Determining How Citations Are Used in Citation Contexts

Michael Färber[1] and Ashwath Sampath[2]

[1]Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany
michael.faerber@kit.edu
[2]University of Freiburg, Freiburg, Germany
ashwath92@gmail.com

Abstract. Citations have been classified based on their textual contexts w.r.t. their worthiness, function, polarity, and importance. To the best of our knowledge, so far citations have not automatically been classified by their grammatical role, that is, whether the citation (1) is grammatically integrated in the sentence, (2) is annotated directly after the occurrence of author names, (3) backs up a concept, (4) backs up a claim, or (5) is not appropriate because the context is incomplete or noisy. We argue that determining such classes for citation contexts is useful for a variety of tasks, such as improved citation recommendation and scientific impact quantification. In this paper, we propose this classification scheme, as well as a machine-learning-based approach to determine the classes automatically. Our evaluation reveals that the classification performance varies significantly between the citation types.

Keywords: citations, scholarly data, bibliometrics, classification

1 Motivation

Citing sources has always been an integral part of academic research. Scientific works need to contain appropriate citations to other works due to several reasons [1]. Most notably, all claims written by an author need to be backed up to ensure transparency, reliability, and truth. Furthermore, mentions of methods, data sets and important domain-specific concepts need to be linked via references to help the reader properly understand the text and to give attribution to the corresponding publications and authors.

Citation contexts have been classified in several respects so far. Noteworthy to mention is the classification of citation contexts (e.g., sentences) concerning the cite-worthiness [2]. Apart from that, citation contexts have been classified according to their citation function using some annotation scheme [1,3] (i.e., determining the "role" of a citation in its pragmatic context; e.g., that the author mentions a weakness of an approach). Similar tasks to the citation function determination are the polarity determination (i.e., if the author speaks in a positive, neutral, or negative way about the cited paper) [4,5]

Table 1: Examples for citation types taken from our data set.

Citation type	Example sentence
IN-TEXT	"The approaches of [7] and [1] allow to plug an advisor system in a scenario from the environments Telos and ExploraGraph."
AUTHOR	"Gibson et al. [12] used hyperlink for identifying communities."
CONCEPT	"To this end, SWRL [14] extends OWL-DL and OWL-Lite with Horn clauses."
CLAIM	"In the traditional hypertext Web, browsing and searching are often seen as the two dominant modes of interaction (Olston & Chi, 2003)."
INCOMPLETE	"see [16, 15]."

and the determination of the citation importance [6]. Redi et al. [7] proposed reasons why citations are used and needed in Wikipedia texts. However, to our knowledge, citation contexts have not been classified automatically by grammatical functions, that is, whether the citation (1) is grammatically integrated in the sentence, (2) is mentioned directly after the occurrence of author names, (3) backs up a concept, (4) backs up a claim, or (5) is not appropriate due to incomplete or noisy context. Determining such classes for citation contexts can be seen as a prerequisite for a variety of tasks, such as for improved citation recommendation and for creating alternative measurements of the scientific impact other than citation counts.

2 Citation Types

Based on the findings gained by manually examining various citation contexts and previous works on citation classification [8], we introduce five classes of citation types (see also the examples in Table 1):

1. IN-TEXT: The citation marker is part of the sentence. Citations are then "weaved" into the grammar of the sentence.
2. AUTHOR: The citation marker is set directly after the occurrence of author names.
3. CONCEPT: The citation marker follows a concept. This can be a named entity (e.g., a specific data set, method, or project) or an abstract scientific concept.
4. CLAIM: A claim made by the author is backed up by the citation.
5. INCOMPLETE: The citation does not need to be considered because the citation context is incomplete or noisy (e.g., containing only formulas or references from the reference section if extracted automatically) and, thus, not fully understandable.

