

Toward Stance Classification Based on Claim Microstructures

Filip Boltužić and Jan Šnajder

University of Zagreb, Faculty of Electrical Engineering and Computing
Text Analysis and Knowledge Engineering Lab
Unska 3, 10000 Zagreb, Croatia
{filip.boltuzic, jan.snajder}@fer.hr

Abstract

Claims are the building blocks of arguments and the reasons underpinning opinions, thus analyzing claims is important for both argumentation mining and opinion mining. We propose a framework for representing claims as *microstructures*, which express the beliefs, judgments, and policies about the relations between domain-specific concepts. In a proof-of-concept study, we manually build microstructures for over 800 claims extracted from an online debate. We test the so-obtained microstructures on the task of claim stance classification, achieving considerable improvements over text-based baselines.

1 Introduction

In online discussions, users express their opinions using more or less well structured arguments. The building blocks of these arguments are *claims*: statements that are in dispute and that we are trying to support with reason Govier (2013). Claims can support or attack other claims, giving rise to complex argumentative structures. Thus, the ability to identify and analyze claims in text is a crucial part of argumentation mining (Moens, 2014; Lippi and Torroni, 2016). Outside the realm of well-structured argumentation, the ability to analyze claims is crucial for tasks such as stance classification (Anand et al., 2011; Hasan and Ng, 2013; Mohammad et al., 2016) and fine-grained opinion analysis (Stoyanov and Cardie, 2008; Yang and Cardie, 2013), as well as the converging task of argument-based opinion mining (Clos et al., 2014; Boltužić and Šnajder, 2014), which aims to uncover the reasons underpinning the opinions.

Previous research has tackled the claim detection task for diverse domains, including legal docu-

ments (Palau and Moens, 2009), microtexts (Peldszus and Stede, 2015), Wikipedia articles (Aharoni et al., 2014; Levy et al., 2014; Rinott et al., 2015), student essays (Stab and Gurevych, 2017), and user-generated web discourse (Habernal and Gurevych, 2015). Boltužić and Šnajder (2015) addressed the task of identifying prominent claims in online debates, while Boltužić and Šnajder (2016) analyzed the implicit premises between two claims. Recently, Bar-Haim et al. (2017) introduced the claim stance classification task, where classification is done at the claim rather than document level.

In this paper, we address the task of claim analysis from a different angle. While prior work has dealt with claims as textual fragments, we study the possibility of a more precise, domain-specific analysis of claims based on their internal logical structure. The work closest to ours is that of Wyner and Van Engers (2010) and Wyner et al. (2016), who explored normalizing claims from the policy making domain by translating them to Attempto Controlled English (Fuchs et al., 2008), and then mapping them to propositions. In contrast, we propose a framework for representing claims as *microstructures*: structures expressing the relations between the domain-specific concepts, reflecting the beliefs, value judgments, or desired policies of the claim author. We present a preliminary proof-of-concept study, where we use the proposed framework to manually create microstructures for over eight hundred claims extracted from an online debate.

We envisage that claim microstructures could play an important role in a variety of opinion mining and argument mining tasks, including stance classification, extraction of argumentative structures, analysis of implicit premises, fine-grained opinion mining, identifying prominent claims, and claim matching. To demonstrate the viability of claim microstructures for a downstream task, we look into supervised claim stance classification

and show that, even with a simple encoding of microstructures as features, we get substantial improvements on this task over text-based baselines.

The contribution of our work may be summed up as follows: (1) we investigate the feasibility of using microstructures for representing claims, (2) we demonstrate the use of microstructures for stance classification, and (3) to promote further research, we make available the dataset annotated with claim paraphrases and microstructures.¹

2 Claim Microstructures

We introduce a framework for representing claims from text using logical *microstructures* whose purpose is to capture the gist of a claim. The initial motivation came from the analysis of our dataset (cf. Section 4), which revealed that a large majority of claims can be conceived of as expressing *relations* between *concepts* using a certain *modality*. Figure 1 shows a claim microstructure bringing together these three elements.

