

## **Toward a Relation Hierarchy for Information Retrieval**

**Sung H. Myaeng**

School of Information Studies  
Syracuse University  
Syracuse, NY 13244  
shmyaeng@mailbox.syr.edu

**Michael L. McHale**

Rome Laboratory  
Griffiss AFB, NY 13441-5700  
mlmchale@sunrise.acs.syr.edu

### **1. INTRODUCTION**

Natural language text can be seen as a symbolic representation of a cognitive state of the creator that comprises concepts and the relations among those concepts. Much work has been done in Information Science, especially within Information Retrieval (IR), concerning the handling of concepts, most notably in the form of keywords. Much less effort has been spent toward the understanding and handling of the semantic relations that contextually bind concepts together. While it has been shown (Wang, et al., 1985) that the use of these semantic relations for query enhancement can increase retrieval effectiveness, the proper handling of semantic relations has a much wider application than just query enhancement. Once relations inherent in text are identified and captured, they can be used to provide contextual information to the concepts in the representations of the text, which otherwise would be treated as if they were independent and separate.

As part of an effort to develop IR systems that exploit such semantic relations existing in documents and information need statements, the research described in this paper is concerned with how those relations could be classified for IR purposes. We are particularly interested in grouping relations and constructing a hierarchy of relations of a homogeneous type. Since ideas and information are expressed at different levels of abstraction in documents and information-need statements, for example, it is obvious that relations classified in a hierarchical manner would prove to be of utmost value. Without such a hierarchy, no structural matching between two rich representations would be sensible.

Since, in our mind, IR is the major application of such a hierarchy, there are several issues that need to be investigated. More specifically, we have the following goals:

- discovering relations that are ontologically appropriate and suitable for IR;
- classifying relations into different groups in such a way that relations in a group are homogeneous enough to be placed in a hierarchy;
- organizing relations into a hierarchy that facilitates matching between document and query representations and, at the same time, proves meaningful for users; and
- devising ways of automatically identifying relations from natural language text.

In effect, the first and the last goals serve as constraints that need to be satisfied by the resulting relation hierarchies.

Accordingly, we have been exploring various approaches to creating a relational schema suitable for IR, which essentially defines the meaning of a set of particular relations useful for IR and their inter-relationships or classificatory structure. In order to be suitable, a schema must be flexible enough to be applicable across varied databases, rich enough to enhance retrieval performance, and intuitive enough that users find it meaningful and comfortable to use in search processes.

As a first step, we have been investigating and comparing various existing schemas for relations that have been developed for different purposes and contexts, including Fillmore (1968), Farradane (1980), Wang et al. (1985), Chaffin & Herrmann (1987), Sowa (1984), Dick (1991), and Sabah and Vilnat (1991). The next section outlines each of these in some detail and their respective value and relevance to our goal of constructing a relation hierarchy for IR. The following section describes our current work, with a report on some results drawn from our attempt to understand the types of knowledge required to identify relations from natural language sentences as we try to construct a relation hierarchy.

## 2. EXAMINATION OF EXISTING RELATION SCHEMAS

While there are many ways relational schemas can be categorized based on a number of factors such as psychological reality, procedures by which they are discovered, syntagmatic vs. paradigmatic, or semantic vs. lexical (Evens, 1988), we view existing relation schemas as belonging to one of two different approaches. The first can be characterized as a **pragmatic** approach, since world knowledge is often required to assign a relation between concepts. The other kind of approach can be seen as **linguistic** in nature in that the main motivation behind it is to understand the nature of relations that connect linguistic constituents and thus reveal semantic structure of the sentences. While this distinction is discernible, it should be noted that some pragmatic schemas include linguistically oriented relations.

### 2.1. Pragmatic Approaches

**Relational Indexing.** Farradane (1980) devised a method of systematically indexing the relations between indexing terms. As shown in Figure 1, this relational indexing scheme consists of a 3x3 matrix of relations along Associative and Discriminatory dimensions. According to Farradane these nine relations have been found in practice to be necessary and sufficient for indexing purposes. Furthermore, it is argued that the schema is capable of representing different levels of relations.