3 Citation Type Classification Approach

To determine the citation types for a given citation context (including the citation markers), we develop a classifier based on a multi-label (one vs. the rest) gradient boosting classifier.[1] We use the following features, each extracted from the given citation context:

1. the number of words;
2. the normalized number of nouns, verbs, proper nouns, numbers, and prepositions;
3. the normalized number of citation markers;
4. the normalized number of mentioned people's names, using the named entity recognition implementation of the Python library Spacy;
5. the citation positions sum, which is calculated by adding the normalized word positions of all the citation markers in the citation context;
6. the sum of the distance of each of the citation markers in the citation context to the nearest noun, nearest proper noun, nearest preposition, and nearest verb, respectively.
7. the sum of the distance of each of the citation markers to the nearest person mentioned;
8. the TF-IDF values of the words in the context;
9. the average of the fastText vectors of all citation context's words.[2]

4 Evaluation

4.1 Data Set

We manually created a ground truth data set for the citation type classification task. To that end, we manually labeled 500 citation contexts that were written in English and extracted from the Microsoft Academic Graph where the cited document is tagged by the Microsoft Academic Graph with "semantic web". To create a test data set, we annotated 100 English citation contexts retrieved from the Microsoft Academic Graph where the cited document is tagged as "natural language processing." The data set is freely available online.[3] Note that we used different domains for the training and test data sets to keep our machine learning model generic enough and to avoid overfitting. The domains were chosen based on the expertise of the authors as assessors.

[1] See `https://github.com/michaelfaerber/citation-type-classifier` for our source code. Note that each citation context can belong to one or several citation types. This makes our classification task a multi-label classification task.

[2] See `https://fasttext.cc/`. The pretrained vectors were trained on Common Crawl and Wikipedia using the CBOW model of fastText. fastText operates at the character level, and therefore can generate vectors for words not seen in the training corpus.

[3] See `https://github.com/michaelfaerber/citation-type-classifier`.

Table 2: Evaluation results.

Class	IN-TEXT	AUTHOR	CONCEPT	CLAIM	INCOMPLETE
precision	0.50	1.00	0.46	0.57	0.50
recall	0.45	0.06	0.59	0.72	0.40
F1	0.47	0.11	0.52	0.64	0.44

4.2 Evaluation Results

Table 2 shows the performance of our classifier. Overall, we can observe that automatically determining all classes for a citation context is a difficult task. However, the classifier seems to work significantly better for some classes than others. The CLAIM and AUTHOR classes lie at the extremes, while the other three classes are around average. We obtain an accuracy over all classes (defined as the proportion of samples in which all the predicted labels match all the true labels) of 0.4. The Hamming loss, a more suitable metric than accuracy for multi-label classification tasks (as it allows partial mismatches), is 0.216.

5 Conclusion

In this paper, we presented a novel classification scheme for in-text citations. We then proposed a machine-learning-based approach to determine these citation types automatically for given citation contexts. In the future, we will analyze the citation types for various scientific disciplines. We also plan to incorporate the citation type classifier into citation recommendation systems.

References

1. Teufel, S., Siddharthan, A., Tidhar, D.: An annotation scheme for citation function. In: Proc. of SIGdial'09. (2009) 80–87
2. Färber, M., Thiemann, A., Jatowt, A.: To Cite, or Not to Cite? Detecting Citation Contexts in Text. In: Proc. of ECIR'18. (2018)
3. Teufel, S., Siddharthan, A., Tidhar, D.: Automatic classification of citation function. In: Proc. of EMNLP'07. (2006) 103–110
4. Abu-Jbara, A., Ezra, J., Radev, D.R.: Purpose and Polarity of Citation: Towards NLP-based Bibliometrics. In: Proc. of NAACL-HLT'13. (2013) 596–606
5. Ghosh, S., Das, D., Chakraborty, T.: Determining sentiment in citation text and analyzing its impact on the proposed ranking index. CoRR **abs/1707.01425** (2017)
6. Valenzuela, M., Ha, V., Etzioni, O.: Identifying Meaningful Citations. In: Proc. of SBD'15. (2015)
7. Redi, M., Fetahu, B., Morgan, J.T., Taraborelli, D.: Citation Needed: A Taxonomy and Algorithmic Assessment of Wikipedia's Verifiability. In: Proc. of WWW'19. (2019)
8. Petrić, B.: Rhetorical functions of citations in high-and low-rated master's theses. Journal of English for Academic Purposes **6**(3) (2007) 238–253