Relations. Many claims can be represented as expressing a relation between two concepts. For example, on the topic of gay rights, the relations may be ‘promotes(GayMarriage, Depopulation)’ or ‘purpose(Love, Procreation)’. There are also comparably fewer claims that can be expressed via higher-order relations, e.g., ‘entails(Constitution, allow(State, GayMarriage))’. Each relation can be negated, e.g., ‘¬promotes(GayMarriage, Depopulation)’ expresses that gay marriage does not cause depopulation.

Concepts. The relations are established between concepts, expressed by noun phrases. For ease of access, these can be arranged into a small, domain-specific taxonomy of concepts. For instance, “gay marriage”, “heterosexual marriage”, and “religious marriage” all belong under the concept of “marriage”. The taxonomic relations could also be useful for later computational processing. Unlike relations, concepts are domain dependent and need to be defined for each new topic.

Modalities. We furthermore observed that the claims express different *modalities*, which can roughly be categorized into *beliefs*, *value judgments*, and *policies*. We formalize this via unary relations ‘believes’, ‘approves’, and ‘desires’, corresponding to beliefs (factual, religious,

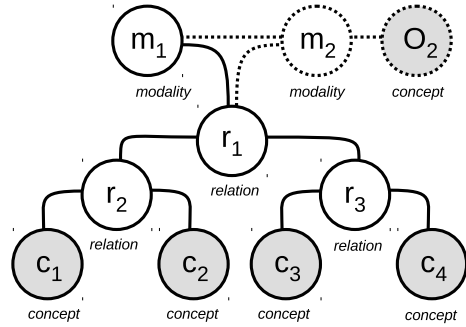


Figure 1: Claim microstructure (2nd-order).

and opinion-based), positive value judgment, and desired policy (desired state of affairs), respectively. The three modalities act as a wrapper on the propositional content of the claim, effectively modulating what is being claimed. For instance, ‘believes(purpose(Love, Procreation))’ expresses the belief that love serves procreation, while ‘desires(¬allow(State, GayMarriage))’ expresses the wish for the state not to allow gay marriages. Finally, we observed that in a fair number of cases the claims are supported by a reference to a second opinion holder (e.g., the Bible, the state). We represent this by introducing one additional modality layer with the opinion holder as an additional modifier. For instance, ‘believes(believes[State](promotes(Marriage, Advancement)))’ corresponds to the belief that the state believes gay marriages lead to an advancement. By convention, the opinion holder of the first modality is always the author of the post.

Let \mathcal{R} , \mathcal{C} , and \mathcal{M} denote the set of relations, concepts, and modalities, respectively. Formally, we define a claim microstructure as a quadruple (m_1, m_2, o_2, r) , where $m_1 \in \mathcal{M}$ and (optionally) $m_2 \in \mathcal{M} \cup \{\epsilon\}$ are the modalities, $o_2 \in \mathcal{C} \cup \{\epsilon\}$ is the (optional) second opinion holder, and $r = (t, c_1, c_2) \in \mathcal{R}$ is the (possibly higher-order) relation between two concepts or relations $c_1, c_2 \in \mathcal{C} \cup \mathcal{R}$, conveyed by the relation type t . Table 1 defines the relation types used in this work.

It should be noted that, unlike Aharoni et al. (2014), who consider as claims only the statements that directly support or contest the debating topic, we consider all statements with propositional content. For example, in the context of gay rights, ‘belief(purpose(Life, Love))’ is a valid claim in our framework, although it does not support nor contest the topic, i.e., the stance of that claim is neutral.

¹<http://takelab.fer.hr/claim-micro>

Relation	Definition
promotes(A, B)	Promoting agent A promotes, fosters, leads, increases likelihood, boosts B.
suppress(A, B)	Suppressing agent A suppresses, decreases likelihood, puts down, vanquishes B
allow(A, B)	Principle A allows, approves, licenses state of affairs B
entails(A, B)	State of affairs A, necessarily, per definition or causally, makes B true.
contradicts(A, B)	State of affairs A, necessarily, per definition or causally, makes B false.
purpose(A, B)	The purpose of A is B.
equal(A, B)	State of affairs A is equal to state of affairs B.
has(A, B)	A has the properties affected by the existence of B.

Table 1: Relations types in claim microstructures.

