Calculating Breadth of Knowledge

Lisa F. Rau
GE Research and Development Center
Schenectady, NY 12301 USA
and
Computer Science Department
University of Exeter, Exeter UK

Abstract

Since the advent of computers, information systems have grown in terms of the quantity of knowledge they deal with. Advances in data management are on the critical path for usability of these systems. This paper reports on a novel approach to an important problem; that of calculating the conceptual *breadth* of knowledge or data in a knowledge base or database. Breadth determination is useful in that ascribing meta-level knowledge of conceptual content can help to predict, for example, the validity of the closed-world assumption or the likelihood of encountering new information of a particular type. The point at which a system determines it is likely to have breadth in a given knowledge area may also serve as the trigger point for calculations that assume relatively complete knowledge in that area. The accurate determination of when a system has complete knowledge in an area is crucial for the accurate application of many AI algorithms.

Introduction

As information systems continue to grow in size and complexity, it becomes more and more critical to develop innovative methods of ascertaining the contents of these information systems. Three features motivate such a construction: (1) someone unfamiliar with the database can pose general questions about the breadth of content of the system (such as "do you know about X?"); (2) the system can now distinguish potentially missing information from information not likely to exist, i.e., the difference between a "no" response and an "I don't know" response and (3) knowing when complete knowledge exists allows for processes that depend on this assumption to operate. For example, the work of Pollack [Pollack, 1986] on plan failure analysis assumes the system has complete knowledge of what the user knows. Knowing when this is assumption is valid is critical in order to correctly apply the methods. The last two motivations address the important problem of knowing when to apply the closed world hypothesis [Reiter, 1978]; that everything not known is *false*.

This paper reports on a theory and implementation of the cognitive notion of conceptual *breadth*, a concept not typically modelled from a computational perspective. While the concept of *breadth of knowledge* has a clear intuitive meaning, this work aims to formalize the notion and provide a computational mechanism for determining breadth. By implementing this computation on an arbitrary knowledge or data base, it is now possible to construct a meta-representation for the scope or limits of what a system knows about. The representation is a "meta-representation" because it is knowledge *about* the knowledge in the system.

This work is one portion of a larger project to provide computational methods for automatically deriving what a system knows, for example, what is salient [Rau, 1992]. Breadth of knowledge is the notion of having a conceptual covering of a subject area. This covering may be shallow or deep. Because virtually all databases necessarily cover an arbitrary knowledge area, it is useful to make a distinction between *absolute* breadth and *relative* breadth. Relative breadth is breadth with respect to the domain of operation of a system. Relative breadth is achieved when a system contains knowledge along all of its possible expected dimensions. For example, consider a database that contains, among other things, information about business occupation for each person in the database. It might be limited to encountering only certain occupations due to the sample of the population in the database, from a socio-economic, regional, age or sex related or functional perspective. For relative breadth, seeing almost all of this subset of occupations is sufficient.

For absolute breadth, the set of occupations encountered should approximate the set of all occupations people have with any regularity. That is, absolute breadth is achieved when the system has seen all the *possible* values a given field can take from the standpoint of an objective measure of human knowledge. Clearly, computing absolute breadth is the more challenging problem.

This paper discusses algorithms for computing absolute and relative breadth with respect to both open class (infinite) categories, and closed class (finite) categories. It details the computational method used to ascertain relative breadth for closed class categories, and provides an experimental validation of the method on a test of a 1,917 record database produced by an NLP system. With these results, users may pose questions such as "What countries of interest does this database cover?"

and obtain, in addition to a list of those countries, the information as to whether these are *all* the countries or just the ***countries the system happens to have seen up to that point.*** Of additional utility, expert reasoning systems can make likely inferences based on the closed world assumption with this type of analysis.

Overview of Methodology

In this section, I describe the expected functionality of a breadth calculation in general terms, followed by the specific algorithms that implement this functionality. The experimental results are also described.

