

Analyzing research papers using citation sentences

Wendy Lehnert, Claire Cardie, and Ellen Riloff
Department of Computer and Information Science
University of Massachusetts
Amherst, MA 01003
E-mail: lehnert@cs.umass.edu

Abstract

By focusing only on the citation sentences in a research document, one can get a good feel for how the paper relates to other research and its overall contribution to the field. The main purpose of a citation is to explicitly link one research paper to another. We present a taxonomy of citation types based upon empirical data and claim that we can recognize these citation types using domain-independent predictive parsing techniques. Finally, an experiment based on a corpus of research papers in the field of machine learning demonstrates that this is a promising new approach for processing expository text.

1 Introduction

One can get a reasonably good understanding of a research paper by merely skimming the sentences that reference other papers. By looking at how the author relates his work to other work in the field, a casual reader can get a good idea of what the paper is all about. As an example, consider the following opening sentence taken from a paper by (Braverman 88):

The methods of explanation-based learning (EBL) [DeJong & Mooney, 1986] and explanation-based generalization (EBG) [Mitchell, Keller, & Kedar-Cabelli, 1986] involve two conceptual phases: explanation and generalization.

There are two important inferences that can be drawn from the sentence: (1) The paper is pigeonholing itself in the explanation-based learning (EBL) and explanation-based generalization (EBG) paradigms and (2) The distinction between the two conceptual phases is likely to be relevant to the paper. These inferences require general knowledge about research and research papers but no domain-dependent knowledge about EBL or EBG per se. Based on this observation, we contend that research documents can be understood at two distinct levels:

1. **The Semantic Level:** the domain-dependent details of the paper
2. **The Paradigmatic Level:** how the paper relates to other papers in the field

Previous work in understanding expository text has concentrated on summarization at the semantic level. Traditional approaches required extensive domain-dependent knowledge as well as an analysis of the surface structure of the text (see [Britton and Black 85] and [Voss and Bisanz 85] for discussions of these approaches). By focusing on the paradigmatic level of research documents, however, we can ignore the domain-dependent details of a paper and thereby make the task of “understanding” the text a tractable problem. Two main claims follow from this approach:

- Research documents can be understood at the paradigmatic level using a set of *conceptual references*.
- *Conceptual references* can be extracted from research documents using domain-independent predictive parsing techniques.

2 Conceptual References

A *conceptual reference* is a relation between a referencing paper and the referenced paper. Conceptual references represent the reasons behind a citation and tell us why an author references another paper. We have identified two distinct levels of conceptual references: conceptual reference categories and conceptual reference structures. In the following sections, we describe these two levels of conceptual references.

2.1 Conceptual Reference Categories

Conceptual reference categories identify the abstract object that the author is pointing to in the referenced paper. For example, the author may reference another paper to refer to a method, example, or result presented in that paper. Based on the observation that the majority of citation sentences rely on a small set of conceptual reference types, we have created a taxonomy of 18 conceptual reference categories:

1. **System:** a system is described in the referenced paper; e.g. in (Silver 88),
"Other approaches were considered, including the use of ID3 [Quinlan]..."
2. **Method:** a method is described in the referenced paper; e.g. in (Shavlik 88),
"Mooney [Mooney 88a] presents an algorithm for generalizing..."
3. **Concept:** a concept is described in the referenced paper; e.g. in (Keller 88),
"Mostow's original definition of operationality [Mostow 81]..."
4. **Result:** a result is claimed in the referenced paper; e.g. in (Cohen 88),
"In [Cohen 87] it is shown that PAs are Turing-equivalent..."
5. **Fact:** a fact is stated in the referenced paper; e.g. in (Ellman 88),
"Standard explanation-based learning (EBL) methods apply only to domains for which a tractable domain theory is available [Mitchell 86]"
6. **Criticism:** a criticism is made in the referenced paper; e.g. in (Hunter 88),
"A detailed criticism of such purely empirical systems can be found in [Schank 86]"
7. **Example:** an example is used in the referenced paper; e.g. in (Shavlik 88),
"An example in [Shavlik 88a] shows that ..."
8. **More details:** the referenced paper has more details; e.g. in (Swaminathan 88),
"For details, the reader is referred to [Swaminathan 88a]"
9. **Attribution:** an item is attributed to the referenced paper; e.g. in (Bylander 88),
"In [Chandrasekaran 87], three types of explanation are ..."
10. **View:** a view is expressed in the referenced paper; e.g. in (Hirsh 88),
"Generalization can be viewed as a search problem ([Mitchell 82] [Simon 74])..."
11. **Model:** a model is presented in the referenced paper; e.g. in (Mahadevan 88),
"While there exist formal models for concept learning [Natarajan 87a] ..."
12. **Research:** research is presented in the referenced paper; e.g. in (Clancey 88),
"Apprenticeship learning research has considered ... [Mitchell 85] [Smith 85]"
13. **Extends:** the referenced paper extends previous work; e.g. in (Cohen 88),
"[Cohen 88a] has extended the system described in this paper ..."
14. **Application:** the referenced paper presents an application; e.g. in (Bennett 88),
"For example, Segre has applied EBL to learning robotics tasks in a simplified blocks world [Segre 87b]"

15. **Merge:** the referenced paper merges two techniques; e.g. in (Prieditis 88),
"[this work] is based on combining partial evaluation with other techniques (see Seki and Furukawa [37])"
16. **Proposal:** the referenced paper proposes an idea; e.g. in (Swaminathan 88),
"[Tadepalli 85] has proposed replacing the original theory with ..."
17. **Problems:** the referenced paper identifies a problem; e.g. in (Dietterich 88),
"[An important problem] ... is the imperfect theory problem [Mitchell 86]"
18. **Argument:** the referenced paper argues a position; e.g. in (Ginsberg 88),
"... Kleer gives a similar argument for the ATMS [DeKleer 86]"

We created this taxonomy of conceptual reference categories based upon an exploratory empirical study of citation sentences. Our corpus contained 372 citation sentences from 40 papers in the Proceedings of the AAAI Spring Symposium Series on Explanation-Based Learning, March 1988.¹ To establish a good set of categories, we set a priori criteria for an acceptable taxonomy: (1) every category must be present in at least 2 sentences from at least 2 different papers (to limit idiosyncracies due to a particular author) and (2) the final set of categories must cover at least 90% of the citation sentences.

First, we arbitrarily selected 208 sentences (the odd-numbered sentences) from our corpus and labelled them by hand with the object being referenced by the citation – its candidate conceptual reference type.² Second, we compiled a list of these candidate conceptual reference types and retained only those that occurred in at least 2 sentences from 2 different papers (satisfying the first criterion). Finally, we measured the coverage of this final set and found that it covered 93.7% of the test set. Since this exceeded our 90% threshold (satisfying the second criterion), this set became our 18 conceptual reference categories.

2.2 Conceptual Reference Structures

The conceptual reference categories describe objects being pointed to in the referenced paper. Conceptual reference structures fit on top of these categories to describe the relationship between the referenced object and the current paper. The structures combine conceptual reference categories to explain exactly how a referenced object is being used by the referencing paper. For example, the author might reference a method in another paper to show how his own method *is similar to* the referenced method. There are currently three types of conceptual reference structures: similarity, difference, and flagship references:

- **SIMILARITY:** an object in the current paper is similar to an object in the referenced paper; e.g. in (Ellman 88): *"The general approach ... is similar to methods described in [Keller 87] and [Mostow and Fawcett 87]"*
- **DIFFERENCE:** an object in the current paper differs from an object in the referenced paper; e.g. in (Minton 88): *"Using a different approach, DeJong and Mooney's GENESIS system [DeJong 86b] ..."*
- **FLAGSHIP:** the current paper is pigeonholed via a group of citations of the same type; e.g. in (Prieditis 88): *"See [9,28,29,27,18,35,21] for examples of EBL systems."*

Since these conceptual reference structures fit on top of many conceptual reference categories (e.g. systems, methods, or models can be similar), we can visualize conceptual reference structures as being on a higher plane than conceptual reference categories. A citation sentence may therefore be represented as a conceptual reference category alone (e.g. if a related system is referenced) or as several conceptual reference categories that are embedded in a conceptual reference structure (e.g. if a method is compared to a related method). It is also possible for a sentence to be mapped into

¹All examples used in the paper are taken from this corpus.

