

Computational Argumentation — Part V

Perspective Assessment

Henning Wachsmuth

<https://ai.uni-hannover.de>



Learning goals

▪ Concepts

- The multi-faceted notion of perspective
- Stance and myside bias of authors/speakers
- Frames and framing of argumentative texts



<https://commons.wikimedia.org>

▪ Methods


- Stance classification contrastiveness, route kernels, and graph convolutional networks
- Frame recognition using different clustering methods



<https://pixabay.com>

▪ Associated research fields

- Natural language processing



<https://pixabay.com>

▪ Within this course

- How to determine what perspectives are conveyed by (previously mined) arguments



Outline

- I. Introduction to computational argumentation
- II. Basics of natural language processing
- III. Basics of argumentation
- IV. Argument mining
- V. Perspective assessment**
- VI. Argument quality assessment
- VII. Argument generation
- VIII. Applications of computational argumentation
- IX. Conclusion

- a) **Introduction**
- b) Stance classification
- c) Frame recognition
- d) Conclusion

Perspective

■ Perspective

- Coverage term for facets that characterize how an argument pictures an issue and how this reflects the view of the argument's author on the issue

■ Common facets of perspective

- **Stance.** Position of the author toward the issue
- **Topic.** Aspects of the issue focused on
- **Frame.** Emphasis of a specific viewpoint on the issue
- **Bias.** Preconceptions leading to a subjective picturing
- **Values.** Reference to human value systems
- **Morals.** Reference to moral behaviors

We focus on stance and frames below, but touch some others.

■ Formalization of facets

- Meta-information that reflects an understanding of some facet
- May be given as labels, scores, additional text fragments, or similar



Perspective assessment: Example

▪ Examples

- Two claims on the issue of banning death penalty

“The death penalty legitimizes an irreversible act of violence by the state and will inevitably claim innocent victims.”

“Capital punishment may brutalise society in a [...] more fundamental way, one that has implications for the state’s relationship with all citizens.”

both taken from https://www.bbc.co.uk/ethics/capitalpunishment/against_1.shtml

▪ Assessing perspectives for the examples

- **Stance.** Both agree in being *pro* banning the death penalty
- **Topic.** 1st claim focuses on *state violence* and *innocent victims*
- **Frame.** 1st claim frames the issue from a *legal* viewpoint
- **Bias.** Both claims have a fully *one-sided* view on the topic
- **Values.** 2nd claim values a *stable society* rather than the *role of individuals*
- **Morals.** Both claims target morals related to the *care/harm* contrast

Perspective assessment

- **Perspective assessment**

- Family of tasks that assess any facet of the perspective of an argument

- **Selected perspective assessment tasks**

Stance classification	Frame identification	Social bias quantification
Agreement classification	Frame classification	Social bias detection
Aspect detection	Myside bias classification	Human value identification
Key point analysis	Political bias classification	Moral foundation classification

- **Why perspective assessment?**

- Argumentative structure alone is not sufficient for many applications.
- Often, some understanding is needed of how an argument relates to an issue.

- **Notice**

- Where mining ends and assessment starts is not defined exactly.

For example, classifying evidence types might be seen as assessment.

Next section: Stance classification

- I. Introduction to computational argumentation
- II. Basics of natural language processing
- III. Basics of argumentation
- IV. Argument mining
- V. Perspective assessment**
- VI. Argument quality assessment
- VII. Argument generation
- VIII. Applications of computational argumentation
- IX. Conclusion

- a) Introduction
- b) Stance classification**
- c) Frame recognition
- d) Conclusion

Stance (recap)

■ Stance

- The overall position held by a person toward some target, such as an object, statement, or issue
- To have/take a stance on a target means to be *pro* or *con* toward it.

Related terms: Viewpoint, view, standpoint, stand, position



Pro toward ban of death penalty

The death penalty must be abolished.

Con toward ban of death penalty

Life-long imprisonment is pricy, though.

■ Observations on stance

- Often but not necessarily conveys sentiment
- Depends on what a speaker claims to be true
- Can be expressed without any reference to the target

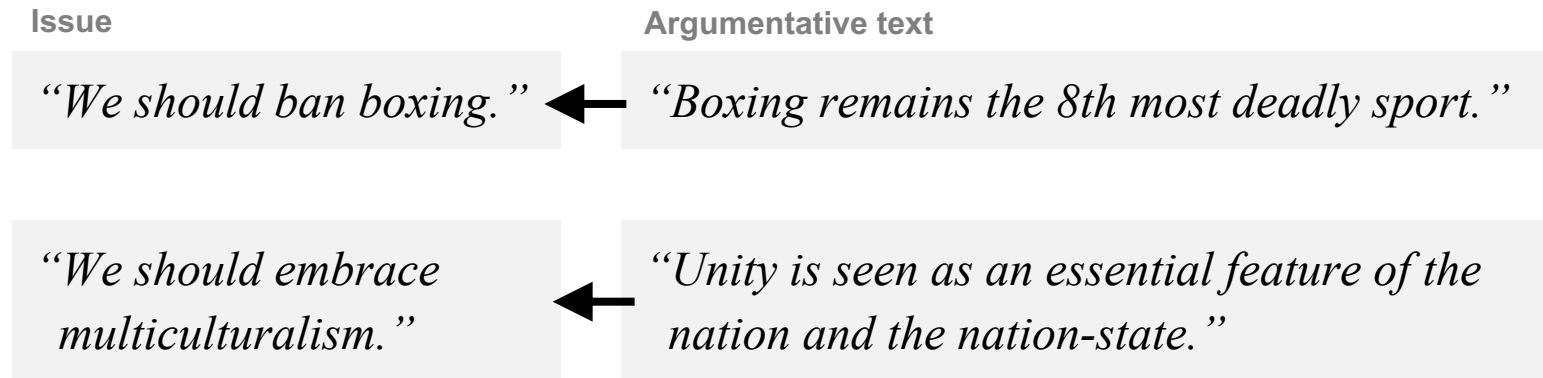
■ Claim

- A claim is a statement that conveys a stance toward a controversial issue, implicitly or explicitly

Stance classification

■ Stance classification

- The classification of the stance of an argumentative text toward a given issue
- **Stance.** Usually *pro* or *con*, sometimes also *neutral* or *neither*
- **Issue.** A controversial topic or claim



minimally shortened examples of Bar-Haim et al. (2017a)

■ Challenges (beyond those on previous slide)

- The *contrastiveness* of targeted concepts needs to be accounted for.
- Positive stance can be expressed with negative sentiment, and vice versa.