Functionality

Intuitively, a system has broad knowledge in a given area when it "knows about" (that is, has instances of) virtually all of the possible *manifestations* of knowledge in that area. Translated into database terms, given a database field $\mathcal{F}$ that contains unique fillers (types, not tokens) $(f_1, f_2, \cdots, f_n)$, the system has breadth of knowledge with respect to $\mathcal{F}$ when it has encountered virtually all the f_i, independent of how frequently each of the f_i appears. The assumption here is that a database field filled with a value is equivalent to the database "knowing" that information[1].

This straightforward observation implies that when a system has seen, in a particular field, as many distinct fillers as total numer of fillers, the expectation is that breadth has not been achieved, and that there is a potentially sizable number of unseen fillers not yet encountered. That is, when every filler of a field is unique, it is more likely that the next filler will also be unique. Conversely, when a system has experienced high redundancy in the distinct fillers, and/or has not seen new distinct fillers for a while, the intuition and expectation is that the system has likely seen most if not all the fillers it is likely to encounter, and breadth has been achieved. This scenario is complicated when it is impossible ever to have breadth of knowledge given an infinite number of possible distinct types of fillers. The next section describes how this and the above functionality is achieved.

Implementation

As we have just seen, the key concept for the computation of breadth is the certainty that the system has seen all possible instances of a given concept. To compute this, the system must first make a judgement as to whether the fillers consist of a closed or open class. That is, the system must determine whether the types of fillers are finite or infinite. A closed class is a class of objects with a fixed set of members, and doesn't change. An open class, on the other hand, has no finite number of possible members. For example, the syntactic category of *determiner* is a closed class consisting of a small set of words (such as **a**, **the**, **this**, etc.). The syntactic category of *noun* is an open class, as just about any word can be a noun, and new nouns are created all the time. This qualitative distinction is useful for natural language processing because a system must be able to predict the part of speech for a new, unknown word. It is important for AI systems in general because it gives information as to when it might be correct to apply a closed world assumption.

In summary, in order to compute breadth, the system must be able to perform two functions: (1) differentiate between cases where there are potentially infinitely many unique instance types (open-class fills) and cases where the instances come from a finite set (closed-class fills) and (2) quantify "virtually all" so that the expectation of when breadth is achieved is satisfied. This is done by providing a numerical threshold based on the probability of the given distribution's containing only those distinct types, under the assumption that the distribution of values already seen provides limits on the nature of the distribution in the database.

Determining Class Boundedness

Although the distinction between an open class and a closed class is not always clear [Collins *et al.*, 1975], from the perspective of a database, fillers of a field can always be so classified.

Some databases may have a data dictionary that gives type information for each field of each record. This information does not always map into a correct open/closed class determination, as in the case of person age. Person age in a data dictionary may simply be a number (an open class) whereas in reality it is a closed class of numbers from 0 to 120 or so. However it is quite easy to specify manually, for each field, whether it is closed or open class, so this is an option as well.

The easiest purely automated method for determining whether values of the field or slot are members of a closed class or an open class is to look at the number of distinct values. There are much larger (order of magnitude) numbers of distinct values in open class fields in any data or knowledge base of reasonable size. A reasonable sized database would contain at least thousands of records. Although it is possible to have a closed-class fill that contains thousands of elements (for example in the case of cities of the world), these "borderline" cases are best assumed to be open-class, as it is likely that the set of possible fills is not exhaustive. It is a safe assumption then that closed-class fills, by their nature, are bounded and the number of choices is dwarfed by the number of instances in large samples. Therefore, computing these values for a sample of the data will neatly classify a field as either closed-class or open-class.

Determining Breadth

The next four sections discuss the breadth computation for each of the four cases of relative or absolute breadth with respect to open or closed classes. Recall that relative breadth is breadth of knowledge with respect to a system's domain of operation, as opposed to some objective, absolute measure of human knowledge, termed absolute breadth.

[1]Note that this only applies to knowing the contents of a category, where category is defined by the fields in a database. For example, in the database used here, country is a field. More complex concepts, such as "big countries" or "things like countries" can be handled in exactly the same manner as long as a class membership function can be defined.