²There were actually 223 references because some of the sentences had multiple citations.

more than one conceptual reference if the author references a paper for several reasons. The next section describes how we recognize conceptual reference categories and build structures on top of them.

3 Parsing into Conceptual References

As stated earlier, we claim that understanding the relationships among research papers hinges on the ability to extract conceptual references from a document. Our second claim extends to the parsing mechanisms needed to understand these conceptual references:

1. *The parser relies strictly on predictive parsing techniques.* Predictive parsers are knowledge-based sentence analyzers that create conceptual representations for sentences. They have been used extensively in understanding narrative texts (e.g. [Lehnert 89], [Riesbeck and Schank 76], [Birnbaum and Selfridge 81], [Dyer 83]).
2. *The memory model underlying the parser contains only domain-independent knowledge.* Because we draw our examples from research literature in the field of machine learning, we distinguish between knowledge about research in general (domain-independent knowledge) and knowledge about machine learning (domain-dependent knowledge). In mapping sentences into conceptual references, the parser applies only general research knowledge.

Relationships among research papers are usually made explicit when one research paper references another. For this reason, we can recognize conceptual references without parsing the entire research document and restrict our attention to only those sentences that cite other papers. Citation analysis has been recognized as an important subfield of information retrieval since the early 1960's (e.g., [Garfield 55], [Garfield 64], [Garfield 79]). For the most part, however, researchers in information retrieval have concentrated on statistical analyses of citations to assess the value of an individual paper or the influence of a particular author (see [Salton and McGill 83], [O'Connor 83], and [Hurt 87]). They examine citations only to find the *existence* of explicit links between papers and make no attempt to attach any semantics to the links.

Because our work requires a deeper understanding of the relationships implied when one paper references another, we analyze citations within the context of the sentences in which they occur. These sentences are easily identifiable by a preprocessor that recognizes the peculiar format of citations. Occasionally, there are sentences that contain conceptual references without explicitly citing other papers. There is nothing about our parsing approach that prohibits mapping these sentences into the appropriate conceptual reference types. Identifying these sentences, however, would entail a scan of every sentence in the document. By limiting our attention to explicit citation sentences, we lose little information and only need to parse a small, easily identifiable portion of the text.

The goal of the parser is to represent a citation sentence as an instantiated form of one or more conceptual reference types. Sometimes it is adequate to map a sentence into a simple conceptual reference category. Often, however, a sentence contains more than one conceptual reference or contains one of the more complicated conceptual reference structures. In these cases, the parser returns multiple or embedded representations of the input.

Thus far, we have concentrated on recognizing 10 of the most common reference types described in section 2 — System, Method, Concept, Result, Criticism, Example, More Details, Attribution, Similarity, and Difference references. In the next section, we discuss our approach for parsing citation sentences into one or more of these types. Section 4 walks through a more detailed parse of a sentence from the machine learning corpus.