		Associative Mechanisms		
		Awareness	Temporary Association	Fixed Association
Discriminatory Mechanisms	Concurrent Conceptualization	Concurrence	Self-Activity	Association
	Not-Distinct Conceptualization	Equivalence	Dimensional	Appurtenance
	Distinct Conceptualization	Distinctness	Action	Causation

**Fig. 1. Farradane's Relations**

One problem with Farradane's relations is a large discrepancy in the granularity of relationships. For instance, the Equivalence relation implies either a synonym or an equivalent expression, and thus it is usable only for a very few items. The Association relation, on the other hand, can be used to express any unspecified relation and thus is a sort of "catch- all." This is a problem for two reasons. The first is that Association will be used more than any other relation, in effect making the other relations "second-class". The second problem is the obvious effect that if we remove the Equivalence relation we could still classify all the relations because we could use Association for the unspecified equivalence. It follows directly from this that if a schema has a catch-all relation such as Association then that single relation is the only one that is necessary, and just by itself it is sufficient. In spite of the difference in specificity levels, the relations are only laid out in a matrix, as opposed to a hierarchy.

Another problem is that the relations as they are, are ambiguous. For example, Concurrence relation is used to refer to two unrelated situations: one between terms that co-occur regularly in text and the other used between those terms that have a mutual temporal relationship. This type of ambiguity is detrimental to retrieval precision. For example, if the query refers to simultaneous events, then documents which are retrieved because the terms are simply co-occurring in them are irrelevant to the query.

**Relational Thesauri.** Wang, Vandendorpe and Evens (1985) describe a series of experiments designed to investigate the effect of using relational thesauri for query enhancement

on retrieval effectiveness. They used 43 relations clustered into five different groups (i.e., Part/Whole, Collocation, Paradigmatic, Taxonomy/Synonymy and Antonymy) and found a significant improvement in retrieval effectiveness especially when queries were ill-defined or lacked normal index terms. The only group of the five that didn't show improvement was the Antonymy relation, and that actually degraded performance.

Their schema is in fact a two-level hierarchy of relations. Since their goal was to determine how (through what relations) new terms should be introduced to expand queries, however, the hierarchical structure of the relations was of little use; the schema was only helpful in concluding that the relation of Antonymy didn't contribute to any performance improvement. Furthermore, the type of relation existing between a query term and a newly introduced term is of little significance as far as the retrieval process is concerned. After all, the relations were not used at all in representations of documents; they were simply treated as two terms in the query.

It is interesting to observe that the Wang, Vandendorpe and EVens schema mixes some syntactically oriented relations (e.g., perfect-infinitive, adjective-noun) with both semantic (e.g., synonym, count mass) and pragmatic relations (e.g., part-whole, reciprocal kinship), resulting in tangled hierarchies. This heterogeneous mixture would probably prove confusing for users, but Wang, et al. were strictly interested in using the schema to enhance query statements and it proved useful in doing that.

**Relation Element Theory.** Chaffin and Herrmann (1987) empirically constructed a taxonomy of relations by having subjects sort examples of the relations. The 31 relations used came from various sources in the psychological and linguistic literature and were organized into five families (i.e., Contrasts, Similar, Class Inclusion, Case Relations, and Part-Whole). This *a priori* classification of relations was found to be compatible with the hierarchical clustering done by the subjects. The intent of the study was to demonstrate that:

The ability of subjects to compare relations for similarity suggests that they decomposed the relations into elements or aspects in which they are the same and elements or aspects in which they differ (p. 226).

In other words, the 31 relations themselves are decomposable into more basic elements. The basic elements they used were taken from Stasio, et al. (1985) and include such items as Bipolarity (items on opposite sides of a midpoint) and Dimension (sharing of a single dimension). Thus Incompatibility, one of their relations (as in *frank - hypocritical*), can be seen as Dimensional bipolarity (written as DIM(BIP)).

This work goes a long way in developing a hierarchy that is meaningful for users. What isn't clear is whether the relations in the hierarchy are sufficiently specific for Information Retrieval. For instance, in a given application, it may be beneficial to differentiate between TOOLS and METHODS, but in their scheme, both fall into the INSTRUMENT relation. If there were another level (or more) of relations then the differentiation could be made when it was helpful to do so. Another question that arises is whether all the relations in the hierarchy are necessary and sufficient to meet users' needs and the need to represent naturally occurring relations that bind together independent concepts in text.

