

Argumentation Technology for Artificial Intelligence

Part 2: Argument Mining and Assessment

Philipp Cimiano, Benno Stein, and **Henning Wachsmuth**

September 24th, 2019



Bauhaus-Universität
Weimar



Argumentation in natural language

- **Real-world arguments**

- Mostly not logically valid
- Leave much implicit
- May be hidden in longer utterances
- May be split over multiple utterances
- May depend on the context

”If you wanna hear my view I think that the death penalty should be abolished. It legitimizes an irreversible act of violence. As long as human justice remains fallible, the risk of executing the innocent can never be eliminated.”

Alice. *I think a university degree is important. Employers always look at what degree you have first.*

Bob. *LOL ... everyone knows that practical experience is what does the trick.*

Alice. *Good point! Anyway, in doubt I would always prefer to have one!*

- **Can we actually use formal approaches?**

- Long story short: To a wide extent, yes
- But we need to *mine* arguments and *assess* their properties before

Argument mining and assessment: Outline

1. Argument mining

- How to find argument units and relations in text?

think that the death penalty should be abolished.
support ↑ support
act of violence. As long as human justice remains

2. Stance classification


- How to determine whether an argument is pro or con?



<https://pixabay.com/>


3. Overall argumentation analysis

- How to analyze longer argumentative structures?



4. Argumentation quality assessment

- How to judge whether an argument is good or bad?



5. Fallacy detection

- How to identify argumentative flaws in discussions?

"Again, how old are you?"

1. Argument Mining

Argument mining: Introduction

- **Argument mining** (aka argumentation mining)
 - Automatic identification of arguments in natural language text
 - Core task for natural language argumentation

” *If you wanna hear my view I think that **the death penalty should be abolished** .*
Premise *It legitimizes an irreversible act of violence .* **Conclusion** *As long as human justice remains*
fallible , the risk of executing the innocent can never be eliminated . **Premise**

support ↗ ↖ *support*

- **Three main argument mining steps**
Different task decompositions found in literature
 - Segmenting a text into argument units and other parts
 - Classifying the type or role of each unit
 - Identifying and classifying relations between units

Unit segmentation

- **Argument units** (aka argumentative discourse units)
 - Text segments with an argumentative function
Usually, the premises and conclusions of arguments
- **Unit segmentation**
 - **Task.** Given a text, segment it into argument units and other parts
 - **Method.** Usually, token-level sequence labeling (more on this below)

” If you wanna hear my view I think that the death penalty should be abolished. “

○ ○ ○ ○ ○ ○ ○ ○ ○ ○ **B** **I** **I** **I** **I** **I** ○

- **Challenges**
 - Unit granularity differs: Anything between clauses and paragraphs
 - Usually the first step: Unclear what are the arguments
- **State of the art** (Ajjour et al., 2017)
 - Rather reliable on narrow genres (F_1 0.72–0.82), unsolved across them

Background: Evaluation measures

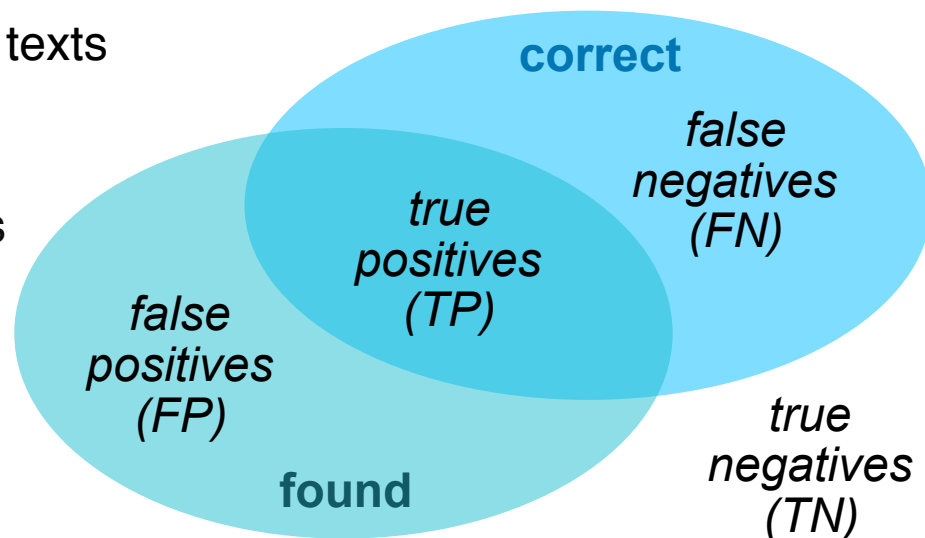
- **Data-driven research**

- Models and methods developed on training texts
- Most methods not fully "correct"
- Effectiveness evaluated on test texts

- **Effectiveness measures**

- **Accuracy**. Used if both positives and negatives important

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$



- **Precision, recall, and F_1 -score**. Used if positives in the focus

$$\text{Precision } (P) = \frac{TP}{TP + FP} \quad \text{Recall } (R) = \frac{TP}{TP + FN} \quad F_1\text{-score} = \frac{2 \cdot P \cdot R}{P + R}$$

Unit classification

- **Unit classes**

- Claim and evidence types
(Rinott et al., 2015; Al-Khatib et al., 2017)
- Roles within argumentation
(Stab and Gurevych, 2014; Habernal and Gurevych 2015)
- Often corpus-specific

assumption statistics other
anecdote common ground testimony

claim premise none
major claim

- **Unit classification**

- **Task.** Given an argument unit, assign one class from a set of classes
- **Method.** Usually, supervised text classification (more on this below)

Conclusion
that *the death penalty should be abolished.*
violence . As long as human justice remains
innocent can never be eliminated . ” **Premise**

- **State of the art**

- Reliable on "explicit" argumentation, such as essays (F_1 0.87) (Stab, 2017)
- Still rather reliable on news editorials (F_1 0.77) (Al-Khatib et al., 2017)
- Minority classes may be problematic, though

