Working with open argument corpora

JOHN LAWRENCE
Centre for Argument Technology, University of Dundee, UK
j.lawrence@dundee.ac.uk

MATHILDE JANIER
Centre for Argument Technology, University of Dundee, UK
m.janier@dundee.ac.uk

CHRIS REED
Centre for Argument Technology, University of Dundee, UK
c.a.reed@dundee.ac.uk

AIFdb Corpora provides a facility to group Argument Interchange Format (AIF) argument maps and search for maps that are related to each other (for example, analyses of related texts.) Users can create and share corpora containing any number of argument maps from within AIFdb. By integrating with the OVA+ analysis tool, AIFdb Corpora allows for the creation of corpora compliant with both AIF and Inference Anchoring Theory, a philosophically and linguistically grounded counterpart to AIF.

KEYWORDS: argument corpora, Online Visualisation of Argument, Argument Interchange Format, Inference Anchoring Theory

1. INTRODUCTION

The number of Argument Interchange Format (AIF) (Chesñevar et al., 2006) argument maps contained within the open and publicly accessible database, AIFdb[1] (Lawrence, Bex, Reed, & Snaith, 2012), now exceeds 4,000 with over 60,000 individual nodes in eleven different languages. These numbers are growing rapidly and, as an increasing number of argumentation tools, such as Arvina[2], the AnalysisWall (Bex, Lawrence, Snaith, & Reed, 2013) and ArguBlogging (Bex, Snaith, Lawrence, & Reed,

[1]http://www.aifdb.org
[2]http://arvina.arg-tech.org

2014) begin using AIFdb to store their data, the rate of growth is set to increase.

Although AIFdb offers a search interface to locate both a given node and the maps within which that node occurs, there is no real ability to group argument maps or to search for maps that are related to each other in some way (for example, being analyses of related texts.)

AIFdb Corpora[3] offers such an ability, allowing a user to create and share a corpus containing any number of argument maps from within the database. By integrating closely with the OVA+ (Online Visualisation of Argument) analysis tool, AIFdb Corpora allows for the rapid creation of large corpora compliant with both AIF and *Inference Anchoring Theory* (IAT) (Budzynska & Reed, 2011), a philosophically and linguistically grounded counterpart to the AIF.

2. ARGUMENT DATA

The continuing growth in the volume of data which we produce has driven efforts to unlock the wealth of information this data contains. Automatic techniques such as Opinion Mining and Sentiment Analysis (Liu, 2010) allow us to determine the views expressed in a piece of textual data, for example, whether a product review is positive or negative. Existing techniques struggle, however, to identify more complex structural relationships between concepts.

Argument Mining[4] is the automatic identification of the argumentative structure contained within a piece of natural language text. By automatically identifying this structure and its associated premises and conclusions, we are able to tell not just *what* views are being expressed, but also *why* those particular views are held.

The desire to achieve this deeper understanding of the views which people express has led to the recent rapid growth in the Argument Mining field (2014 saw the first ACL workshop on the topic in Baltimore[5] and meetings dedicated to the topic in both Warsaw[6] and Dundee[7]). One of the challenges faced by current approaches to argument mining however, is the lack of large quantities of appropriately annotated arguments to serve as training and test data. Several recent efforts have been made to improve this situation by the creation of corpora across a range of different domains.

For example, (Green, 2014) aims to create a freely available

[3]http://corpora.aifdb.org
[4]Sometimes also referred to as Argumentation Mining
[5]http://www.uncg.edu/cmp/ArgMining2014/
[6]http://argdiap.pl/argdiap2014
[7]http://www.arg-tech.org/swam2014/

corpus of open-access, full-text scientific articles from the biomedical genetics research literature, annotated to support argument mining research. However, there are challenges to creating such corpora, such as the extensive use of biological, chemical, and clinical terminology in the BioNLP domain.

In (Houngbo & Mercer, 2014), a straightforward feature of co-referring text using the word "this" is used to build a self-annotating corpus extracted from a large biomedical research paper dataset. This is achieved by collecting pairs of sequential sentences where the second sentence begins with "This method...", "This result...", or "This conclusion...", and then categorising the first sentence in each pair respectively as Method, Result or Conclusion sentences. The corpus is annotated without involving domain experts and in a 10-fold cross-validation, gives an overall F-score of 0.97 with naïve Bayes and 0.987 with SVM.

