To appear as:

Schvaneveldt, R. W, & Cohen, T. A. (in press). Abductive reasoning and similarity: Some computational tools. In D. Ifenthaler, P. Pirnay-Dummer, & N. M. Seel (Eds.), *Computer based diagnostics and systematic analysis of knowledge.* New York: Springer.

Chapter 11

ABDUCTIVE REASONING AND SIMILARITY

Some Computational Tools

Roger W. Schvaneveldt & Trevor A. Cohen

Arizona State University *University of Texas at Houston*

Abstract

Abductive reasoning includes discovering new hypotheses or explanations. This chapter identifies several factors involved in abductive reasoning in an effort to move toward a theory of such reasoning. The chapter has an ultimate focus on the nature and influence of similarity. A major goal of our work is to develop computational tools that provide intelligent abductive suggestions to people engaged in discovering new knowledge. Novel abductive inferences often exhibit interesting similarities to the phenomena under investigation. Such similarities are not strong or direct but rather are often only obvious once the inference has been drawn. Some of our research is directed at discovering indirect similarities from text by using measures that are sensitive to indirect relations between terms in the text. By focusing on terms that are related but do not co-occur, potentially interesting indirect relations can be explored. Our work employs Random Indexing methods and Pathfinder networks to identify and display relations among terms in a text corpus. These displays are provided to individuals to suggest possible abductive inferences. We explore a variety of methods for identifying indirect similarities. The degree to which surprising and interesting inferences are suggested is the primary measure of success. Several examples are presented to illustrate the method. An analysis showing a positive relationship between (a) the strength of indirect similarity in one period of time and (b) the likelihood that the terms involved become directly related in future time. This correlation supports the hypothesis that discoveries may be latent in such indirect similarities. Presumably, noticing such similarity brings indirectly related concepts together suggesting a new idea.

Keywords: abductive reasoning; discovery; hypothesis generation; problem solving; indirect similarity; computational tools

1. INTRODUCTION

This chapter outlines a psychological theory of certain aspects of creative thinking, specifically abductive reasoning, a term coined by the philosopher and logician, C. S. Peirce (1839-1914). Peirce held that the hypothetico-deductive method in science required a logic underlying the generation of hypotheses in addition to the inductive and deductive logic involved in testing hypotheses. Given some observations that are surprising or unexpected, abductive reasoning is concerned with generating hypotheses about the observations or with reasoning to the best explanation. Problem solving, in general, can often be seen to fit the abductive reasoning framework. The problem motivates a search for a solution, and abductive reasoning produces potential solutions to the problem. Peirce suggested that people have an impressive ability to formulate promising hypotheses on the basis of only a few observations.

Issues concerning novelty, evaluation, optimality, consilience, aesthetics, and pragmatics among others arise in the study of abductive reasoning. While these issues will be briefly addressed in the paper, the primary focus is on the involvement of similarity relations in generating potential abductive inferences. In other words, the focus is on one possible explanation of how new ideas arise. We propose methods for identifying potential new connections among ideas and for displaying connections using Pathfinder networks to assist experts in searching for such promising connections. While reasoning by analogy is a form of abductive reasoning, not all abductive inferences are analogies. We return to this point later.

Similarity-based abduction is proposed as a theory for generating ideas as hypotheses or problem solutions. Abductive reasoning begins by activating a goal state characterized by a problem to be solved with no immediate solution found. Essentially, no available solution means that none are directly associated with the problem. However, a process of spreading activation would lead to the activation of other ideas related to the problem. Over time, continuing to think about the problem or engaging in still other activities would lead to the activation of other ideas together with patterns of connections among the ideas. Interconnections among the activated ideas could lead to an enhancement of the connections of ideas to the elements of the problem in two ways. First, activation among the connections could simply increase the activity in existing weak links between the problem and other ideas. Second, indirect connections of between newly activated ideas and the problem could be detected by means of similar patterns of connections. Such newly activated ideas might be indirectly or implicitly related to the problem. These new promoted weak connections and newly identified indirect connections provide links to potential solutions to the problem. They constitute potential hypotheses.

Developing models of similarity-based abduction involves developing methods of generating activation of ideas on the basis of activation of existing connections among ideas. Examples of such methods can be found in GeneRanker (Gonzalez, et al., 2007), Hyperspace Analog of Language or HAL (Burgess, Livesay, & Lund, 1988), Latent Semantic Analysis or LSA (Landauer & Dumais, 1997), and Random Indexing (Kanerva, Kristofersson, & Holst, 2000). Cohen (2008) has shown how identifying new connections can lead to novel hypotheses concerning potential treatments for medical conditions. Also, developing tools to assist users in identifying fruitful new ideas pertinent to hypothesis discovery and problem solving requires generating possible ideas, ranking the ideas, and providing informative displays of connections for users to examine and evaluate for their potential utility. Examples of models and tools are also presented in the paper.

1.1 Abductive Reasoning

C. S. Peirce wrote extensively about logic and scientific method. Several important pieces were published in 1940 under the editorship of Justus Buchler (Peirce, 1940a, 1940b). Peirce proposed that there were three essential types of reasoning including the familiar deductive and inductive reasoning. The testing, confirming, and rejecting of hypotheses is covered by deduction and induction. In contrast with many logicians, Peirce also thought there was a logic underlying the origin of new hypotheses. He called this logic variously “abduction”, “retroduction”, and “hypothesis” in his writings over the years. The kind of reasoning he envisions proceeds something like the following:

I make some observations (O) that are surprising, unusual, or puzzling in some way. It occurs to me that if a particular hypothesis (H) were true, then O would follow as a matter of course. In other words, H implies O so we could say that H explains O. Thus, H is plausible and should be considered further. Abductive reasoning is illustrated by Figure 1.

Consider Figure 1 to be a set of observations (O). Now ask, “What is this?” or “How could these observations be explained?” Now we are seeking hypotheses (H) that would explain the diagram (O). We might come up with such conjectures as:

- H1: “It’s olives on toothpicks.”
- H2: “It’s barbeque spits with tomatoes.”
- H3: “It’s two pair of spectacles.”
- Etc.

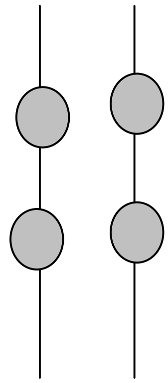

Figure 1. What is this?

Notice that each of these conjectures (Hi) has the property that the O would follow if H were true. This is the abductive form of logic. Cast in the form of a syllogism, abductive logic would appear as in Table 1. In the example above, the arrangement of the lines and circles constitute the observations (O). The various suggestions are potential hypotheses.

Table 1.
Abductive Inference

Major Premise	O
Minor Premise	If H then O
Conclusion	H is plausible

Obviously this is not a deductive argument which requires that the conclusion necessarily follows from the premises. In abductive inference, H does not follow with certainty so the conclusion about H only reaches plausibility. The observations could have resulted from H so H is a reasonable conjecture about why O is as it is. As such, H deserves further consideration as a possible explanation of O. Abductive reasoning bears a strong similarity to inductive inference which is illustrated in Table 2.

Table 2.
Inductive Inference

Major Premise	If H then O
Minor Premise	O
Conclusion	H is confirmed

Induction, too, does not carry certainty. In the deductive realm, H does not necessarily follow from the premises given so at best we can say that the observation confirms or supports the hypothesis. Finding confirming evidence for a hypothesis simply allows us to continue entertaining it, perhaps with increased confidence, but confirming evidence does not *prove* a hypothesis. The difference between abduction and induction is due to the temporal relations of the premises. The major premise precedes the minor premise in time so the hypothesis occurs as an explanation in abduction while the observations occur as a test of the hypothesis in induction. Tests of hypotheses do not always lead to confirmation, however, which leads to the third type of inference, deduction (modus tollens) as in Table 3.