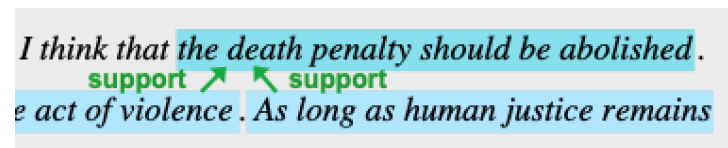
Relation identification and classification

- **Argumentative relations**

- Within arguments (premises to conclusion) or between them (arg to arg)
- **Types.** Support or attack, partly more fine-grained

- **Relation identification and classification**

- **Task.** Given two units /arguments, what relation holds between them, if any
- **Method.** Various, e.g., computing the minimum spanning tree
(Peldszus and Stede, 2015)



- **State of the art**

- Semi-reliable on narrow genres, such as essays (F_1 0.73) (Stab, 2017)
- Identification works better than classification
- Relations hard to agree on for "hidden" arguments, such as in editorials

- **Related tasks**

- Given an argument, classify its argumentation scheme (Feng and Hirst, 2011)
- Given an argument, find the best counterargument (Wachsmuth et al., 2018)

2. Stance Classification

Stance classification: Introduction

- **Stance**

- Overall position of a person towards some target, such as an issue or statement
- To have/take a stance on a target means to be *pro* or *con* towards it
(Somasundaran and Wiebe, 2010)



<https://pixabay.com>

Con towards "death penalty"
" If you wanna hear my view I think that *the death penalty should be abolished* .
Pro towards claim above
It legitimizes an irreversible act of violence . As long as human justice remains fallible , the risk of executing the innocent can never be eliminated . " **Pro** towards claim above


- **Stance classification**

- Determination of the stance towards a target encoded in a text span
- **In argumentation**. Conceptual overlap with relation classification, but usually stance refers to the issue at discussion

Background: Supervised text classification

- **Text classification**

- **Task.** Given a text, assign one class from a set of classes
Stance classification is a text classification problem
- Usually done with supervised machine learning



- **Feature-based classification**

- Map text to feature vector, map feature vector to class label
Features engineered manually or semi-automatically
- **Models.** Support vector machines, random forest, ...

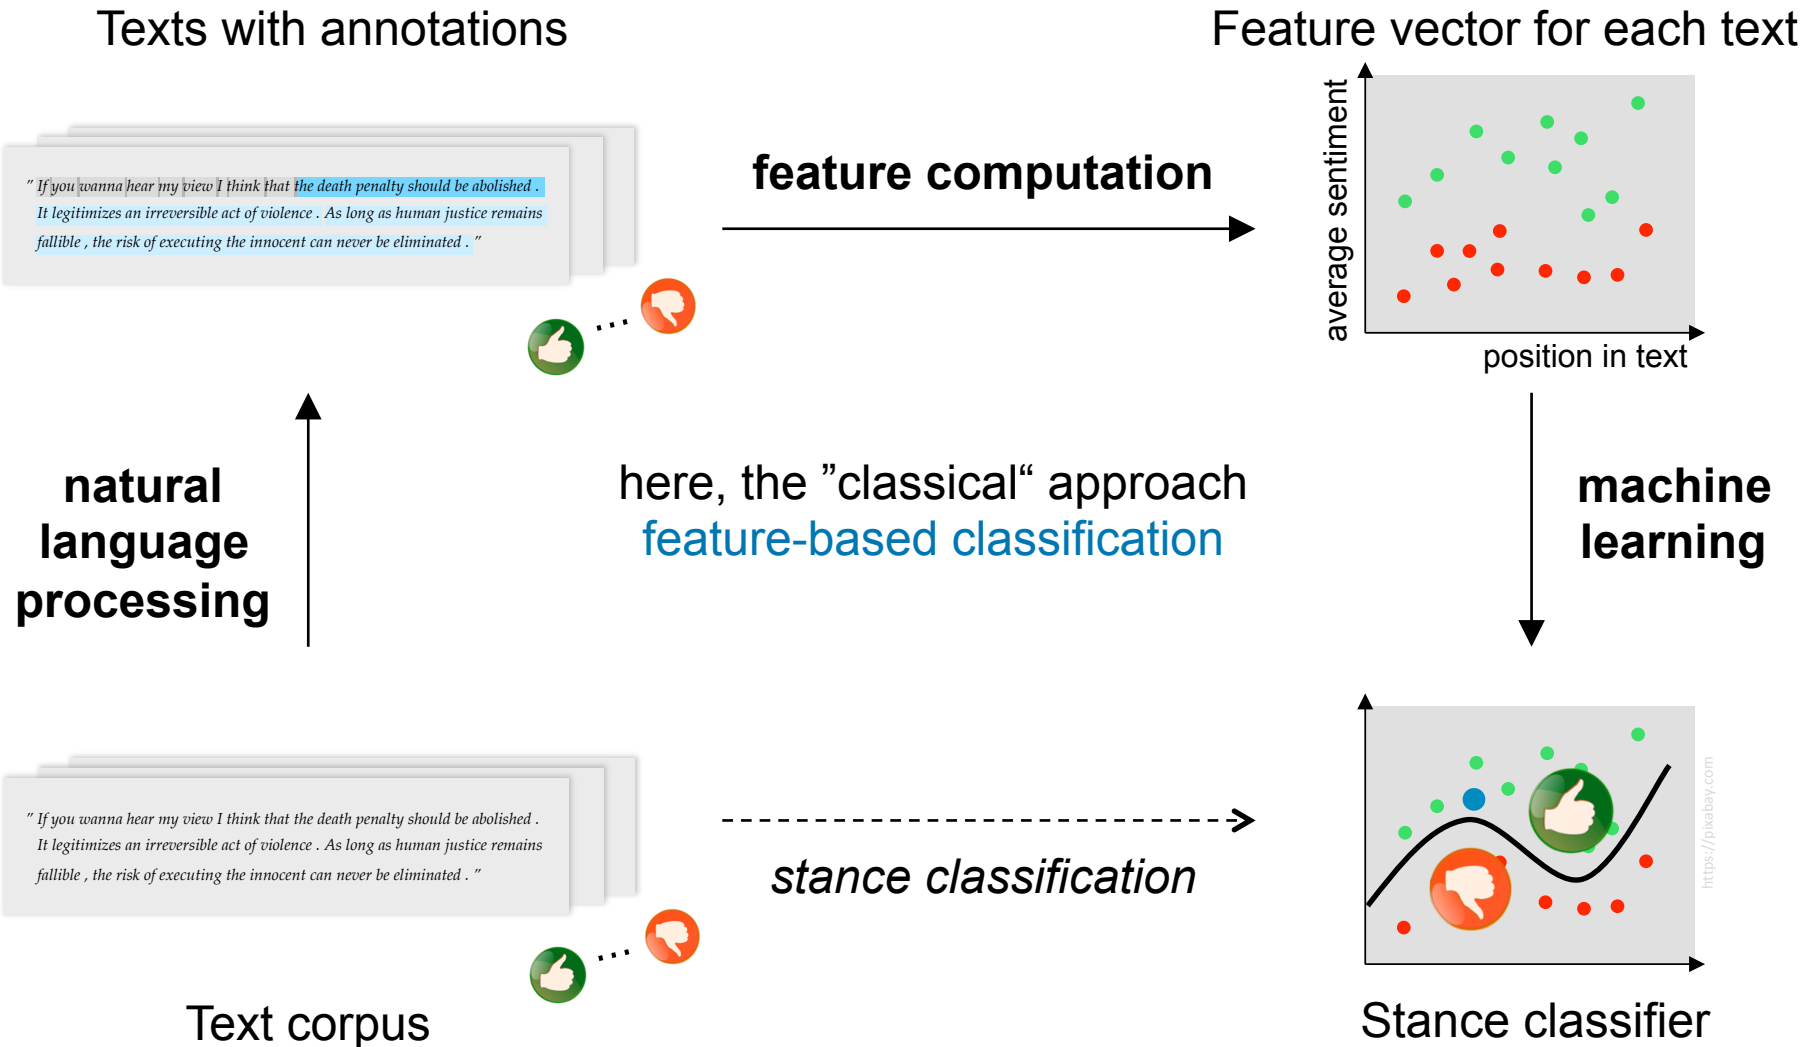
- **Neural classification** (usually works better, given enough data)

