

Book Review

Argumentation Mining

Manfred Stede and Jodi Schneider

(University of Potsdam, University of Illinois at Urbana-Champaign)

Morgan & Claypool (Synthesis Lectures on Human Language Technologies, edited by Graeme Hirst, volume 40), 2018, xvi+175 pp; paperback, ISBN 978-1-68173-459-0; ebook, ISBN 978-1-68173-460-6; doi:10.2200/S00883ED1V01Y201811HLT040

Reviewed by
Henning Wachsmuth
Paderborn University

Argumentation is pervasive—in everyday life as well as in politics and media. People exchange arguments to persuade each other, to achieve agreement, to make decisions, and more. Argumentation has been studied since ancient times (Aristotle 2007), and it is an active research topic across disciplines today, from logic to rhetoric to linguistics (van Eemeren et al. 2014). In a time of alternative facts and filter bubbles, arguments are of ever-increasing importance; when the truth of facts is unclear, we need to compare reasons for opposing claims, and we should do so beyond our own view.

Computational research on argumentation is still young. It first evolved in the AI community, oriented toward formal argumentation (Dung 1995). Except for some early pioneering works, the natural language side is getting attention since the publication of the first approaches to mine arguments from text (Palau and Moens 2009). Since then, computational linguistics research on argumentation grows constantly, and impressive industrial applications such as IBM's *Debater* start to appear. Since 2014, the *ArgMining* workshop series exists, annually taking place at ACL, EMNLP, or NAACL.

Argumentation Mining is the first textbook on the topic. In line with ongoing research activities, it does not only tackle the identification and classification of claims, reasons, and their relations, but also the assessment and generation of argumentation. Although the focus is on natural language processing (NLP) techniques, the book covers connections to formal argumentation and some underlying techniques. The complementary pair of authors was a smart choice in this regard: Manfred Stede is one of the leading computational linguists for discourse processing, and Jodi Schneider represents the AI community and has specific expertise on scholarly communication and knowledge organization. Both have co-authored several papers on argumentation, and they obviously have a comprehensive overview of the topic; I couldn't think of more than a few papers that I missed being discussed.

The book contains ten chapters. After definitions and relevant basics (Chapters 1–2), it outlines common argument models and existing corpora (Chapters 3–4). Chapters 5 to 7 present computational aspects of argumentation mining, followed by assessment and generation approaches (Chapters 8–9), and a final conclusion (Chapter 10).

Chapter 1 begins with a very compressed definition of argumentation from the literature. It then explicates each single term in the definition one after another, directly

<https://doi.org/10.1162/coli.r.00358>

confronting the reader with what argumentation is all about—a very elegant opening. After illustrative examples for different aspects, the chapter provides a first idea of the full process of argumentation mining, and why this process is worth being studied.

The nature of argumentative language is reviewed in Chapter 2. Based on the notion of public and private states, the reader learns about various facets of subjectivity; from opinions and stance to verifiability, from speech acts to the usage of causal relations in argumentation, explanation, and justification; and from discourse modes to rhetoric and its linguistic reflection in style, figures of speech, and text organization. I found this to be an excellent introduction for any researcher dealing with subjective language.

Chapter 3 treats the modeling of argumentative structure. The authors clarify that the best model depends on the given task and text genre. They first discuss types of arguments and their components, advising the reader not to confuse the function of a component in an argument with its relevance. The chapter goes on with a concise discussion of defeasible reasoning, argumentation schemes, and types of argument models before it presents ideas of and tools for argument mapping. In my view, the latter is perhaps a bit extensive; I would have rather preferred some other details (e.g., on the role of critical questions for argumentation schemes or on the actual model of the *New Rhetoric*).

In Chapter 4, major corpus annotation schemes and selected corpora related to argumentation mining are briefly summarized. Visual examples make the annotation schemes easily understandable. The chapter explicitly does not strive for completeness, but it was not fully clear to me what governed the covered selection.

Chapter 5 opens the discussion of computational methods. As the authors point out, choosing mining methods to use depends on the goal. They begin with the distinction of argumentative from non-argumentative text spans. Due to an important observation, they keep the summary of relevant work short: Very few approaches go beyond standard features so far, hardly capturing argumentation-specific properties. Later, the authors go more into depth on identifying claims. The reader gets to know what text features work within several domains, but also learns how context-dependent the notion of a claim is and that many claims actually remain implicit.

Chapter 6 continues with the mining of premises that support or object a claim. It provides nice insights into several approaches to this task, some of which exploit knowledge about the claim, others not. The authors reveal that premise identification works better than claim identification in most studies, while support is easier to detect than objections. Also, they clarify the close relation of argumentation mining and stance classification, to then summarize noteworthy findings on stance, such as the existence of ideological patterns (e.g., “pro marijuana” correlates with “pro gay rights”).