Legal texts are the focus of (Walker, Vazirova, & Sanford, 2014), where a type system is developed for marking up successful and unsuccessful patterns of argument in U.S. judicial decisions. Building on a corpus of vaccine-injury compensation cases that report factfinding about causation, based on both scientific and non-scientific evidence and reasoning, patterns of reasoning are identified and used to illustrate the difficulty of developing a type or annotation system for characterising these patterns. A further example of legal material is the ECHR corpus (Mochales & Ieven, 2009), a set of documents extracted from legal texts of the European Court of Human Rights (ECHR). The ECHR, over the years, has developed a standard type of reasoning and structure of argumentation resulting in material which, although not specifically annotated for argumentative content is easily adapted to serve as data for argument mining.

Such efforts add to the volume of currently available data for which at least some elements of the argumentative structure have been identified. The most comprehensive and completely annotated existing collection of such data is the openly accessible database, AIFdb[8] (Lawrence et al., 2012), containing over 4,000 Argument Interchange Format (AIF) argument maps with over 60,000 individual nodes in ten different languages. These numbers are growing rapidly, thanks to both the increase in analysis tools interacting directly with AIFdb and the ability to import analyses produced with the Rationale and Carneades tools (Bex, Gordon, Lawrence, & Reed, 2012).

Additionally, several online tools such as DebateGraph[9],

[8]http://www.aifdb.org
[9]http://debategraph.org

TruthMapping[10], Debatepedia[11], Agora[12], Argunet[13] and Rationale Online[14] allow users to create and share argument analyses. Although these tools are helping to increase the volume of analysed argumentation, they generally do not offer the ability to access this data and each use their own formats for its annotation and storage. In order to help overcome this challenge, AIFdb currently offers the facility to import and convert Rationale analyses into AIF and development is underway to allow for conversion of the DebateGraph and Argunet formats.

In addition to the previously discussed corpora of structured argument data, there are large corpora of unstructured data available that are rich in argumentative structure, from, for example, Wikipedia, Google Books, meeting data from the AMIDA Meeting Corpus[15] annotated using the Twente Argumentation Scheme (Rienks, Heylen, & Weijden, 2005) and product reviews from websites such as Amazon and epinions.com. Whilst these corpora may be useful for certain argument mining techniques, such as those using unsupervised learning methods, the full utilisation of these resources is limited by their lack of annotation. Despite the lack of marked argument structure, Wikipedia, in particular, represents a considerable amount of data rich in argumentative content. IBM's recently announced Debater project,[16] is an argument construction engine utilising a corpus of unstructured Wikipedia text. Debater can respond to a given topic by automatically constructing a set of relevant pro/con arguments phrased in natural language. For example, when asked for responses to the topic "The sale of violent video games to minors should be banned", Debater scanned approximately 4 million Wikipedia articles and determined the ten most relevant articles, scanned all 3,000 sentences in those articles, detected sentences which contain candidate claims, assessed their pro and con polarity and then presented three relevant pro and con arguments.

Although Debater is able to extract simple pro and con reasons from Wikipedia articles, it falls short of being able to offer a detailed understanding of the argumentative structure. In (Aharoni et al., 2014), work towards annotating articles from Wikipedia using a meticulously monitored manual annotation process is discussed. The result is 2,683 argument elements, collected in the context of 33 pre-defined controversial topics, and organised under a simple structure detailing a

---

[10]https://www.truthmapping.com
[11]http://www.debatepedia.org
[12]http://agora.gatech.edu/
[13]http://www.argunet.org/
[14]https://www.rationaleonline.com/
[15]http://corpus.amidaproject.org/
[16]http://www.kurzweilai.net/introducing-a-new-feature-of-ibms-watson-the-debater

claim and its associated supporting evidence.

Another possible avenue for increasing the volume of annotated argument is crowdsourcing, as discussed in (Ghosh, Muresan, Wacholder, Aakhus, & Mitsui, 2014), where a two-tiered approach is proposed to determine which portions of texts are argumentative and what is the nature of argumentation. The first step suggested adopts a coarse-grained annotation scheme based on Pragmatic Argumentation Theory (van Eemeren, Grootendorst, Jackson, & Jacobs, 1993) and asked expert annotators to label entire threads using this scheme. A clustering technique is then used to identify which pieces of text were easier or harder to annotate and it is shown that crowdsourcing is a feasible approach to obtain annotations, particularly on those text segments that were identified as being easier for the Expert Annotators.