- Features (weights in neural networks) learned automatically
- **Models.** Convolutional neural networks, bi-directional LSTMs, ...

- **Sequence labeling** (applicable, when a sequence of texts is classified)

- Like other techniques, but considering previous classifications
- **Models.** Conditional random fields, recurrent neural networks, ...

How to develop a stance classification algorithm



Modeling stance

- **Candidate features of the text** (Somasundaran and Wiebe, 2010, Hasan and Ng, 2013)

- **Bag-of-words.** Distribution of words or word n-grams
 - **Core vocab.** Terms from subjectivity lexica
 - **Discourse.** Connectives + relations between units
 - **Sentiment.** Aspect or topic-directed polarity
- ... and many more...

→ accuracy up to 0.70

stance tend to be the same

- **Candidate features of the context**

- Exploit author knowledge in dialog → up to 0.74
(Ranade et al., 2013)
- Exploit opposing views in dialog → up to 0.75
(Hasan and Ng, 2013)
- Connections between topics of claim and target
(Bar-Haim et al., 2017)

→ 0.84 for most confident 10%, 0.65 overall (3 classes)

Alice: I think a **university degree** is important. Employers always look at what degree you have first.

↓ stance tend to be opposite

Bob: LOL ... everyone knows that **practical experience** is what does the trick.

Alice: Good point! Anyway, in doubt I would always prefer to have one!

3. Overall Argumentation Analysis

(Wachsmuth et al., 2017c)



Giovanni
Da San Martino



Dora
Kiesel



Benno
Stein

Overall argumentation analysis: Introduction

The death penalty is a legal means that as such is not practicable in Germany.


For one thing, inviolable human dignity is anchored in our constitution,

and further no one may have the right to adjudicate upon the death of another human being.

Even if many people think that a murderer has already decided on the life or death of another person,

this is precisely the crime that we should not repay with the same.

(Peldszus and Stede, 2016)



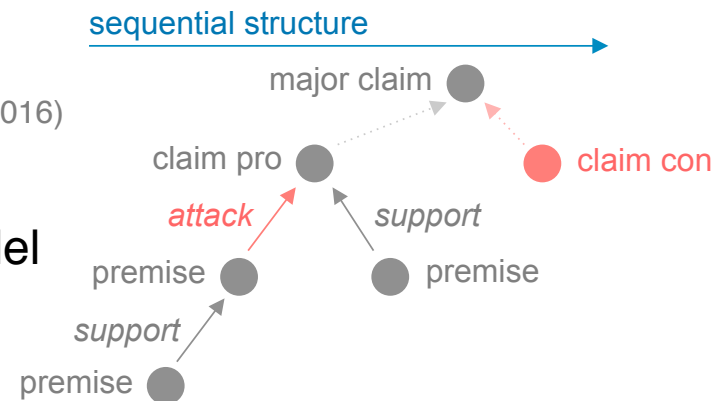
- **Hypothesis**

- Overall structure is decisive for downstream analysis tasks

Analysis tasks and corpora

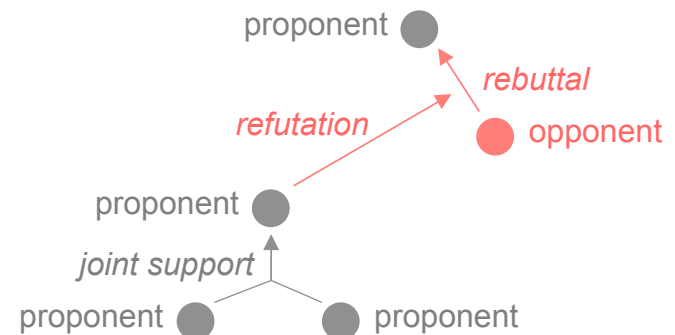
- **Myside bias on AAE corpus** (Stab and Gurevych, 2016)

