Hedging in Scientific Articles as a Means of Classifying Citations

Chrysanne Di Marco and **Robert E. Mercer**
University of Waterloo, Waterloo, Ontario, Canada
The University of Western Ontario, London, Ontario, Canada

Abstract

Citations in scientific writing fulfil an important role in creating relationships among mutually relevant articles within a research field. These inter-article relationships reinforce the argumentation structure intrinsic to all scientific writing. Therefore, determining the nature of the exact relationship between a citing and cited paper requires an understanding of the rhetorical relations within the argumentative context in which a citation is placed. To automatically determine these relations, we have suggested that various stylistic and rhetorical cues will be significant. One such cue that we are studying is the use of hedging to modify the affect of a scientific claim. We have previously shown that hedging occurs more frequently in citation contexts than in the text as a whole. With this information we conjecture that hedging is a significant aspect of the rhetorical structure of citation contexts and that the pragmatics of hedges may help in determining the rhetorical purpose of citations.

Introduction

Scientific writing and the need for affect

Since the inception of the formal scientific article in the seventeenth century, the process of scientific discovery has been inextricably linked with the actions of writing and publishing the results of research. Rhetoricians of science have gradually moved from a purely descriptive characterization of the science genre to full-fledged field studies detailing the evolution of the scientific article. During the first generation of rhetoricians of science, e.g., (Myers 1991), (Gross 1996), (Fahnestock 1999), the persuasive nature of the scientific article, how it contributes to making and justifying a knowledge claim, was recognized as the defining property of scientific writing. Style (lexical and syntactic choice), presentation (organization of the text and display of the data), and argumentation structure were noted as the rhetorical means by which authors build a convincing case for their results. Recently, second-generation rhetoricians of science (e.g., (Hyland 1998), (Gross, Harmon, & Reidy 2002)) have begun to methodically analyze large corpora of scientific texts with the purpose of cataloguing specific stylistic and rhetorical features that are used to create the pragmatic effects that contribute to the author's knowledge claim. One particular type of pragmatic effect, *hedging*, is especially common in scientific writing and can be realized through a wide variety of linguistic choices.

Citation analysis in scientific writing

We believe that pragmatic attitudes such as hedging (Hyland 1998), politeness (Myers 1989), and persuasion play an essential role in building the argumentative structure of the scientific article, and in conveying the nuances that help to support the author's knowledge claims. Moreover, we believe that these pragmatic effects work together with both global discourse structure—e.g., the traditional 'IMRad'[1] design of scientific discourse—and local text structure, including lexical choice, syntactic arrangement, citation placement and other aspects of scientific presentation, to create the overall rhetorical effect of a research article. In particular, we are studying the pragmatic function of citations in providing a textual means of relating articles in the space of documents which defines a research community. Studies in citation analysis indicate that the author's intent in including a citation at a particular point in the text reflects the pragmatic purpose of the citation, whether, for example, it indicates supporting or contrasting work to the topic under discussion. Our basic hypothesis is that the specific pragmatic function of citations may be determined through the analysis of fine-grained linguistic cues in the surrounding text.

We are presently studying the analysis of hedging cues in scientific writing as a means of classifying the purpose of citations in scientific texts. Hedging analysis seems well-suited as a means of approaching this problem: hedging in scientific writing is both pervasive and often readily detectable by surface textual features, while hedging cues have been well-studied (e.g., (Hyland 1998)) in terms of their pragmatic function.

Hedging in scientific writing

The term "hedging" was introduced by Lakoff (1972) in describing "words whose job it is to make things more or less fuzzy." As Hyland (1998) elaborates, "[Hedging] has subsequently been applied to the linguistic devices used to qualify a speaker's confidence in the truth of a proposition, the kind

[1] Introduction, Methods, Results, and Discussion.

of caveats like *I think*, *perhaps*, *might*, and *maybe* which we routinely add to our statements to avoid commitment to categorical assertions. Hedges therefore express tentativeness and possibility in communication, and their appropriate use in scientific discourse is critical (p. 1)".