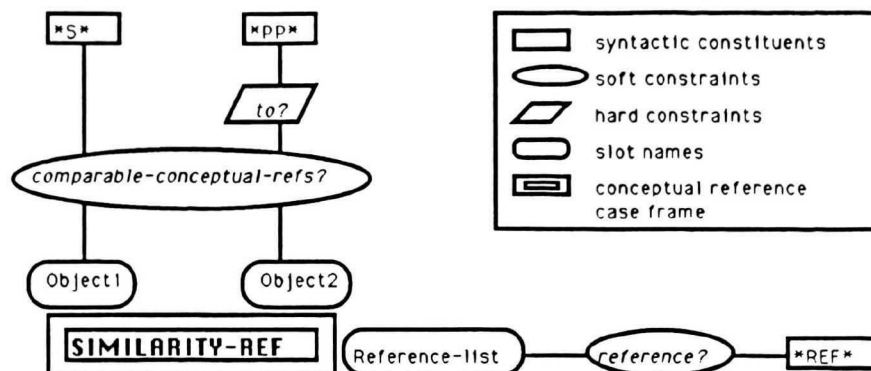


Figure 1: Case Frame Definition for a Similarity Reference

3.1 The CIRCUS Parser

We currently use the semantically-oriented CIRCUS parser [Lehnert 89] for understanding citation sentences. As a conceptual sentence analyzer, CIRCUS represents the meaning of sentences in terms of semantic case frames. Within CIRCUS, a conceptual reference is therefore represented as a case frame structure. Before sentence analysis begins, however, case frame definitions for each conceptual reference type must be hand-coded for the predictive semantics module that performs the slot-filling task. A Similarity reference, for example, has a case frame definition illustrated by Figure 1. Similarity references have three slots: Object1 and Object2 hold the two objects being compared; the Reference-list slot contains a list of citations. In addition, the case frame definition specifies that Object1 will be located in the subject of the sentence, Object2 will be in a prepositional phrase, and the special *REF* syntactic constituent³ will contain the Reference-list.

Before filling a case frame slot, a syntactic constituent must satisfy the slot's semantic constraints.⁴ For a Similarity reference, the prepositional phrase filling Object2 should begin with the preposition "to"⁵, the head nouns in the subject and prepositional phrase constituents should be conceptual references that address research at the same level of generality, and the contents of *REF* should be a list of citations.⁶ Below we present an example to illustrate how CIRCUS maps a sentence into a Similarity reference using the case frame definition described above.

4 An Example

Consider the following citation sentence:

³The sole purpose of the *REF* syntactic buffer is to hold citations.

⁴CIRCUS allows both *hard* and *soft* constraints. A hard slot constraint is a predicate that *must* be satisfied. In contrast, a soft constraint defines a preference for a slot filler rather than a predicate that blocks slot-filling when it is not satisfied.

⁵The Similarity reference case frame in Figure 1 recognizes citation sentences of the form "... {object1}...{ "to be" verb}...{any synonym for "similar"}...to {object2}...".

⁶Because we employ only domain-independent knowledge in mapping citation sentences into conceptual references, the semantic constraints access knowledge about research in general, but do not use any knowledge about the machine learning domain.




Figure 2: Desired Case Frame

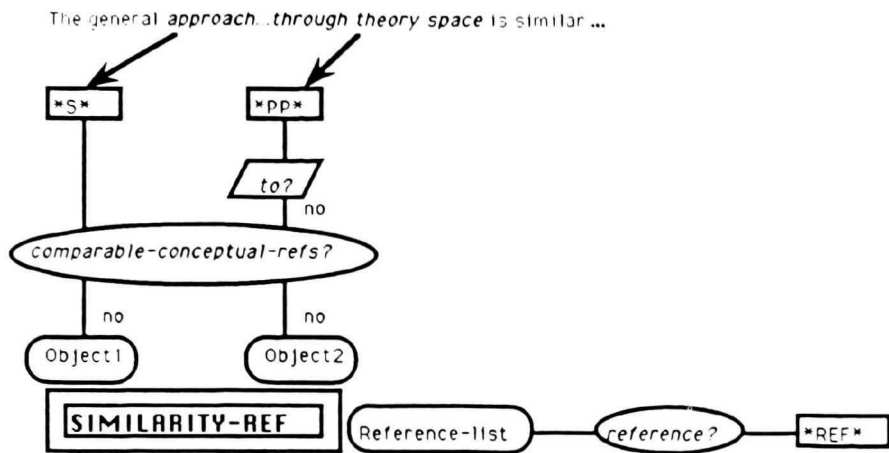


Figure 3: Case Frame Status After the Word "similar"

"The general approach of using examples to guide search through an approximate theory space is similar to methods described in [Keller 87] and [Mostow and Fawcett 87]" (Ellman 88).