- 402 persuasive student essays
- 15.1 units/text, proprietary argument model
- 251 one-sided, 151 two-sided




- **Stance on Arg-Microtexts** (Peldszus and Stede, 2016)

- 112 short argumentative texts
- 5.1 units/text, model of Freeman (2011)
- 46 pro stance, 42 con stance, 24 unlabeled



- **Genre on Web Discourse** (Habernal and Gurevych, 2015)


- 340 argumentative web texts
- 3.4 units/text, modified model of Toulmin (1958)
- 216 comments, 46 blog posts, 73 forum posts, 5 articles



Study: The impact of overall argumentation

- **Research questions**

1. How to jointly model sequential and hierarchical overall structure?
2. How important is overall structure in analysis tasks?



- **Background: Kernel methods in machine learning**

- Representation of instances in implicit feature space
- Similarity function used by classifier (e.g., SVM)
- Strong when good features unknown and/or data limited

- **Kernels for structured data**

- Subsequence kernel for sequential structure


(Mooney and Bunescu, 2006)

- Tree kernel for hierarchical structure

(Collins and Duffy, 2001)


- **Route kernel: Tree kernel with positions**

(Aiolli et al., 2009)




Experiments for each analysis task

- Overall argumentation approaches




frequencies

linear kernel




sequences

subsequence kernel




hierarchies

tree path kernel



routes


route kernel



- Baseline approaches

majority

always majority class



pos

linear kernel



tokens

linear kernel




- Experiments on ground-truth argument corpora


- SVM for each kernel in repeated 10-fold cross-validation
- Hyperparameter tuning, fairness in training

Accuracy results


- Myside bias on AAE-v2



- Stance on Arg-Microtexts



- Genre on Web Discourse



4. Argumentation quality assessment

Argumentation quality assessment: Introduction

- **Argumentation quality**

- Natural language arguments rarely logically valid
- Need to quantify how *strong* an argument or argumentation is.

” Is a strong argument an effective argument which gains the adherence of the audience, or is it a valid argument, which ought to gain it? “

(Perelman and Olbrechts-Tyteca, 1969)

- **Argumentation quality assessment**

- Absolute or relative judgment of specific quality dimensions
- Identification of flaws and fallacies

- **Critical for any application**

- **Argument search.** What argument to rank highest?
- **Writing support.** What argumentative flaws does a text have?
- **Decision making.** Which arguments outweigh others?



What to assess and how to assess it

- **What to assess**

- Several, partly very subjective quality dimensions
- Different granularity levels

cogent?

effective?

reasonable?

” *If you wanna hear my view I think that the death penalty should be abolished.*

It legitimizes an irreversible act of violence . As long as human justice remains

fallible , the risk of executing the innocent can never be eliminated. ” **acceptability:**
4 out of 5

acceptable?

clear?

relevant?

more acceptable than

- **How to assess**

- Absolute or relative?
- How *should* we argue vs. how *do* we argue?
- Based on manual assessments or ”objective“ properties?
- Include model of audience?

”*Human beings never act freely and thus should not be punished for even the most horrific crimes. “*

Three aspects of argumentation quality

$$\frac{A \quad A \rightarrow B}{B}$$

Logic

"A dialectical discussion derives its reasonableness from a dual criterion: problem validity and intersubjective validity."

van Eemeren (2015)

$$\frac{A \quad A \rightarrow B}{B}$$
$$\frac{B \rightarrow C}{C}$$

Dialectic



<https://de.wikipedia.org>

"An argument is cogent if its premises are relevant to its conclusion, individually acceptable, and together sufficient to draw the conclusion."

Blair (2012)