Modeling stance classification

Input to stance classification?

- Beyond text, issue always needed (at least implicitly)
- Knowledge about debate structure may help
- Further information may include user profiles etc.

Modeling stance classification

- **Instance-level.** Classify stance of individual texts
- **Collection-level.** Classify all stances in a debate
- **Agreement.** Classify if two texts match in stance

Common stance classification approaches

- Standard text classification trained for specific issues
- Relation-like classification with the issue as input
- Graph-based analysis over all arguments in a debate

Target: Rescue boats

Alice: *The EU should allow rescue boats in the Mediterranean Sea, to save the innocent refugees.*

↓ **stance tend to be opposite**

Bob: *So naïve... having such boats makes even more people die trying.*

↓ **stance tend to be the same**

Alice: *Well, I actually read that rescue boats haven't led to an increase yet.*

Stance classification: Overview of approaches

▪ Selected instance-level approaches

- Author information as features for classification in dialogues (Ranade et al., 2013)
- Opposing views as feature for classification in dialogues (Hasan and Ng, 2013)
- [Stance as sentiment and contrast of claim and issue targets](#) (Bar-Haim et al., 2017a)
- [Route kernels based on overall structure of texts](#) (Wachsmuth et al., 2017d)
- Mixture-of-experts transformer for cross-domain stance (Hardalov et al., 2021)
- Few-shot selection for LLM-based stance classification (Ajjour and Wachsmuth, 2025)

▪ Selected collection-level approaches

- [Graph convolutional network on whole debate structure](#) (Barrow et al., 2021)
- Derive stance from statistical evidence based on stance tree (Saha et al., 2024)

▪ Selected agreement approaches

- Bi-LSTM-based twin network for classifying argument pairs (Xu et al., 2019)
- Cross-domain comparison of several transformer approaches (Körner et al., 2021)

Stance as sentiment and contrast

▪ Task

- Given a claim relevant to a given issue, classify the claim's stance on the issue.
- The issue is supposed to be a claim itself.

Issue. *“Advertising is harmful.”*

Claim. *“Marketing creates consumerism and waste.”*

▪ IBMSC data

- 55 issues from iDebate with and 2394 claims from Wikipedia
- 25 issues for training (1039 claims), 30 for testing (1355 claims)
- The *target* and *sentiment* of each claim were annotated manually.

▪ Presented approach (Bar-Haim et al., 2017a)

1. Identify the target of the issue and the claim.
2. Classify the sentiment toward each target.
3. Determine whether the targets are contrastive.
4. Derive stance from sentiment and contrast.

claim target polarity
× contrastiveness
× issue target polarity

≈ **stance**

Stance as sentiment and contrast: Approach steps 1+2

▪ Sentiment toward targets

- **Target.** The topic/concept that is addressed by an issue or claim
- **Sentiment.** The polarity expressed toward the target (positive or negative)

1. Identify targets x_c and x_i of claim and issue

- **Candidates.** Any noun phrase
- **Features.** Position in parse tree, relation to sentiment, Wikipedia title?, ...
- **Approach.** Logistic regression classifier

Issue. “*Advertising* is harmful.”

Claim. “*Marketing* creates consumerism and waste.”

2. Score polarities $p(x_c)$ and $p(x_i)$ in $[-1,1]$

- **Features.** Terms from sentiment lexicon and a “polarity shifter” lexicon
- **Approach.** Weigh values based on token distance to targets

Issue. “*Advertising* is harmful.”
-0.9

Claim. “*Marketing* creates consumerism and waste.”
-0.7


Stance as sentiment and contrast: Approach steps 3+4

▪ Contrastiveness of targets

- Two targets x_c and x_i are called *contrastive* (-1), if a pro stance on x_c makes a con stance on x_i likely, and vice versa.
- Otherwise, x_c and x_i are called *consistent* (+1).

▪ Score contrastiveness $c(x_c, x_i)$ in $[-1, 1]$

- **Features.** Polarity shifters, relatedness measures, Wikipedia headers, ...
- **Approach.** Random forest classifier



▪ Score stance s

- **Feature.** $s := p(x_c) \cdot c(x_c, x_i) \cdot p(x_i)$
- **Approach.** Classify as pro if $s \geq \tau$ and as con if $s \leq -\tau$, otherwise discard.

τ is a threshold that can be understood as a confidence level.

$$\begin{aligned} s &= -0.7 \cdot +1 \cdot -0.9 \\ &= 0.63 \end{aligned}$$

Say, $\tau = 0.5$. then stance is **pro**.

Stance as sentiment and contrast: Results

▪ **Evaluation** (Bar-Haim et al., 2017a)

- **Baseline.** Support vector machine (SVM) with unigram and sentiment features
- **Measure.** Accuracy@coverage depending on threshold for s (here 20–100%)

Approach	20%	40%	60%	80%	100%
SVM baseline	0.717	0.709	0.691	0.668	0.632
Approach w/o contrast	0.770	0.749	0.734	0.632	0.632
Full approach	0.847	0.793	0.740	0.632	0.632

▪ **Observations**

- Reliable for confident cases, but does not beat baseline if all are classified
- The hardest cases are those where stance is expressed without sentiment.

▪ **Extended approach** (Bar-Haim et al., 2017b)

- Expansion of sentiment lexicon using bootstrapping techniques.
- Use of sentiment from neighboring sentences if claim/issue lack sentiment.

Extended approach	0.935	0.856	0.776	0.734	0.691
--------------------------	--------------	--------------	--------------	--------------	--------------

Stance classification: Relation to myside bias

▪ Myside bias

- An argumentative text is said to have myside bias iff. it does not consider opposing viewpoints.
- In other words, it only supports its own stance.

▪ Myside bias classification

- The classification of a text as having myside bias or not

▪ Modeling myside bias classification

- Conceptually, a standard text classification task

▪ Approaches to myside bias classification

- Supervised classification using various features (Stab and Gurevych, 2016)
- [Route kernels based on overall structure of texts](#) (Wachsmuth et al., 2017d)



Stance and myside bias: Impact of overall structure?