The discussion of argumentation mining ends in Chapter 7 with methods for argumentative relations between claims and for complete hierarchical argument structures. It becomes clear when relations actually need to be mined, and how this task relates to textual entailment and relationship extraction. What particularly impressed me is the second half: Starting from the observation that there is no consensus on how to describe complete structures, the book convincingly opposes various ways of building up the hierarchical structure induced by argumentative relations, including minimum spanning trees, integer linear programming, and neural networks. Afterward, the authors emphasize an important aspect that is well known in rhetoric but still understudied in NLP: the *sequential* structure of arguments and its effect on downstream tasks.

Chapter 8 turns its attention to the assessment of the mined structure: reconstructing implicit components (enthymemes), classifying argumentation schemes, and analyzing quality. Only a few approaches to enthymemes exist, hence the chapter just

sketches challenges and findings briefly. Next, it outlines insightful empirical results on the distribution of schemes and their varying complexity, and also accounts for critical questions. The authors then make a strong claim, which I tend to agree with, although it may be debatable: “the question about the quality of an argument is the ‘ultimate’ one for argumentation mining.” They detail three facets of quality assessment: (1) the consistency checking of arguments based on abstract argumentation frameworks; (2) the absolute rating and relative comparison of logical, rhetorical, and dialectical quality dimensions; and (3) the role of emotions. As they argue, emotions may be seen as flaws in some argumentative contexts, but they notably impact the perceived quality of arguments in practice.

Finally, Chapter 9 closes the discussion of computational methods by looking at the two general techniques to generate argumentation: transforming knowledge into text, and rewriting text into other text. Research on the former started in the 1990s, whereas the latter came up more recently with statistical NLP research on argumentation. For both, the chapter explains important approaches in detail, giving the reader a concrete idea of how to proceed step by step. If anything, I found the evaluation of how well the approaches work a bit short, leaving implicit how hard the problem is. The authors mention a few results, though, and a comparison of the two techniques nicely concludes the chapter.

Chapter 10 first concludes the book’s contents in an unusual but charming way: It delineates algorithm for how to design an argumentation mining algorithm, depending on the given task and data. After an overview of several downstream applications of the presented methods, the authors then end with valuable last thoughts, from which I pick out one here: Technology that analyzes or synthesizes argumentation has the potential to benefit society—but also to create harm, for example, by influencing humans in a deceptive way. “Researchers must be alert to such risks.”

I really enjoyed reading the book. The authors found elegant ways to explain everything in a clear and easy-to-understand manner. They give many examples to illustrate the introduced concepts and methods. Although I’ve been studying argumentation for several years, I learned many new details, filling gaps in my knowledge. I particularly liked the summaries at the end of many sections, which concisely analyze and interpret the state of the art. Here and there, I would have liked to see the authors taking a stance on some approaches, but I see that this may have reduced objectivity.

The book does not use many formalizations, avoiding going into technical detail. This makes sense, given that it does not aim to teach NLP basics. To learn implementing approaches, a closer look at the respective papers may be needed (i.e., the book alone does not fully serve as a self-contained textbook for courses). Nevertheless, it is a great overview of basics and the current state of the art in the realm of argumentation mining. I often struggled in the past with what literature to give to students who are new to the topic. The book solves this problem. It is the best entrance point to the computational analysis and synthesis of natural language argumentation available so far.

References

- Aristotle. 2007. *On Rhetoric: A Theory of Civic Discourse* (George A. Kennedy, Translator), Clarendon Aristotle series. Oxford University Press.
- Dung, Phan Minh. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357.
- van Eemeren, Frans H., Bart Garssen, Erik C. W. Krabbe, A. Francisca Snoeck Henkemans, Bart Verheij, and Jean H. M. Wagemans. 2014. *Handbook of Argumentation Theory*. Springer Netherlands.

Palau, Raquel Mochales and Marie-Francine Moens. 2009. Argumentation mining: The detection, classification and structure of arguments in text. In *Proceedings of the*

12th International Conference on Artificial Intelligence and Law, pages 98–107, New York, NY.

Henning Wachsmuth is a junior professor at the Paderborn University, Paderborn, Germany. The computational analysis and synthesis of natural language argumentation is his main research focus. So far, he has co-authored more than 25 scientific papers on computational argumentation. In collaboration with the Bauhaus-Universität Weimar, his group develops the first search engine for arguments on the Web, `args.me`. E-mail: `henningw@upb.de`.