3 Data Annotation

We adopted the dataset of Hasan and Ng (2014), which contains user posts from online two-sided debates on a number of topics. For reasons of feasibility, in this study we consider only one topic: “Gay rights”. We sampled 100 posts (50 for and 50 against) from this topic. The manual annotation was carried out in two steps. In the first step, the annotators segmented out the individual claims from user posts and paraphrased them into well-articulated claims. In the second step, the annotators translated each paraphrased claim into the corresponding logical microstructure.

While in principle the claim microstructures could have been built directly from segments, we chose to introduce the additional step of claim paraphrasing for three reasons. First, we assumed that paraphrasing would help in identifying the segments corresponding to individual claims, since paraphrasing demonstrates understanding. In that respect, our work is similar to that of Wyner and Van Engers (2010), who used a controlled language for paraphrasing the claims. Second, we assumed that paraphrases will make overt the logical structure of claims, making their translation into microstructures easier. Lastly, we assumed that paraphrases could help in identifying the prominent concepts for the domain-specific taxonomy.

3.1 Claim Segments and Paraphrases

The purpose of the this step was to extract claim segments from user posts, thus separating argumentative from non-argumentative content, and to

paraphrase the claims into simple, well-articulated statements. This obviously involves two non-trivial tasks: segmentation and paraphrasing. Arguably, there are many ways in which a post can be segmented into claims, and even more ways in which each segment could be paraphrased. We hypothesize that much of this ambiguity can be reduced by considering these two tasks jointly, and by adopting certain paraphrasing principles aimed at obtaining *simplifying paraphrases* – paraphrases that expresses the essence of the claims devoid of superfluous words and phrases. To this end, we adopted the following nine paraphrasing principles: (1) **Argumentativeness** – Only argumentative text should be paraphrased; (2) **Atomicity** – A claim should convey a single thought; (3) **Authority** – Experts in claims from expert opinion should be made explicit in the paraphrase; (4) **Brevity** – Paraphrases should keep only the relevant argumentative content; (5) **Canonicity** – Canonical terms and phrases are preferred over idiomatic language; (6) **Contextuality** – Claims should be paraphrased by considering their local and topical context as well as their context; (7) **Declarativity** – paraphrases should be in declarative form, and (8) **Dereferencing** – Pronouns and nominal references should be resolved; and (9) **Explicitness** – Only explicitly stated information should be paraphrased, and not whatever might be implied by the claim.

The annotation was carried out by one trained annotator and took 25 hours. The 100 user posts yielded 920 claim segments and the same number of paraphrases. Table 2 gives an example. Note that generally the claim segments may overlap, though this is not the case in this example. Overall, the segments covered 79.6% of the text, while the remaining 20.4% may be considered non-argumentative.

3.2 Claim Microstructures

In the second step, we asked two annotators (A1 and A2) to translate each of the 920 paraphrases into claim microstructures. The annotators were provided with a domain-specific taxonomy on “Gay rights”, compiled based on a manual analysis of claim paraphrases. The taxonomy consists of 150 concepts arranged into a tree of a maximum depth of four. The annotators were instructed to use the existing concepts from the taxonomy, and introduce new ones only if they could not find a suitable one in the taxonomy. They were also instructed not to use microstructures of order higher than two.

User post	Claim segment	Claim paraphrase	Claim microstructure
<i>Men should fall in love with women that's why they where created and women should get married to men because it makes everything easier.</i>	<i>Men should fall in love with women.</i>	<i>People of opposite sex should fall in love.</i>	desires(entails(OppositeSex, FallingInLove))
	<i>that's why they where created</i>	<i>Men and women are created to pair.</i>	believes(purpose(MenAndWomen, Procreation))
	<i>women should get married to men because it makes everything easier.</i>	<i>Heterosexual marriages make everything easier.</i>	believes(entails(HeterosexualMarriage, Normal))

Table 2: An example of a user post segmented into three claim segments, each paraphrased and translated into the corresponding claim microstructure.

Out of 920 claim paraphrases, annotator A1 managed to translate 882 claims into 707 distinct microstructures, while annotator A2 translated 842 claims into 767 distinct microstructures. The average annotation effort was 33 hours. The number of claims for which both annotators provided a microstructure is 819 (89%), while the number of claims for which both provided an identical microstructure is only 58 (6.3%). The annotators introduced a total of 157 new concepts, indicating that the initial taxonomy was of too limited a scope. The low annotator agreement and the relatively large number of newly added concepts suggest that a fair amount of ambiguity exists in translating paraphrases to microstructures. Our analysis revealed that, in the majority of cases, the ambiguity is genuine and in such cases having more candidate microstructures for a single claim can be considered advantageous.