Argumentation quality

Rhetoric

$$\frac{A \quad A \rightarrow B}{B}$$

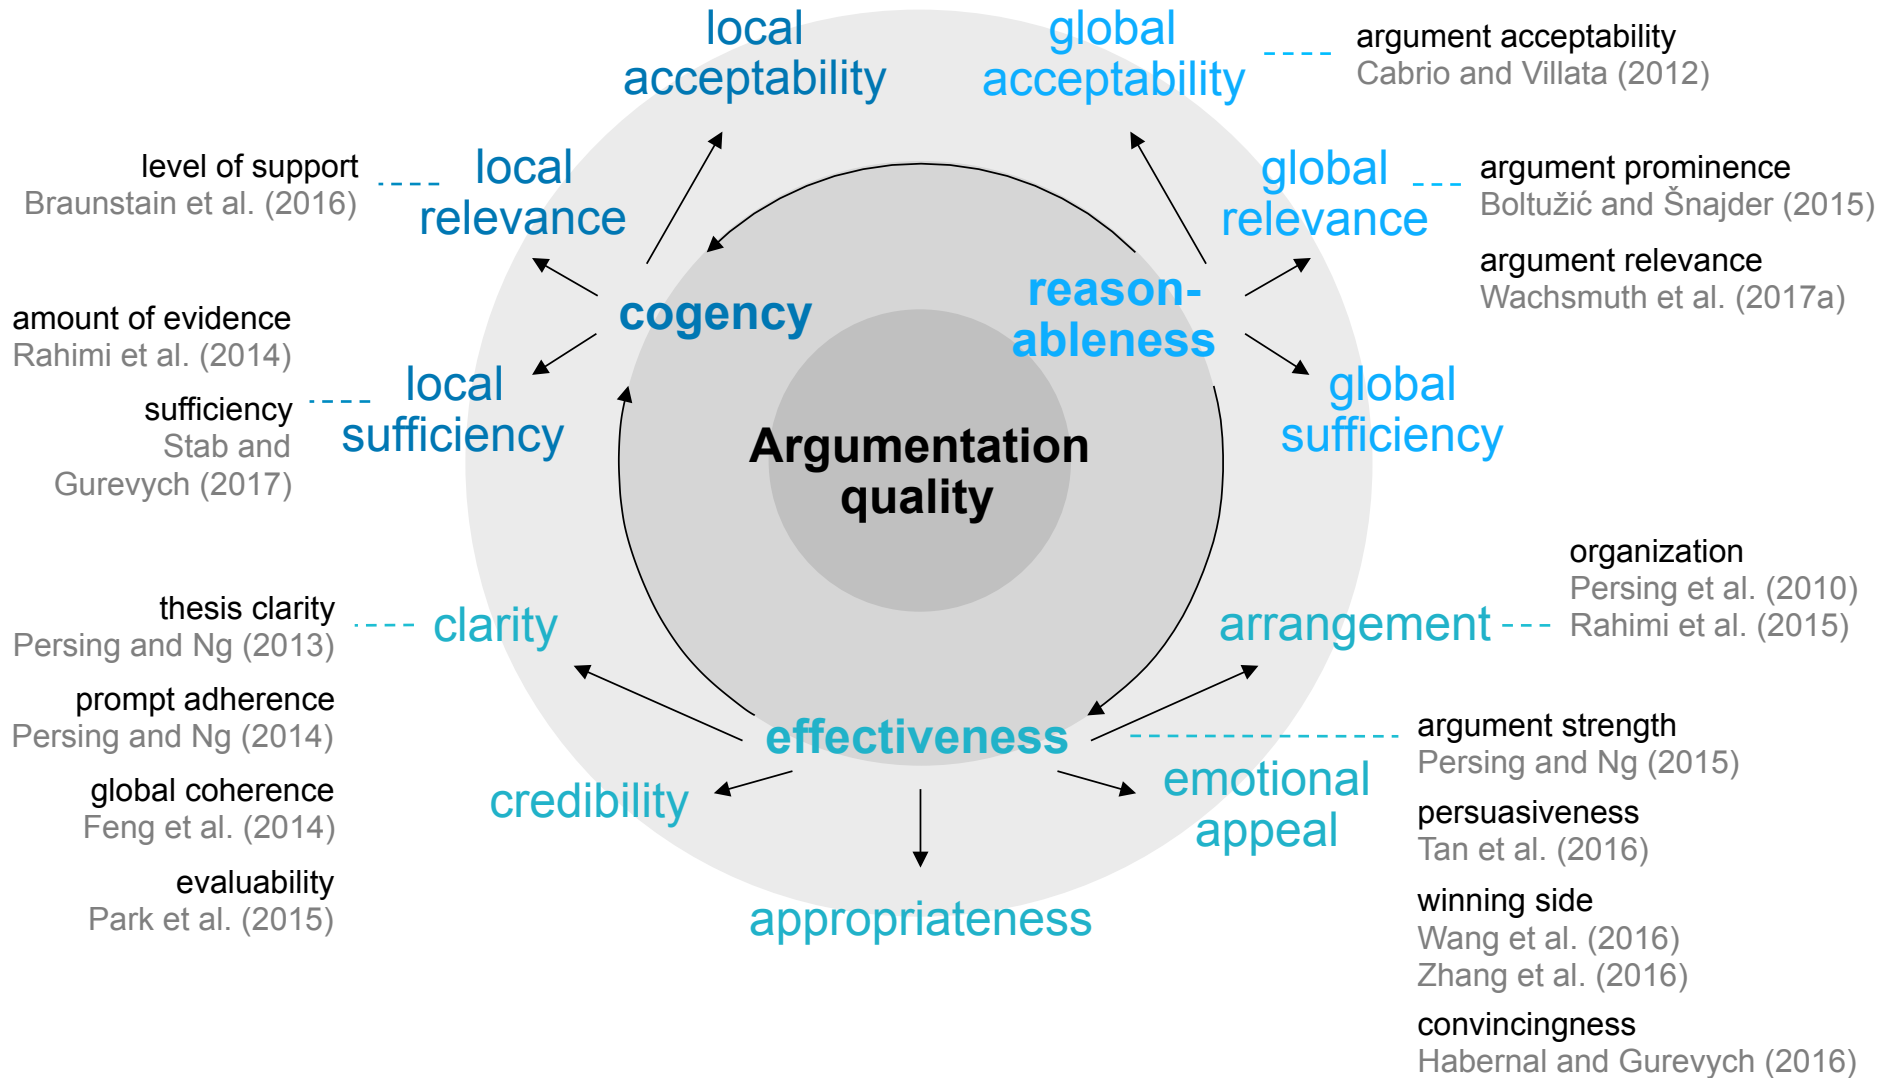


<https://commons.wikimedia.org>

"In making a speech, one must study three points: the means of producing persuasion, the style or language to be used, and the proper arrangement of the various parts."

Aristotle (2007)

A taxonomy of argumentation quality (Wachsmuth et al., 2017b)



The role of participants in argumentation

Author (or speaker)

- Argumentation is connected to the person who argues
- The same argument is perceived differently depending the author

Reader (or audience)

- Argumentation often targets a particular audience
- Different arguments and ways of arguing work for different persons

”University education must be free. That is the only way to achieve equal opportunities for everyone.“

”According to the study of XYZ found online, avoiding tuition fees is beneficial in the long run, both socially and economically.“



<https://pixabay.com>




<https://commons.wikimedia.org>




<https://pixabay.com>

Selected assessment approaches


- **Absolute assessment** (Wachsmuth et al., 2016)
 - Regression of 4 essay quality dimensions
 - Features based on argument mining




- **Relative assessment** (Zhang et al., 2016)
 - Classification of the winner of a debate
 - Modeling own and attack of opponent's points



- **”Objective“ assessment** (Cabrio and Villata, 2012)
 - Graph analysis to determine acceptability (Dung, 1995)
 - Textual entailment to obtain attacking arguments



- **Audience-specific assessment** (El Baff et al., 2018)
 - Analysis of editorial effectiveness for audience
 - Model of audience's ideology and personality