The following examples illustrate some of the ways in which hedging may be used to deliberately convey an attitude of uncertainty or qualification. In the first example, the use of the verb *suggested* hints at the author's hesitancy to declare the absolute certainty of the claim:

(1) The functional significance of this modulation is suggested by the reported inhibition of MeSo-induced differentiation in mouse erythroleukemia cells constitutively expressing c-myb.

In the second example, the syntactic structure of the sentence, a fronted adverbial clause, emphasizes the effect of qualification through the rhetorical cue *Although*. The subsequent phrase, *a certain degree*, is a lexical modifier that also serves to limit the scope of the result:

(2) Although many neuroblastoma cell lines show a certain degree of heterogeneity in terms of neurotransmitter expression and differentiative potential, each cell has a prevalent behavior in response to differentiation inducers.

Hedging may be used in different rhetorical contexts within a scientific article to convey persuasive effect and enhance the knowledge claims of the author. For example, hedging may be realized through various linguistic cues in the Introduction, Results section, a controversial Discussion section, or generally throughout the research paper.

Within the Introduction to a scientific article, the use of hedging may serve both to establish the results within a wider research context and highlight the significance of this new work. In the extract below, the authors repeatedly use the key phrase *is/are consistent with* to first establish the reliability of their results, and then turn to more-hesitant cues (*provide circumstantial evidence*, *may be responsible*, *Regardless of the validity of this specific proposal*) to support, yet not overreach, their assertions. Nevertheless, the authors do still manage to get their claims across through a number of subtle but signficant cues: *not appear to*, *we reasoned*, *would*.

(3) Transgenic Arabidopdis seedlings over expressing phytochrome B exhibit enhanced sensitivity to Rc but wild-type responsiveness to FRc (Wagner et al, 1991; McCormac et al, 1993). This result is consistent with the behaviour of endogenous phytochrome B deduced from the *hy* 3 mutant studies... By contrast, transgenic Arabidopdis over expressing phytochrome A exhibits enhanced sensitivity to FRc (Whitelam et al, 1992; McCormac et al, 1993). Together these results are consistent with the possibility, although do not prove, that the capacity to mediate the FR-HIR may be an intrinsic property of phytochrome A.

Accumulated biochemical and physiological data also provide circumstantial evidence that phytochrome A may be responsible for the FR-HIR... [the data] are consistent with the possibility that this photolabile phytochrome pool may be responsible for the FR-HIR.

Regardless of the validity of this specific proposal, however, because phytochrome B does not appear to be involved in the FR-HIR, we reasoned that mutants defective in the activity of the phytochrome mediating this response would retain phytochrome B, and, therefore, retain responsiveness to Rc...

The Results section of a scientific paper, whether implicit or set off as a formal structure, tends to be lengthy and subdivided according to topic (Hyland 1998) (p. 193). The topics present the paper's findings, while associated hedges may be used to enhance the persuasive effects of the authors' interpretations of the findings and the resulting claims.

In the following example, the authors appear to be hedging certainty, putting forth their claim, but tempering the persuasive effect. They have chosen a modal verb, *would*, rather than a strong positive verb, such as *indicates*, so that the effect of the claim is restrained. Then, the following sentence seems to signal the possibility of a strong contrast by the explicit discourse marker, *However*, and use of a negative phrase, *cannot be ruled out*. Overall, the rhetorical effect is one of hesitance and tentativeness on the author's part.

(4) The faint 21-kD band observed in the PBM lane (Figure 2) would reflect the transient passage of this protein across the PBM from the plant cell cytoplasm to the bacteroids. However, the opposite is also possible, and it cannot be ruled out that the 21-kD polypeptides seen in the bacteroid lane and in the soluble proteins lane are totally different proteins with the same apparent molecular weight.