## 2.2. Linguistic Approaches

**Case Grammar.** Fillmore's Case Grammar (1968) and other subsequent works (Cook, 1989) have enjoyed much popularity with computational linguists and others building systems that process natural language texts. Hirst gives the following summary of Case Grammar.

In its most basic form, case theory views a sentence as an assertion whose predicate is denoted by the verb of the sentence and whose arguments are denoted by the noun phrases (Hirst 1987).

The five or so Cases represent semantic primitive relations joining the concepts of a sentence. For instance, in the sentence

John loves Mary.

John has the AGENT relation to the act of loving and Mary the EXPERIENCER relation. Fillmore's Cases consist of very simple primitives and thus capture practically all relations in some sense. In fact, Cook (1989) argues that five of the Cases (i.e., AGENT, EXPERIENCER, BENEFACTIVE, OBJECT and LOCATIVE) have proven to be both necessary and sufficient in relating concepts to the verb in a sentence. As support for this statement Cook gives a mapping of Peter Roget's classification schema (Roget, 1977) (which was designed to cover the semantic world) to Cases (Cook 1989: 190).

However, since these Cases are very general in nature, associating them to detailed semantic relations and constructing a hierarchy is neither straightforward nor unambiguous. Also, there is no guarantee that the resulting hierarchy would be meaningful for users. It is expected that a combination of Case relations and concepts have to define different detailed relations in different contexts and that the result will form multiple hierarchies built with the common base or children in a tree.

**Relations in Conceptual Graphs.** Sowa's Conceptual Graphs (1984) were inspired by Pierce's Existential Graphs, and as such, are a combination of semantic networks and predicate calculus. One of the main motivations for the creation of Conceptual Graphs was the desire to have a useful knowledge representation for natural language sentences, providing relatively easy mapping between natural language sentences and Conceptual Graphs representations. A sentence would be represented as a graph consisting of concept nodes and relation nodes.