<https://flickr.com>

5. Fallacy detection

(Habernal et al., 2018)



Ivan
Habernal



Iryna
Gurevych



Benno
Stein



Fallacy detection: Introduction

- **What is a fallacy?** (Tindale, 2007)
 - An argument with some (often hidden) flaw in its reasoning, i.e., it has a failed or deceptive scheme
- **Example types of fallacies**
 - **Ad-hominem.** Attacking the opponent instead of his or her arguments
 - **Red herring.** Reasoning based on an unrelated issue
 - **Appeal to ignorance.** Taking lack of evidence as proof for the opposite

*My girlfriend **won't** give me a gift for my birthday. I have received no indication to the contrary from her.*

*My flight tomorrow **won't** be delayed. I have received no indication to the contrary from the airline.*

(credit to Mario Treiber for this example)

- **Fallacies are hard to detect**
 - Structure identical to other arguments
 - Understanding and context knowledge needed

A study of ad-hominem arguments on the web

- **Ad-hominem arguments**

- Attacking the opponent instead of his or her arguments
- 20% of all online news comments uncivil (Coe et al., 2014)

That's an ad hominem fallacy
Calvin!!

"YOU'RE FACE IS AN
AD HOMINEM!!!"



- **Research questions**

- Can we identify ad-hominem automatically?
- What are triggers of ad-hominem?

- **Data**

- 2M posts from Reddit ChangeMyView
 - 3866 posts (0.2%) contain an ad-hominem argument
- Ad-hominem is deleted by moderators, but we obtained all comments from them



- **Reddit ChangeMyView (CMV)**

- An opinion poster (OP) states a view
- Others argue for the opposite
- OP gives Δ to convincing posts

Deltas(s) from OP **CMV: Trump has done nothing of substance since being elected to office.**

This is kind of a counter to the other post made recently about Trump being a great president.

He pointed out things like the economy, which was growing before he even took office and it actually seems to be

Ad-hominem on CMV

"Reading comprehension is your friend"

"Ever have discussions with narcissistic idiots on the internet? They are so tiring"

"You still refuse to acknowledge that you used a strawman argument against me"

"little buddy"

"Thank you so much for all your pretentious explanations"

"To say that people intrinsically understand portion size is idiotic."

"You started with a fallacy and then deflected."

"Please dont waste peoples time pretending to know what you're talking about"

"boy"

"Did you even read this?"

"Read what I posted before acting like a pompous ass"

"Do you even know what you're saying?"

"Again, how old are you?"

"You're making the claims, it's your job to prove it. Don't you know how debating works?"

"You're obviously just Nobody with enough brains to operate a computer could possibly believe something this stupid"

"Your second paragraph is fairly idiotic"

"You have no capability to understand why"

"Wow. Someone sounds like a bit of an anti-semite"

"you dumb fuck"

"Your just an asshole"

"Possible lie any harder?"

"How can you explain that? You can't because it will hurt your feelings to face reality"

"You're just a dishonest troll"

"You are just a liar."

"Can you also use Google?"

"Willful ignorance is not something I can combat"

"You're trash at debating."

"You're too dishonest to actually quote the verse because you know it's bullshit"


"You're using troll tactics"

"sir"

"If you can't grasp the concept, I can't help you"

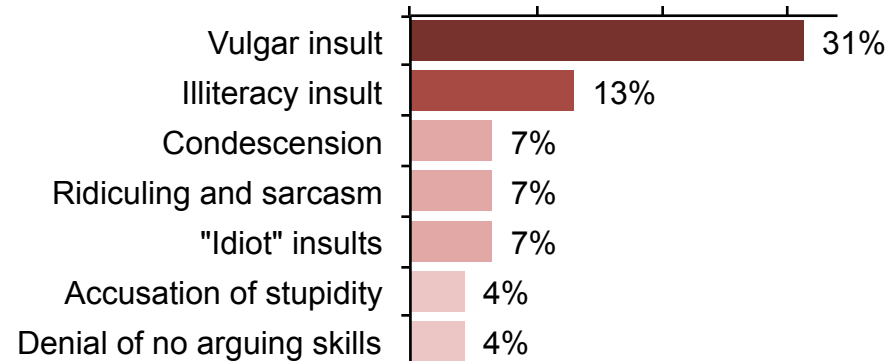
Identification of ad-hominem

- **Distribution of ad-hominem**




- **Types of ad-hominem**

- Ad-hominem annotated in 400 arguments by 7 crowdworkers
- 15 types derived manually from their annotations



- **Identification of ad-hominem**

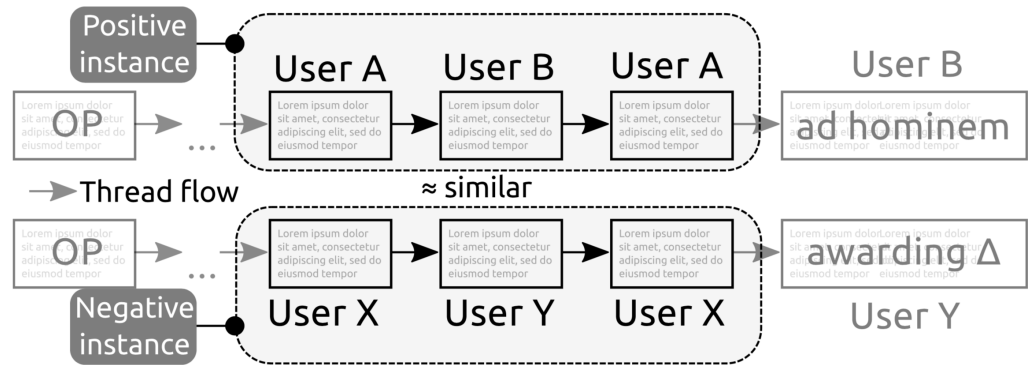
- **Manual.** 100 balanced arguments (50 ad-hominem) classified by 6 workers
- **Automatic.** 7242 balanced arguments classified by 2 neural classifiers (Bi-LSTM & CNN)



Analysis of triggers of ad-hominem

- **Prediction of ad-hominem**

- **Attentive LSTM** trained on 2852 argument 3-tuples
- Accuracy 0.72
- Manual attention analysis



(OOV_comment_begin) If only you would n't rely on [fallacious] (http : OOV) [arguments] (http : OOV) to make your point. So no , I do n't realize how stupid and naive I am. All I 've realized is that you are n't actually prepared to have an actual discussion .