Table 3.
Deductive Inference

Major Premise	If H then O
Minor Premise	O is false (not O)
Conclusion	H is disproved (not H)

Finding that the predictions of a hypothesis fail to hold leads to the certain conclusion that the hypothesis is false. This asymmetry between induction and deduction was the basis of Popper's (1962) philosophy of science. Because disproving hypotheses is more conclusive than confirming them, Popper thought that scientists should make great efforts to disprove their favorite hypotheses rather than seeking more and more confirmatory evidence. This comparison of abduction, induction, and deduction helps to understand the relative roles of these logic forms in certain aspects of forming, confirming, and rejecting hypotheses. Let's return to the abductive case.

Because Peirce sought to characterize abduction as a form of logic, he sought some "rules" of abduction. Harman (1965) characterizes abduction as "inference to the best explanation." We are somewhat more comfortable thinking about abduction in terms of certain kinds of "constraints" rather than rules. For one thing, constraints operate to influence a process without completely determining it. Abduction is concerned with generating plausible ideas, not proving them, so relaxing the requirements of "rules of logic" to constraints seems more appropriate for a theory of abductive reasoning. Peirce proposed that certain conditions associated with testing hypotheses might figure into their value for scientific research. For example, other things being equal, hypotheses that are easy to test might be favored over those that require more time, effort, or money to test. This brings economic criteria to bear in selecting hypotheses.

There are several other criteria or constraints that affect our judgments about the quality of hypotheses. Returning to the example presented in Figure 1, consider the hypothesis that the figure depicts a bear climbing up the other side of a tree. Do you like it? Most people like this suggestion more than the others advanced earlier. Why? One characteristic of the bear hypothesis is that it explains the *entire* figure. It explains not only the existence of the lines and the circles, but it explains the number of lines and circles. It explains why the lines are roughly parallel and why the circles are spaced in just the way they are. In other words, the bear hypothesis makes the most of the observations provided, perhaps even expanding the observations over what they were originally taken to be. Certain features that might have been considered arbitrary or coincidental become necessary and meaningful by virtue of the hypothesis. This is what makes a *good* hypothesis. We might call these constraints *coverage* and *fruitfulness*. Coverage refers to the extent of the coverage of the facts by the hypothesis. Fruitfulness refers to the information added by virtue of the interpretation afforded by the hypothesis.

Syllogisms are usually applied in the realm of deductive reasoning where we say that a syllogism is valid if the conclusion follows necessarily from the premises. When we add qualifications such as “plausible” to the conclusion, we may question the value of presenting the argument as a syllogism. The syllogistic form may tempt one to seek forms of certainty in the realm of abduction, but such an endeavor is fruitless because abduction does not yield certainty. A better quest may be to clarify what it means for a hypothesis to be plausible, and then identify methods that would help to achieve plausibility.

Peirce held that inquiry serves to relieve doubt. If one’s beliefs are up to the task of accounting for experience, there is little motivation to examine those beliefs. Thus, surprise leads to a search for explanations to relieve doubt. Such explanations may lead to a change in beliefs either by adding new beliefs or by modifying established beliefs.

1.2 The Importance of Novelty

There are some reasons to think that there are different forms of abduction. Eco (1998) discusses the distinction in terms of the prior availability of the hypothesis. Two types of explanation differ in the status of the hypothesis before the abductive step. One form of explanation amounts to providing the general rule under which an observed case falls. This kind of reasoning occurs in medical diagnosis, for example, where a set of presenting symptoms (O) is “explained” by diagnosing the patient as carrying a certain “disease” (H). In this case, the disease was known as a possibility beforehand, and it provides an explanation of the symptoms in a particular case. This form of abduction amounts to determining which of a set of known explanations is to be adduced in a particular case. This might be called *selective* abduction.

A second form of abduction is at play when a new hypothesis is proposed as an explanation. This is where true creativity is at work. Historical examples in science are found in the Copernican heliocentric theory of the solar system, Pasteur's germ theory of disease, Darwin's theory of evolution, and Einstein's relativity theory. More common examples are to be found in problem solving and other creative activities in which novel ideas are generated to solve problems. This might be called *generative* abduction.

In its various forms, abductive reasoning is actually quite commonplace. Peirce, himself, proposed that perception is fundamentally abductive inference. Sherlock Holmes, notwithstanding, detective work also seems to be better characterized as primarily abductive rather than deductive. Solving a crime involves finding an explanation for the facts of the case (O) by postulating a perpetrator (H). The degree to which detection involves selective as opposed to generative abduction is an open question. It may depend upon the details of a particular case.

By using a variety of constraints in the generation process, the distinction between generation and evaluation may be obscured, but it may be of value to distinguish between cases where abduction leads to new knowledge in a system as opposed to calling up old knowledge. In actuality, novelty may come in degrees as knowledge is modified by abductive inference. Stuart Kauffmann (2000) develops the interesting idea of the "adjacent possible" by which he means that a system may take on a number of novel states that are "adjacent" in some sense to the prior state of the system. Thus, novelty for a system is relative to the state of a system at a given point in time. Still, some state changes may represent larger steps than others. It may be useful in distinguishing different abductive procedures and/or abductive outcomes by the magnitude of the change brought about by the abduction.

Novelty can be introduced at several levels including revising or expanding existing concepts, creating new concepts and categories, forming new propositions in the form of hypotheses or laws, or applying a system of relations to a new situation as in reasoning by analogy. Often abductive reasoning is triggered by a failure of expectation or a conflict between current beliefs and new observations.

1.3 Approaches to Understanding the Generation of Hypotheses

Methods for generating new knowledge generally depend in some way on *similarity*. Similarity can take many forms and includes both superficial and relational similarity. New concepts and categories depend on similarity of features or functions. Often some deep similarity is revealed by creative thought as illustrated by Arthur Koestler (1990) in his book, *The Act of Creation*, with the concept of bisociation. Koestler points out how creativity in humor, art, and science often involves bringing two distinct ideas together to reveal a deep similarity. This is illustrated in the following joke:

> A woman observes her friend in apparent deep distress. She asks, "Vats da matta, Millie?" She responds, "Oye ve, it's our son Sammy; the doctor says he has an Oedipus complex." She replies, "Oh Oedipus, Schmedipus, vats da matta as long as he's a good boy and loves his mama."

Here the creative juxtaposition of two ways of loving one's mother (bisociation) produces a humorous result. The similarity of oedipal love and the love for one's mother can be exploited to bring together two quite incompatible ideas.

Similarity is also involved in creating new propositions in Coombs, Pfeiffer, & Hartley's (1992) e-MGR system by combining parts of older propositions located by similarity to the data to be modeled (see also Coombs & Hartley, 1987). Gentner (1983) uses relational similarity as the basis of identifying analogies in her structure mapping system. Case-based reasoning systems (Kolodner, 1993) are related to analogical reasoning systems that attempt to find analogous past cases to use to analyze a current case. Similarity is at the heart of finding cases.