The availability of large scale corpora of annotated argument data has a wide range of possible applications, for example, allowing for the comparison of analysis to real world dialogue (Goodwin & Cortes, 2010), determining the validity of argument coding schemes (Pallotta, Seretan, Ailomaa, Ghorbel, & Rajman, 2007), and providing insight into patterns of argument in discourse (O'Halloran, 2011). By building a diverse range of corpora spanning different times and domains, it is possible to perform comparative research into argument usage in a discourse field over time and across discourse fields.

3. AIFDB

The Argument Web (Bex et al., 2013) is a vision for a large-scale Web of inter-connected arguments posted by individuals on the World Wide Web in a structured manner. As such it is necessary to provide a service which not only allows for the storage and retrieval of this structured argument data, but is compatible with the widest possible range of currently existing argumentation software and provides a stable and flexible platform around which future software can be developed. AIFdb is a database implementation of the Argument Interchange Format (AIF) (Chesñevar et al., 2006), allowing for the storage and retrieval of AIF compliant argument structures. AIFdb offers a wide range of web service interfaces for interacting with stored argument data, as well as offering its own search and argument visualisation features all consistent with the formal ontology of the AIF.

At the lowest level, AIFdb's web services allow for the insertion and querying of the basic components of an AIF argument such as nodes, edges and schemes. These components are represented by tables in the database as seen in Figure 1. Building upon these lower level interactions, AIFdb also offers a "middle layer" which groups these simple

queries to allow more complex interactions to be easily performed. For example it is possible, with a single query, to determine all of the statements made by a particular person in support of a given information node (I-node). At the highest level of interaction, AIFdb supports modules handling the import and export of numerous formats such as SVG, DOT, RDF-XML and the formats of the Carneades, Rationale and Araucaria tools.

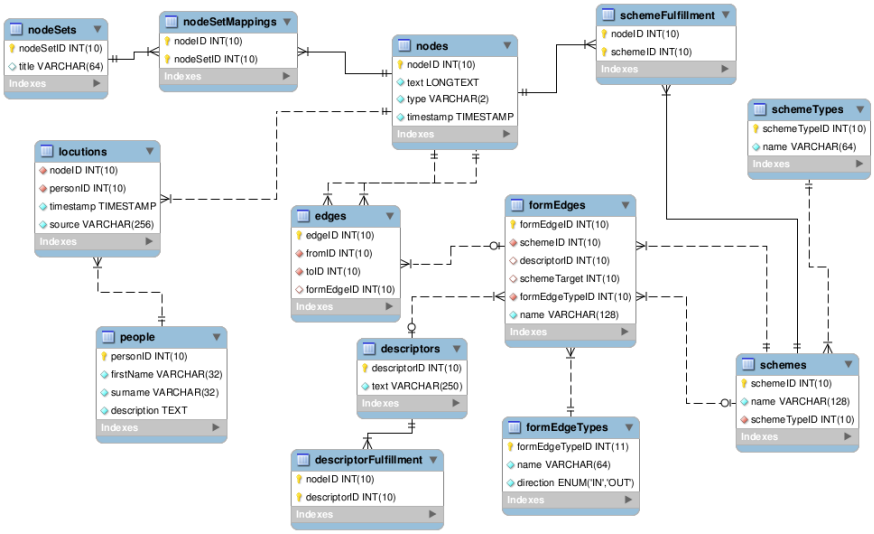

Figure 1 – AIFdb Data Structure Diagram

4. AIFDB CORPORA

AIFdb Corpora extends the functionality offered by AIFdb, allowing a user to create and share a corpus containing any number of argument maps from within the database. Users are able to create their own corpora for specific projects or themes, grouping together argument maps that are related in some way, and enabling these to be viewed together or downloaded in a variety of formats.

4.1 The AIFdb Corpora Interface

In order to create a corpus, the user must specify simple details, including the title of the corpus, a shortname used in the corpus URL, and a brief description, shown in the list of corpora. Once these details are entered, the user is given a unique link to a page where they can edit their corpus. The edit page, as shown in Figure 2, allows the user to update these details, as well as providing text corresponding to the corpus as a whole, for example, if the corpus consists of analyses of different parts of