(OOV_comment_begin) What god do you believe in ? And it 's not a fallacy when it 's very comparable to the most popular gods .

(OOV means out-of-vocabulary)

- **Terms with much attention**

- Mostly topic-independent rhetorical devices
- A few loaded keywords, such as "racist"
- Partly meta about argumentation

vulgar intensifiers
"... the fuck..."

direct imperatives
"You should..."

missing evidence
"unsupported claims!"

bad argumentation
"You're grasping at straws" ...

Conclusion

Conclusion

- **Argument mining and assessment**

- Finding arguments in natural language text
- Classifying stance, schemes, and similar
- Assessing quality dimensions and flaws

Conclusion
think that **the death penalty should be abolished**.
support ↗ ↖ **support**
t of violence. As long as human justice remains
Premise
e innocent can never be eliminated.”

- **State of the art**

- Most tasks now tackled with neural approaches
- In narrow domains, reasonable effectiveness achievable
- Robust "off-the-shelf" algorithms rare so far



- **Role within argumentation technology**

- Builds on argumentation theory and computational linguistics
- Needed to process natural language arguments
- Converts arguments to (semi-) structured information



References

- **Aioli et al. (2009)**. Fabio Aioli, Giovanni Da San Martino, and Alessandro Sperduti. 2009. Route kernels for trees. In Proceedings of the 26th Annual International Conference on Machine Learning, pages 17–24.
- **Ajjour et al. (2017)**. Yamen Ajjour, Wei-Fan Chen, Johannes Kiesel, Henning Wachsmuth, and Benno Stein. Unit Segmentation of Argumentative Texts. In Proceedings of the Fourth Workshop on Argument Mining, pages 118–128, 2017.
- **Al-Khatib et al. (2017)**. Khalid Al-Khatib, Henning Wachsmuth, Matthias Hagen, and Benno Stein. Patterns of Argumentation Strategies across Topics. In Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, pages 1362–1368, 2017.
- **Aristotle (2007)**. Aristotle (George A. Kennedy, Translator). On Rhetoric: A Theory of Civic Discourse. Clarendon Aristotle series. Oxford University Press, 2007.
- **Bar-Haim et al. (2017)**. Roy Bar-Haim, Indrajit Bhattacharya, Francesco Dinuzzo, Amrita Saha, and Noam Slonim. Stance Classification of Context-Dependent Claims. In Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers, pages 251–261, 2017.
- **Blair (2012)**. J. Anthony Blair. Groundwork in the Theory of Argumentation. Springer Netherlands, 2012.
- **Boltužic and Snajder (2015)**. Filip Boltužic and Jan Snajder. Identifying Prominent Arguments in Online Debates using Semantic Textual Similarity. In Proceedings of the 2nd Workshop on Argumentation Mining, pages 110–115, 2015.
- **Braunstain et al. (2016)**. Liora Braunstain, Oren Kurland, David Carmel, Idan Szpektor, and Anna Shtok. Supporting Human Answers for Advice-seeking Questions in CQA Sites. In Proceedings of the 38th European Conference on IR Research, pages 129–141, 2016.

References

- **Cabrio and Villata (2012)**. Elena Cabrio and Serena Villata. Combining Textual Entailment and Argumentation Theory for Supporting Online Debates Interactions. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 208–212, 2012.
- **Collins and Duffy (2001)**. Michael Collins and Nigel Duffy. 2001. Convolution kernels for natural language. In *Advances in Neural Information Processing Systems 14*, pages 625–632.
- **Dung (1995)**: Phan Minh Dung. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. *Artificial Intelligence*, 77(2):321–357, 1995.
- **El Baff et al. (2018)**. Roxanne El Baff, Henning Wachsmuth, Khalid Al-Khatib, and Benno Stein. Challenge or Empower: Revisiting Argumentation Quality in a News Editorial Corpus. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 454–464, 2018.
- **Feng and Hirst (2011)**. Vanessa Wei Feng and Graeme Hirst. Classifying Arguments by Scheme. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, pages 987–996, 2011.
- **Feng et al. (2014)**. Vanessa Wei Feng, Ziheng Lin, and Graeme Hirst. The Impact of Deep Hierarchical Discourse Structures in the Evaluation of Text Coherence. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 940–949. Dublin City University and Association for Computational Linguistics, 2014.
- **Freeman (2011)**. *Argument Structure: Representation and Theory*. Springer, 2011.
- **Habernal and 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, 2015.

References

- **Habernal and Gurevych (2016a).** Ivan Habernal and Iryna Gurevych. 2016. Which Argument is More Convincing? Analyzing and Predicting Convincingness of Web Arguments using Bidirectional LSTM. In Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 1589–1599.
- **Habernal et al. (2018a).** Ivan Habernal, Henning Wachsmuth, Iryna Gurevych, and Benno Stein. Before Name-calling: Dynamics and Triggers of Ad Hominem Fallacies in Web Argumentation. In Proceedings of the 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 386–396, 2018.
- **Hasan and Ng (2013).** Kazi Saidul Hasan and Vincent Ng. Stance Classification of Ideological Debates: Data, Models, Features, and Constraints. In Proceedings of the Sixth International Joint Conference on Natural Language Processing, pages 1348--1356, 2013.
- **Mooney and Bunescu (2006).** Raymond J. Mooney and Razvan C. Bunescu. 2006. Subsequence kernels for relation extraction. In Advances in Neural Information Processing Systems 18, pages 171–178. MIT Press.
- **Park et al. (2015).** Joonsuk Park, Cheryl Blake, and Claire Cardie. Toward Machine-assisted Participation in eRulemaking: An Argumentation Model of Evaluability. In Proceedings of the 15th International Conference on Artificial Intelligence and Law, pages 206–210, 2015.
- **Peldszus and Stede (2015).** Andreas Peldszus and Manfred Stede. 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, 2015.
- **Peldszus and Stede (2016).** Andreas Peldszus and Manfred Stede. 2016. An annotated corpus of argumentative microtexts. In Argumentation and Reasoned Action: 1st European Conference on Argumentation.