Relative Breadth, Closed Class: The problem of determining relative breadth for a closed class category is equivalent to the problem of making a determination as to when the system has seen almost all, if not all, of the members of that category. The threshold for this determination can be manipulated according to how certain you would like the system to be in this judgement. Relative breadth of a closed class category translates into the point at which the ratio of the number of distinct fillers seen to the total number of fillers seen (call this $\mathbf{d} \ / \ \mathbf{n}$) is around 1 in $\mathcal{T}$, where $\mathcal{T}$ is the threshold. In the case where the number of distinct fillers is equal to the number of fillers, this ratio is equal to 1. In this case, the system has seen only unique instances and intuitively would expect to see more; it has no reason to believe that it has breadth. In cases where this ratio is small and $\mathbf{n}$ is large, the system has seen many instances of the finite set of fillers. Taking this distribution to be a good approximation to the expected distribution across the database, the system is unlikely to see any additional distinct fillers. In this case, we can assume the system has breadth. The point at which this cutoff is made depends on how certain you would like the system to be in its breadth determination.

In particular, at any given point in time, the system computes either:

1. Breadth has been achieved. All the most frequently occurring kinds of fillers have been seen, and the only outliers are likely to be anomalous or both extremely infrequent and small in number compared to the total number of distinct fillers.
2. Breadth has not been achieved. The frequency of new fillers is such that it is likely that there are additional fillers that the system has yet to encounter, or there is insufficient data to judge.

We have experimented with more complicated formulations involving information theory and statistics, but this simple heuristic captures the intuition as well as the data. The intuition is that breadth is achieved when you are likely to have seen all the instances there are. The results are described in the next section.

Absolute Breadth, Closed Class: Absolute breadth with a closed class category takes advantage of a generic conceptual hierarchy. This hierarchy was built to support generic text processing, and contains concepts representing the most frequent words of English and appropriate super-categories. Because it has broad coverage of the more frequently used concepts that are used in language, it is a good domain-independent starting point for experiments comparing system knowledge to general human knowledge. Absolute breadth is achieved when over some percentage of the members of the closed-class category, as dictated by the hierarchy, have been seen. The exact percentage depends on what percentage of the knowledge area one assumes needs to be covered before breadth is achieved. Certainly more than 50%, probably more than 75%, and less than 100%. Working with the value of 98%, a simple example is breadth of knowledge of the states in the U.S. Given that there are absolutely and precisely 50 states, the system would have breadth of knowledge of what the states in the U.S. were if it knew of any 49.

Absolute Breadth, Open Class: It is very difficult to determine an absolute measure of breadth for an open class category. For example, consider the problem of ascertaining breadth of knowledge about the different shapes that snowflakes can be; something that at least folk science believes to be infinite. Since it is not possible even for an expert to know all the members of this conceptual category, it is impossible to determine when someone is likely to have breadth in this category, from an absolute perspective.

However if the elements of the open-class can be correctly assigned to superclasses (conceptual parents), the simple percentage of the total possible elements as defined by a generic conceptual hierarchy can give an approximation of when absolute breadth of an open-class category has been achieved. Continuing with the snowflake example, if snowflakes can be mapped into superclasses (one such assignment may be by the shape of the perimeter; hexagonal, round, five-pointed star, etc.), then some breadth is achieved when all the possible shapes have been seen.

There are two obstacles to overcome in order for this approach to succeed. The first is correctly determining the superclass categorization for each individual. In addition to choosing the correct parent from multiple possible parents of each individual, the hierarchy may not contain a parent whose elements correspond to exactly the set of possible individuals. For example, the set of possible shapes of snowflakes in the world may not contain shapes with less than four sides. Second, when the superclass itself is infinite, as the category of "shape" surely is, the superclass must be recursively partitioned into *its* superclass. If no finite partition exists, it is not possible to determine absolute breadth of this open class. These are non-trivial problems requiring some theory of superclasses, still under formulation.