Perhaps an alternative to the use of similarity to guide the formation of new knowledge units is the use of some random process. Genetic algorithms (Holland, 1992) provide a good example of the successful use of randomness in creating new units. Of course, there are other important constraints at work in genetic algorithms besides randomness. Total randomness would hold little value in the search for effective new knowledge. Selective reproduction according to "fitness" helps direct genetic algorithms toward more "fit" units. In his paper, *The Architecture of Complexity*, Simon (1962) suggested that evolution depends on the formation of stable intermediate forms. The following quote makes this point and relates the process of evolution to problem solving:

> "A little reflection reveals that cues signaling progress play the same role in the problem-solving process that stable intermediate forms play in the biological evolutionary process. In problem solving, a partial result that represents recognizable progress toward the goal plays the role of a stable subassembly. "

In other words, if fruitful steps toward finding a solution to a problem can be recognized, the probability of finding a solution by trial and error can be greatly increased over the probability of generating a complete solution all at once which may be so small as to be nearly impossible. The importance of stable intermediate forms is further analyzed in Simon's 1981 book, *The Sciences of the Artificial.* Several additional constraints at work in abductive reasoning will be discussed in a later section.

1.4 Optimizing vs. Satisficing

Several approaches to abduction have been proposed and analyzed by researchers in cognitive science (Aliseda, 2000; Charniak & Shimony, 1990; Fann, 1970; Flach & Kakas, 2000; Josephson & Josephson, 1994; Kakas, Kowalski, & Toni, 1998; Konolige, 1996; Levesque, 1989; Peng & Reggia, 1990; Poole, 2000; Prendinger & Ishizuka, 2005; Senglaub, Harris, & Raybourn, 2001; Shrager & Langley, 1990; Walton, 2004). Many of these researchers have investigated the computational complexity of various algorithms associated with abductive reasoning. Such algorithms often exhaustively search some space of possibilities to optimize some measure. The algorithms are generally found to have complexity beyond reasonable computability which means they cannot scale up to the demands in most real applications. For example, Thagard and Verneurgt (1998) showed that deciding about the consistency of a set of propositions is NP hard which is generally believed to be intractable. Bylander, Allemang, Tanner, & Josephson (1991) reached the same conclusion in another analysis of the computational complexity of abduction. Santos & Santos (1996) showed that linear programming leads to good solutions for some abduction problems using relaxation of integer program formulations. Thagard and Verneurgt also report on a connectionist (neural net) approximation algorithm which gives good results in reasonable time. Reggia and Peng (1993) proposed a connectionist solution to diagnostic problem solving. Adaptive resonance theory (Carpenter & Grossberg, 2003) is still another approach to discovery in the framework of dynamical systems theory. Juarrero (1999) presents a compelling account of the connections between dynamical systems theory and intentional action. Such ideas appear to have a great deal to contribute to the development of theories of abductive reasoning.

Because abduction produces only plausible and not certain conclusions, it seems unnecessary to approach the problem through optimization. There are heuristic methods that arrive at very good solutions in reasonable time. Such methods seem particularly appropriate for abduction. Heuristic solutions amount to what Simon (1947) called satisficing, finding a satisfactory solution rather than an optimal one.

2. GENERATING & EVALUATING HYPOTHESES

Several factors influence the evolution of hypotheses. To varying degrees, the factors affect the generation or the evaluation of hypotheses. Generation and evaluation are not necessarily completely distinct processes. There is likely continually interplay between generating ideas and evaluating them in the search for an acceptable hypothesis. The following section enumerates several of the factors at work in terms of constraints operating in abductive reasoning. The criteria are characterized as constraints, in part, because each criterion is defeasible. That is, useful abductions may result by discounting any of the criteria.

2.1 Constraints on Abduction

Although abductive reasoning does not carry the certainty of deduction, there are constraints on what characterizes good hypotheses. A general account of abduction can proceed by identifying the constraints satisfied by the abduction. Abduction systems can be analyzed in terms of the constraints they embody. Different prospective hypotheses can also be compared by the extent to which they meet the constraints. An ordering of the hypotheses by preference follows from such comparisons. An important avenue for research is to determine the proper weighting of the various constraints. *A principled method for varying the weighting of the constraints would produce a variety of hypotheses according to different assumptions.* Here is a summary of some constraints to be considered as contributing to abductive reasoning.

- **The Observations.** Providing an explanation is a primary constraint on abduction. That the observations follow from the hypothesis is a first condition of plausibility of the hypothesis. At first blush, the observations appear to be primarily involved in evaluation as opposed to generation. However, the observations are also the starting point of the whole process. As discussed later, the observations also enter into similarity relations which are critical in generating potential abductive inferences.

- **Reliability of the Observations.** While observations provide primary constraints, the possibility of error in part or all of the observed data must also be considered. More reliable data should be weighted more heavily. If discounting some aspects of the data lead to a coherent account of the remaining data, the discounted data may be submitted to closer scrutiny.

- **Surprise.** Surprising or unexpected observations point to the need for a new hypothesis. When existing explanations of events fail to cover a newly observed event, abductive inference is called into play. While this is generally true, there may also be value in generating new hypotheses even while the current ones seem to be adequate to the task. Such hypotheses might provide for novel perspectives suggesting new ways to evaluate existing hypotheses.

- **Novelty of Hypotheses.** For observations to be considered surprising, it should be the case that ready explanations for the observations are not available. Thus, by this criterion, the novelty of a hypothesis counts in its favor. Novel hypotheses emphasize generation rather than a search among existing hypotheses.
- **Economics.** Hypotheses that are easier (less expensive) to test should be entertained before those that entail more difficult (more expensive) means of testing. This is one of the criteria suggested by Peirce in his work on abductive reasoning.
- **Parsimony**. Simpler hypotheses are preferred over more complex ones (Occam's razor). Parsimony would appear to be primarily an evaluative criterion, but it is also possible that simpler hypotheses would be easier to generate than more complex ones.
- **Aesthetics**. Beauty, elegance, symmetry, and appeal figure into the value of a hypothesis. Again, this constraint seems to be evaluative, but aesthetic factors could also influence certain characteristics of the hypotheses generated.
- **Plausibility and Internal Consistency**. Hypotheses consistent with each other and with background knowledge are preferred over ones that lead to contradictions. This constraint can also be seen to have both evaluative and generative dimensions. Generation might be expected to be strongly influenced by what is already known, and the acceptability of a generated hypothesis may well affect the likelihood of its survival.
- **Explanatory Power (Consilience).** This criterion is primarily evaluative in the sense that it applies to a hypothesis in hand where its explanatory power can be seen. There are various aspects of consilience such as:
    - **Coverage**. The extent to which a hypothesis explains the details of observations including incidental, in addition to central, details – the greater the coverage, the better the hypothesis.
    - **Fruitfulness**. The information added to the observations by virtue of the interpretation afforded by a hypothesis including providing meaning to features that were previously seen as incidental – the more information added, the better the hypothesis.
    - **Organization of the observations**. Hypotheses that reveal connections among the observations that were not obvious

before are of particular value. For example, a hypothesis may suggest related clusters of observations.