References

- **Perelman and Olbrecht-Tyteca (1969).** Chaïm Perelman and Lucie Olbrechts-Tyteca. 1969. *The New Rhetoric: A Treatise on Argumentation* (John Wilkinson and Purcell Weaver, translator). University of Notre Dame Press.
- **Persing and Ng (2013):** Isaac Persing and Vincent Ng. Modeling Thesis Clarity in Student Essays. In: *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, pages 260–269, 2013.
- **Persing and Ng (2014):** Isaac Persing and Vincent Ng. Modeling Prompt Adherence in Student Essays. In: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, pages 1534–1543, 2014.
- **Persing and Ng (2015):** Isaac Persing and V. Ng. Modeling Argument Strength in Student Essays. In: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 543–552, 2015.
- **Persing et al. (2010).** Isaac Persing, Alan Davis, and Vincent Ng. Modeling organization in student essays. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 229–239, 2010.
- **Rahimi et al. (2014).** Zahra Rahimi, Diane J. Litman, Richard Correnti, Lindsay Clare Matsumura, Elaine Wang, and Zahid Kisa. Automatic Scoring of an Analytical Response-to-Text Assessment. In *Proceedings of the 12th International Conference on Intelligent Tutoring Systems*, pages 601–610, 2014.
- **Rahimi et al. (2015).** Zahra Rahimi, Diane Litman, Elaine Wang, and Richard Correnti. Incorporating Coherence of Topics as a Criterion in Automatic Response-to-Text Assessment of the Organization of Writing. In *Proceedings of the Tenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 20–30, 2015.
- **Ranade et al. (2013):** Sarvesh Ranade, Rajeev Sangal, and Radhika Mamidi. Stance Classification in Online Debates by Recognizing Users' Intentions. In: *Proc. of the SIGDIAL 2013*, 61–69, 2013.

References

- **Rinott et al. (2015).** Ruty Rinott, Lena Dankin, Carlos Alzate Perez, M. Mitesh Khapra, Ehud Aharoni, and Noam Slonim. 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, 2015.
- **Somasundaran and Wiebe (2010):** Swapna Somasundaran and Janyce Wiebe. Recognizing Stances in Ideological On-Line Debates. In: Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text, pages 116–124, 2010.
- **Stab and Gurevych (2016).** Christian Stab and Iryna Gurevych. 2016. Recognizing the absence of opposing arguments in persuasive essays. In Proceedings of the Third Workshop on Argument Mining (ArgMining2016), pages 113–118.
- **Stab (2017).** Christian Stab. Argumentative Writing Support by means of Natural Language Processing, Chapter 5. PhD thesis, TU Darmstadt, 2017. <http://tuprints.ulb.tu-darmstadt.de/6006/1/PhD-Thesis-ChristianStab.pdf>
- **Stab and Gurevych (2017).** Christian Stab and Iryna Gurevych. Recognizing Insufficiently Supported Arguments in Argumentative Essays. In Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, pages 980–990, 2017.
- **Tan et al. (2016).** Chenhao Tan, Vlad Niculae, Cristian Danescu-Niculescu-Mizil, and Lillian Lee. Winning Arguments: Interaction Dynamics and Persuasion Strategies in Good-faith Online Discussions. In Proceedings of the 25th International Conference on World Wide Web, pages 613–624, 2016.
- **van Eemeren (2015).** Frans H. van Eemeren. Reasonableness and Effectiveness in Argumentative Discourse: Fifty Contributions to the Development of Pragma-Dialectics. Argumentation Library. Springer International Publishing, 2015.
- **Wachsmuth et al. (2016).** Henning Wachsmuth, Khalid Al-Khatib, and Benno Stein. Using Argument Mining to Assess the Argumentation Quality of Essays. In: Proceedings of the 26th International Conference on Computational Linguistics, pages 1680–1692, 2016.

References

- **Wachsmuth et al. (2017a).** Henning Wachsmuth, Benno Stein, and Yamen Ajjour. "PageRank" for Argument Relevance. In Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, pages 1116–1126, 2017.
- **Wachsmuth et al. (2017b).** Henning Wachsmuth, Nona Naderi, Yufang Hou, Yonatan Bilu, Vinodkumar Prabhakaran, Tim Alberdingk Thijm, Graeme Hirst, and Benno Stein. Computational Argumentation Quality Assessment in Natural Language. In Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, pages 176–187, 2017.
- **Wachsmuth et al. (2017c).** Henning Wachsmuth, Giovanni Da San Martino, Dora Kiesel, and Benno Stein. The Impact of Modeling Overall Argumentation with Tree Kernels. In Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, pages 2369–2379, 2017.
- **Wachsmuth et al. (2018).** Henning Wachsmuth, Shahbaz Syed, and Benno Stein. Retrieval of the Best Counterargument without Prior Topic Knowledge. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, pages 241–251, 2018.
- **Walton et al. (2008).** Douglas Walton, Christopher Reed, and Fabrizio Macagno. Argumentation Schemes. Cambridge University Press, 2008.
- **Wang et al. (2017).** Lu Wang, Nick Beauchamp, Sarah Shugars, and Kechen Qin. Winning on the Merits: The Joint Effects of Content and Style on Debate Outcomes. In: Transactions of the Association for Computational Linguistics 5, pages 219--232, 2017.
- **Zhang et al. (2016).** Justine Zhang, Ravi Kumar, Sujith Ravi, and Cristian Danescu-Niculescu-Mizil. Conversational Flow in Oxford-style Debates. In Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 136–141, 2016.