

An Informatics Perspective on Argumentation Mining

Jodi Schneider*

INRIA Sophia Antipolis France

jschneider@pobox.com

Abstract

It is time to develop a community research agenda in argumentation mining. I suggest some questions to drive a joint community research agenda and then explain how my research in argumentation, on support tools and knowledge representations, advances argumentation mining.

1 Time for a community research agenda

This year, argumentation mining is receiving significant attention. Five different events from April to July 2014 focus on topics such as arguing on the Web, argumentation theory and natural language processing, and argumentation mining. A coordinated research agenda could help advance this work in a systematic way.

We have not yet agreed on the most fundamental issues:

- Q1 What counts as ‘argumentation’, in the context of the argumentation mining task?
- Q2 How do we measure the success of an argumentation mining task? (e.g. corpora & gold standards)

“Argumentation mining, is a relatively new challenge in corpus-based discourse analysis that involves automatically identifying argumentative structures within a document, e.g., the premises, conclusion, and argumentation scheme of each argument, as well as argument-subargument and argument-counterargument relationships between pairs of arguments in the document.”¹ (Green et al., 2014)

¹This work was carried out during the tenure of an ERCIM “Alain Bensoussan” Fellowship Programme. The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 246016.

An informatics perspective (i.e. concerned with supporting human activity) could help us understanding how we will apply argumentation mining; this should sharpen the definition of the argumentation mining task(s). Given such an operationalization, we can then use the standard natural language processing approach: define a corpus of interest, make a gold standard annotation, test algorithms, iterate...

For instance, to operationalize the definition of argumentation mining (Q1), we need to know:

- Q1a How do we plan to use the results of argumentation mining?
- Q1b What domain(s) and human tasks are to be supported?
- Q1c What is the appropriate level of granularity of argument structures in a given context? Which models of argumentation are most appropriate?

This can be challenging because argumentation has a variety of meanings and uses, in fields from philosophy to rhetoric to law; some of the purposes for using argumentation are shown in Figure 1. Understanding how we will use the results of argumentation mining can help address important questions related to Q2, such as measuring the success of algorithms and support tools for identifying arguments. In particular:

- Q2a How accurate does argumentation mining need to be?
- Q2b In which applications are algorithms for automatically extracting argumentation most appropriate?
- Q2c In which applications are support tools for semi-automatically extracting argumentation more appropriate?

In my work I have tried to bring applications of argumentation mining to the forefront. My work falls into three main areas: supporting human argumentation with computer tools (CSCW), rep-

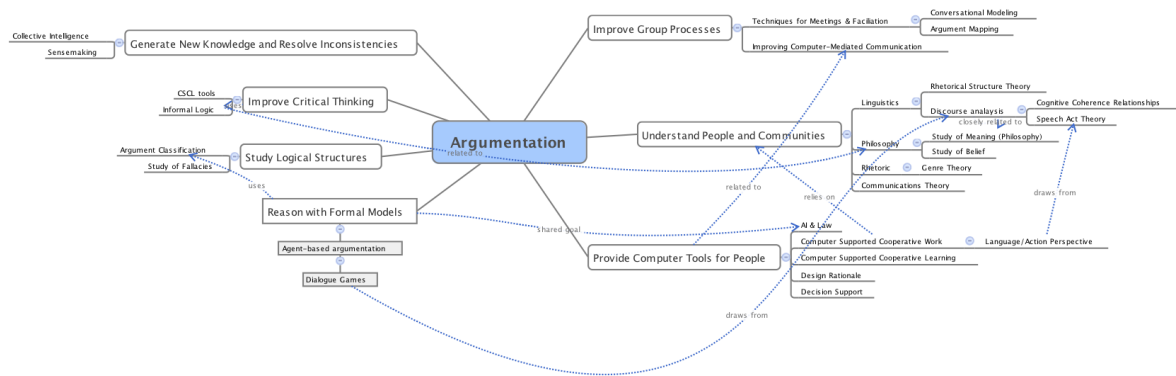


Figure 1: Argumentation can be used for many purposes.
 Download an editable version of this figure from FigShare DOI <http://dx.doi.org/10.6084/m9.figshare.1149925>

representing argumentation in ontologies (knowledge representation), and mining arguments from social media (information extraction using argumentation theory).

2 Computer-Supported Collaborative Work

Arguing appears throughout human activity, to support reasoning and decision-making. The application area determines the particular genres and subgenres of language that should be investigated (Q1b). The appropriate level of granularity (Lawrence et al., 2014) depends on whether we are in a literary work or a law case or a social media discussion (Q1c). The acceptable error rate (Q2a) follows from human tolerances, which we expect to depend on the area; this in turn determines whether we completely automate argumentation mining (Q2b) or merely provide semi-automatic support (Q2c). This is why I emphasize looking at application areas to determine which problems to focus our attention on, for argument mining.