- **Pragmatics**. Pragmatics emphasizes the influence of goals and the context in which reasoning occurs. Goals and context are additional constraints on abductions. Pragmatics can operate both to direct generation and evaluation of hypotheses.
- **Analogy Formation**. Analogy formation works by finding sets of relations found in a source domain that can be applied to a target domain. An often cited example is the analogy between the solar system and an atom where parallels can be drawn between the sun and the nucleus of an atom and between planets and electrons of an atom. Once an analogy is drawn on the basis of known relations, characteristics from the source domain can be hypothesized to hold for the target domain.
- **Random Variation**. Hypotheses may be found by some random variation in older hypotheses. A system that constantly seeks for better hypotheses might be expected to occasionally find particularly good hypotheses that had not been considered before. Constraints in addition to randomness are probably necessary. Random variation alone is unlikely to lead to good results. Genetic algorithms employ randomness with other constraints. Genetic algorithms come primarily from the work of John Holland (1992, 1995, 1998). These methods are inspired by genetic reproduction where such processes as crossover and mutation lead to increases in "fitness" of individuals in populations. The methods are used in a variety of optimization problems. Genetic algorithms include a degree of randomness in the selection of mates and in mutation. Mate selection is also controlled by fitness which constrains the influence of random selection.
- **Similarity and Associations**. Similarity at various levels is a weak constraint on abductive reasoning, but similarity, at some level, is often involved in suggesting abductive inferences. Similarity may guide the search for commonalities among features, objects, and rules. In analogical reasoning, similarity of relations is a critical feature. Koestler (1990) proposed bisociation as a prominent feature of creative endeavors. Bisociation is the bringing together of unrelated ideas in a way that draws out a relation between them. In analogical reasoning, patterns of similarities provide constraints on abduction, but with analogies, the similarity is found at the level of relations – similar relations suggest analogies (Gentner, 1983; Gentner, Holyoak, & Kokinov, 2001; Gentner & Markman, 1997;

Holyoak & Thagard, 1995). In a study of insight in problem solving, Dayton, Durso, and Shepard (1990) found that critical associative connections underlying the solution of the problem often appeared in Pathfinder networks before the problem was solved suggesting that arriving at a solution may be mediated by establishing the critical connections. A similar process may be at play in some cases of abductive inference.

2.2 Similarity in Abductive Inference

Novel abductive inferences are not strongly associated with the phenomena to be explained. Rather such strong associations would make the inference obvious rather than novel or surprising. However, similarity or association of some kind may well be involved. The similarity may be indirect or the association weak, but the connection is often obvious in hindsight. Bruza, et al. (2006) discuss the value of identifying indirect associations in discovering a novel medical treatment involving the use of fish oil to treat Raynaud's Syndrome (intermittent blood flow in the extremities). Swanson (1986, 1987) proposed the treatment by connecting ideas from two unconnected literatures regarding the syndrome and dietary fish oil. Bruza, et al. suggest that such connections can be generated from textual sources by identifying concepts (terms) that do not occur together, but they do tend to co-occur with the same other concepts. The HAL system (Burgess, et al., 1998; Lund & Burgess, 1996) and the LSA method (Landauer & Dumais, 1997; Landauer, Foltz, & Laham, 1998) lead to identifying high degrees of similarity for terms that have similar patterns of co-occurrence with other terms in the database. They use a similarity measure based on the cosine of the angle between the vectors for each of the terms where the vectors represent the co-occurrence patterns of each of the terms. Bruza, et al. show the connection of fish oil and Raynaud's Syndrome discovery using such methods.

There is a longstanding interest in the role of geometric models or conceptual spaces in cognition (Gärdenfors, 2000; Kruskal, 1964a, 1964b; Shepard, 1962a, 1962b; Widdows, 2004). Gärdenfors proposes an important role for a geometric level of representation, distinct from both low-level connectionist processes and higher-level symbolic representations. An important role of the geometric level is to provide a basis for establishing similarity relations by virtue of the relations among concepts in conceptual space. Much of this work sees concepts as corresponding to regions of low-dimensional conceptual space. In other models, such as HAL and LSA, concepts correspond to vectors in high-dimensional conceptual space. Both views support the idea that similarity can be derived from spatial information.

The use of cosine measures on vectors representing the distribution of terms in text provides a way of assessing similarities between terms. Such similarity reflects both the co-occurrence of terms and similarities in the patterns of co-occurrences across all of the terms in a corpus. By eliminating pairs of terms that occur together in the corpus, one can focus on "indirect" similarity, similarity that derives from similar patterns of co-occurrence rather than direct co-occurrence. These indirect similarities may suggest possible abductive inferences. Not all indirect similarities can be expected to constitute such inferences. Synonyms rarely occur together in text, but they could be expected to have similar patterns of co-occurrence with other terms. While these would not qualify as novel inferences, they should be relatively easy to identify. Often we can characterize the *type* of thing that would qualify as a useful inference. For example, if we are looking for possible treatments of a disease or syndrome, only indirect similarity with things that could be treatments would be entertained as potential abductive inferences pertinent to the disease or syndrome. At this stage of our work, we rely on human judgment to determine which, if any, of the terms with indirect similarity to a target of interest constitute interesting potential abductive inferences.

For human evaluation, it is useful to view collections of terms indirectly related to a target term as Pathfinder networks (McDonald, Plate, & Schvaneveldt, 1990; Schvaneveldt, Dearholt, & Durso, 1988; Schvaneveldt, Durso, & Dearholt, 1989; Schvaneveldt, 1990) which depict patterns of relationships among the terms via patterns of links among the terms. Such networks show the strongest similarities among the terms, often revealing interesting paths among the terms as a way of identifying intermediate relationships of interest in addition to showing terms of interest.

3. RANDOM VECTORS AND PATHFINDER NETWORKS AS AIDS FOR ABDUCTION FROM TEXT

In this section, we present some findings from our work on developing computational tools to support abductive inference from textual corpora. Here we provide only a brief look at the work which will appear in more detail in Cohen, Schvaneveldt, & Widdows (under review).

The ability of methods such as LSA and HAL to find meaningful connections between terms (such as "raynaud", "fish" and "oil") that do not co-occur directly in any text passage can be considered as a sort of inference.

Landauer and his colleagues describe this as an "indirect inference" and estimate that much of LSA's human-like performance on tasks such as the TOEFL synonym test relies on inferences of this sort (Landauer and Dumais, 1997). In Figure 2 we illustrate the ability of LSA to identify interesting similarities. These indirect inferences are abductive in nature. They arise from similar patterns of occurrence across the corpus in the absence of co-occurrence.

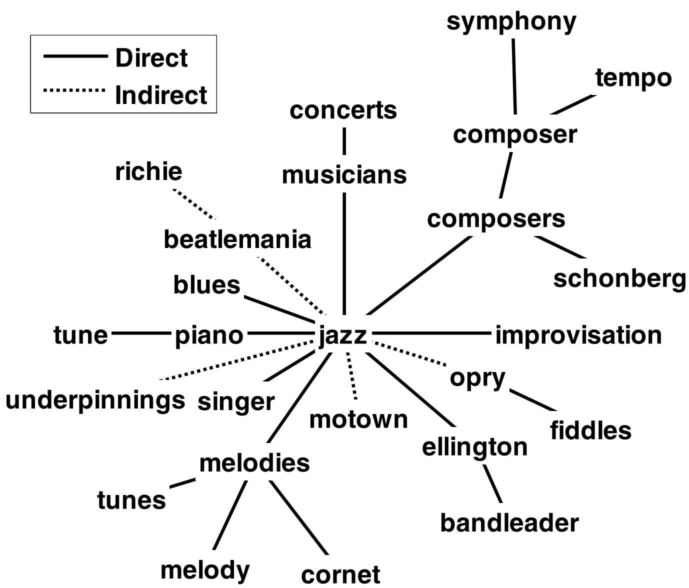

Figure 2. PFNET of nearest indirect LSA neighbors of "beatlemania"

This figure shows a Pathfinder network (PFNET) of the 20 *nearest indirect neighbors* of the term "beatlemania" in a semantic space derived from the Touchstone Applied Sciences (TASA) corpus using the General Text Parser software package (Giles, Lo, & Berry, 2003), obtained by screening out all terms that occur directly with the term "beatlemania" in any document in this corpus. The links in the PFNET are determined by the cosine similarities between all pairs of the terms, but after the application of Pathfinder network scaling only those links representing the most significant pairwise similarities are preserved[1]. Dashed links illustrate indirect

1 The PFNETs presented here were all computed with parameters, $r = \infty$ and $q = n - 1$, where n is the number of nodes in the network. The links preserved with these parameters

connections between terms that do not co-occur directly in any document. Many of these connections make intuitive sense, as they refer to musical forms and performers more commonly associated with musical genre other than pop. The figure also reveals a number of other interesting indirect neighbors of the term “jazz”, such as “motown” and (the grand 'ol) “opry”.