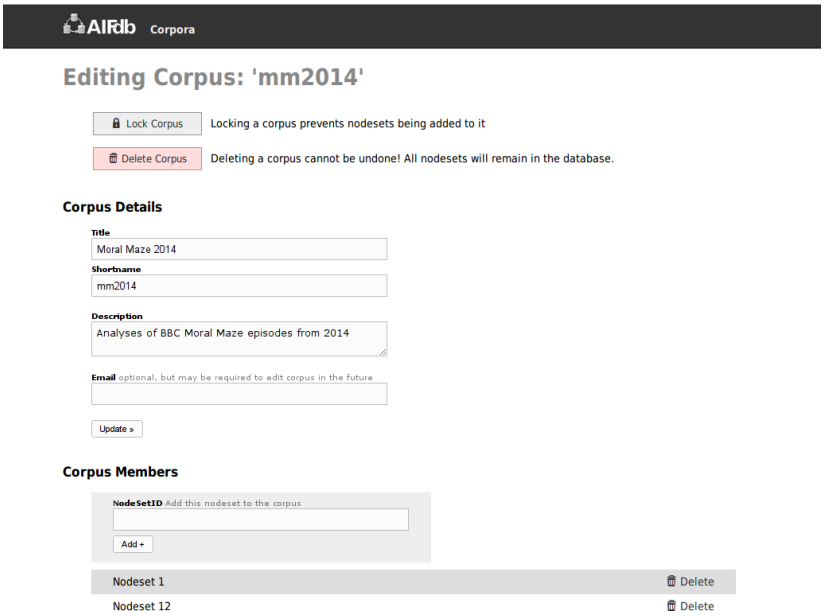

Figure 2 – Corpora Admin Interface

the same text, it can be provided here and will be made available as part of the full corpus download. The edit page also allows the user to manually add individual argument maps to the corpus, and to add sub-corpora; existing corpora that form a part of the corpus being edited. Any subcorpora then act like a part of the parent corpus, for example, when an argument map is added to a sub-corpus, it is added to the parent corpus as well. Additionally, the user may lock the corpus, preventing any other applications from adding maps to it.

AIFdb Corpora also offers interfaces to list and search corpora as seen in Figure 3, and to display corpora, as seen in Figure 4. This allows a user to share a link to their corpus with others and allows for easy viewing and downloading of the corpus contents, either as individual files in SVG, PNG, DOT, JSON, RDF-XML, Prolog and the formats of the Carneades (Gordon, Prakken, & Walton, 2007) and Rationale (van Gelder, 2007) tools, or as an archive file containing JSON format representations of each argument map as well as the original text of the entire corpus. The display interface also provides links to view, evaluate and edit any of the argument maps contained within the corpus in OVA or OVA+.

Although AIFdb does not allow for the storage of the text corresponding to an argument map, an additional database is available to store these texts and tools such as OVA are able to store the original text for analyses in this database. When the archive file for a particular corpus is generated this database is queried and the text for each individual