Hedging may be used not only in enhancing or mitigating the persuasive effects of an author's specific knowledge claims, but in setting up a strong 'protective' position from which to defend a highly controversial position. Hyland (1998) (p. 196) describes a text in which the writer has proposed a radical explanation for a process that is a core issue in her research area. As he analyzes the text, he points out how the writer goes even further, in making serious challenges to current theories. Not only is the writer concerned about supporting her own scientific claim, Hyland observes, but with protecting her position in her research community: "In making this proposal, the writer implicitly attributes serious inadequacies in current theories in their interpretations of critical data. She therefore runs the very real risk of having the claim rejected by a community of peers who, she perceives, have a great deal invested in the existing view and who are likely to defend it without giving serious consideration to her work" (p. 196).

How then does this writer manage to simultaneously put forth her own claim, challenge established theory, and protect her position in the community? Not surprisingly, the paper is thick with hedges: modal verbs and adverbs, epistemic lexical verbs, indefinite quantifiers, and admissions of limiting conditions, all contriving to "[create] a rhetorical and interpersonal context which seeks to pre-empt the reader's

rejection" (Hyland 1998) (p. 196).

As these examples illustrate, hedging effects are commonly used throughout scientific articles, while the ways in which hedging may be realized are both varied and easy to recognize. These characteristics suggested to us that the detection of hedging effects might be used as the basis for locating linguistic cues in scientific texts that might then help to determine the intended communicative effect of citations placed in the surrounding text.

Application: Classifying citations[2]

The need for citation classification

Scientific citations play a crucial role in maintaining the network of relationships among mutually relevant articles within a research field. Customarily, authors include citations in their papers to indicate works that are foundational in their field, background for their own work, or representative of complementary or contradictory research. But, determining the nature of the exact relationship between a citing and cited paper is often difficult to ascertain. To address this, the aim of formal citation analysis has been to categorize and, ultimately, automatically classify scientific citations.

A *citation* may be formally defined as a portion of a sentence in a citing document which references another document or a set of other documents collectively. For example, in sentence (5) below, there are two citations: the first citation is *Although the 3-D structure... progress*, with the set of references (Eger et al., 1994; Kelly, 1994); the second citation is *it was shown... submasses* with the single reference (Coughlan et al., 1986).

(5) Although the 3-D structure analysis by x-ray crystallography is still in progress (Eger et al., 1994; Kelly, 1994), it was shown by electron microscopy that XO consists of three submasses (Coughlan et al., 1986).

The primary purpose of scientific citation indexing is to provide researchers with a means of tracing the historical evolution of their field and staying current with on-going results. Citations link researchers and related articles together, and allow navigation through a space of mutually relevant documents which define a coherent academic discipline. Citation statistics play an important role in academic affairs, including promotion and tenure decisions and research grant awards. Scientific citations are thus a crucial component in the research and administrative life of the academic community. However, with the huge amount of scientific literature available, and the growing number of digital libraries, standard citation indexes are no longer adequate for providing precise and accurate information. What is needed is a means of better judging the relevancy of related papers to a researcher's specific needs so that only those articles most related to the task at hand will be retrieved. In previous work, Garzone and Mercer (Garzone 1996), (Garzone & Mercer 2000) presented a system for citation classification that relied on characteristic syntactic structure to determine citation category. We are now extending this idea to develop a method for using fine-grained rhetorical cues within citation sentences to provide such a stylistic basis for categorization (Mercer & DiMarco 2003), (DiMarco & Mercer 2003), (Mercer, DiMarco, & Kroon 2004).

[2]This section contains some material from (Mercer & DiMarco 2003).