PFNETS preserve the most significant links between nodes in a network, and consequently reveal the semantic structure underlying this group of near neighbors, such as the western classical music related connection between “Schonberg”, “composers”, “composer” and “symphony”. It is also possible to use PFNETS to attempt to uncover the similarities that lead to an interesting indirect connection. For example, if we combine the representations of “jazz” and “beatlemania” by simply adding their corresponding vectors together and generate the nearest neighbors of this combined representation, we derive the PFNET shown in Figure 3.

consist of the union of the links in all minimum spanning trees or, in terms of similarities, the union of the links in all maximum spanning trees. The sum of the similarities associated with the links in such trees is the maximum over all possible spanning trees. See “Pathfinder Networks” in Wikipedia for additional information.

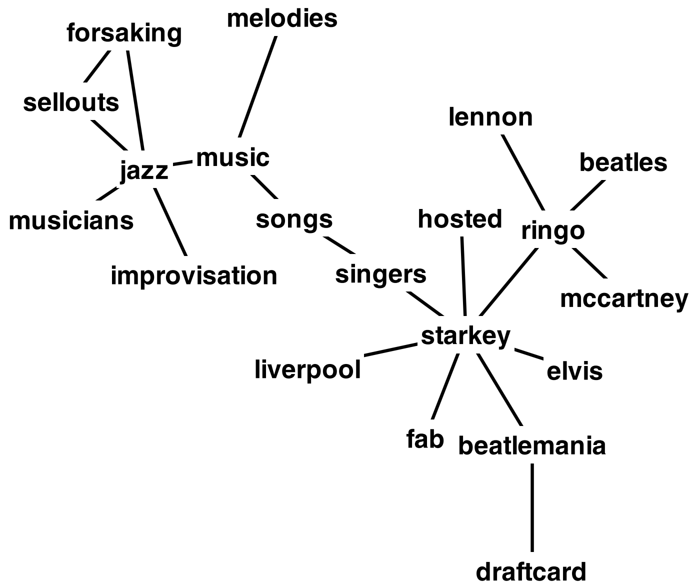

Figure 3. PFNET of the nearest RRI neighbors of "beatlemania + jazz"

Pathfinder has revealed a chain of significant links leading from "jazz" through "music" (jazz is a musical genre), "songs" (a musical form), "singers" (of songs) and "starkey" to beatlemania. This track of relations between "jazz" and "beatlemania" might be seen as a form of bisociation (Koestler, 1990) where the intermediates explain the indirect connection. Starkey here refers to Richard Starkey, the name on the birth certificate of Beatles member Ringo Starr. Although Starr was the group's drummer, he was also a backing vocalist as well as lead vocalist on several well-known tunes such as "Yellow Submarine" and "With a Little Help from My Friends".

Figure 3 was generated using as a basis semantic distances estimated by the Random Indexing model (Kanerva, et al., 2000; Karlgren & Sahlgren, 2001) using the Semantic Vectors Package (Widdows & Ferraro, 2008). Random Indexing is similar in concept and underlying assumptions to Latent Semantic Analysis in that terms are represented in a vector space according to their distribution across a large set of documents. However, unlike LSA, Random Indexing does not depend upon computationally demanding methods of dimension reduction to generate a condensed vector representation for each term. Rather, it achieves this end by projecting terms directly into a vector space of predetermined dimensionality (usually > 1000) by assigning to each document a randomly-generated index vector in this subspace that is close-to-orthogonal to every other assigned index vector. While more investigation is needed to determine which aspects of LSA's performance this method as able to reproduce, initial investigations show it is possible to use this method of dimension reduction without degrading the model's performance on synonym tests (Kanerva, et al., 2000). Unlike LSA, the model scales comfortably to large corpora, as we illustrate below with networks derived from the MEDLINE corpus of abstracts. In our studies, we have found Random Indexing using a term-document approach to be somewhat limited in its ability to generate meaningful indirect inferences. Consequently the remaining diagrams, with the exception of the "thrombophilia" example, were produced using Reflective Random Indexing (RRI), a method that creates term vectors by an iterative construction (Cohen, et al., under review).

The PFNET in Figure 4 was created with a similar approach, however in this case the semantic distance between terms was generated from the abstracts of articles in the MEDLINE database of biomedical literature occurring from 1980 to 1986. An indirect connection between the terms "pneumocystis" (*pneumocystis carinii pneumonia* occurs in immunocompromised patients) and "promiscuity" (promiscuity amongst the homosexual community was implicated in the transmission of the recently discovered Acquired Immune Deficiency Syndrome) was retrieved amongst the 20 nearest indirect neighbors of the term "pneumocystis". Figure 4 shows the 20 nearest neighbors of the combined representation of "pneumocystis" and "promiscuity" using the RRI method of creating the index.

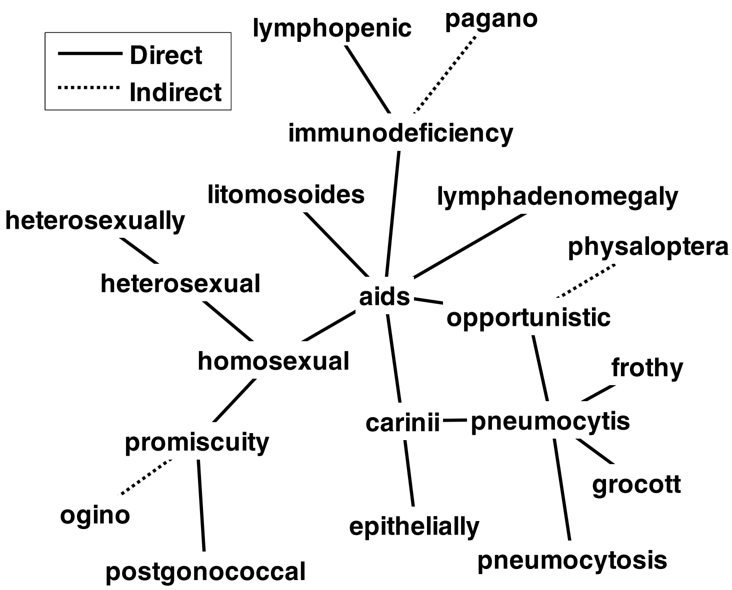

Figure 4. PFNET of nearest neighbors of "pneumocystis + **p**romiscuity"

Again, the PFNET reveals a plausible line of reasoning connecting these two terms. Pneumocystis is connected through "carinii", "aids" and "homosexual" to promiscuity. This PFNET illustrates an inferred relationship between pneumocystis carinii pneumonia and promiscuity, which was not explicitly stated in any MEDLINE abstract used to build this model. Interestingly, the "ogino" indirectly connected to "promiscuity" in this diagram refers to Kyusaku Ogino, who measured the fertile period of the female menstrual cycle. While Ogino did not believe this method could be used as a reliable form of contraception, the Rhythm Method of contraception is nonetheless referred to as the "Ogino Method" in the occasional MEDLINE record.

