is not structurally similar. It is rejected because of the inexactness of the functional match. Since the match evaluation indicates that DIGESTIVE-SYSTEM is structurally dissimilar from the target, the search control heuristics allow the search to explore more concepts which are structurally similar to the digestive system.

(5) Spreading activation searches more in the biological system domain, and activates CIRCULATORY-SYSTEM. It is accepted by the matching component as a good analogy for invention, because (1) the functions of CIRCULATORY-SYSTEM and SPRINKLER-IRRIGATION-SYSTEM have a strong, albeit abstract, mapping, and (2) the behavioral and structural components of the two concepts do not map very well.

4 Discussion

The major difference between KDSA and indexing retrieval systems is that KDSA is well-suited to retrieving far-flung analogies that meet the requirements for a large range of tasks, while most indexing structures are optimized for a single task in a homogeneous case library. But in addition to this difference, we also expect qualitative differences in the computational behavior between KDSA and most indexing schemes under varying knowledge base conditions. Two of these expected differences are listed here as *distinguishing hypotheses* about the computational behavior of KDSA which differentiate it from other approaches:

Hypothesis 1 *As the size of the knowledge base grows (in a random fashion), the time it takes KDSA to retrieve an acceptable far-flung analogy should decrease.*

KDSA uses other concepts in the knowledge base to guide it to a final analogy, so having more concepts in the KB increases the number of potential "beacons" leading to a final analogy. The concepts added to the KB which are not relevant to a given target concept should be either (1) never reached by the activation search, or (2) quickly deactivated by the search control component. Thus KDSA should spend less time exploring blind alleys in its search for analogies, and should be able to use the increased number of beacons to focus on promising areas of the semantic network more quickly. The time taken to retrieve a case by indexing, by contrast, should at best remain constant as new cases are added to the case library.

Hypothesis 2 *As the size of the knowledge base grows, the quality of analogies retrieved by KDSA should increase, both in terms of their likelihood to yield useful results and their likelihood to yield creative results.*

Indexing systems will see an increase in solution quality as the case base grows [Golding and Rosenbloom, 1991] because of the increased probability that a case closely matching the target will be found in the library. KDSA's solution quality should see this same source of improvement: the more interdisciplinary knowledge a KB has, the more likely that a creative useful analogy exists for a given problem. However, KDSA should also see an additional source of improvement: the increased number of potential beacon concepts will make it possible for the spreading activation search to reach creative useful analogies which were previously present in the KB but unretrievable.

These hypotheses seem to agree with previous studies of inventors: the more interdisciplinary knowledge an inventor has, the easier and more quickly he can brainstorm to find analogies, and the more insightful those analogies are likely to be. We have implemented KDSA, and are currently implementing a diverse medium-scale knowledge base to be used as a testbed to examine these hypotheses and other aspects of the mechanism's performance under various KB conditions.

5 Related Work

There is a large body of AI literature on information retrieval in semantic networks. Among the more recent work, SCISOR [Rau, 1987] and GRANT [Cohen and Kjeldsen, 1987] both use heuristic information to direct a spread of activation in semantic networks. KDSA's search control component is similar to the relatedness condition which controls the spread of markers in SCISOR and the path endorsements which direct spreading activation in GRANT. KDSA differs from these systems, however, in that it uses information from previous match evaluations to dynamically adjust the direction of the spread of activation. KDSA in effect runs a *series* of SCISOR-like or GRANT-like searches, starting each sub-search from the near-misses it has encountered in previous sub-searches, and using the evaluations of those near-misses to formulate its search control for the next sub-search.

There is other research in the areas of design and creativity which is relevant to KDSA. Researchers in case-based design (e.g., [Goel and Chandrasekaran, 1989]) have recognized the importance of function in retrieving useful design histories. CADET [Sycara and Navinchandra, 1991] represents a different approach to cross-contextual analogy in design: it retrieves cases of devices that are seemingly functionally dissimilar from the target device by matching only subparts of larger cases, and it uses transformations of the design goal to make the cross-contextual leap. MINSTREL [Turner, 1992] also uses a form of index transformation in its case-based approach to creativity. As with KDSA, MINSTREL's heuristics are task-specific (its task is story-telling); unlike KDSA, though, MINSTREL uses a standard indexing approach to retrieval.

6 Conclusion

We have presented knowledge-directed spreading activation, a task-independent method for retrieving analogous concepts in a multi-domain knowledge base. KDSA overcomes the shortcomings that indexing methods may have in large knowledge bases by applying task-specific knowledge to a general semantic network search technique. And it extends previous models of semantic network retrieval by using evaluations of failed analogies encountered in the earlier stages of the search to influence the direction of the search in later stages. Future directions of research on KDSA include validating the hypotheses that performance will improve as interdisciplinary knowledge in the KB grows, and demonstrating the task-independence of the method by identifying heuristics for other tasks and degrees of creativity.

References

[Anderson, 1983] J. R. Anderson. *The Architecture of Cognition.* Harvard University Press, 1983.

[Carbonell, 1983] J. G. Carbonell. Learning by analogy: Formulating and generalizing plans from past experience. In R. S. Michalski, J. G. Carbonell, and T. M. Mitchell, editors, *Machine Learning: An Artificial Intelligence Approach*, pages 137–162. Tioga Press, Palo Alto, 1983.

[Cohen and Kjeldsen, 1987] P. R. Cohen and R. Kjeldsen. Information retrieval by constrained spreading activation in semantic networks. *Information Processing and Management*, 23(4):255–268, 1987.

[Goel and Chandrasekaran, 1989] A. Goel and B. Chandrasekaran. Functional representation of designs and redesign problem solving. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI-89)*, Detroit, Michigan, August 20-25 1989.

[Golding and Rosenbloom, 1991] A. R. Golding and P. S. Rosenbloom. Improving rule-based systems through case-based reasoning. In *Proceedings of the Ninth National Conference on Artificial Intelligence (AAAI-91)*, pages 22–27, July 1991.

[Hayes-Roth, 1990] B. Hayes-Roth. Architectural foundations for real-time performance in intelligent agents. *The Journal of Real-Time Systems*, 2:99–125, 1990.

[Hughes, 1971] T. P. Hughes. How did the heroic inventors do it? *American Heritage of Invention and Technology*, 1(2):22–23, Fall 1971.

[Kock, 1978] W. Kock. *The Creative Engineer.* Plenum Press, 1978.

[Koestler, 1965] A. Koestler. *The Act of Creation.* Macmillan, 1965.

[Pearl and Korf, 1987] J. Pearl and R. E. Korf. Search techniques. *Annual Review of Computer Science*, 2:451–467, 1987.

[Rau, 1987] L. F. Rau. Knowledge organization and access in a conceptual information system. *Information Processing and Management*, 23(4):269–283, 1987.

[Sycara and Navinchandra, 1991] K. P. Sycara and D. Navinchandra. Index transformation techniques for facilitating creative use of multiple cases. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, 1991.

[Thagard *et al.*, 1990] P. Thagard, K. J. Holyoak, G. Nelson, and D. Gochfeld. Analog retrieval by constraint satisfaction. *Artificial Intelligence*, 46:259–310, 1990.

[Turner, 1992] S. Turner. MINSTREL: A model of storytelling and creativity. Technical Report UCLA-AI-N-91, Artificial Intelligence Laboratory, Computer Science Department, University of California, Los Angeles, 1992. forthcoming.