The analysis also revealed that ‘believes’ is the most frequent modality, used for about 79% of claims. For A1, *entails* is by far the most common relation (61%), while A2 made a more balanced use of relations, with the top two being *has* (21%) and *entails* (15%). The concepts most frequently used by A1 are *homosexuality*, *homosexual people* and *marriage*, while for A2 these are *The Bible*, *homosexual people*, and *government interest*.

4 Stance Classification

4.1 Setup

Annotation. We consider claim stance classification as one potential application of claim microstructures. To this end, we asked two annotators to label the stance of each of 819 claim paraphrases on a five-point scale: *strong favor* (F), *likely favor* (f), *neither* (N), *likely against* (a), and, *strong against* (A). We adopt the definition of F, N, and A stance from Mohammad et al. (2016), with the

Claim paraphrase	Stance	
	A1	A2
<i>Gay couples should be able to experience parenting.</i>	F	F
<i>Gay couples don't have children.</i>	N	N
<i>By natural means, infertility is wrong.</i>	a	a
<i>A homosexual relationship lacks the ability to procreate.</i>	a	A

Table 3: Claim stance annotations.

Features	Regression	Classification		
	5-way	5-way	3-way-N	3-way-E
<i>seg-w2v</i>	0.084	0.230	0.259	0.383
<i>seg-tfidf</i>	0.133	0.170	0.248	0.297
<i>par-w2c</i>	0.133	0.170	0.248	0.297
<i>par-tfidf</i>	0.250	0.290	0.487	0.377
<i>ms-onehot</i>	0.316	0.320	0.507	0.473
<i>ms-path</i>	0.331	0.315	0.501	0.462

Table 4: Stance classification macro-averaged F1-score using segments (*seg*), paraphrases (*par*), and microstructures (*ms*) as features. The best result in each group is shown in boldface.

addition of the in-between options (f and a) for indicating implicit or indirect stance. Table 3 shows some examples. On the five classes, we observe a moderate inter-annotator agreement of 0.53 Cohen’s κ (Cohen, 1960). The aggregation was done by first removing 16 instances on which the annotators disagreed in stance polarity, and then averaging and rounding the two labels by treating them as numbers from the $[-2, +2]$ interval.

Baselines. We compare against two baselines, which, to the best of our knowledge, are considered state of the art for stance classification (Sobhani et al., 2016): (1) a sum of skip-gram vec-

tors (Mikolov et al., 2013)² for each word and a (2) tf-idf unigram and bigram representation of a claim. For baselines, we use these representations on claim segments (*seg*). For the sake of completeness, we also run the baselines on claim paraphrases (*par*), but note that this serves only as a reference, as obtaining paraphrases is arguably a task that is more difficult to automate than obtaining microstructures.

Microstructures. To represent the claim microstructures (*ms*), we adopted a simple one-hot encoding scheme: we use one one-hot vector for each of the modalities, relations, relation negations, concepts, and opinion holders concatenating the vectors into a single feature vector (*onehot*). In addition, to leverage the taxonomical relations between concepts, we experimented with encoding for each concept its ancestors in the taxonomy, by encoding the nodes along the path leading from the root to the concept (*path*).

Models. We used support vector machine (SVM) classifier and regression models with an RBF kernel, as implemented in the LibSVM library of Chang and Lin (2011). We trained and evaluated the models on 803 claim instances (either segments, paraphrases, or microstructures) using a 5×3 nested cross-validation, using grid search to optimize hyperparameters C and γ .

Tasks. We considered four task: (1) a 5-way regression setup, in which the model is trained to predict the numeric stance score, but afterwards the predictions are rounded and mapped to labels, (2) a 5-way classification task, (3) a 3-way classification task in which the implicit labels (a and f) are mapped to neutral (3-way-N), and a (4) 3-way classification task in which the implicit labels are mapped to explicit for and against labels (3-way-E). The last two tasks are easier, so we expected the models to perform better on these tasks.