Relative Breadth, Open Class: The solution to this case is similar to the absolute breadth, open class case. Where there are potentially infinitely many unique instances (open-class), it is sufficient to have seen at least one of every *class* of instance of the possible values in that area in order to have achieved breadth. That is, breadth with respect to open class values is calculated by using the conceptual parent of each filler as opposed to the filler itself in the same calculation as for closed class values. The parent of each filler is computed by determining the conceptual category that the unique values belong to. For another example, given a unique person name, we transform that name to the category of `human`. The same problems with correct classification still remain. For example, a given person may be subcategorized as a `military officer` or `government official`, each of these subcategories in turn are `human`. Thus determining closed-class "clusters" of the open-class values requires additional mechanics. The methods to perform this clustering accurately are currently under experimentation. After the clusters are determined, they will be treated in the same manner as the truly closed-class values to provide a threshold for the breadth decision.

Experimental Results

This section details some of the results of running the computations just described on a database. The database used contained almost 2,000 database records of information automatically extracted from texts reporting on terrorist activities in Latin America. These records were all created with a natural language text processing program, described elsewhere [Jacobs and Rau, 1990; Jacobs and Rau, 1993; Krupka *et al.*, 1991]. Using news stories as a source suggests that this work has the potential to operate on arbitrary and general knowledge, as well as specific databases. Figure 1 shows a sample message and template from this set.

DEV-MUC4-0351

BOGOTA, 18 AUG 89 (EFE) -- [TEXT] SENATOR LUIS CARLOS GALAN, LIBERAL PARTY PRESIDENTIAL HOPEFUL, WAS SHOT THIS EVENING WHEN HE WAS ABOUT TO GIVE A SPEECH AT MAIN SQUARE OF SOACHA, 15 KM SOUTH OF BOGOTA, IT WAS CONFIRMED BY POLICE AND HEALTH AUTHORITIES.

ACCORDING TO THE FIRST REPORTS, AT LEAST ONE MAN FIRED ON THE SENATOR FROM AMONG THOSE GATHERED. THE SENATOR IS CURRENTLY AT THE EMERGENCY ROOM OF A HOSPITAL IN BOSA, CLOSE TO SOACHA. TWO OTHER PERSONS WERE WOUNDED DURING THE ATTACK.

0. MESSAGE: ID	DEV-MUC3-0351
1. MESSAGE: TEMPLATE	1
2. INCIDENT: DATE	18 AUG 89
3. INCIDENT: LOCATION	COLOMBIA: SOACHA (CITY)
4. INCIDENT: TYPE	ATTACK
5. INCIDENT: STAGE OF EXECUTION	ACCOMPLISHED
6. INCIDENT: INSTRUMENT ID	-
7. INCIDENT: INSTRUMENT TYPE	GUN: "-"
8. PERP: INCIDENT CATEGORY	TERRORIST ACT
9. PERP: INDIVIDUAL ID	"AT LEAST ONE MAN" / "ONE MAN"
10. PERP: ORGANIZATION ID	-
11. PERP: ORGANIZATION CONFIDENCE	-
12. PHYS TGT: ID	•
13. PHYS TGT: TYPE	•
14. PHYS TGT: NUMBER	•
15. PHYS TGT: FOREIGN NATION	•
16. PHYS TGT: EFFECT OF INCIDENT	•
17. PHYS TGT: TOTAL NUMBER	•
18. HUM TGT: NAME	"LUIS CARLOS GALAN"
19. HUM TGT: DESCRIPTION HOPEFUL" /	"LIBERAL PARTY PRESIDENTIAL "SENATOR": "LUIS CARLOS GALAN" "TWO OTHER PERSONS"
20. HUM TGT: TYPE CARLOS	GOVERNMENT OFFICIAL: "LUIS GALAN" CIVILIAN: "TWO OTHER PERSONS"
21. HUM TGT: NUMBER	1: "LUIS CARLOS GALAN" 2: "TWO OTHER PERSONS"
22. HUM TGT: FOREIGN NATION	-
23. HUM TGT: EFFECT OF INCIDENT	INJURY: "LUIS CARLOS GALAN" INJURY: "TWO OTHER PERSONS"
24. HUM TGT: TOTAL NUMBER	

Figure 1: Example Text and Data Extracted

Closed/Open Class Determination

First, the fields in the records are divided into closed and open classes, according to gross frequency of occurrence of distinct values, as discussed previously. This results in the closed class slots illustrated in Figure 2.

