

Report on STSM 1868, Agreement Technologies COST Action ICO801

Jodi Schneider, with Adam Wyner, Katie Atkinson, and Trevor Bench-Capon

February 8, 2012

My COST internship at Liverpool has been very productive, and we expect this research to continue towards three near-term outcomes:

1. Building and using an argument exploration tool.
2. Analysing argumentation in online product reviews using a domain ontology and the argument exploration tool.
3. Finding rhetorical templates for argumentation schemes

We discuss these three lines of research, and our progress and next directions, in turn.

1 Argument Exploration Tool

Our argumentation exploration tool aims to highlight areas of interest in written texts in order to direct an analyst’s attention to areas likely to contain argumentative content. The theory underlying the tool is that rhetorical relations are frequently signaled by cue words such as ‘but’, ‘finally’, and ‘therefore’. These words have been extensively studied in linguistics, and after surveying a number of lists of cue words [9, 12, 2, 14, 11, 10], we have settled on the Penn Discourse Treebank 2.0 [15]. We think that this is the right starting place for looking at argumentation structure through rhetorical cues—it’s a thoroughly annotated, well-researched sample from prominent linguists.

Towards building an argumentation exploration tool, we have implemented a General Architecture for Text Engineering (GATE) [5] text analysis pipeline, using a gazetteer and JAPE [6] rules that annotate and highlight each discourse relation found in list of Senses and Explicit Connectives from the Appendix B of the Penn Discourse Treebank 2.0 Annotation Manual. This pipeline will benefit from further development. There are certain general linguistic problems that will need to be addressed, ideally with existing GATE modules, or otherwise with our own work. First of all, certain multiword phrases (e.g. “either..or”, “if..then”, etc.) are currently not handled correctly. Second, certain discourse relations admit of variability; thus generalisation is needed to correctly handle

the precedence (e.g. “a day or two before”, “two years before”, etc.) and succession “18 months after”, “two days after”, ...) cues. Similar generalisation will be needed, if continuation is used, (e.g. “first(ly)”, “second(ly)”). By writing additional JAPE rules, we should be able to handle all of these. Areas for further near-term exploration include the relationship between sentence boundaries and rhetorical relations; do we need to restrict relations to within a single sentence or paragraph?

Further, we have identified several corpora of interest, for use with the argument exploration tool. Our current work on analysing product reviews (discussed in the next section) focuses on a small corpus of Amazon¹ reviews. For a popular camera, we selected all the 3-star (5 reviews, 3 A4 pages) and 5-star reviews (10 reviews, 6 A4 pages) for analysis.

We envision future work using our pipeline on the TREC Twitter corpus² as the starting point for a submission to The 18th International Conference on Knowledge Engineering and Knowledge Management (EKAW). We also have the Open American National Corpus³ (chosen since it is a large, well-known corpus, and since it is freely available), but have no current projects in mind for its use.

One particularly interesting corpus is unavailable to us. Although we have also investigated acquiring access to the Penn Discourse Treebank Corpus 2.0, in order to check our pipeline against its carefully hand-annotated examples, its cost to non-members is prohibitive (\$1000 for the stand-off annotations⁴; but the Wall Street Journal section of Treebank-2⁵ (\$3150 for non-members) is needed for surrounding context).

2 Analysing Argumentation in Online Product Reviews

Reviews are a social space with purposeful discussions that include argumentative content. Our work on the online review domain is inspired in part by Stella Heras’ work on critical questions on Amazon [8] from COMMA 2010. Rather than focusing on critical questions, we want to model argumentation schemes in online reviews. The objective is to develop an analysis to argue about camera buying along the lines of the practical reasoning scheme [1]:

Result Action: You should *Action1* camera X.

Properties Premise: X has properties A, B, C.

Value Premise: *Action1* (or Property B) promotes value V1.⁶

¹<http://amazon.co.uk>

²<http://trec.nist.gov/data/tweets/>

³<http://americannationalcorpus.org/OANC/>

⁴<http://www ldc.upenn.edu/Catalog/CatalogEntry.jsp?catalogId=LDC2008T05>

⁵<http://www ldc.upenn.edu/Catalog/CatalogEntry.jsp?catalogId=LDC95T7>

⁶While strictly speaking, only actions promote values, in the context of product purchase recommendations, there is an implication that a property can also promote a value, in so

We expect reviews to discuss actions, properties, and values, and are interested in investigating the concrete features (properties) and abstract ones (values), as well as what actions promote which values for which agents. We note that value promotion depends on the person: the action of Katie buying a camera promotes the value of *portability* for her since it fits in her purse, but the action of Trevor buying the same camera does *not* promote this value for him, since the camera doesn't fit in his trouser pocket. That is, arguments about properties are objective while arguments about values are inherently subjective. What counts as promoting a value depends on the particular agent; further, any particular value may or may not be important to that agent. Understanding this agent-specific worldview is crucial for successful persuasion of a particular agent.

To look for such argument scheme patterns in reviews, we are using our argument exploration tool and a domain ontology. Thus far, we have analysed our Amazon camera reviews corpus (discussed above) by hand, to identify actions, objects, properties, and values. For properties, both binary values (such as 'has a flash') and ranges (such as the number of megapixels or scope of the zoom) are common. Similarly, values promotion can be either effectively absolute ('promotes value V1') or scalar ('promotes portability'). For properties that range, and values that can be promoted to varying degrees, it is important to know which direction is better. Further we note that acceptable and appropriate ranges vary over time, as products change.