CIRCUS should parse this sentence into the instantiated Similarity reference case frame shown in Figure 2. The parser scans a sentence from left to right, using its stack-oriented control to assign words/phrases to syntactic constituents until it notices a trigger for one of the predefined case frames.⁷ Once a case frame is active, CIRCUS' predictive semantics module uses a marker passing algorithm to fill slots in the frame.

In our example, the presence of a "to be" verb followed by the adjective "similar" activates the Similarity reference case frame. In addition, the words "approach" and "methods" are a subset of the phrases that trigger Method references. Figure 3 illustrates the Similarity reference case frame just after CIRCUS has scanned the word "similar". The subject contains a Method reference for "approach" and the most recent prepositional phrase is "through theory space". Although CIRCUS places these constituents in the Object slots (as specified by the case frame definition), the Reference-list slot remains empty. In addition, the semantic constraints associated with Object1 and Object2 have not yet been satisfied. The prepositional phrase filling Object2 should begin with "to" and both Object slots should point to conceptual references.

⁷ Some Similarity reference triggers are: "Similarly, ...", "In the same way, ...", "X-like" where X can be any noun, "...is the same as...", "...is similar...".

CIRCUS continues scanning the sentence until all slots of the active frame(s) are filled without any semantic failures or until it reaches the end of the sentence. After picking up the references, the parser successfully returns the instantiated Similarity reference case frame of Figure 2 because each slot is filled with an object that satisfies the slot’s semantic constraints: *REF* contains legitimate references, the preposition in the prepositional phrase constituent is “to”, and Object1 and Object2 point to conceptual references that address research at the same level of generality (i.e., both are methods).

5 Evaluation

Our system currently recognizes 10 conceptual reference types — 8 conceptual reference categories (System, Method, Concept, Result, Criticism, Example, More Details, Attribution) and 2 of the higher level conceptual reference structures (Similarity and Difference). It correctly parses 69 sentences from papers in our machine learning corpus and contains over 450 lexicon entries.

To evaluate our progress, we ran an informal experiment. The goal of the experiment was to test the generality of our current set of conceptual reference case frame definitions. We selected two papers from the field of machine learning that were not part of the original corpus⁸ and parsed the 28 citation sentences from those papers. We allowed the addition of lexicon entries for any new words occurring in the citation sentences, but did not define any new conceptual reference case frames or case frame triggers.

The system correctly parsed 75% of all citation sentences in the two papers. However, two of the sentences (7%) contained conceptual references whose case frames had not been predefined for CIRCUS. (They were not one of the 10 reference types listed above.) Modifying existing conceptual reference frame definitions and adding new triggers allowed 3 more of the remaining sentences to be parsed. With minor modifications to our parser definitions, we could therefore cover 86% of the test sentences. Discounting the 7% covered by undefined case frames, our success rate was then 93%.

6 Conclusion

This paper introduces an original strategy for parsing research documents using conceptual references. The work began in conjunction with the RA document summarization project which aims to summarize scientific research papers in terms of underlying research trends [Swaminathan 90]. RA models domain-independent structural relations in a research field in terms of research schemas and conceptual references. The RA system can use the conceptual references produced by our system to construct a summary of a corpus of research papers.

In closing, we emphasize that our approach to processing expository text is unique in two distinct ways. First, we model domain-independent structural relations between research papers in terms of conceptual references. This is a fundamental departure from other knowledge-based systems that emphasize the semantic content of their domains. Second, we demonstrate that expository text can be processed at the paradigmatic level using strongly predictive techniques. Our results have demonstrated that these two aspects of our system make it a promising new approach for processing expository texts.