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

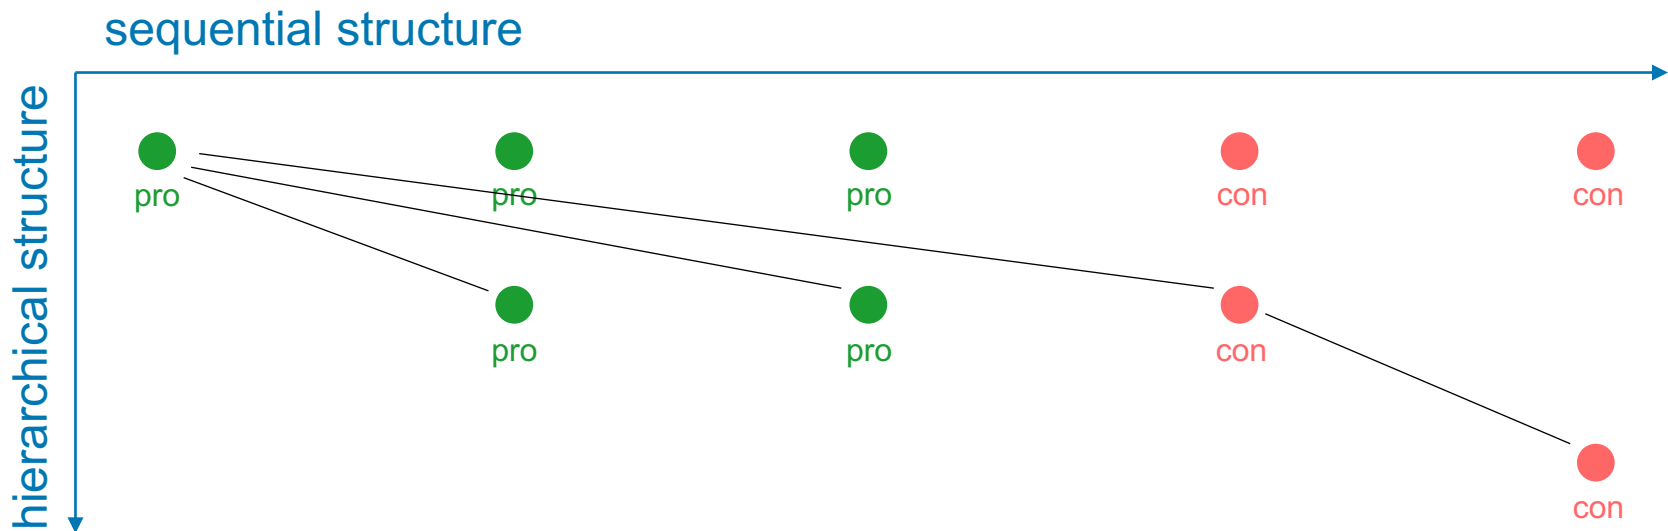
For one thing, inviolable human dignity is anchored in our constitution,

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)




Route kernels for stance and bias

▪ Task

- Given a monological argumentative text, classify stance and myside bias (without knowing the issue discussed).


▪ Idea

- Argumentative texts with similar overall structure may also be similar in myside bias and stance.



▪ Research questions

1. How to jointly model sequential and hierarchical structure?
2. What structure is important for myside bias and stance?



▪ Approach in a nutshell (Wachsmuth et al., 2017d)

- Start from argumentative structure of a text.
- Model overall structure using *route kernels*.
- Classify text based on kernels.

Route kernels for stance and bias: Tasks and data

▪ Myside bias on AAE-v2

(Stab and Gurevych, 2016)

- 402 persuasive student essays
- Essay-specific argument model
- 251 myside bias, 151 no myside bias

▪ Stance on Arg-Microtexts




(Peldszus and Stede, 2016)

- 112 short argumentative texts
- Freeman model (Freeman, 2011)
- 46 pro stance, 42 con stance, 24 unlabeled

▪ Genre on Web Discourse (for comparison)

(Habernal and Gurevych, 2015)

- 340 argumentative web texts
- Modified Toulmin model (Toulmin, 1958)
- 216 comments, 46 blog posts, 73 forum posts, 5 articles



Route kernels for stance and bias: Unification

▪ A unified model

- Order nodes according to position.
- Encode stance toward parent as node label.
- Model relations between node *pairs* only.
- The root implicitly defines the thesis.



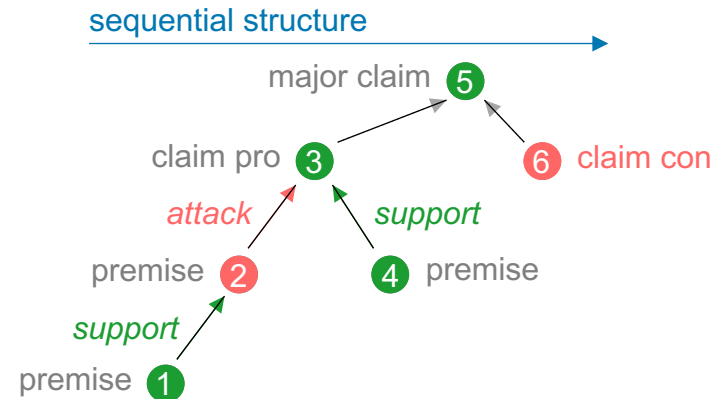
▪ Pros and cons

- + Sequential structure captured
- + Same analyses on all corpora
- + Comparisons across corpora
- + Simpler argument mining (presumably)
- Partly less expressive

▪ In this lecture, only unified model

- For experiments with specific models, see paper.

(Wachsmuth et al., 2017d)



Background: Kernel methods

Kernel methods in machine learning

- Kernel methods classify instances by comparing them to known instances.
- Strong when good features are unknown and/or when data is limited

Often used for structured input data, such as trees

Kernel method: Three building blocks

- **Kernel.** Represents an instance in a task-specific implicit feature space

Different kernels can be combined mathematically.

- **Similarity function.** Quantifies the similarity of any two kernels

- **Classifier.** Distinguishes classes based on similarities

A typical kernel-based classifier is the SVM.