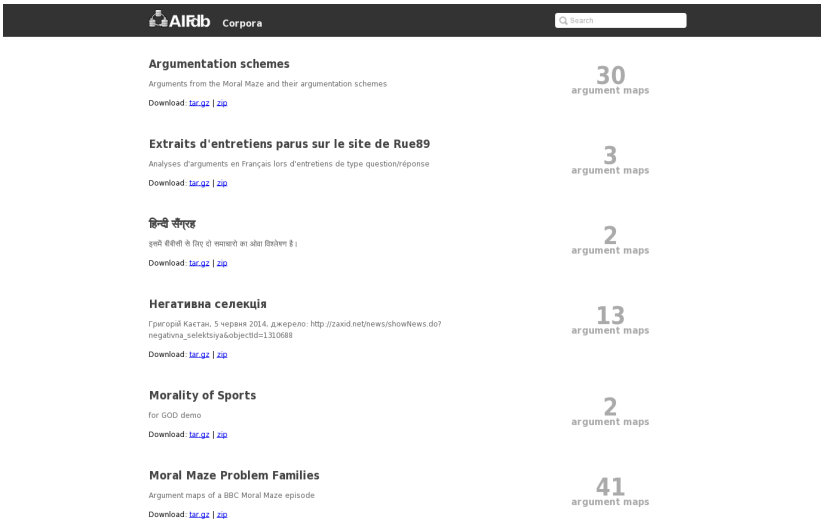

Figure 3 – Corpus Listing

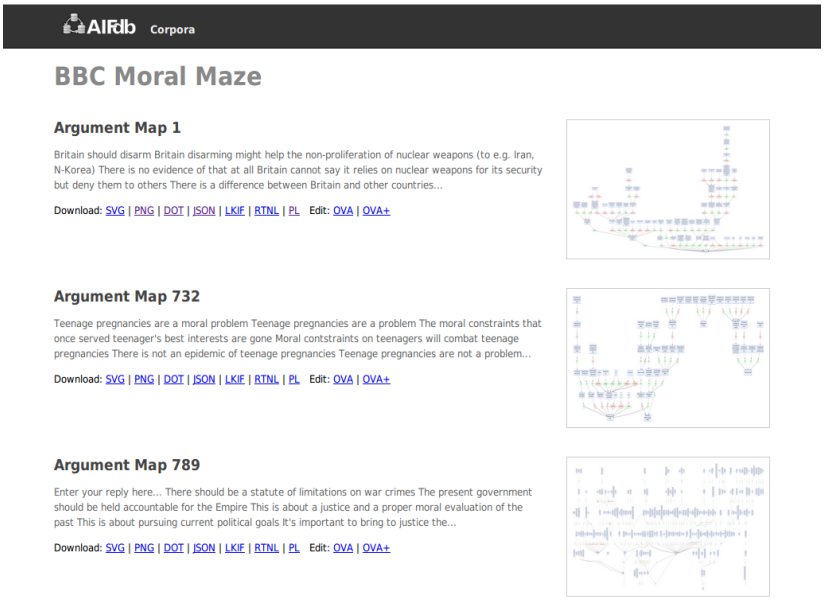

Figure 4 – Corpus Details

argument map added to the archive file. If no text has been provided for the corpus as a whole, then the individual texts for each argument map are concatenated and provided together. In cases where a corpus contains subcorpora, again, if no text has been provided for the corpus as a whole, the text for any individual maps is joined with any text that has been given for the subcorpora.

4.2 Integration with OVA+

OVA+ offers a web based interface for the analysis of arguments and is accessible from any web browser at http://ova.arg-tech.org/. The tool was built as a response to the AIF theory and allows for the creation and representation of argument structure, combined with the ability to exchange, share and reuse the resultant argument maps. The system relies on the Inference Anchoring Theory (IAT), with the schemes provided allowing for a graphical representation of the argumentative structure of a text and, more interestingly, of dialogues.

OVA+ handles texts of any type and any length. To begin the analysis, the first relevant utterance for argumentation must be extracted in order to create an information node (I-node). Then it is possible to create the locution node associated (L-node) and to specify the name of the speaker (participant); the locution appears, preceded by the name of the participant assigned to it, and arrows link the L-node to the I-node via a YA-node.

YA-nodes are the illocutionary forces of locutions, and can be given a scheme provided by the IAT model. Each following utterance can be annotated accordingly. According to the AIF, it is possible to evidence supports or attacks between arguments. An RA-node (application of a rule of inference) should connect two I-nodes. To elicit an attack between arguments, RA-nodes can be changed into CA-nodes, namely applications of a pattern of conflict. Linked arguments can be established by connecting all the arguments to the proper scheme-node (RA or CA). According to IAT, it is also possible to indicate the transitions (TA-nodes) between locutions by linking two L-nodes. Finally, it is possible for the analyst to assign the illocutionary forces anchored in the transitions. This can be done thanks to the set of IAT schemes which are proposed when a TA-node has been linked to its corresponding scheme-node.

At the end of the analysis, OVA+ permits saving the work on the user's computer as either an image file or as a JSON format file representing the AIF structure. Most interestingly, however, OVA+ offers the possibility of saving an analysis to AIFdb and its further addition to any corpus in AIFdb Corpora. This ability allows for analyses to be reused via AIFdb, and consequently any of the growing number of argument web

tools which use AIFdb as their data store; or loaded in OVA+ for consultation or correction, as well as allowing for the rapid and collaborative creation of large AIF and IAT compliant corpora.

4.3 AIFdb Corpora Usage

AIFdb Corpora already collects over 1,000 analyses into a range of corpora, the largest of which are described below:

AraucariaDB An import of 667 argument analyses produced using Araucaria(Reed & Rowe, 2004) and stored in the Araucaria database (Reed, 2006).

AraucariaDBpl A selection of over 50 Polish language analyses created using the Polish version of Araucaria (Budzynska, 2011).

Digging By Debating Argument Study Collection of analyses of 19th century philosophical texts from the Hathi Trust collection, created for the Digging by Debating project[17].