Another interesting indirect connection to emerge from the TASA corpus through Random Indexing is an association between Picasso and impressionism shown in Figure 5. Deriving a PFNET from the combined vector for "picasso" and "impressionism" reveals a pathway from Picasso through "cubism, " "cubist," and "carafe," an important cubist work of Picasso, to "manet" to "impressionists" to "impressionism." Manet's work, particularly *Le déjeuner sur l'herbe,* is considered to by many critics to have exerted a seminal influence on the evolution of cubism through the work of Picasso (and George Braque). Picasso also painted variations of several works of Manet, including *Le déjeuner* which at the time of this writing is featured in the exhibition "*Picasso / Manet: Le déjeuner sur l'herbe"* at the Musée d'Orsay in Paris.

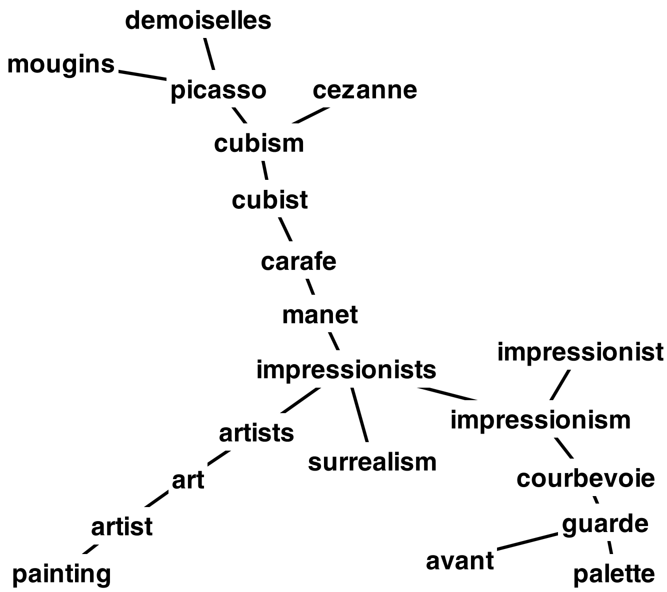

Figure 5. PFNET of nearest RRI neighbors of "picasso + impressionism"

We have further investigations underway to produce and evaluate indirect connections obtained from the MEDLINE database. An interesting indirect association was observed between "spongiform" and "cannibalism." This association was noted in a subset of MEDLINE abstracts occurring between 1980 and 1985. The spongiform encephalopathies, such as bovine spongiform encephalopathy in cattle (BSE, aka Mad Cow Disease), scrapie in sheep and Creutzfeld-Jacob Disease (CJD) in humans are degenerative neurological disease that are currently thought to be caused by prions, infectious protein agents that replicate in the brain. While Prusiner's prion hypothesis (Prusiner 1982) was contested at this time, he was later awarded the Nobel Prize for this work. "Kuru" is an encephalopathy that was transmitted by cannibalistic practice in New Guinea.

A PFNET for the combined terms "spongiform+cannibalism" is shown in Figure 6. This PFNET reveals a pathway (CJD via "kuru" to "cannibalism"). This pathway reveals a plausible line of reasoning connecting cannibalism through kuru to other spongiform encephalophathies, and the prion hypothesis which was first proposed in the context of scrapie. A similar line of reasoning was explored by Prusiner during the course of his research, in which he developed an experimental model of the transmission of scrapie using the natural cannibalistic activity of hamsters (Prusiner 1985). No direct connection between "prion" and "kuru" exists in the 1980-1985 corpus of abstracts, and while the notion that kuru may also be caused by a prion protein is unlikely to have been novel at the time, this example provides an interesting illustration of how the exploration of one meaningful indirect inference can reveal another.

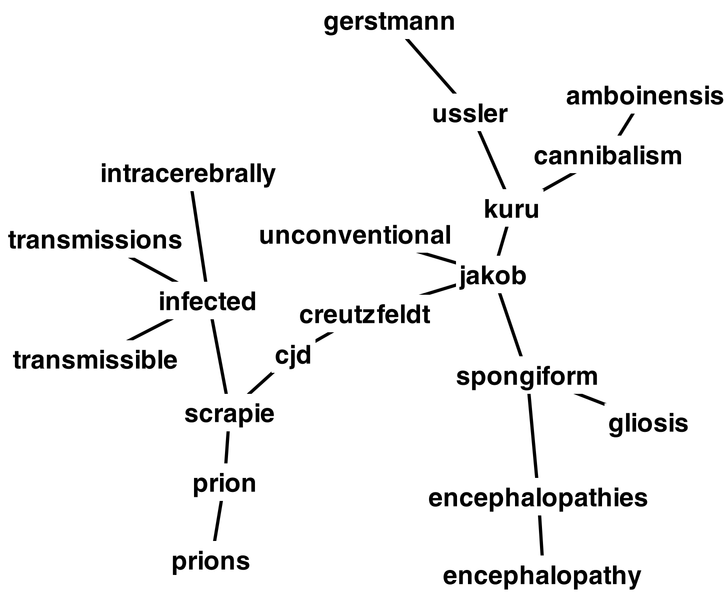

Figure 6. PFNET of nearest RRI neighbors of spongiform+cannibalism

The discovery of a hypothesis is illustrated in Figure 7, the term "rppgf" was returned as a near neighbor to a cue term "thrombophilia" (Cohen, 2008). RPPGF is the protein sequence Arg-Pro-Pro-Gly-Phe, the sequence of an inhibitor of platelet aggregation that could be therapeutically useful in thrombophilia. However, a PubMed search (conducted on 06/06/2008) for "rppgf AND thrombophilia" does not retrieve any results. Further examination of the MEDLINE corpus shows that these terms do not directly co-occur in any of the abstracts in MEDLINE. However, despite this lack of direct co-occurrence, the indirect similarity between these two terms in the RI space derived from these abstracts was sufficient for "rppgf" to be among the nearest neighbors of "thrombophilia." Discoveries of this sort are the focus of our present research including work on using random indexing methods to encode and retrieve the types of relations that exist between concepts (Cohen, Schvaneveldt, & Rindflesch, 2009)

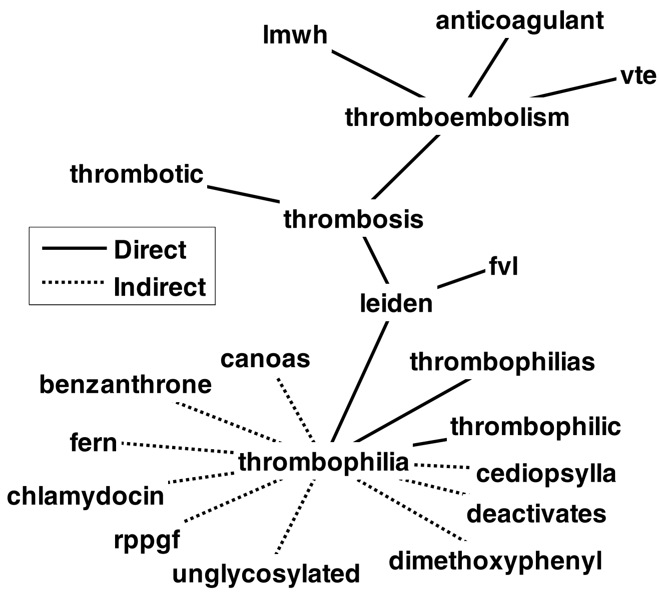

TFigure 7. PFNET of nearest RI neighbors of “thrombophilia”

4. PREDICTING “DISCOVERIES”

If the indirect similarity of terms is a harbinger of an undiscovered relationship between the concepts corresponding to the terms, we might expect that indirect neighbors would tend to become direct neighbors over time. By indirect neighbors, we mean items that are similar to a target item but do not co-occur with the target. Direct neighbors are similar items that do co-occur with the target. Using the MEDLINE database, we assessed the proportion of nearest indirect neighbors between 1980 and 1986 that became direct neighbors after 1986 (“discoveries”). In this experiment we investigated two different method for creating term vectors, standard random indexing (RI) developed by Kanerva, Kristofersson, and Holst (2000) and a new reflective random indexing (RRI) method adjusted to improve indirect similarity (Cohen, et al., under review). The reflective method involves iteratively creating term and document vectors starting with random vectors. The full MEDLINE index of abstracts contains 9,003,811 documents and 1,125,311,210 terms of which 3,948,887 terms are unique. Our index consists of about 300,000 unique terms which excludes terms occurring less than 10 times and terms that contain non-alphabetic characters.