Selected kernels for structured data

- **Linear kernels** capture distributions only

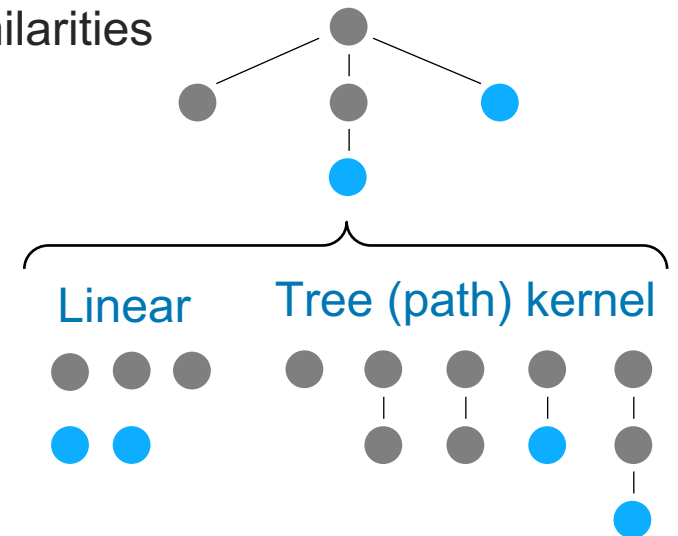
The correspondent of standard feature vectors

- **Subsequence kernels** for sequential structure

(Mooney and Bunescu, 2006)

- **Tree kernels** for hierarchical structure


(Collins and Duffy, 2001)



Background: Route kernels


- **Route kernel** (Aiolli et al., 2009)

- Captures both sequential and hierarchical structure
- Tree kernel with edge labels, indicating node positions relative to siblings
- Models all paths starting from the root of a tree



- **Adapted route kernel for arguments**


- Positions are relative to parent node.
- A polynomial kernel “combines” paths to capture full overall structure.



Route kernels for stance and bias: Approach


Overall structure as a positional tree

- A tree $T = (V, E)$ where nodes in V represent argumentative units and edges in E relations between two units
- Node labels.** Each node labeled as *pro* or *con*
- Edge labels.** Node position in a text relative to parent node



Kernel function for overall structure

- Let two trees $T = (V, E)$ and $T' = (V', E')$ be given.
- The similarity of the trees is defined as:



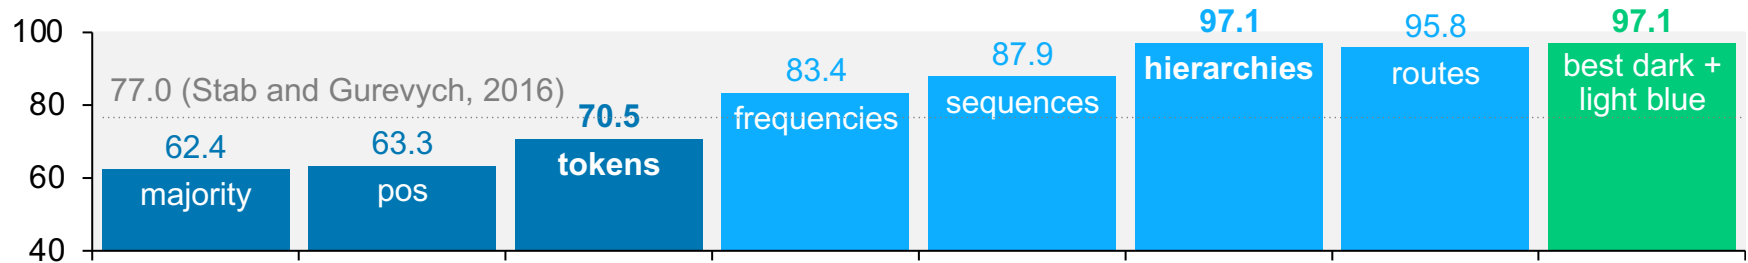
$$K_{\xi\pi}(T, T') = \left(\sum_{v \in V} \sum_{v' \in V'} \frac{\delta(\xi(v), \xi(v')) \cdot \delta(\pi(v), \pi(v'))}{(|V| \cdot |V'|)^2} \right)^d$$

1 for identical paths, 0 otherwise Node label path from root to v Edge label path from root to v Degree of polynomial ($d = 2$ best in experiments)


Sum over all pairs of paths of the two trees Normalization over maximum possible score

Route kernels for stance and bias: Results


■ Myside bias accuracy on AAE-v2



■ Stance accuracy on Arg-Microtexts



■ Genre accuracy on Web Discourse



Stance classification: Discussion

▪ **Effective stance classification**

- Approaches achieve an accuracy < 0.8 in many settings.
- Stance is subjective, so a notably higher accuracy may not be feasible.
- Myside bias is conceptually simpler, but not studied extensively.

▪ **Impact of argumentative structure**

- At least for entire argumentative texts, modeling overall structure is important.
- Theoretically, modeling hierarchical structure “solves” myside bias.
- Practically, the impact depends on the effectiveness of argument mining.

▪ **Stance classification, an independent task**

- Stance classification is also studied apart from computational argumentation.
- Not in all literature on the topic, arguments are considered explicitly.
- Still, the notion of stance implies an argumentative context.

Next section: Frame recognition

I. Introduction to computational argumentation

II. Basics of natural language processing

III. Basics of argumentation

IV. Argument mining

V. Perspective assessment

VI. Argument quality assessment

VII. Argument generation

VIII. Applications of computational argumentation

IX. Conclusion

- a) Introduction
- b) Stance classification
- c) Frame recognition**
- d) Conclusion

Frames and framing (recap)

- **Frame** (Entman, 1993)

- A frame highlights an aspect under which an issue may be considered.
- Both topic-specific and generic sets of frames have been proposed.



- **Framing**

- The selection of specific aspects of an issue to make them more *salient*, i.e., more noticeable, meaningful, and/or memorable.
- The same issue framed in a different way may be perceived entirely different.
- Selecting the right frames is decisive to achieve persuasion.

Frame recognition

▪ Frame recognition

- The determination of the frames encoded in argumentative texts
- **Identification.** Given a set of texts, identify all frames.
- **Classification.** Given a text, assign it a frame from a given set of frames.

“I support legalization of marijuana since it can be taxed for revenue gain.”

“Legalizing prostitution would increase government revenue. A tax on the fee charged by a prostitute, and the imposition of income tax on the earnings of prostitutes would generate revenue.”

**The two claims tackle different topics,
but both emphasize the *economic* frame.**

both taken from Ajjour et al. (2019)

▪ Challenges

- Frames are usually not fully distinct, but overlap in what they emphasize.
- An argumentative text may mix up more than one frame.
- What frames to distinguish may depend on the given text corpus.