⁸In order to remove any bias in the selection of papers, someone outside the parsing development group chose the papers for the experiment, [Keller 87] and [Kedar-Cabelli 87].

7 Acknowledgements

The authors would like to thank Kishore Swaminathan for commenting on earlier drafts of this paper and Stefan Wermter for his work on the project. This research was supported by the Advanced Research Projects Agency of the Department of Defense, monitored by the Office of Naval Research under contract #N00014-87-K-0238, the Office of Naval Research, under a University Research Initiative Grant, Contract #N00014-86-K-0764 and NSF Presidential Young Investigators Award NSFIST-8351863.

References

- [Birnbaum and Selfridge 81] Lawrence Birnbaum and Mallory Selfridge, 1981. Conceptual Analysis of Natural Language. In R. Schank and C. Riesbeck, (Eds.), *Inside Computer Understanding*, pp. 318-353. Hillsdale, NJ; Lawrence Erlbaum Associates.
- [Britton and Black 85] Bruce Britton and John Black, 1985. Understanding Expository Text: From Structure to Process and World Knowledge. In B. Britton and J. Black, (Eds.), *Understanding Expository Text*, pp. 1-9. Hillsdale, NJ; Lawrence Erlbaum Associates.
- [Dyer 83] Michael Dyer, 1983. *In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension*. Cambridge, MA; MIT Press.
- [Garfield 55] E. Garfield, 1955. Citation Indexes for Science. *Science*, 122, pp. 198-111.
- [Garfield 64] E. Garfield, 1964. Can Citation Indexing be Automated? In *Proc. of Symposium on Statistical Association Methods for Mechanized Documentation*, Washington, D.C.
- [Garfield 79] E. Garfield, 1979. *Citation Indexing — Its Theory and Applications in Science, Technology and Humanities*. New York; John Wiley.
- [Hurt 87] C.D. Hurt, 1987. Conceptual Citation Differences in Science, Technology, and Social Sciences Literature. *Information Processing & Management*, Vol. 23, No. 1, pp. 1-6.
- [Kedar-Cabelli 87] S. Kedar-Cabelli, 1987. "Formulating Concepts According to Purpose" *Proc. of AAAI*, Seattle, WA.
- [Keller 87] R. Keller, 1987. "Defining Operationality for Explanation-Based Learning". *Proc. of AAAI*. Seattle, WA.
- [Lehnert 89] Wendy Lehnert, 1989. Symbolic/Subsymbolic Sentence Analysis: Exploiting the Best of Two Worlds. In: Barnden, J. and Pollack, J. (Eds.): *Advances in Connectionist and Neural Computation Theory*, Vol. 1, Ablex Publishers, in press. (Also available as COINS Technical Report 88-99, Department of Computer and Information Science, University of Massachusetts, Amherst, MA.)
- [O'Connor 83] John O'Connor, 1983. Biomedical Citing Statements: Computer Recognition and Use to Aid Full-text Retrieval. *Information Processing & Management*, Vol. 19, No. 6, pp. 361-368.
- [Riesbeck and Schank 76] Chris Riesbeck and Roger Schank, 1976. Expectation-Based Analysis of Sentences in Context. Research Report No. 78. New Haven, CT; Department of Computer Science, Yale University.
- [Salton and McGill 83] Gerard Salton and Michael J. McGill, 1983. *Introduction to Modern Information Retrieval*. New York, NY; McGraw-Hill Book Company.
- [Swaminathan 90] Kishore Swaminathan, 1990. "Knowledge Evolution and Memory Organization for Scientific Domains" — a PhD dissertation in preparation. Department of Computer & Information Science. University of Massachusetts, Amherst, MA.
- [Voss and Bisanz 85] James Voss and Gay Bisanz, 1985. Knowledge and the Processing of Narrative and Expository Texts. In B. Britton and J. Black, (Eds.), *Understanding Expository Text*, pp. 173-198. Hillsdale, NJ; Lawrence Erlbaum Associates.