My thesis described a general, informatics approach to supporting argumentation in collaborative online decision-making (Schneider, 2014b):

1. Analyze requirements for argumentation support in a given situation, context, or community.
2. Consider which argumentation models to use; test their suitability, using features such as the appropriate level of granularity and the tasks to be supported.
3. Build a prototype support tool, using a model of argumentation structures.
4. Evaluate and iterate.

In this approach, argumentation mining supports scalability, by providing automatic or semi-automatic identification of the relevant arguments.

I have applied this methodology to Wikipedia information quality debates, which are used to determine whether to delete a given topic from the encyclopedia (Schneider, 2014b). We tested two argumentation models: Walton’s argumentation schemes (Schneider et al., 2013) and the theory of factors/dimensions (Schneider et al., 2012c), and our annotated data is available online.² Whereas Walton’s argumentation schemes could have provided support for writing arguments, we instead chose to use domain-specific decision factors to filter the overall debate in the prototype support tool we built. One difference is that Walton’s argumentation schemes are at the micro-level—structuring the premises and conclusions of a given argument—whereas decision factors are at the macro-level, identifying the topics important to discuss; this distinction may be relevant for argumentation mining (Schneider, 2014a).

3 Knowledge Representation

Argumentation mining assumes a way to package arguments so that they can be exchanged and shared. Structured representations of arguments allow “evaluating, comparing and identifying the relationships between arguments” (Rahwan et al., 2011). And the knowledge representations most commonly used for the Web are ontologies.

To investigate the existing ontologies for structuring arguments on the social web, we wrote “A Review of Argumentation for the Social Semantic Web” (Schneider et al., 2012b).

²<http://purl.org/jspshd>

The review compares:

- 13 theoretical models for capturing argument structure (Toulmin, IBIS, Walton, Dung, Value-based Arg. Frameworks, Speech Act Theory, Language/Action Perspective, Pragma-dialectic, Metadiscourse, RST, Coherence, and Cognitive Coherence Relations).
- Applications of these theoretical models.
- Ontologies incorporating argumentation (including AIF, LKIF, IBIS and many others).
- 37 collaborative Web-based tools with argumentative discussion components (drawn from Social Web practice as well as from academic researchers).

Thus the argumentation community can choose from a number of existing approaches for structuring argumentation on the Web.

Still, new approaches continue to be suggested. Peldszus and Stede have suggested a promising proposal for annotating arguments using Freeman’s argumentation macrostructure (Peldszus and Stede, 2013). And for biomedical communications, Clark et al have proposed a micropublications ontology based on Toulmin’s model for pay-as-you-go construction of claim-argument networks from scientific papers (Clark et al., 2014). We are using this ontology—the micropublications ontology³—to model evidence about pharmacokinetic drug interactions (Schneider et al., 2014a) in a joint project organized by Richard Boyce.

We have also developed two ontologies related to argumentation. First, WD, the Wiki Discussion ontology⁴ (Schneider, 2014b) was alluded to in Section 2: WD is used for argumentation support for decision-making discussions in ad-hoc online collaboration, applying factors/dimensions theory. Second, ORCA is an Ontology of Reasoning, Certainty and Attribution⁵ (de Waard and Schneider, 2012). Based on a taxonomy by de Waard, ORCA is motivated by scientific argument. ORCA allows distinguishing completely verified facts from hypotheses: it records the certainty of knowledge (lack of knowledge; hypothetical; dubitative; doxastic) as well as its basis (reasoning, data, unidentified) and source (author or other, explicitly or implicitly; or none).

³<http://purl.org/mp/>

⁴<http://purl.org/wd/>

⁵<http://vocab.derii.ie/orca>

4 Mining from Social Media

The third strand of our research is in mining arguments from social media.

4.1 Characteristics of social media

To identify arguments in social media, we need to know where to look. The intention of the author might be relevant, for instance we can expect different types of argument in messages, depending on whether they are recreation, information, instruction, discussion, and recommendation (Schneider et al., 2014b). In (Schneider et al., 2012a), we suggested that relevant features for argumentation in social media may include the genre, metadata, properties of users, goals of a particular dialogue, context and certainty, informal and indirect speech, implicit information, sentiment and subjectivity.

4.2 Information extraction based on argumentation schemes

In a corpus of camera reviews, we examine the argument that consumers give in reviews, focusing on rationales about camera properties and consumer values.

In collaboration with Liverpool researchers including Adam Wyner (Wyner et al., 2012), we describe the argumentation mining task in consumer reviews as an information extraction task, where we fill slots in a predetermined argumentation scheme, such as:

Consumer Argumentation Scheme:

Premise: Camera X has property P .

Premise: Property P promotes value V for agent A .

Conclusion: Agent A should *ActionI* camera X .

Further details of the information extraction are given in (Schneider and Wyner, 2012). In particular, we developed gazetteers for the camera domain and user domain, and selected appropriate discourse indicators and sentiment terminology. These form part of an NLP pipeline in the General Architecture for Text Engineering framework. Resulting annotations can be viewed on a document or searched with a corpus indexing and querying tool, informing an argument analyst who wishes to construct instances of the consumer argumentation scheme.