4.2 Results

Table 4 shows the classification results in terms of macro-averaged F1-score. As expected, the 3-way classification tasks are easier than 5-way classification tasks. Furthermore, the 5-way regression model performs better than 5-way classifier, suggesting that using distance-sensitive loss is beneficial for this task. In all four tasks, the

²We use the pre-trained vectors available at <https://code.google.com/p/word2vec/>

claim microstructures considerably outperform both segment-based baselines, yielding between 9 and 25 points of improvement in F1-score, depending on the task. All differences between the baseline and the microstructure model are statistically significant at $p < 0.05$ (tested using a two-tailed permutation test (Yeh, 2000)). By comparing with claim paraphrases as a reference, we find that microstructures give comparable performance for 5-way and 3-way-N classification tasks (the differences are not statistically significant at $p < 0.05$), while for 5-way regression and 3-way-E classification tasks the microstructures outperform paraphrase representations. Finally, the performance difference between one-hot encoded microstructures and microstructures with path-encoded concepts are not statistically significant at $p < 0.05$, suggesting that stance classification did not profit from encoding taxonomical relations.

In the above experiments, the results for microstructures were obtained on annotations of A1. The models trained on annotations of A2 gave consistently lower performance, albeit still better (and statistically different) than the baseline.

We conclude the experimental section by noting that microstructures improve claim stance classification performance over a segment-based baseline by a maximum of 50.7% F1-score for a 3-way classification setup.

5 Conclusion and Future Work

We presented a framework for representing the microstructures of claims. A microstructure expresses the relations between domain-specific concepts, and is intended to capture the beliefs, value judgments, and desired policies conveyed by claims. In the proof-of-concept study, we manually annotated microstructures for one debating topic. The annotators were able to translate 89% of claims into microstructures, thus proving the viability of the approach. We next demonstrated the usefulness of microstructures on the task of claim stance classification, where a simple encoding of microstructures yielded notable performance improvements over segment-based baselines. This in turn suggests that a claim microstructure does a good job in capturing the argumentative gist of the claim.

We note, however, that this is a preliminary study, which has left aside some important practical issues. While our results are promising, the major question now is how to automatically extract the

microstructures from text. In our study, the claims were segmented and paraphrased by human annotators; an end-to-end system would need to both segment out the claims and extract the corresponding microstructures. We believe that one way to tackle this problem might be to frame it as an information extraction task.

In our preliminary study, the annotators managed to translate most of the claims into microstructures. However, the low agreement rate (6.3%) suggests that the annotation workflow could perhaps be improved.

Another issue worth investigating is the application of the framework to a new domain: the tedious work of deriving a domain-specific taxonomy of concepts and the microstructures could perhaps be alleviated using active learning methods.

Finally, it would of course be interesting to investigate the use of microstructures in other opinion mining and argument mining tasks, including tasks that could profit from analyzing the logical links between claims. We intend to pursue some of these directions in future work.