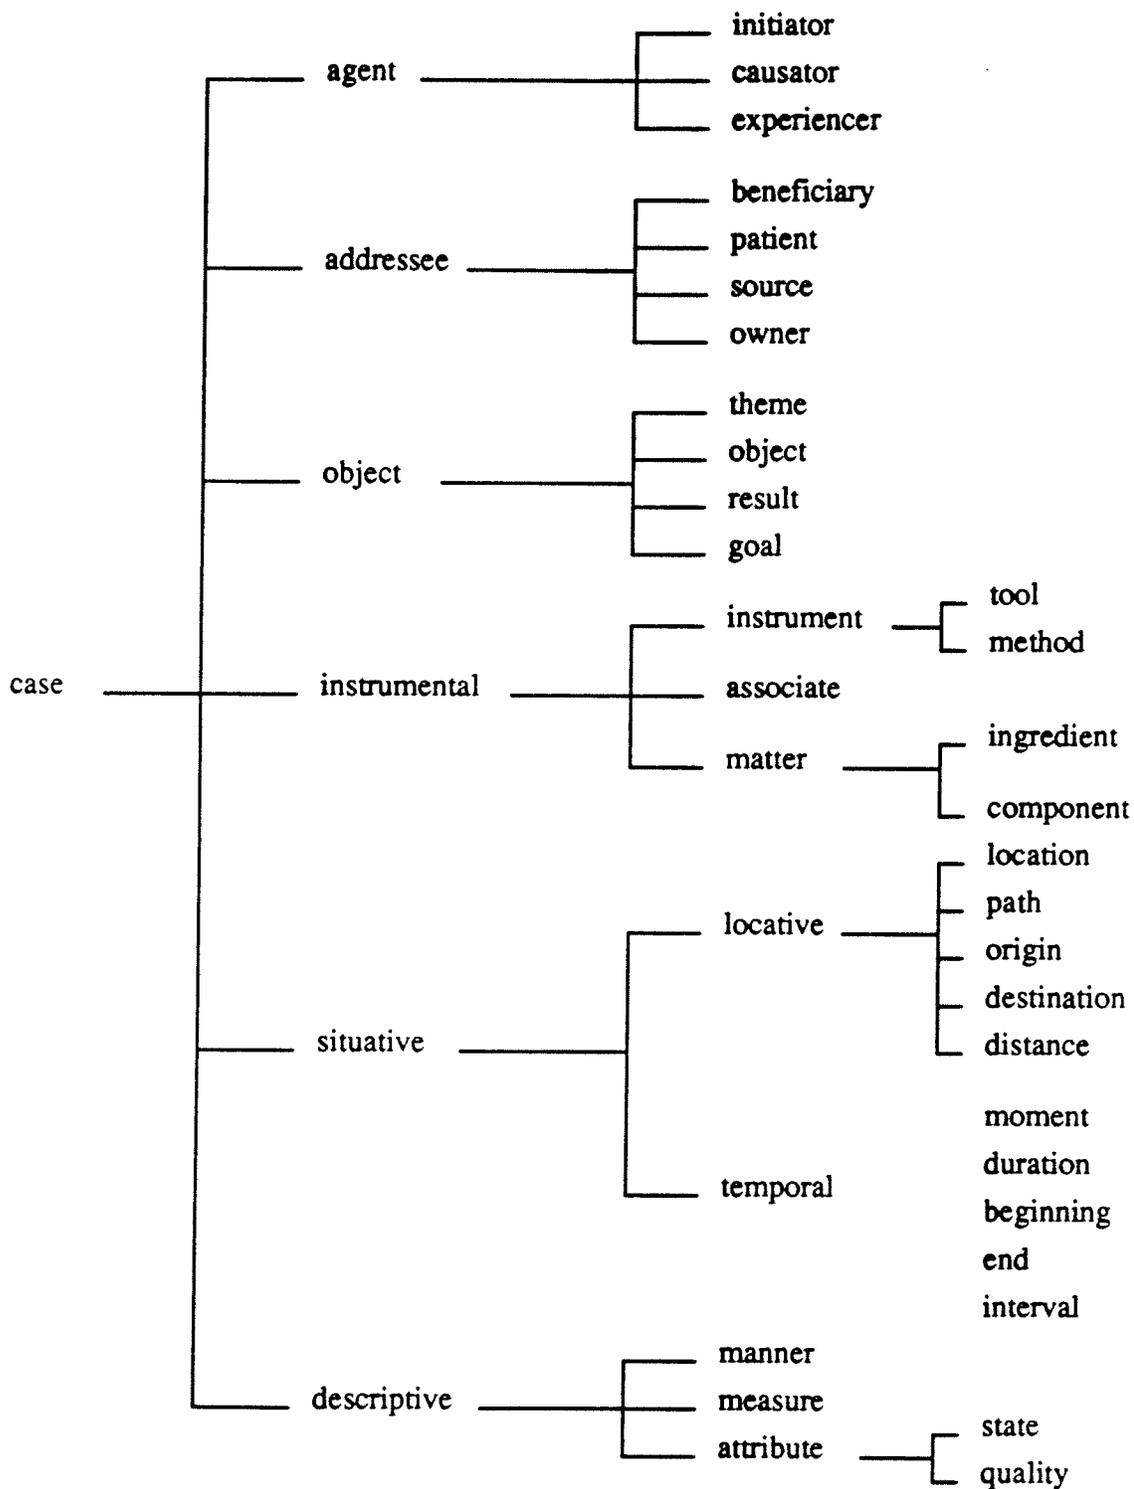
While the formalism puts considerable emphasis on relations, Sowa doesn't specify which relations to use; he introduced in his book (Appendix B.3) a catalogue of 37 relations as an example of the way relations can be formed. These relations include all of the Cases listed under Fillmore above and many more at different levels of granularity. The basic primitive relation in this framework is LINK from which all other relations can be defined (i.e., LINK is necessary and sufficient) as a Conceptual Graph. For instance, the AGENT relation can be defined as the LINK between a concept for an action and a concept for an animate object. In this way, relations with any level of specificity can be introduced in terms of existing relations and concepts and used as an

abstraction of the definition graph. While the notion of concept hierarchy is discussed and used heavily in the framework, the relation hierarchy is introduced only as a possibility and left completely unspecified.