Two thousand (2,000) target terms were randomly selected, and the nearest indirect neighbors (NINs) of each of the targets were found in the database between 1980 and 1986. Then each of the indirect neighbors was checked to determine whether it co-occurs with its target after 1986. The ones that did co-occur were dubbed "discoveries." The results are shown in Figure 8.

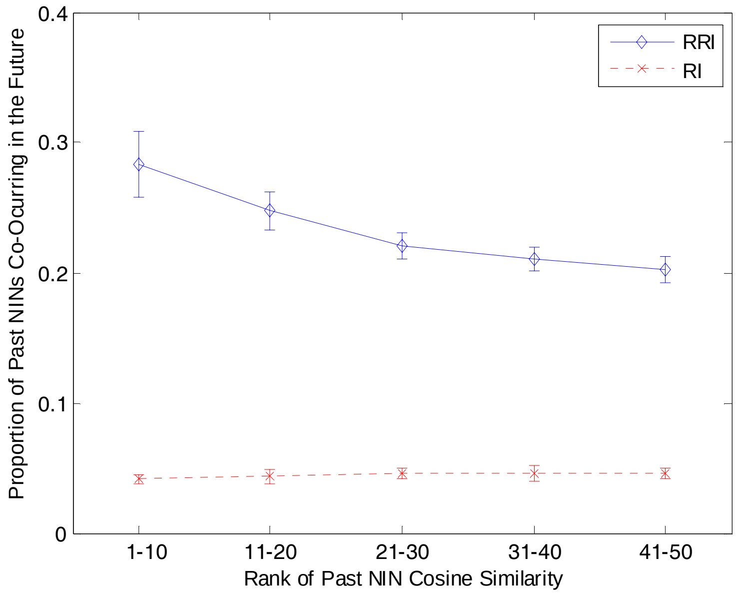

Figure 8. Future "discoveries" from past indirect neighbors (NIN).

The RI index did not produce many "discoveries," a maximum of 4.5% while the rate of discoveries with the RRI index reached 28.4% for the 10 nearest indirect neighbors. The difference between the two indexes is statistically significant, t (1999) = 53.11, p < .0001. There was also a significant decrease in the rate for the RRI index from 28.4% to 22.0% for nearest indirect neighbors 11-50, t (1999) = 17.81, p < .0001. This decrease shows that stronger indirect similarity leads to a greater rate of "discoveries" which suggests that the indirect similarity measure does reflect the importance of the relation between the terms. For the less successful RI index, decreasing similarity lead to a slightly increased rate of "discoveries," 4.5% compared to 4.2%.

These findings suggest that indirect similarity may well be a precursor to the future realization of the relations between concepts. Clearly, there is more work to be done to explore and evaluate these findings. At this point, we find some clear support for continuing this line of work.

5. CONCLUSIONS

New ideas may be sparked by noticing indirect similarities. The spark is essential in leading to novel possibilities in the abductive reasoning found in problem solving and hypothesis generation. We have shown the value of tracing indirect similarities through examples and an analysis of the fate of indirectly similar terms. Our initial efforts at understanding the nature and role of indirect similarity encourage us to continue to pursue the development of this approach and the tools to support it. Although the efforts reported in this paper have concentrated on finding indirect similarities in textual corpora, it could be argued that analogous processes operate in cognition generally. Exciting work lies ahead to elaborate on such possibilities.

6. REFERENCES

Aliseda, A. (2000). Abduction as epistemic change: A Peircean model in artificial intelligence. In Flach, P. & Kakas, A. (Eds.), *Abduction and induction: Essays on their relation and integration.* Boston: Kluwer Academic Publishers.

Bruza, P., Cole, R., Song, D., & Bari, Z. (2006). Towards operational abduction from a cognitive perspective. *The Author.* Oxford: Oxford University Press.

Burgess, C., Livesay, K, & Lund, K. (1998). Explorations in context space: Words, sentences, discourse. *Discourse Processes,* 25, 2&3, 211-257.

Bylander, T., Allemang, D., Tanner, M. C., & Josephson, J. R. (1991). The computational complexity of abduction. *Artificial Intelligence,* 49, 25-60.

Carpenter, G.A. & Grossberg, S. (2003). Adaptive Resonance Theory. In M.A. Arbib (Ed.), The *Handbook of brain theory and neural networks*, *2nd Ed*., Cambridge, MA: MIT Press, 87-90.

Charniak, E., & Shimony, S. E. (1990). Probabilistic semantics for cost based abduction. In *Proceedings of the National Conference on Artificial Intelligence*, (pp. 106-111), Boston.

Cohen, T. (2008). Exploring MEDLINE space with random indexing and Pathfinder networks. *American Medical Informatics Association Symposium*, Washington, DC.

Cohen, T., Schvaneveldt, R. W., & Rindflesch, T. C. (2009). Predication-based semantic indexing: Permutations as a means to encode predications in semantic space. *American Medical Informatics Association Symposium*, Washington, DC.

Cohen, T., Schvaneveldt, R., & Widdows, D. (under review). Reflective random indexing and indirect Inference: A scalable method for discovery of implicit connections.

Coombs, M., & Hartley, R. (1987). The MGR algorithm and its application to the generation of explanations for novel events. *International Journal of Man-Machine Studies*, 20, 21-44.

Coombs, M. J., Pfeiffer, H. D., & Hartley, R. T. (1992). e-MGR: An architecture for symbolic plasticity. *International Journal of Man-Machine Studies*, 36, 247-263.

Dayton, T., Durso, F. T., & Shepard, J. D. (1990). A measure of the knowledge organization underlying insight. In R. Schvaneveldt, (Ed.) *Pathfinder associative networks: Studies in knowledge organization.* Norwood, NJ: Ablex.

Eco, U. (1998). Hooves, horns, insteps: Some hypotheses on three types of abduction. In U. Eco & T. A. Sebeok (Eds.), *The sign of three*. Bloomington, IN: Indiana University Press.

Fann, K. T. (1970). *Peirce's theory of abduction.* The Hague: Martinus Nijhoff.

Flach, P. A., & Kakas, A. C. (2000). *Abduction and induction: Essays on their relation and integration.* Boston: Kluwer Academic Publishers.

Gärdenfors, P. (2000). *Conceptual spaces: The geometry of thought.* Cambridge, MA: MIT Press.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science, 7*, 155-170.

Gentner, D., Holyoak, K. J., & Kokinov, B. N. (Eds.). (2001). *The analogical mind: Perspectives from cognitive science.* Cambridge, MA: MIT Press.

Gentner, D., & Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist, 52,* 45-56.