Frame recognition: Generic vs. topic-specific frames

- **Generic frames** (Boydston et al., 2014)

- The most widely used taxonomy distinguishes 15 frames

Economic	Health & safety
Capacity & resources	Quality of life
Legality, constitutionality & jurisprudence	External regulation & reputation
Morality	Political
Fairness & equality	Cultural identity
Crime & punishment	Public opinion
Policy prescription & evaluation	Security & defense
	Other

- **Definition: Economic**

“The costs, benefits, or monetary/financial implications of the issue (to an individual, family, community, or to the economy [...]).”

- **Topic-specific frames** (Naderi and Hirst, 2016)

- Seven frames in the discussion of the issue *gay marriage*

Fiscal & legal benefits
Discrimination
Marriage vs. procreation
Union without marriage
Man and woman
Major world religions
Institution of marriage

- **Definition: Fiscal & legal benefits**

“Gay couples should be able to take advantage of the fiscal and legal benefits of marriage.”

Frame recognition: Overview of approaches

▪ **Modeling frame recognition**

- Identifying frames implies a clustering problem.
- Classifying frames rather constitutes a standard classification task.
- Some works focus on topic-specific or generic frames, some cover both.

▪ **Selected identification approaches**

- Latent Dirichlet allocation (LDA) for topic-specific frames (Tsur et al., 2015)
- [Topic-removed frame clustering using latent semantic analysis](#) (Ajjour et al., 2019)
- [Debate-level viewpoint identification using community detection](#) (Barrow et al., 2021)

▪ **Selected classification approaches**

- SVM using word/syntax embeddings for topic-specific frames (Naderi and Hirst, 2016)
- Embedding-based gated recurrent units for generic frames (Naderi and Hirst, 2017)
- Cross-genre multitask LSTMs for quality and generic frames (Hartmann et al., 2019)

Topic-removed frame clustering


▪ Task

- Given a collection of arguments, partition it into k sets of frames (either topic-specific or generic ones).
- k is not predefined, but is to be determined.

▪ Debatepedia data (Ajjour et al., 2019)

- 12,326 arguments from 465 debates on debatepedia.org
- 1623 different frames in total (1293 topic-specific, 330 “generic”)

Generic is here seen as: the frame appears in >1 topics.



▪ Approach in a nutshell (Ajjour et al., 2019)

1. **Topic clustering.** Cluster all arguments into l topics.
2. **Topic removal.** Remove topic-specific features from all arguments.
3. **Frame clustering.** Re-cluster all arguments into k frames.

Background: TF-IDF and latent semantic analysis

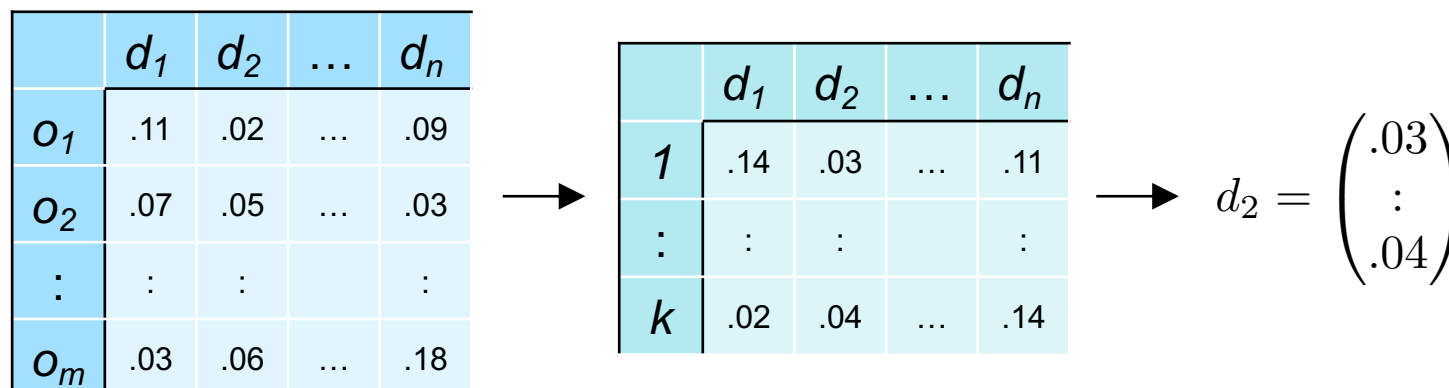
▪ Term Frequency – Inverse Document Frequency (TF-IDF)

- Refine weighting of bag-of-words features in a corpus D .
- A term o in a text $d \in D$ gets a high weight, if o is frequent in d but rare in D .

$$tf - idf(o, d) := \frac{\#(o, d)}{\sum_{o_i \in d} \#(o_i, d)} \cdot \log \left(\frac{\#(d, D)}{\#(d \in D : o \in d)} \right)$$

▪ Latent semantic analysis (LSA) (Deerwester et al., 1990)


- Create term-text frequency matrix for all m terms in a corpus $D = \{d_1, \dots, d_n\}$.
- Use singular-value decomposition to reduce the number of rows from m to k .
- Each column i is then used as the representation of d_i .



Topic-removed frame clustering: Representation

Input

A set of arguments



Legend:

 argument

 topics

 frames

▪ Argument model

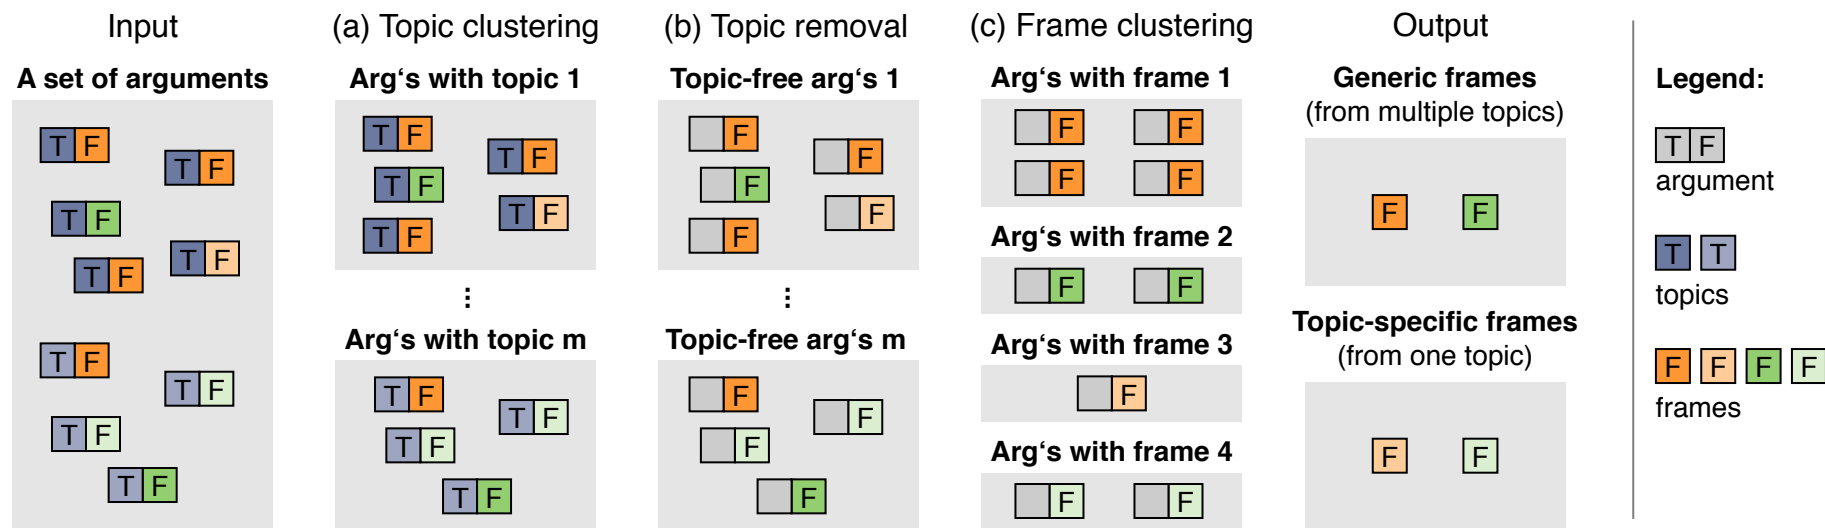
- Arguments are given as pairs of premise and conclusion.
- Each argument is assumed to be a combination of topic and frame.

▪ Three alternative representations

1. **TF-IDF.** Each argument mapped to TF-IDF weights of the top-5000 words
2. **LSA Argument.** LSA with each argument as a text, with $k = 1000$
3. **LSA Debate.** LSA with each debate as a text, with $k = 1000$

In all cases, stop words (determiners etc.) are removed before.

Topic-removed frame clustering: Approach



1. Topic clustering

- l -means clustering using Euclidean TF-IDF/LSA distance, $l \in \{1, \dots, 1000\}$

2. Topic removal (alternatives)

- E_1 . Remove all words with TF-IDF score > 0.005 (determined experimentally).
- E_2 . Remove all words that are part of the conclusion.

3. Frame clustering

- k -means clustering using Euclidean TF-IDF/LSA distance, $k \in \{1.0k, \dots, 1.6k\}$

Topic-removed frame clustering: Experiments

Measures for an argument a

- **B³ precision.** Proportion of arguments in the cluster of a with the frame of a
- **B³ recall.** Proportion of arguments with the frame of a found in the cluster of a
- **B³ F₁-score.** Harmonic mean of B³ precision and B³ recall

These values are averaged over all arguments.

Topic clustering results

Representation	B ³ F ₁
TF-IDF	0.45
LSA Argument	0.44
LSA Debate	0.52

Frame clustering results

Repres.	Removal	Generic	Topic-sp.	All
TF-IDF	None	0.19	0.48	0.26
	E_1	0.28	0.45	0.28
	E_2	0.17	0.45	0.27
LSA Deb.	None	0.44	0.39	0.21
	E_1	0.40	0.47	0.26
	E_2	0.25	0.46	0.25

Discussion

- Different representations best for topic vs. frame
- Highest tested number of frames (1600) turned out best overall

Tested ranges had been chosen based on an initial inspection of the data.

Debate-level viewpoint identification

▪ Motivation

- Claims reflect different viewpoints on a discussed topic, where a viewpoint is a triple $V = (topic, frame, stance)$

The presented work speaks of *aspects* rather than *frames*.

▪ Tasks

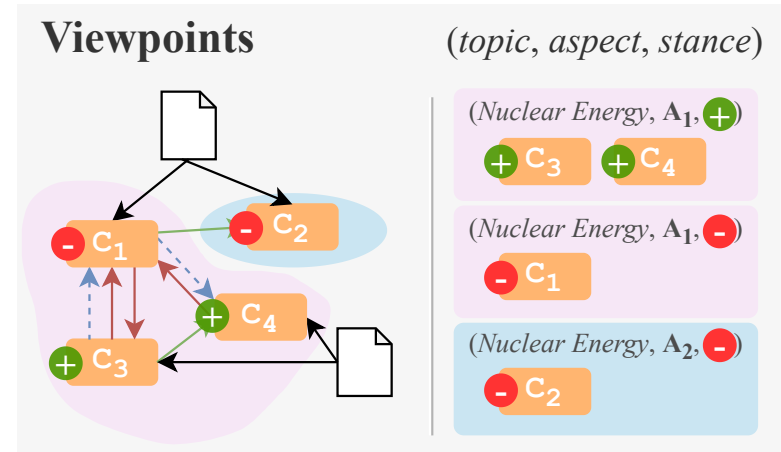
- Given all claims from a set of debates, classify the stance of each claim and group the claims into a set of frames.

▪ Idea (for either task)

- Treat all claims on collection level to exploit their interactions.
- Consider various interactions between claims (and their source texts).