Randomization of Data

After this determination, records are randomly selected from the database and the closed-class fillers are tabulated. Random selection is critical to overcome time-dependent patterns of values typical in any real database.

	Slot Type	Status	Total Distinct Values
3.	LOC OF INCIDENT	CLOSED	(23 distinct values)
4.	INCIDENT TYPE	CLOSED	(6 distinct values)
5.	INCIDENCT STAGE	CLOSED	(3 distinct values
7.	INSTRUMENT TYPE	CLOSED	(19 distinct values)
8.	INCIDENT CATEGORY	CLOSED	(4 distinct values)
11.	PERP ORG-CONF	CLOSED	(7 distinct values)
13.	PHYS TGT TYPE	CLOSED	(17 distinct values)
15.	PHYS TGT NATION	CLOSED	(15 distinct values)
16.	PHYS TGT EFFECT	CLOSED	(8 distinct values)
20.	HUMAN TGT TYPE	CLOSED	(12 distinct values)
22.	HUMAN TGT NATION	CLOSED	(29 distinct values)
23.	HUMAN TGT EFFECT	CLOSED	(11 distinct values)

Figure 2: Closed Class Categories

Threshold Determination

After experimentation, a threshold value T of 50 was chosen. This was the minimal value that captured almost every distinct type, with any types remaining being anomalous. This number could be larger and still include these same distinct values. Any arbitrary number, even if experimentally determined, is open to suspicions of "picking a number out of a hat". However if breadth is viewed as a binary predicate, then such an arbitrary threshold is required.

An alternative view of breadth is as a continuum, whereby you have more or less breadth, and it monotonically increases the more distinct types encountered. Under this view, the threshold can be a range, and it maps onto the breadth scale linearly. However for ease of exposition, I assume here the binary predicate model. This means that breadth is achieved when $d/n < 1/50$. The number of distinct types seen so far (`d`), the total number of fillers seen to this point (`n`), and the threshold (`50`) are the numbers that estimate the breadth. It is possible to graph `d` *vs.* `n` and see the point at which breadth is achieved, as well as the theoretical projection of where breadth would be achieved if no additional distinct values were seen past the point at which the data ends.

A graphical form of presentation does not capture the frequency with which each distinct value occurs in the sample. In all cases where a cutoff (breadth determination) was made when not all distinct values had been seen, the remaining distinct values were anomalous. Anomalous data is data that occurs only once or twice and is almost always a result of mistyping or other errors. This is to be predicted, in that any values not seen after $T \times$ `d`, where `d` is the number of distinct values, is likely not to occur with frequency greater than 1 in that amount. For example, suppose as is true in the data presented here, that a breadth determination for the field of `Physical-Target-Type` was made after seeing 1134 instances. In this case, $n = 1134$, $d = 17$ and no additional distinct fill exists in the remainder of this set of data. As such, breadth is still achieved if only truly anomalous data has not been seen. If a significant portion of the data is anomalous, the system would not conclude breadth in the area, as the prediction is that more unseen values would be likely to occur.

3.	LOCATION OF INCIDENT	546 410 133 88 39 32 15 11 7 5 3 2 2 2 2 1 1 1 1 1 1 1 1(no breadth)
4.	INCIDENT TYPE	734 349 145 51 20 7 (0)
5.	INCIDENCT STAGE	1204 60 48 (0)
7.	INSTRUMENT TYPE	550 194 192 145 79 78 50 38 14 12 11 10 8 6 4 3 2 (2 1)
8.	INCIDENT CATEGORY	967 219 130 1 (0)
11.	PERPETRATOR ORG-CONFIDENCE	617 353 148 103 102 80 (2)
13.	PHYSICAL TGT TYPE	580 160 133 121 118 114 77 58 46 42 41 31 23 17 17 1 1 (0)
15.	PHYSICAL TGT FOREIGN NATION	1092 162 25 11 6 5 5 3 2 1 1 1 1 1 (1)
16.	PHYSICAL TGT EFFECT	765 346 162 156 45 12 7 (2)
20.	HUMAN TGT TYPE	1019 382 165 143 99 61 39 37 17 14 13 (1)
22.	HUMAN TGT FOREIGN NATION	1179 49 32 10 10 9 8 6 5 5 4 4 4 4 3 2 2 2 2 2 2 2 1 1 1 1 1 1 1 (no breadth)
23.	HUMAN TGT EFFECT	856 457 280 122 65 31 10 10 7 5 1 (0)