We have also presented additional argumentation schemes that model evaluative expressions in reviews, focusing in (Wyner and Schneider, 2012) on user models within a context of hotel reviews.

5 Conclusions

We have described our work related to argumentation mining, which uses CSCW, knowledge representation, argumentation theory and information extraction. As we noted, different approaches are appropriate for identifying and modeling arguments in online debates (Schneider, 2014b) versus scientific papers (Schneider et al., 2014a), so different application areas need to be considered. We hope that our questions about argumentation mining—starting with *What counts as ‘argumentation’*, *in the context of the argumentation mining task?* and *How do we measure the success of an argumentation mining task?*—drive the community towards establishing shared tasks. Shared corpora and well-defined tasks are needed to propel argumentation mining beyond a highly discussed area into an agreed upon research challenge.

References

- Tim Clark, Paolo N. Ciccarese, and Carole A. Goble. 2014. Micropublications: a semantic model for claims, evidence, arguments and annotations in biomedical communications. *Journal of Biomedical Semantics*, 5(27), July.
- Anita de Waard and Jodi Schneider. 2012. Formalising uncertainty: An Ontology of Reasoning, Certainty and Attribution (ORCA). In *Joint Workshop on Semantic Technologies Applied to Biomedical Informatics and Individualized Medicine (SATBI+SWIM 2012) at International Semantic Web Conference (ISWC 2012)*.
- Nancy Green, Kevin Ashley, Diane Litman, Chris Reed, and Vern Walker. 2014. Workshop description, first workshop on argumentation mining at the association for computational linguistics.
- John Lawrence, Chris Reed, Colin Allen, Simon McAlister, and Andrew Ravenscroft. 2014. Mining arguments from 19th century philosophical texts using topic based modelling. In *Proceedings of the First Workshop on Argumentation Mining*, pages 79–87. Association for Computational Linguistics.
- Andreas Peldszus and Manfred Stede. 2013. From argument diagrams to argumentation mining in texts: a survey. *International Journal of Cognitive Informatics and Natural Intelligence (IJCINI)*, 7(1):1–31.
- Iyad Rahwan, Bitan Banihashemi, Chris Reed, Douglas Walton, and Sherief Abdallah. 2011. Representing and classifying arguments on the Semantic Web. *The Knowledge Engineering Review*, 26(04):487–511.
- Jodi Schneider and Adam Wyner. 2012. Identifying consumers’ arguments in text. In *SWAIE 2012: Semantic Web and Information Extraction. In conjunction with EKAW 2012*.
- Jodi Schneider, Brian Davis, and Adam Wyner. 2012a. Dimensions of argumentation in social media. *Knowledge Engineering and Knowledge Management*, pages 21–25.
- Jodi Schneider, Tudor Groza, and Alexandre Passant. 2012b. A review of argumentation for the Social Semantic Web. *Semantic Web-Interoperability, Usability, Applicability*, 4(2):159–218.
- Jodi Schneider, Alexandre Passant, and Stefan Decker. 2012c. Deletion discussions in Wikipedia: Decision factors and outcomes. In *Proceedings of the International Symposium on Wikis and Open Collaboration, WikiSym 2012*, pages 17:1–17:10.
- Jodi Schneider, Krystian Samp, Alexandre Passant, and Stefan Decker. 2013. Arguments about deletion: How experience improves the acceptability of arguments in ad-hoc online task groups. In *Proceedings of the ACM conference on Computer Supported Cooperative Work, CSCW 2013*, pages 1069–1080.
- Jodi Schneider, Carol Collins, Lisa Hines, John R Horn, and Richard Boyce. 2014a. Modeling arguments in scientific papers. In *The 12th ArgDIA Conference: From Real Data to Argument Mining*.
- Jodi Schneider, Serena Villata, and Elena Cabrio. 2014b. Why did they post that argument? Communicative intentions of Web 2.0 arguments. In *Arguing on the Web 2.0 at the 8th ISSA Conference on Argumentation*.
- Jodi Schneider. 2014a. Automated argumentation mining to the rescue? Envisioning argumentation and decision-making support for debates in open online collaboration communities. In *Proceedings of the First Workshop on Argumentation Mining*, pages 59–63, Baltimore, Maryland, June. Association for Computational Linguistics.
- Jodi Schneider. 2014b. *Identifying, Annotating, and Filtering Arguments and Opinions in Open Collaboration Systems*. Ph.D. dissertation, Digital Enterprise Research Institute, National University of Ireland, Galway. Corpus and supplementary material also available online at <http://purl.org/jsp/d>.
- Adam Wyner and Jodi Schneider. 2012. Arguing from a point of view. In *First International Conference on Agreement Technologies*.
- Adam Wyner, Jodi Schneider, Katie Atkinson, and Trevor Bench-Capon. 2012. Semi-automated argumentative analysis of online product reviews. In *Computational models of argument: Proceedings of COMMA 2012*.