**Relations in Legal Document Representation.** Working within the Conceptual Graph framework, Dick (1991) investigated the use of Case Grammar for IR purposes as a way of representing argument structure in legal documents. Her particular choice of Cases was based on Somers (1987). Somers approach is intensely verb-based and consists of a two-dimensional Case grid. Along one axis are the Cases ACTIVE, OBJECTIVE, DATIVE, LOCATIVE, TEMPORAL and AMBIENT. SOURCE, PATH, GOAL and LOCAL are on the other axis. In the two sentences, "John sold a house to Mary" and "Mary bought a house from John," for example, both John and Mary are agents and categorized to have ACTIVE cases. However, a finer distinction can be made between the roles played by John and Mary; while John has ACTIVE SOURCE case (animate instigator of action) as the agent of "sell," Mary has ACTIVE GOAL case (animate active recipient) as the agent of "buy." Dick found the Case grid to be generally usable within her domain (legal reasoning), but often it lacked the degree of specificity required. In those instances she enhanced the Cases as required for the domain by successfully combining two different relation schemas (Somers' and Sowa's).

Unlike the work done by Wang, et al., the intent of this work is to use the verb-centered relations as the basis for rich representation of the text. With the ability to perform inferences, this representation would facilitate conceptual retrieval of legal documents. However, one limitation with this approach is that while the relations used have a strong linguistic justification, they lack in empirical justification in IR. It's not clear whether the relations have the right level of specificity and exhaustivity for IR and whether they are compatible with the users' cognitive demand. The use of a flat, as apposed to hierarchical, relational structure was acceptable for Dick's purpose, but has the disadvantage for our purposes of not considering the possible ontological mismatch between users' information need representation and document representation.

**Hierarchy of Case Relations.** Also working in the Conceptual Graph framework, Sabah and Vilnat (1991) have developed a relational hierarchy that uses Case Grammar. As shown in Figure 2, the top of the hierarchy is simply CASE, and at the next level are AGENT, ADDRESSEE, OBJECT, INSTRUMENTAL, SITUATIVE, and DESCRIPTIVE. These Cases can be further refined as needed for increased specificity. For instance, the INSTRUMENTAL CASE can be refined as either INSTRUMENT, ASSOCIATE or MATTER. INSTRUMENT can be further refined into TOOL or METHOD. The lowest level contains about 40 Cases.



**Fig. 2. Sabah & Vilnat's Relation Hierarchy**

Their hierarchy was developed by analysis of a novel and from linguistic work on French prepositions. It was constructed in a bottom-up fashion in that all the necessary detailed relations were first determined and then grouped to form the next level. With the main goal of obtaining a semantic representation of sentences using Conceptual Graphs, this hierarchy allowed for joining two partial Conceptual Graphs derived from a sentence that may contain different yet compatible relations, thereby forming a complete Conceptual Graph for the sentence.

There seems to be very little concern in their work for the naturalness of the terms used for the grouped Cases. OBJECT, for instance, refines as OBJECT, THEME, RESULT or GOAL. The use of the same term at different levels is confusing. Also there seems to be some arbitrariness in both terms and in grouping. The former may be more a problem with translation from the original French but the latter isn't so easily dismissed (e.g., referring to DISTANCE as a LOCATIVE). Nonetheless, the bottom-up, text-driven approach to developing a Case hierarchy is a valid one on pragmatic grounds and expected to prove useful in our work.

### 2.3. Summary and Problems

A problem that surfaces for all three of the schemas in the category of pragmatic approaches is related to the desired ability to automatically extract the relations from text. Since many of the relations that are used are pragmatic in nature, they require either fairly large amounts of world knowledge or rather extensive inferencing to determine (for instance, that a *rake* and a *fork* have similar attributes or that *food* is often *tasty*). It is possible to generate taxonomies of concepts to aid in this type of inferencing, but in general, this in itself remains a very difficult problem for automatic methods.