Related work in citation classification

The usefulness of citation categorization for other applications is directly related to the comprehensiveness (breadth and granularity) of the citation classification scheme. Garzone and Mercer (Garzone 1996), (Garzone & Mercer 2000) proposed a citation classification scheme with 35 categories. This classification scheme is more comprehensive than the union of all of the previous schemes: it has a finer granularity than the often-used scheme of Garfield and Weinstock (Garfield 1965), (Weinstock 1971) and the one which previously had the most categories, Duncan, Anderson, and McAleese (1981), and it includes the full breadth of the other schemes (Cole (1975), Finney (1979), Frost (1979), Lipetz (1965), Moravcsik and Murugesan (1975), Peritz (1983), Small (1982), and Spiegel-Rösing (1977)). The Garzone and Mercer scheme and its relationship to the previous ones is discussed in detail in (Garzone 1996).

To demonstrate the kinds of citation categories, we list a few in a compressed format (the slashes indicate separate categories):

- Citing work disputes/corrects/questions some aspect of cited work.
- Citing work confirms/illustrates some aspect of cited work.
- Use of materials, equipment, or tools/methods, procedures, and design/theoretical equation/definition/numerical data.

We have a prototype citation classification system that takes journal articles (currently only biochemistry and physics) as input and maps each citation into one of the 35 citation categories. The prototype system relies on a large number of cue words (for example, discourse cues, nouns, and verbs which are closely related to the science and its methodology), some simple syntactic relationships, and knowledge about the IMRaD structure.

In direct contrast to Garzone and Mercer, which we take as our own starting-point, Teufel (1999) questions whether fine-grained discourse cues do exist in citation contexts, and states that "many instances of citation context are linguistically unmarked." (p. 93). She adds that while "overt cues" may be recognized if they are present, the problems of detecting these cues by automated means are formidable (p. 125). Teufel thus articulates the dual challenges facing us: to demonstrate that fine-grained discourse cues can play a role in citation analysis, and that such cues may be detected by automated means. While Teufel represents a counterposition to our approach, her work does complement ours in a number of ways. Teufel's research has a different goal to ours—it is aimed at generating summaries of scientific articles—but she does acknowledge the importance of a recognizable discourse structure in scientific articles, the IM-

RaD structure, and she also relies on local rhetorical structure to help determine where to find specific types of information to construct her 'fixed-form' summaries. However, Teufel voices her concern about the "potentially high level of subjectivity" (p. 92) inherent in judging the nature of citations, a task made more difficult by the fine granularity of her model of argumentation and the absence, she claims, of reliable means of mapping from citations to the author's reason for including the citation. As a consequence, Teufel confines her classification of citation categories to only two clearly distinguishable types: the cited work either provides a basis for the citing work or contrasts with it.

Nanba et al. ((Nanba & Okumura 1999), (Nanba, Kando, & Okumura 2000)) also present work in automated citation classification that is complementary to ours: their aim is to automatically generate review articles in a specific subject domain using citation types as the basis for the classification of papers. As does Teufel, they rely on two primary citation categories (works that provide a supporting basis for the citing paper, works that have a contrasting or 'negative' relationship), but also add a third 'others' category to indicate some form of unspecified relationship exists between the citing and cited papers. Collections of 'cue phrases' (including discourse markers, lexical usage, specific phrases), are used to classify citations into the different categories but these cues are heuristically motivated rather than theoretically based. In contrast, the types of cues we are using to detect the purpose of a citation are based in discourse analysis (Mercer & DiMarco 2003) and the rhetoric of science (Mercer, DiMarco, & Kroon 2004).

We can thus summarize the differences between our approach to citation categorization and that of Teufel and Nanba as follows:

- Our aim is a literature indexing tool using the rhetoric of science.
- We use a fine-grained citation categorization scheme with a greater number and variety of categories.
- We rely on cue phrases derived from formal linguistic theories as the basis for the detection and classification of citations.