Towards creating a review-oriented camera ontology, we have made an initial candidate feature list. We have also explored existing ontologies; two are outdated toy ontologies [4, 3]. A third photography ontology [7] (which developed from [3]) is worth further consideration, and comparison with our feature list. We have made the observation that ontologies for argumentation are purpose-dependent. Different ontologies are needed for evaluation versus construction versus summarising a debate.⁷ So the key question is what do you need in an ontology to support purchase decisions? We need to clarify and explore this question. Our domain ontology will be used to identify important terms and relations in the reviews, which may stand in certain 'grammatical' positions in their instantiations of argumentation schemes. This needs further exploration, in combination with the argumentation explorer.

far as that property is conveyed with product ownership. For example, a camera with a fast shutter speed can be construed to promote the value of ease, because owning such a camera would allow a person to take pictures rapidly.

⁷As we have discussed, two arguments may have the same premises and conclusions but different intermediate argumentation strategies, which opens them to different attacks—making them the same for some purposes, but different for others.

3 Finding Rhetorical Templates for Argumentation Schemes

Using our argumentation explorer, we hope to find the template forms and sentence relationships that signal the rhetorical moves associated with several specific argumentation schemes such as practical reasoning, credible source, or expert testimony. This seems plausible, based around the idea that the rhetorical cue words that signal discourse relations are commonly used in writing arguments. We may want to describe some notion of ‘signaled argumentation schemes’ to describe the argumentation schemes found near the rhetorical cue words that signal discourse relations.

4 Future Work: Exploring the Empirical Basis of Argumentation Schemes

We are also interested in exploring the empirical basis of signaled argumentation schemes in written texts. From a small set of instantiated argumentation schemes evaluated against a particular corpus (Amazon camera reviews), we could broaden scope in two respects. First, in that corpus, are there instances of the schemes that lack signalling indicators? Second, in other, larger corpora, are there instances of schemes, signalled or not? Both will further help us understand the use of argumentation schemes in practice.

5 Outcomes

The outcomes so far are an early draft of a COMMA submission, a first pass at a GATE pipeline using discourse relations, an initial candidate feature list for a review-oriented camera ontology, and some statistics about prevalence of the original list of cue words in our corpus.

Personal outcomes include great advances on a theoretical and practical level. I have received extensive feedback on revisions for a submitted (now accepted) journal article, as well as comments on a workshop paper for the 2012 ACM Conference on Computer Supported Cooperative Work (CSCW). I have increased my familiarity with relevant argumentation theory literature. Even some fundamentals were new to me in some cases. For example, that there are three ways of attacking an argument: rebuttal–i.e. ‘the conclusion is false’, undercutting–i.e. ‘this argument is generally ok, but doesn’t work here’ by undermining the premise, or undermining–i.e. Since P, then Q. Not P. Another example is the distinction between justification and explanation: one justifies one’s own actions, but explains others’ actions – based on Joseph Raz’ practical reasoning [13].

On a practical level, I have gained sufficient comfort with GATE to imple-

ment and run a pipeline and to use GATE's Lucene⁸ search (ANNIC) to explore the results. I have also had significant support in focusing my work.

References

- [1] K. Atkinson, T. Bench-Capon, and P. McBurney. Computational representation of practical argument. *Synthese*, 152(2):157–206, 2006.
- [2] G. Bordea. First steps toward semi-automatic extraction of claims from scientific publications. Master's thesis, Politechnic University of Bucharest, Faculty of Automation and Computer Science, 2009.
- [3] Co-ode project from Bio Health Informatics Group at The University of Manchester Department of Computer Science. Photography ontology. <http://www.co-ode.org/ontologies/photography/>.
- [4] R. L. Costello and D. B. Jacobs (Mitre Corporation). Camera OWL Ontology. <http://www.xfront.com/camera/camera.owl>.
- [5] H. Cunningham, D. Maynard, K. Bontcheva, V. Tablan, N. Aswani, I. Roberts, G. Gorrell, A. Funk, A. Roberts, D. Damljjanovic, T. Heitz, M. A. Greenwood, H. Saggion, J. Petrak, Y. Li, and W. Peters. *Text Processing with GATE (Version 6)*. 2011.
- [6] H. Cunningham, D. Maynard, and V. Tablan. JAPE: a Java Annotation Patterns Engine (Second Edition). Technical report CS-00-10, Department of Computer Science, University of Sheffield, 2000.
- [7] N. Drummond. Photography ontology v2. <http://www.photographyontology.org.uk/>.
- [8] S. Heras, K. Atkinson, V. Botti, F. Grasso, V. Juliana, and P. McBurney. How argumentation can enhance dialogues in social networks. In *Computational Models of Argument - Proceedings of COMMA 2010*, number 216 in *Frontiers in Artificial Intelligence and Applications*, pages 267–274. IOS Press, 2010.
- [9] A. Knott. *A Data-Driven Methodology for Motivating a Set of Coherence Relations*. PhD thesis, University of Edinburgh, 1996.
- [10] A. Lascarides. List of discourse markers. Personal communication, February 2, 2012.
- [11] W. C. Mann and S. A. Thompson. Rhetorical Structure Theory: Toward a functional theory of text organization. *Text*, 8(3):243–281, 1988.
- [12] D. Marcu. *The rhetorical parsing of natural language texts*. PhD thesis, Department of Computer Science, University of Toronto, 1997.

⁸<http://lucene.apache.org/>

- [13] J. Raz. *Practical reason and norms*. Oxford University Press, Oxford; New York, 1999.
- [14] C. Sporleder and A. Lascarides. Using automatically labelled examples to classify rhetorical relations: A critical assessment. *Natural Language Engineering*, 14(3):369–416, 2008.
- [15] The PDTB Research Group. The Penn Discourse Treebank 2.0 Annotation Manual. Technical report IRCS-08-01, Institute for Research in Cognitive Science, University of Pennsylvania, 2008.