While linguistically oriented relations are easier to construct and hence more likely to be identifiable and extractable from natural language text, they are not without any problems. It is not clear how useful the actual Case hierarchy or a set of Case relations would be for IR. The level of specificity of the relations may not be appropriate for both users and the final text representations since they are all simply Case relations rather than other semantic relations, such as, for instance, "cause -- effect.". Linguistically driven Case relations themselves may not be the ones users would recognize from text or use to express their information needs naturally. In addition, the fact that Case relations are verb-centered implies that a relation between any two nouns, for example, is not a Case and needs to be inferred or determined in a different way. This in turn implies the need for relations other than Cases for IR purposes.

## 3. CONSTRUCTING A RELATION HIERARCHY FOR IR

Our goal is to create a hierarchy of relations that is useful for retrieval, extractable from text in a principled way, and natural for users. From the linguistic approaches above, we have seen that by starting with Case relations, it may be possible to construct hierarchies that are usable in an information processing system. Our intention is to produce not only a hierarchy useful for the retrieval process but to also have one that is not a cognitive load on users. As such, our concern is not so much in the mechanical procedure *per se*, by which relations can be grouped and a hierarchy

is constructed, as in the degree to which the resulting groupings and the hierarchy of relations satisfy activities necessary in IR and in the degree to which the relations in the hierarchy can be recognized and extracted from text by automatic means.

The hierarchical structure would be useful at least in two specific ways. One is that even if two relations are different at the surface level (e.g. INSTRUMENT and TOOL), they can match as long as they have a hierarchical relationship. A less obvious reason is that when users are asked to define the relation between two concepts existing in their queries, the hierarchy could help them decide the exact relation in the same way a thesaurus assists users in formulating queries.

### 3.1. Our Approach

We take both top-down and bottom-up approaches. As indicated earlier, the top-down approach involves our examination of existing relation schemas to see how relevant they are for constructing our own hierarchy. This not only provides us with a starting point as to what relations we need to consider and look after when the bottom-up approach is taken, but also gives an insight on what categories of relations exist and how relations could be grouped together. In other words, we can easily produce our own schema just based on our intellectual analysis of the existing ones and on our intuitive understanding of what IR processes entail. This schema is then validated and verified by the bottom-up approach which involves actual text, people, and retrieval tasks.

A complementary approach we are taking at the same time is driven bottom-up by actual text. We are going through an intellectual analysis of abstracts and query statements in an IR test collection to see what relations, both implicit and explicit, are used in the text and can be recognized. By analyzing the two different types of text, we are attempting to understand whether there is any discernible difference between producers and consumers of text in terms of types of relations and their classificatory structure.

While this process is necessarily guided and influenced by our knowledge of various schemas acquired through the top-down task, care must be taken to ensure that we do not limit ourselves to the existing schemas and go beyond what's included in them if necessary (i.e., if they are not sufficient). For that purpose, we are using dictionary definitions as a cross-check, as illustrated in the following sub-section. Since each of the senses associated with a particular lexical item in the dictionary can reveal a relation that might exist in the text segment that includes the lexical item, by including these relations we achieve the same effect as adding relations to those in the existing relation schemas. An important consideration being made is to what extent the relation recognition process can be automated.

To that effect, we are using Sowa's Conceptual Graphs for the representation framework of our IR system. This allows us an easy way of using those linguistic relations that we find and also provides a sound way of handling pragmatic relations when we can extract them. An effective connection between linguistically based relations and pragmatically oriented relations can be accomplished in the Conceptual Graph framework.

The relations resulting from the above processes will be classified into groups to form a schema and examined by people to ensure it matches the cognitive structure of people who are engaged in

the process of retrieving information. We plan to use the Repertory Grid (Shaw & Gaines, 1986) as a methodological tool by which distinctions and similarities of relations as viewed by users are systematically recorded and organized not only to produce a hierarchy but also to reveal the set of discriminatory dimensions being used.

### 3.2. Preliminary Result

As a first step toward understanding what relations should be included in the hierarchy and could be automatically extractable, we analyzed the various relations possible for the preposition FOR. This analysis gave us an insight into what kind of knowledge would be necessary for automatically identifying relations that are classifiable into a hierarchy. It also provided us with evidence that existing relation schemas are probably inadequate for IR purposes.