Figure 3: Frequency Distribution of Values

Breadth Determination

Figure 3 illustrates the frequency with which each distinct value occurs. For example, there were 734 instances of the first distinct value in the **Type** field, 349 instances of the next distinct value, etc. Numbers in parenthesis indicate those anomalous values that are missed by an early cutoff; **no breadth** indicates that a likely breadth determination cannot be made. Note, for example, how the preponderance of singular values in the **human-target-foreign-nation** slot precludes a breadth determination, as the expectation here is towards more such values. When low frequency values contribute significantly towards the total number of distinct values, breadth of knowledge must include even these low frequency values as well. As can be seen from these data, it is possible to answer the question "Do you have breadth?" applied to the fields in a database. Moreover this question can be answered so that in only some of the fields, only a small fraction of the total number of distinct values is unseen; these values are likely to be anomalous at the threshold value of 50. Breadth determinations even more likely to cover 100% of all values can be made at larger thresholds. With these results, a user may get different answers to questions posed to a database where the field referenced had complete breadth of knowledge from one which did not have breadth. For example, all data in which every member of a field participates in some relationship must be qualified as to whether the relationship hold for all-currently-known, as opposed to all-possible.

Probabilistic Analysis

A probabilistic analysis shows that after 50 trials, if only 1 distinct value has surfaced, the chances of seeing a new value next (i.e., on the 51th trial) are:

$$\left(\frac{50}{51}\right)^{51}\left(\frac{1}{51}\right) \simeq .007$$

under the assumption that the *a posteriori* probability equals the *a priori* probability; a good assumption with large n. This only decreases for other values of **d**, so that if 2 distinct values have occurred in $50 \times 2 = 100$ trials, the likelihood of seeing a pattern consisting of these 2 values in an arbitrary internal distribution, followed by a new value, is .003.

The best case scenario is that all distinct values have been seen. And although taking the pattern of historical values to predict the probability during the next trial assumes a 1/50 chance the next value will be distinct (in the worst case of having seen only one distinct value), there is only a .007 chance that with this weighting we would have seen the pattern just seen, thus lending confidence that any distinct values yet unseen occur with smaller probabilities. This analysis indicates that after this many values, almost all, if not all, distinct values have been seen with good certainty, and any values not seen are likely to be anomalous or extremely infrequent.

The results presented in this section indicate that fairly simple calculations can accurately predict when a database or knowledge base has breadth in closed-class subject areas.

Related Work

A great deal of research has addressed the problem of what a system might know or believe [Halpern, 1986; Vardi, 1988]. This work primarily concerns itself with calculating what a system might be expected to believe, given some set of assertions. The work described here contributes to that body of research by adding a new metric that is calculated from what is known, the breadth of knowledge. This work is related to recent work in the area of knowledge discovery in databases [Piatetsky-Shapiro and Frawley, 1991] that attempts to learn new knowledge from the structure and content of databases. However, the particular problem of computing breadth of knowledge has not been directly addressed in this new research area.