Moral Maze 2012 Analyses of episodes of the BBC Moral Maze Radio 4 programme from 2012. These analyses are split into two sub-corpora **MM2012a** (containing the first four programmes) and **MM2012b** (containing the last four). **MM2012** comprises the full run of 8 programmes from the 2012 summer season. These corpora of analysed arguments have already been used in different projects, in particular the one described in (Budzynska, Janier, Kang, et al., 2014) and (Budzynska, Janier, Reed, et al., 2014).

Argumentation Schemes Examples of occurrences of Walton's argumentation schemes (Walton, Reed, & Macagno, 2008) found in episodes of the BBC Moral Maze Radio 4 programme

Expert Opinion and Positive Consequences Examples of the Expert Opinion and Positive Consequences argumentation schemes taken from online news sources

Dispute Mediation Argument maps of mediation session transcripts.

This corpus was created to facilitate the retrieving of analysed excerpts of mediation transcripts, as part of the DrEAMS poject (Dialogue-based Exploration of Arguments and Mediation Space) in the University of Dundee. The research project indeed aims at exploring discourse in dispute mediation through analyses of the dialogues between disputants and mediators. A repository of the

[17]http://diggingbydebating.org/

argument maps (currently 65) revealed necessary in order to keep trace of the phenomena proper to mediation highlighted by these argument maps. This corpus is actually composed of 2 subcorpora: **Dispute mediation: Excerpts taken from publications** and **Mock mediation**. Each corpus of AIFdb Corpora can indeed be integrated to a larger one. In addition to their practical aspect for the project in itself, those corpora offer a valuable data set for the argumentation community, as the analysed mediation dialogues can be freely consulted and downloaded.

Language Opposition Corpus Argument maps of online multi-party interactions

This is currently the largest corpus (1946 argument maps distributed through five sub corpora). It is used in Rutgers for the SALTS project[18], the goals of which are to advance the understanding of how expression of argumentation shapes the ebb and flow of online interactions, and to develop computational models capable of identifying and characterizing the expression of argumentation in multi-party interactions.

ECC - Bank of America and ECC - The Coca Cola Company . Argument maps of Earning Conference Calls.

The argument maps are the result of a project at Universita della Svizzera Italiana which explores the dialogues between CEOs or managers of big companies and financial analysts during Earning Conference Calls.

Негативна селекція (Negative Selection) Analyses of a journal article discussing the Dean's election at Lvov University and the effects of the post soviet system on modern Ukranian academia.

5. CONCLUSION

It is hoped that by making the process of creating and updating a corpus as simple as possible, usage will continue to grow and that AIFdb Corpora will prove to be a useful tool for collecting and sharing AIF argument maps. By close integration with analysis tools such as OVA+, AIFdb Corpora allows for the rapid, collaborative creation of AIF and IAT compliant corpora, and so will offer a valuable resource in areas such as argumentation mining which have a demand for large quantities of such annotated material.

[18]http://salts.rutgers.edu/

ACKNOWLEDGEMENTS: The authors would like to acknowledge that the work reported in this paper has been supported in part by the RCUK Lifelong Health and Wellbeing Programme grant number EP/K037293/1 - BESiDE: The Built Environment for Social Inclusion in the Digital Economy, and in part by the British Leverhulme Trust under grant RPG-2013-076.

REFERENCES

Aharoni, E., Polnarov, A., Lavee, T., Hershcovich, D., Levy, R., Rinott, R., … Slonim, N. (2014, June). A benchmark dataset for automatic detection of claims and evidence in the context of controversial topics. In *Proceedings of the first workshop on argumentation mining* (pp. 64–68). Baltimore, Maryland: Association for Computational Linguistics.

Bex, F., Gordon, T. F., Lawrence, J., & Reed, C. (2012). Interchanging arguments between Carneades and AIF – Theory and practice. In *Proceedings of the fourth international conference on computational models of argument (comma 2012)* (pp. 390–397). Vienna: IOS Press.

Bex, F., Lawrence, J., Snaith, M., & Reed, C. (2013, Oct). Implementing the argument web. *Communications of the ACM*, *56*(10), 66–73.

Bex, F., Snaith, M., Lawrence, J., & Reed, C. (2014). Argublogging: An application for the argument web. *Web Semantics: Science, Services and Agents on the World Wide Web*.

Budzynska, K. (2011). Araucaria-PL: Software for teaching argumentation theory. In *Proceedings of the third international congress on tools for teaching logic (ticttl 2011)* (pp. 30–37).