Detecting hedges: Surface cues in science texts

A catalogue of surface features

The surface features through which hedging is realized in scientific texts have been copiously catalogued, in particular by Hyland. Using several corpora, both scientific and general academic, Hyland (1998) carried out a detailed analysis of hedging at several levels of linguistic description, including surface-level cataloguing of hedges and pragmatic analysis of their functions (pp. 98–99). The results of the study yielded a detailed catalogue of hedging cues including a large number of modal auxiliaries, epistemic lexical verbs (most commonly, *suggest*, *indicate*, *predict*), epistemic adjectives, adverbs, and nouns (representing half the major grammatical classes expressing hedging), as well as a variety of non-lexical, discourse-based hedges.

From hedging cues to citation classification

We believe that hedging cues may provide a prime source of fine-grained discourse cues that can be used to determine the intent of citations in the surrounding text. Hedging cues seem ideally suited for this purpose for the following reasons:

- The various types of hedging in scientific discourse have been extensively studied and catalogued by rhetoricians of science, Hyland (1998), in particular.
- The surface cues that give rise to hedging are readily recognizable by linguistic analysis, e.g., modal auxiliaries, specific lexical choice, and the use of discourse markers.

In our inital study (Mercer & DiMarco 2003), we analyzed the frequency of discourse cues in a set of scholarly scientific articles. We reported strong evidence that these cue phrases are used in the citation sentences and the surrounding text with the same frequency as in the article as a whole. We noted in this study that citations appeared to occur quite often in sentences marked by hedging cues. For example, the sentence below (shown earlier as example (1)), contains the hedging verb *suggested*, but also a citation about earlier work by other authors:

(6) The functional significance of this modulation is suggested by the reported inhibition of MeSo-induced differentiation in mouse erythroleukemia cells constitutively expressing c-myb (REF).

We may assume that the hedge and the citation are linked in some way: hesitancy in the current work may be offset by the support of earlier related research.

In the second example (shown earlier as (2)), the lexical and syntactic cues (*Although*, *a certain degree*) express qualification of the claim, but now the accompanying use of several citations serves to bolster the authoritative nature of the underlying argument. (Indeed, two of the citations refer to papers published more than five years earlier, and the third reference is 17 years old.)

(7) Although many neuroblastoma cell lines show a certain degree of heterogeneity in terms of neurotransmitter expression (REF) and differentiative potential (REF), each cell has a prevalent behavior in response to differentiation inducers (REF).

In (Mercer, DiMarco, & Kroon 2004), we followed up on our hypothesis that hedging cues tend to occur in citation contexts by doing a frequency analysis of hedging cues in citation contexts in a corpus of 985 peer-reviewed recent biology journal articles from the BioMed Central corpus. We obtained statistically significant results indicating that hedging is indeed used more frequently in citation contexts than the text as a whole. Given the presumption that writers make stylistic and rhetorical choices purposefully, we propose that we have further evidence that hedging cues are an important aspect of the rhetorical structure of citation contexts and that the pragmatic functions of hedges may help in determining the purpose of citations.

Summary of ongoing work

We are presently developing a biomedical literature indexing tool to automate the classification of citations using the rhetoric of science through the following tasks:

- Adapting existing computational linguistic tools (e.g., online lexicons, part-of-speech taggers, discourse marker analyzers) for the detection of hedging cues and other cue phrases within citation contexts.
- Building test corpora of citation sentences from biomedical and scientific articles.
- Developing methods and tools for automatically classifying the pragmatic functions of hedging cues and other cue phrases in the citation corpora.

Our goal in studying the effects of hedging in scientific writing is to identify linguistic cues that may be used as a means of determining the pragmatic function of citations. Ultimately, we can expect to be able to associate hedging cues and other pragmatic cues with rhetorical relations as determiners of citation function.

Acknowledgements

We thank Fred Kroon for carrying out the frequency analysis of hedging cues in citation contexts. Our research has been financially supported by the Natural Sciences and Engineering Research Council of Canada and by the Universities of Western Ontario and Waterloo.

References

Cole, S. 1975. The growth of scientific knowledge: Theories of deviance as a case study. In *The idea of social structure: Papers in honor of Robert K. Merton*. New York: Harcourt, Brace Jovanovich. 175–220.