Future Directions

The primary area for future work is in the development and testing of the theory of superclasses to determine breadth in the open-class cases. Another critical step currently on-going is the choice of application area to illustrate the utility of knowing when breadth has been achieved in a real AI/database system. Also of importance is automating the process of threshold determination. This value appears to vary as a function of the raw number of distinct fills as well as the inherent frequency distribution of a given field. Even more accurate determinations may be made if the threshold can be dynamically computed for each slot. Another activity currently on-going is embedding the breadth calculation in a suite of tools under development to construct a meta-profile of the contents of an arbitrary knowledge or database. With this profile, users of an application will be able to get a feeling for the contents of a database that can aid in their judgements of the appropriateness of the database for their information needs, as well as in constructing appropriate and answerable queries. In addition, the system will be able to distinguish negative searches from the result of information requests outside an area of expertise from responses due to closed world assumptions.

Conclusions

This paper began with an analysis of the notion of conceptual breadth, with respect to absolute and relative measures, and with respect to closed class and open class categories. A specific computation was detailed for automatically determining the conceptual breadth of a knowledge or database in the case of computing relative breadth of closed class fields. The point at which breadth is achieved is computed by taking a random sample from the database, and keeping track of the percentage of distinct types to the total number of fillers seen. When this percent falls below a threshold, it is expected that most if not all the distinct types have been seen. This method accurately identifies a point at which breadth is very likely to have been achieved, and a probabilistic analysis supports this expectation.

This work is important not only for the methods and computations described here, but for investigating new questions we would like large knowledge based system to be able to answer - questions such as "what do you know?" and "how complete is your knowledge?". It is critical in determining when to apply the closed-world hypothesis, and when there is not enough information for this assumption to apply. Looking at areas traditionally reserved for the purely cognitive realm, such as meta-questions of knowledge scope and extent, offers a new perspective from which to develop computational answers.

Acknowledgements. The results presented here have benefited from discussions with my colleagues at GE, including Piero Bonissone, Fred Faltin, Paul Jacobs and Barbara Vivier. I am also grateful for the perspectives and suggestions of Noel Sharkey (University of Exeter), and Gregory Piatetsky-Shapiro (GTE Labs). David McDonald also provided comments on a draft of this paper.

References

[Collins *et al.*, 1975] Allan Collins, Eleanor H. Warnock, Nelleke Aiello, and Mark L. Miller. Reasoning from incomplete knowledge. In Daniel G. Bobrow and Allan Collins, editors, *Representation and Understanding*, pages 383–415. Academic Press, Inc., Orlando, FL, 1975.

[Halpern, 1986] J. Y. Halpern, editor. *Proceedings of the First Conference on Theoretical Aspects of Reasoning About Knowledge.* Morgan Maufmann, Los Altos, CA, March 1986.

[Jacobs and Rau, 1990] P. S. Jacobs and L. F. Rau. The GE NLToolset: A software foundation for intelligent text processing. In *Proceedings of the Thirteenth International Conference on Computational Linguistics*, volume 3, pages 373–377, Helsinki, Finland, 1990.

[Jacobs and Rau, 1993] P. S. Jacobs and L. F. Rau. Innovations in text interpretation. *Artificial Intelligence (Special Issue on Natural Language Processing)*, 48, To Appear 1993.

[Krupka *et al.*, 1991] G. R. Krupka, P. S. Jacobs, L. F. Rau, and L. Iwańska. Description of the GE NLToolset system as used for MUC-3. In *Proceedings of the Third Message Understanding Conference (Muc-3)*, San Mateo, Ca, May 1991. Morgan Kaufmann Publishers.

[Piatetsky-Shapiro and Frawley, 1991] G. Piatetsky-Shapiro and W. J. Frawley. *Knowledge Discovery in Databases.* MIT Press, Cambridge, MA, 1991.

[Pollack, 1986] M. Pollack. *Simulating Plan Failure.* PhD thesis, University of Pennsylvania, 1986.

[Rau, 1992] L. F. Rau. Calculating salience of knowledge. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, Hillsdale, NJ, July 1992. Cognitive Science Society, Lawrence Erlbaum Associates.

[Reiter, 1978] R. Reiter. On closed world data bases. In H. Gallaire and J. Minker, editors, *Logic and Data Bases*, pages 55–76. Plenum Press, New York, New York, 1978.

[Vardi, 1988] M. Y. Vardi, editor. *Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge.* Morgan Maufmann, Los Altos, CA, March 1988.