Budzynska, K., Janier, M., Kang, J., Reed, C., Saint-Dizier, P., Stede, M., & Yaskorska, O. (2014). Towards argument mining from dialogue. In *Proceedings of the fifth international conference on computational models of argument (comma 2014)* (pp. 185–196). IOS Press.

Budzynska, K., Janier, M., Reed, C., Saint-Dizier, P., Stede, M., & Yaskorska, O. (2014). A model for processing illocutionary structures and argumentation in debates. In *Proceedings of the 9th edition of the language resources and evaluation conference (lrec)* (pp. 917–924).

Budzynska, K., & Reed, C. (2011). *Whence inference* (Tech. Rep.). University of Dundee.

Chesñevar, C., Modgil, S., Rahwan, I., Reed, C., Simari, G., South, M., … others (2006). Towards an argument interchange format. *The Knowledge Engineering Review*, *21*(04), 293–316.

Ghosh, D., Muresan, S., Wacholder, N., Aakhus, M., & Mitsui, M. (2014, June). Analyzing argumentative discourse units in online interactions. In *Proceedings of the first workshop on argumentation mining* (pp. 39–48). Baltimore, Maryland: Association for Computational Linguistics.

Goodwin, J., & Cortes, V. (2010). Theorists' and practitioners' spatial metaphors for argumentation: A corpus-based approach. *Verbum*, *23*, 163–78.

Gordon, T. F., Prakken, H., & Walton, D. (2007). The carneades model of argument and burden of proof. *Artificial Intelligence*, *171*(10), 875–896.

Green, N. (2014, June). Towards creation of a corpus for argumentation mining the biomedical genetics research literature. In *Proceedings of the first workshop on argumentation mining* (pp. 11–18). Baltimore, Maryland: Association for Computational Linguistics.

Houngbo, H., & Mercer, R. (2014, June). An automated method to build a corpus of rhetorically-classified sentences in biomedical texts. In *Proceedings of the first workshop on argumentation mining* (pp. 19–23). Baltimore, Maryland: Association for Computational Linguistics.

Lawrence, J., Bex, F., Reed, C., & Snaith, M. (2012). AIFdb: Infrastructure for the argument web. In *Proceedings of the fourth international conference on computational models of argument (comma 2012)* (pp. 515–516).

Liu, B. (2010). Sentiment analysis and subjectivity. *Handbook of natural language processing*, *2*, 627–666.

Mochales, R., & Ieven, A. (2009). Creating an argumentation corpus: do theories apply to real arguments?: a case study on the legal argumentation of the echr. In *Proceedings of the 12th international conference on artificial intelligence and law* (pp. 21–30).

O'Halloran, K. (2011). Investigating argumentation in reading groups: Combining manual qualitative coding and automated corpus analysis tools. *Applied linguistics*, *32*(2), 172–196.

Pallotta, V., Seretan, V., Ailomaa, M., Ghorbel, H., & Rajman, M. (2007). Towards an argumentative coding scheme for annotating meeting dialogue data. In *Proceedings of the 10th international pragmatics conference, ipra* (pp. 8–13).

Reed, C. (2006). Preliminary results from an argument corpus. In E. M. Bermudez & L. R. Miyares (Eds.), *Linguistics in the twenty-first century* (pp. 185–196). Cambridge Scholars Press.

Reed, C., & Rowe, G. (2004). Araucaria: Software for argument analysis, diagramming and representation. *International Journal on Artificial Intelligence Tools*, *13*(4), 961–980.

Rienks, R., Heylen, D., & Weijden, v. d. E. (2005). Argument diagramming of meeting conversations. In *Multimodal multiparty meeting processing, workshop at the 7th intl. conference on multimodal interfaces.* IOS Press.

van Eemeren, F. H., Grootendorst, R., Jackson, S., & Jacobs, S. (1993). *Reconstructing argumentative discourse.* University of Alabama Press.

van Gelder, T. (2007). The rationale for rationale. *Law, probability and risk*, *6*(1-4), 23–42.

Walker, V., Vazirova, K., & Sanford, C. (2014, June). Annotating patterns of reasoning about medical theories of causation in vaccine cases: Toward a type system for arguments. In *Proceedings of the first workshop on argumentation mining* (pp. 1–10). Baltimore, Maryland: Association for Computational Linguistics.

Walton, D., Reed, C., & Macagno, F. (2008). *Argumentation schemes*. Cambridge University Press.