Giles, J. T, Wo. L, & Berry, M. W. (2003). GTP (General Text Parser) Software for Text Mining. In H. Bozdogan (Ed.), *Statistical Data Mining and Knowledge Discovery,* Boca Raton: CRC Press, pp. 455-471.

Gonzalez, G., Uribe, J.C., Tari, L., Brophy, C. & Baral, C. (2007). Mining Gene-Disease relationships from Biomedical Literature: Incorporating Interactions, Connectivity, Confidence, and Context Measures. *Proceedings of the Pacific Symposium in Biocomputing*, Maui, Hawaii.

Harman, G. H. (1965). The inference to the best explanation. *The Philosophical Review*, 74, 88-95.

Holland, J. H. (1992). *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence*. Cambridge, MA: MIT Press.

Holland, J. H. (1995). *Hidden order: How adaptation builds complexity*. New York: Addison Wesley.

Holland, J. H. (1998). *Emergence: From chaos to order.* Cambridge, MA: Perseus.

Holyoak, K. J., & Thagard, P. (1995). *Mental leaps.* Cambridge, MA: MIT Press.

Josephson, J. R., & Josephson, S. G., Eds. (1994). *Abductive inference: Computation, philosophy, technology.* New York: Cambridge University Press.

Juarrero, A. (1999). *Dynamics in action: Intentional behavior as a complex system*. Cambridge, MA: MIT Press.

Kakas, A. C., Kowalski, R. A., & Toni, F. (1998). The role of abduction in logic programming. In D. M. Gabbay, C. J. Hogger, & J. A. Robinson (Eds.), *Handbook of Logic in Artificial Intelligence and Logic Programming: Vol. 5, Logic Programming*. Oxford: Clarendon Press.

Kanerva, P., Kristofersson, J., and Holst, A. (2000). Random Indexing of text samples for Latent Semantic Analysis. In L. R. Gleitman and A. K. Josh (Eds.), *Proceedings of the 22nd Annual Conference of the Cognitive Science Society*, p. 1036. Mahwah, New Jersey: Erlbaum.

Karlgren, J., & Sahlgren, M. (2001). From words to understanding. *Foundations of Real-World Intelligence*, 294-308.

Kauffmann, S. (2000). *Investigations*. Oxford: Oxford University Press.

Koestler, A. (1990). *The act of creation*. New York: Penguin.

Kolodner, J. (1993). *Case-Based Reasoning*. San Mateo: Morgan Kaufmann.

Konolige, K. (1996). Abductive theories in artificial intelligence. In B. Brewka (Ed.), *Principles of knowledge representation*. Stanford, CA: CSLI Publications.

Kruskal, J. B. (1964a). Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika, 29***,** 1-27.

Kruskal, J. B. (1964b). Nonmetric multidimensional scaling: a numerical method. *Psychometrika, 29***,** 115-129.

Landauer, T. K, Foltz, P. W. & Laham, D. (1998). An introduction to Latent Semantic Analysis. *Discourse Processes*, *25*, 2&3, 259-284.

Landauer, T. K. & Dumais, S. T. (1997). A solution to Plato's problem: the Latent Semantic Analysis theory of acquisition, induction and representation of knowledge. *Psychological Review, 104(2)*, 211-240.

Levesque, H. J. (1989). A knowledge-level account of abduction. In *Proceedings of the International Conference on Artificial Intelligence.* Detroit, MI.

Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments & Computers,* 28(2), 203-208.

McDonald, J. E., Plate, T. A., & Schvaneveldt, R. W. (1990). Using Pathfinder to extract semantic information from text. In R. Schvaneveldt (Ed.), *Pathfinder associative networks: Studies in knowledge organization.* (pp. 149-164). Norwood, NJ: Ablex.

Peirce, C. S. (1940a). Abduction and induction. In J. Buchler (Ed.), *Philosophical writings of Peirce*. New York: Routledge.

Peirce, C. S. (1940b). Logic as semiotic: The theory of signs. In J. Buchler (Ed.), *Philosophical writings of Peirce*. New York: Routledge.

Peng, Y. , & Reggia, J. A. (1990). *Abductive inference models for diagnostic problem-solving.* New York: Springer-Verlag.

Poole, D. (2000). Abducing through negation as failure: Stable models within the independent choice logic. *The Journal of Logic Programming*, 44, 5-35.

Popper, K. (1962). *Conjectures and refutations.* London: Routledge.

Prendinger, H., & Ishizuka, M. (2005). A creative abduction approach to scientific and knowledge discovery. *Knowledge-Based Systems*, 18(7), 321-326.

Prusiner, S. B. (1982). Novel proteinaceous infectious particles cause scrapie. *Science*, *216*(4542), 136-44.

Prusiner, S. B., Cochran, S. P., & Alpers, M. P. (1985). Transmission of scrapie in hamsters. *The Journal of Infectious Diseases*, *152*(5), 971-8.

Reggia, J. A., & Peng, Y. (1993). A connectionist approach to diagnostic problem solving using causal networks. *Information Sciences*, 70, 27-48.

Santos, E., Jr., & Santos, E. S. (1996). Polynomial solvability of cost-based abduction. *Artificial Intelligence*, 86, 157-170.

Schvaneveldt, R. W. (Editor) (1990). *Pathfinder associative networks: Studies in knowledge organization.* Norwood, NJ: Ablex.

Schvaneveldt, R. W., Dearholt, D. W., & Durso, F. T. (1988). Graph theoretic foundations of Pathfinder networks. *Computers and Mathematics with Applications*, *15*, 337-345.

Schvaneveldt, R. W., Durso, F. T., & Dearholt, D. W. (1989). Network structures in proximity data. In G. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory*, Vol. 24 (pp. 249-284). New York: Academic Press.

Senglaub, M., & Harris, D., & Raybourn, E. M. (2001). ***Foundations for reasoning in cognition-based computational representations of human decision making.*** Technical Report SAND2001-3496, Sandia National Laboratories, Albuquerque.

Shepard, R. N. (1962a). The analysis of proximities: multidimensional scaling with unknown distance function Part I. *Psychometrika, 27*, 125-140.

Shepard, R. N. (1962b). The analysis of proximities: multidimensional scaling with unknown distance function Part II. *Psychometrika, 27*, 219-246.

Shrager, J., & Langley, P. (Eds.) (1990). *Computational models of scientific discovery and theory formation.* Palo Aalto, CA: Morgan Kaufmann.

Simon, H. A. (1947). *Administrative Behavior*. New York: The Free Press.

Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society, 106,* 467-482.

Simon, H. A. (1981). ***The Sciences of the Artificial.*** (2nd Edition). Cambridge, MA: MIT Press.

Swanson, D. R. (1986). Fish oil, Raynaud's Syndrome, and undiscovered public knowledge. *Perspectives in Biology and Medicine,* 30(1), 7-18.

Swanson, D. R. (1987). Two medical literatures that are logically but not bibliographically related. *Journal of the American Society for Information Science*, 38, 228-233.

Thagard, P., & Verbeurgt, K. (1998). Coherence as constraint satisfaction. *Cognitive Science*, 22, 1-24.

Walton, D. N. (2004). ***Abductive reasoning.*** Tuscaloosa, Alabama: University of Alabama Press.

Widdows, D. (2004). ***Geometry and meaning.*** Stanford: CSLI Publications.

Widdows, D., & Ferraro, K. (2008). Semantic Vectors: A Scalable Open Source Package and Online Technology Management Application. *Sixth International Conference on Language Resources and Evaluation (LREC 2008).*