DiMarco, C., and Mercer, R. 2003. Toward a catalogue of citation-related rhetorical cues in scientific texts. In *Proceedings of the Pacific Association for Computational Linguistics (PACLING) Conference*.

Duncan, E.; Anderson, F.; and McAleese, R. 1981. Qualified citation indexing: Its relevance to educational technology. In *Information Retrieval in Educational Technology: Proceedings of the First Symposium on Information Retrieval in Educational Technology*, 70–79. University of Aberdeen.

Fahnestock, J. 1999. *Rhetorical figures in science*. Oxford University Press.

Finney, B. 1979. The reference characteristics of scientific texts. Master's thesis, The City University of London.

Frost, C. 1979. The use of citations in literary research: A preliminary classification of citation functions. *Library Quarterly* 49:399–414.

Garfield, E. 1965. Can citation indexing be automated? In Stevens, M., et al., eds., *Statistical Association Methods for Mechanical Documentation (NBS Misc. Pub. 269)*. Washington, DC: National Bureau of Standards.

Garzone, M., and Mercer, R. 2000. Towards an automated citation classifier. In *Proceedings of the 13th Biennial Conference of the CSCSI/SCEIO (AI'2000)*, 337–346.

Garzone, M. 1996. Automated classification of citations using linguistic semantic grammars. Master's thesis, The University of Western Ontario.

Gross, A.; Harmon, J.; and Reidy, M. 2002. *Communicating science: The scientific article from the 17th century to the present*. Oxford University Press.

Gross, A. 1996. *The rhetoric of science*. Harvard University Press.

Hyland, K. 1998. *Hedging in scientific research articles*. John Benjamins Publishing.

Lakoff, R. 1972. The pragmatics of modality. In *Papers from the Eighth Regional Meeting, P. Peranteau, J. Levi, and G. Phares (eds.)*, 229–46. Chicago Linguistics Society.

Lipetz, B. 1965. Problems of citation analysis: Critical review. *American Documentation* 16:381–390.

Mercer, R., and DiMarco, C. 2003. The importance of fine-grained cue phrases in scientific citations. In *Proceedings of the 16th Conference of the CSCSI/SCEIO (AI'2003)*.

Mercer, R.; DiMarco, C.; and Kroon, F. 2004. The frequency of hedging cues in citation contexts in scientific writing. In *submission to the 17th Conference of the CSCSI/SCEIO (AI'2004)*.

Moravscik, M., and Murugesan, P. 1975. Some results on the function and quality of citations. *Social Studies of Science* 5:86–92.

Myers, G. 1989. The pragmatics of politeness in scientific articles. *Applied Linguistics* 10(1):1–35.

Myers, G. 1991. *Writing biology*. University of Wisconsin Press.

Nanba, H., and Okumura, M. 1999. Towards multi-paper summarization using reference information. In *Proceedings of the 16th International Joint Conferences on Artificial Intelligence (IJCAI-99)*, 926–931.

Nanba, H.; Kando, N.; and Okumura, M. 2000. Classification of research papers using citation links and citation types: Towards automatic review article generation. In *Proceedings of the American Society for Information Science (ASIS)/The 11th SIG Classification Research Workshop, Classification for User Support and Learning*, 117–134.

Peritz, B. 1983. A classification of citation roles for the social sciences and related fields. *Scientometrics* 5:303–312.

Small, H. 1982. Citation content analysis. *Progress in Communication Sciences* 3:287–310.

Spiegel-Rösing, I. 1977. Science studies: Bibliometric and content analysis. *Social Studies of Science* 7:97–113.

Teufel, S. 1999. *Argumentative zoning: Information extraction from scientific articles*. Ph.D. thesis, University of Edinburgh.

Weinstock, M. 1971. Citation indexes. In *Encyclopaedia of Library and Information Science*, volume 5. New York: Marcel Dekkar. 16–40.