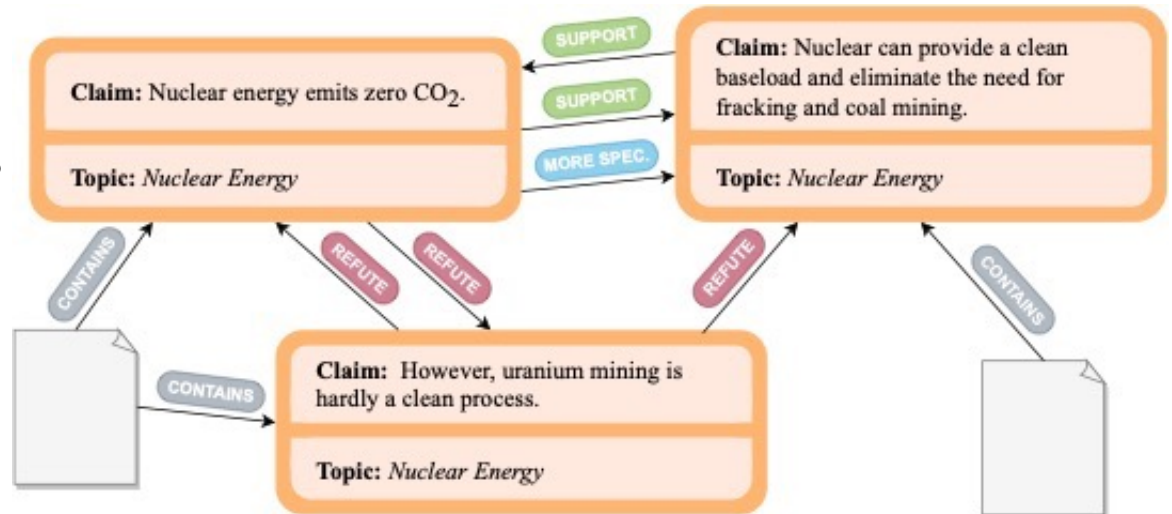
▪ Approach (Barrow et al., 2021)

- Model interactions between claims and texts in a typed multigraph.
- Classify stance using a relational graph convolutional network.
- Identify frames using modularity-based community detection.



Debate-level viewpoint identification: Representation

- **Multigraph $G = (V, E)$ with $V = V_t \cup V_c$ and $E = E_{tc} \cup E_{cc}$**
- **Nodes** represent argumentative texts ($v_t \in V_t$) or claims ($v_c \in V_c$).
- **Edges** indicates that a text contains a claim ($e_{tc} \in E_{tc}$) or that claims have one of various relations ($e_{cc} \in E_{cc}$).
- Each e_{cc} is weighted.



- **Relation types/weights**

- **Topic similarity.** LDA cosine score
- **Term similarity.** TF-IDF cosine score
- **Paraphrase.** $P(\text{yes}) - P(\text{no})$
- **Natural language inference.** $P(\text{entailment}) - P(\text{contradiction})$
- **Relative stance.** $P(\text{support}) - P(\text{refute})$
- **Relative specificity.** $P(\text{more specific}) - P(\text{less specific})$


All similarities and probabilities are determined with existing NLP methods.

Debate-level viewpoint identification: Stance classifier

▪ Relational graph convolutional network (R-GCN)


(Kipf and Welling, 2017; Schlichtkrull et al., 2018)

- **GCN.** A neural network where each layer defines a graph of units (nodes) and relations (edges)
- Related units affect the activation of a unit.
- **R-GCN.** Different input/output relation types



▪ R-GCN illustration on the right

- Computation of the output of the red unit
- For each relation type rel_i , gather activations of all input units and of all output units.
- Apply ReLU activation to weighted sum.



▪ R-GCN-based stance classifier

- **Input.** Contextualized embeddings of texts and claims
- **R-GCN.** Per claim, activations of related nodes, weighted by graph relations
- **Output.** Pro/Con stance prediction for each claim


Debate-level viewpoint identification: Frame identifier

▪ **Modularity-based community detection (MCD)** (Clauset et al., 2004)

- Hierarchical agglomerative clustering for graphs $G = (V, E)$ that maximizes the *modularity* of clusters C (communities)
- **Adjacency matrix.** $A_{vw} = 1$ if $(v_i, v_j) \in E$, else 0
- **Match function.** $\delta(C_i, C_j) = 1$ if $C_i = C_j$, else 0
- **Modularity.** Quantifies how many more edges are within communities than across:

$$m := \frac{1}{2|E|} \sum_{i,j} \left(A_{ij} - \frac{\sum_k A_{ik} \cdot \sum_k A_{jk}}{2|E|} \right) \cdot \delta(C_i, C_j)$$

1. Start with one node per community C .
2. Repeatedly merge C_i and C_j with max. m increase.



A	v_1	v_2	v_3	v_4
v_1	0	1	1	0
v_2	1	0	1	0
v_3	1	1	0	1
v_4	0	0	1	0

▪ **MCD-based frame identifier**

- **Input.** Undirected graph edges (weight: sum of positive directed weights); keep only those with score > 0.6 .
- **Output.** Each community C at maximum m becomes a frame.

Debate-level viewpoint identification: Experiments

▪ Stance classification

- **Data.** IBMCS (2.4 claims), ArgMin (11.1k claims) (Bar-Haim et al., 2017a; Stab et al., 2018)
- **Metric.** Accuracy

Approach	IBMCS	ArgMin
Majority	0.517	0.511
BertGCN (Lin et al., 2021)	0.663	0.587
R-GCN (ablation w/o content)	0.443	0.527
R-GCN (ablation w/o text nodes)	0.830	0.683
R-GCN (approach)	0.834	0.680

▪ Frame identification

- **Data.** Debatepedia (only topic-specific frames) (Ajjour et al., 2019)
- **Metric.** B³ F₁-score

Approach	Debatepedia
Frame clustering (Ajjour et al., 2019)	0.470
MCD (approach)	0.554

Frame recognition: Discussion

▪ **Effective frame recognition**

- Unsupervised nature of frame *identification* makes high effectiveness hard.
- Frame *classification* tends to achieve higher (but not super high) results.
- Frames are generally not clearly defined and fully distinct classes.

▪ **Frames vs. aspects**

- Some research models the aspects of a topic that an arguments addresses.
- Aspects and topic-specific frames are conceptually similar.
- Still, frames encode the idea to emphasize certain aspects (and hide others).

▪ **Framing and persuasion**

- The stance to be argued for often affects what frames to choose.
- Framing ultimately has a rhetorical function: to increase persuasive effect.
- Sometimes, framing is used deceptively: to manipulate.