References

- Ehud Aharoni, Anatoly Polnarov, Tamar Lavee, Daniel Hershcovich, Ran Levy, Ruty Rinott, Dan Gutfreund, and Noam Slonim. 2014. 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*, pages 64–68.
- Pranav Anand, Marilyn Walker, Rob Abbott, Jean E Fox Tree, Robeson Bowmani, and Michael Minor. 2011. Cats rule and dogs drool!: Classifying stance in online debate. In *Proceedings of the 2nd workshop on computational approaches to subjectivity and sentiment analysis*, pages 1–9. Association for Computational Linguistics.
- Roy Bar-Haim, Indrajit Bhattacharya, Francesco Dinuzzo, Amrita Saha, and Noam Slonim. 2017. Stance classification of context-dependent claims. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, pages 251–256. Association for Computational Linguistics.
- Filip Boltužić and Jan Šnajder. 2014. Back up your stance: Recognizing arguments in online discussions. In *Proceedings of the First Workshop on Argumentation Mining*, pages 49–58. Association for Computational Linguistics.
- Filip Boltužić and Jan Šnajder. 2015. Identifying prominent arguments in online debates using semantic textual similarity. In *Proceedings of the 2nd Workshop on Argumentation Mining*, pages 110–115. Association for Computational Linguistics.
- Filip Boltužić and Jan Šnajder. 2016. Fill the gap! Analyzing implicit premises between claims from online debates. In *Proceedings of the 3rd Workshop on Argument Mining*, pages 124–133. Association for Computational Linguistics.
- Chih-Chung Chang and Chih-Jen Lin. 2011. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):27.
- Jérémie Clos, Nirmalie Wiratunga, Joemon Jose, Stewart Massie, and Guillaume Cabanac. 2014. Towards argumentative opinion mining in online discussions. In *Proceedings of the SICSA Workshop on Argument Mining (the Scottish Informatics & Computer Science Alliance)*, page 10.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- Norbert E Fuchs, Kaarel Kaljurand, and Tobias Kuhn. 2008. Attempto controlled English for knowledge representation. In *Reasoning Web*, pages 104–124. Springer.
- Trudy Govier. 2013. *A Practical Study of Argument, Enhanced Edition*, 7th edition. Cengage Learning.
- Ivan Habernal and Iryna Gurevych. 2015. Exploiting debate portals for semi-supervised argumentation mining in user-generated web discourse. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2127–2137.
- Kazi Saidul Hasan and Vincent Ng. 2013. Stance classification of ideological debates: Data, models, features, and constraints. In *International Joint Conference on Natural Language Processing*, pages 1348–1356.
- Kazi Saidul Hasan and Vincent Ng. 2014. Why are you taking this stance? Identifying and classifying reasons in ideological debates. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 751–762. Association for Computational Linguistics.
- Ran Levy, Yonatan Bilu, Daniel Hershcovich, Ehud Aharoni, and Noam Slonim. 2014. Context dependent claim detection. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*.
- Marco Lippi and Paolo Torrioni. 2016. Argumentation mining: State of the art and emerging trends. *ACM Transactions on Internet Technology (TOIT)*, 16(2):10.

- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *Proceedings of ICLR*, Scottsdale, AZ, USA.
- Marie-Francine Moens. 2014. Argumentation mining: Where are we now, where do we want to be and how do we get there? In *Post-proceedings of the forum for information retrieval evaluation (FIRE 2013)*.
- Saif M. Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. SemEval-2016 Task 6: Detecting stance in tweets. In *Proceedings of SemEval-2016*, pages 31–41. Association for Computational Linguistics.
- Raquel Mochales Palau 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. ACM.
- Andreas Peldszus and Manfred Stede. 2015. Joint prediction in mst-style discourse parsing for argumentation mining. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 938–948. Association for Computational Linguistics.
- Ruty Rinott, Lena Dankin, Carlos Alzate Perez, Mitesh M Khapra, Ehud Aharoni, and Noam Slonim. 2015. Show me your evidence—an automatic method for context dependent evidence detection. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 440–450. Association for Computational Linguistics.
- Parinaz Sobhani, Saif M Mohammad, and Svetlana Kiritchenko. 2016. Detecting stance in tweets and analyzing its interaction with sentiment. In *Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics (*SEM 2016)*, pages 159–168. Association for Computational Linguistics.
- Christian Stab and Iryna Gurevych. 2017. Parsing argumentation structures in persuasive essays. *Computational Linguistics*, (in press).
- Veselin Stoyanov and Claire Cardie. 2008. Topic identification for fine-grained opinion analysis. In *Proceedings of the 22nd International Conference on Computational Linguistics-Volume 1*, pages 817–824. Association for Computational Linguistics.
- Adam Wyner, Tom van Engers, and Anthony Hunter. 2016. Working on the argument pipeline: Through flow issues between natural language argument, instantiated arguments, and argumentation frameworks. *Argument & Computation*, 7(1):69–89.
- Adam Wyner and Tom Van Engers. 2010. A framework for enriched, controlled on-line discussion forums for e-government policy-making. In *Proceedings of eGov 2010*, pages 1–11.
- Bishan Yang and Claire Cardie. 2013. Joint inference for fine-grained opinion extraction. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, pages 1640–1649. Association for Computational Linguistics.
- Alexander Yeh. 2000. More accurate tests for the statistical significance of result differences. In *Proceedings of the 18th conference on Computational linguistics-Volume 2*, pages 947–953. Association for Computational Linguistics.