We chose a preposition as our initial focus because most work with Case Grammar allows only the verbs to have Case frames. Thus a verb such as “study” would have the frame [AGENT, OBJECT]. These frames are context free, that is, they are not modified in any way by either the agent or the object. However, this argument becomes clouded when we try to see what relations exist among concepts in the following example sentences:

- 1) I studied for the exam.
- 2) I studied for five hours.

In the first sentence there is an object, but there isn't one in the second sentence because “for five hours” is a modifier (i.e., optional) not an object. This doesn't seem to be an attribute of the verb but rather an effect of the various meanings of the preposition FOR, which are in turn determined by the context in which they occur. It can be seen in the two example sentences above that it is the nature of the noun phrases occurring after FOR that decides between purpose and duration. There are as many as 19 senses for the preposition FOR (25 including phrasal uses) in *Longman's Dictionary of Contemporary English* (LDOCE), and these are very much context dependent, relying on the nature of the cooccurring noun phrases as well as on the verbs.

An empirical question arises whether it is important to differentiate among these relations and to what extent they occur in an IR context. To better understand this question we looked at the uses of FOR in a test collection used for a variety of IR experiments, which consists of the abstracts of the Communications of the Association of Computing Machinery (CACM) and a set of queries together with predetermined relevance judgments on the abstracts. Within the natural language query statements produced by actual users, we found only three of the senses for FOR (i.e., (A) “intended to belong to” or “be given to,” (B) “to show purpose,” and (C) “in order to improve the condition of”). Although there is no statistical evidence for this pattern at this point, this encouraging result indicates that this data-driven approach would prove useful in identifying relations meaningful for IR processes.

The task then becomes that of finding a method to automatically extract these meanings from text. From the data set we analyzed, the following patterns were found.

### **Meaning Pattern**

- (A) for + noun (state or concrete noun)
- (B) for + verb or nominalized verb (action verbs)
- (C) for + achieving, containing, handling... (statives)

Since this type of information on semantic restrictions of nouns and verbs is available in LDOCE it would indicate that the correct selection of meaning for FOR can be automatically accomplished within this database. This is significant because lexical ambiguity is a very difficult problem in general. If for any given database we can start with the prepositions and determine their relational roles by the selectional restrictions of the nouns, which in turn limit the senses of the verb, then we've come a long way toward a pragmatic solution of the problem. Outside of restricted databases, however, the solution is much more difficult. As can be seen in the following sentences, it is clear that some additional information needs to be brought to bear in analyzing text in areas that have not been encountered.

- (3) I brought a friend home for lunch.
- (4) I brought a chicken home for lunch.

To ascertain how much of a problem it would be to recognize additional meanings, we looked at the usage of FOR in search statements prepared for a database whose domain is in education, which were collected as part of the ASK project (Oddy, et al., 1991). These statements contain a number of uses (72 occurrences, to be exact) of FOR, but only 40 percent of those are of the three senses identified from the aforementioned query statements in the CACM collection. In terms of the senses in LDOCE, the most common is still "intended to belong to," but the second most common was for "duration of time." This additional richness would be nice to exploit but the complexity of the patterns involved has made developing an algorithm for extraction difficult. However, one pattern that emerged was that the use of the time sense of FOR was always in a section dealing with background information on the ASK statements and not in the actual topic-related part of the statement. Collapsing of senses in the background sections may be innocuous. In order to determine if the statement is part of the background it will be necessary to automatically extract discourse structure as attempted in Oddy, et al (1991). This remains for further work.

One final, yet important, point that can be drawn from our initial data analysis is that none of the relation schemas we examined so far contain relations that correspond faithfully to the three most commonly occurring senses for FOR. The closest were BENEFICIARY Case in Sabah and Vilnat (1991) and RECIPIENT relation in Sowa (1984). In other words, since 19 out of 64 query statements in the CACM collection contain one of the three relations, a great portion of the query representations that rely on the existing relation schemas only would be either inaccurate or incomplete.