Next section: Conclusion

- I. Introduction to computational argumentation
- II. Basics of natural language processing
- III. Basics of argumentation
- IV. Argument mining
- V. Perspective assessment**
- VI. Argument quality assessment
- VII. Argument generation
- VIII. Applications of computational argumentation
- IX. Conclusion

- a) Introduction
- b) Stance and myside bias classification
- c) Frame recognition
- d) Conclusion**

Conclusion

■ Perspective assessment

- Perspective refers to how an argument pictures an issue
- Covers stance, topic, frames, bias, values, morals, ...
- Assessment of perspective facets usually via classification

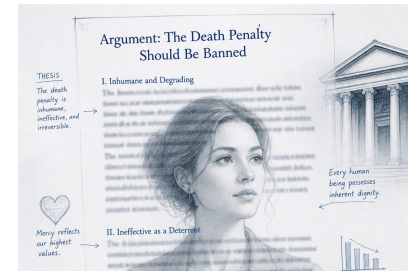



Image generated by DALL-E (via ChatGPT)

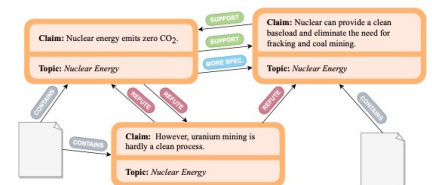
■ Stance classification

- Classifying an argument as pro or con toward an issue
- Stance can be modeled as target sentiment and contrast
- Route kernels enable joint modeling of overall structure



■ Frame recognition

- Identifying and/or classifying the framing of arguments
- Both generic and topic-specific frame models exist
- Frame identification suggests use of clustering methods



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. (2019).** Yamen Ajjour, Milad Alshomary, Henning Wachsmuth, and Benno Stein. Modeling Frames in Argumentation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 2922–2932, 2019.
- **Ajjour and Wachsmuth (2025).** Yamen Ajjour and Henning Wachsmuth. Exploring LLM Priming Strategies for Few-Shot Stance Classification. In *Proceedings of the 12th Argument mining Workshop*, pages 11–23, 2025.
- **Al Khatib et al. (2020).** Khalid Al Khatib, Michael Völske, Shahbaz Syed, Nikolay Kolyada, and Benno Stein. Exploiting Personal Characteristics of Debaters for Predicting Persuasiveness. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7067–7072, 2020.
- **Bar-Haim et al. (2017a).** 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.
- **Bar-Haim et al. (2017b).** Roy Bar-Haim, Lilach Edelstein, Charles Jochim, and Noam Slonim. Improving Claim Stance Classification with Lexical Knowledge Expansion and Context Utilization. In *Proceedings of the 4th Workshop on Argument Mining*, pages 32–38, 2017.
- **Barrow et al. (2021).** Joe Barrow, Rajiv Jain, Nedim Lipka, Franck Deroncourt, Vlad Morariu, Varun Manjunatha, Douglas Oard, Philip Resnik, and Henning Wachsmuth. Syntopical Graphs for Computational Argumentation Tasks. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1583–1595, 2021.

References

- **Clauset et al. (2004).** Aaron Clauset, Mark E.J. Newman, and Cristopher Moore. Finding community structure in very large networks. *Physical review E*, 70(6):066111, 2004.
- **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.
- **Deerwester et al. (1990).** Susan T. Dumals, George W. Furnasand, Thomas K Landauer, and Richard Harshman. 1990. Indexing by Latent Semantic Analysis. 41(6):391–407.
- **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.
- **Hardalov et al. (2021).** Momchil Hardalov, Arnav Arora, Preslav Nakov, and Isabelle Augenstein. Cross-Domain Label-Adaptive Stance Detection. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9011–9028, 2021.
- **Hartmann et al. (2019).** Mareike Hartmann, Tallulah Jansen, Isabelle Augenstein, and Anders Søgaard. Issue framing in online discussion fora. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1401–1407, 2019.
- **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.
- **Kipf and Welling (2017).** Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *Proceedings of the 5th International Conference on Learning Representations*, 2017.

References

- **Körner et al. (2021)**. Erik Körner, Gregor Wiedemann, Ahmad Dawar Hakimi, Gerhard Heyer, and Martin Potthast. On Classifying whether Two Texts are on the Same Side of an Argument. In Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, pages 10130–10138, 2021.
- **Lin et al. (2021)**. Yuxiao Lin, Yuxian Meng, Xiaofei Sun, Qinghong Han, Kun Kuang, Jiwei Li, and Fei Wu. 2021. BertGCN: Transductive text classification by combining GCN and BERT. arXiv preprint arXiv:2105.05727.
- **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.
- **Naderi and Hirst (2016)**. Nona Naderi and Graeme Hirst. Argumentation mining in parliamentary discourse. In Principles and Practice of Multi-Agent Systems, pages 16–25, 2016.
- **Naderi and Hirst (2017)**. Nona Naderi and Graeme Hirst. Classifying Frames at the Sentence Level in News Articles. In Proceedings of the International Conference Recent Advances in Natural Language Processing, pages 536–542, 2017.
- **Peldszus and Stede (2016)**. Andreas Peldszus and Manfred Stede. An annotated corpus of argumentative microtexts. In Argumentation and Reasoned Action: 1st European Conference on Argumentation. 2016.
- **Saha et al. (2024)**. Rudra Ranajee Saha, Laks V. S. Lakshmanan, and Raymond T. Ng. Stance Detection with Explanations. Computational Linguistics, 50(1):193–235, 2024.
- **Schlichtkrull et al. (2018)**. Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In European semantic web conference, pages 593–607, 2018.
- **Stab and Gurevych (2016)**. Christian Stab and Iryna Gurevych. Recognizing the Absence of Opposing Arguments in Persuasive Essays. In Proceedings of the Third Workshop on Argument Mining (ArgMining2016), pages 113–118, 2016.

References

- **Stab et al. (2018).** Christian Stab, Tristan Miller, Benjamin Schiller, Pranav Rai, and Iryna Gurevych. Cross-topic argument mining from heterogeneous sources. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pages 3664–3674, 2018.
- **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.
- **Tsur et al. (2015).** Oren Tsur, Dan Calacci, and David Lazer. A frame of mind: Using statistical models for detection of framing and agenda setting campaigns. In Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 1629–1638, 2015.
- **Xu et al. (2019).** Chang Xu, Cecile Paris, Surya Nepal, and Ross Sparks. Recognising Agreement and Disagreement between Stances with Reason Comparing Networks. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pages 4665–4671, 2019.
- **Wachsmuth et al. (2017d).** 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.

Prompts

- **Image on Slide 4.** Please create a sketch-like professional drawing in portrait format that use black/gray and blue-ish color tones on white background to illustrate the perspective that an author puts into their argument. The drawing should show a blurred textual argument on the issue of banning death penalty. The bust of the author, who is a good female rhetor, should be visible behind the argument, slightly shining through.