## **4. SUMMARY AND DISCUSSION**

We have described our ongoing work on constructing a relational hierarchy to be used for Information Retrieval purposes. For an IR system to utilize semantic information, i.e. concepts and relations, that exists in both documents and users' information need statements, a set of relations meaningful to the users and useful for IR processes needs to be identified and organized in such a

way that the semantic information is maximally used. As an effective way to deal with the disparity in terms of levels of abstraction between document representations and information need statements, relations need to be organized in a hierarchical manner. Based on our examination, we feel that existing relation schemas are not adequate for IR and that much work needs to be done.

As an attempt to exemplify our overall approach to the task, we presented our work with FOR and some preliminary results, which have given us some insight into why existing relation schemas are inadequate and, at the same time, how much and what types of knowledge would be needed to extract relations from text. It also shed some light on how a connection can be made between linguistically oriented relations and pragmatic relations, both of which need to be in a hierarchy.

Our work with the relation hierarchy is just one component of our overall goal of making an Intelligent Information Retrieval system. A main difficulty here is that Information Retrieval is concerned with vast amounts of data on heterogeneous data bases. If semantic relations are ever to be truly useful for Information Retrieval, we must be able to process these large amounts of text. In order to avoid the need to employ full-scale Natural Language Processing systems that in turn often require manual encoding of extensive lexicons and thesauri, we need to focus on more inexpensive ways of extracting such semantic information, perhaps relying on existing resources such as machine readable dictionaries and thesauri, such as *Longman's Dictionary of Contemporary English* (Longman 1987) and *Roget's International Thesaurus* (Roget 1977), which explicitly represent many of the relations used in the above schemas (e.g., synonyms, antonyms, taxonomic relations).

## 5. REFERENCES

- Chaffin, R. and D.J. Herrmann (1987) "Relation Element Theory: A New Account of the Representation and Processing of Semantic Relations," in D. Gorfein and R. Hoffman, (Eds.), *Memory and Learning: The Ebbinghaus Centennial Conference*, Lawrence Erlbaum, Hillsdale, NJ.
- Cook, W. A. (1989). *Case Grammar Theory*, Georgetown University Press, Washington, D.C.
- Dick, J.P. (1991) "On the usefulness of conceptual graphs in representing knowledge for intelligent retrieval". *Proceedings of the Sixth Annual Workshop on Conceptual Graphs*. July 1991, Binghamton, NY.
- Evens, M. (1988). *Relational Models of the Lexicon: Representing Knowledge in Semantic Networks*. Cambridge University Press, Cambridge.
- Farradane, J. (1980) "Relational Indexing." *J. of Information Science* 1 (5,6), 267-276, 313-324.
- Fillmore, C. J. (1968) "The Case for Case." In Bach, Emmon Werner and Harms, Robert Thomas (Eds.). *Universals in Linguistic Theory*, New York: Holt, Rinehart & Winston.
- Hirst, G. (1987) *Semantic interpretation and the resolution of ambiguity*. Cambridge: Cambridge University Press.
- Longman (1987) *Longman Dictionary of Contemporary English*. Harlow, UK: Longman.
- Roget (1977) *Roget's International Thesaurus*. New York: Harper & Row.
- Sabah, G. and A. Vilnat (1991) "Flexible Case Structure Implemented into a Deterministic Parser," in *Proceedings of the Sixth Annual Workshop on Conceptual Graphs*: Binghamton, NY.

- Shaw & Gaines (1988) "An interactive knowledge elicitation technique using personal construct technology." In Kidd, A., Ed., *Knowledge Elicitation for Expert Systems: A Practical Handbook*. Plenum Press.
- Somers, H.L. (1987) *Valency and Case in Computational Linguistics*. Edinburgh: Edinburgh University Press.
- Sowa, J.F. (1984) *Conceptual Structures: Information Processing in Mind and Machine*. Reading, MA: Addison-Wesley.
- Stasio T., D.J. Herrmann and R. Chaffin (1985) Predictions of Relation Similarity according to Relation Definition Theory. *Bulletin of Psychonomic Society*.
- Wang, Y-C, Vandendorpe, J. and Evens, M. (1985) "Relational Thesauri in Information Retrieval." *JASIS* 36 (1), 15-27.

