A Roadmap of Persuasive Argumentation

Christopher Hidey

April 21, 2017

Roadmap for Persuasive Argumentation

Goals of persuasive argumentation:

1. Providing knowledge
2. Convincing

Roadmap for Persuasive Argumentation

Goals of persuasive argumentation:

- 1) Providing knowledge
  - Argumentation structure
  - Causal relations
- 2) Convincing
  - Personal
  - Emotionally moving

Outline

1. Introduction
2. Persuasion
3. Causal Relations
4. Generation

Roadmap for Persuasive Argumentation

1. Persuasion
   1. What makes an argument more persuasive than a logical sequence of reasons?
   2. How are persuasive arguments structured?
2. Causal Relations
3. Generation

Roadmap for Persuasive Argumentation

1. Persuasion
2. Causal Relations
    1. How can we better represent and model causal relations?
    2. How can we model sequences of reasoning?
3. Generation

Roadmap for Persuasive Argumentation

1. Persuasion
2. Causal Relations
3. Generation
    1. How can we customize generation to emphasize persuasion?
    2. How can we generate goal-oriented and globally coherent arguments?

Persuasion

1. What makes an argument more persuasive than a logical sequence of reasons?
2. How are persuasive arguments structured?

Social Media

Tan et al. (2016)

Habernal and Gurevych (2016)

Das et al. (2016)

Rosenthal et al. (2017)

Walker et al. (2012)

Persuasive Essays

Peldszus and Stede (2015)

Ghosh et al. (2016)

Somasundaran et al. (2016)

Forbes-Riley et al. (2016)

Persuasion

1. *What makes an argument more persuasive than a logical sequence of reasons?*
2. How are persuasive arguments structured?

Social Media

Tan et al. (2016)

Habernal and Gurevych (2016)

Das et al. (2016)

Rosenthal et al. (2017)

Walker et al. (2012)

Persuasive Essays

Peldszus and Stede (2015)

Ghosh et al. (2016)

Somasundaran et al. (2016)

Forbes-Riley et al. (2016)

Tan et al. (2016)

Goal: Predict persuasion
Data: Change My View
Method: Logistic Regression
Features: Sentiment, Style, Interplay

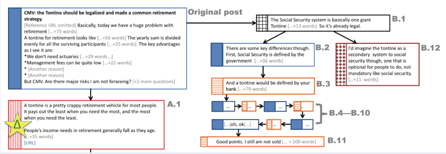

- (+) Naturally labeled open-domain data
- Balanced prediction controlled for topic but (-) assumes persuasion

Winning Arguments: Interaction Dynamics and Persuasion Strategies in Good-faith Online Discussions

Tan et al. (2016)

Goal: Predict persuasion
Data: Change My View
Method: Logistic Regression
Features: Sentiment, Style, Interplay

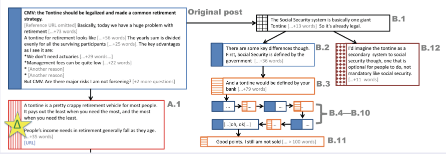

- (+) Naturally labeled open-domain data
- Balanced prediction controlled for topic but (-) assumes persuasion

---

Winning Arguments: Interaction Dynamics and Persuasion Strategies in Good-faith Online Discussions

Persuasion ⟶ Social Media

Tan et al. (2016)

Goal: *Personal* persuasion

- (+) Naturally labeled open-domain data
- Balanced prediction controlled for topic but (-) assumes persuasion

Habernal and Gurevych (2016)

Goal: Ranking arguments
Data: CreateDebate and Procon
Method: SVM and LSTM
Features: Sentiment, Readability

physical education should be mandatory cuhz 112,000 people have died in the year 2011...

- (+) Objective ranking for quality
- (-) May just reveal which arguments are bad

Which argument is more convincing? Analyzing and predicting convincingness of Web arguments using bidirectional LSTM

Persuasion ⟶ Social Media

Tan et al. (2016)

Goal: *Personal* persuasion

- (+) Naturally labeled open-domain data
- Balanced prediction controlled for topic but (-) assumes persuasion

Habernal and Gurevych (2016)

Goal: Ranking arguments
Data: CreateDebate and Procon
Method: SVM and LSTM
Features: Sentiment, Readability

physical education should be mandatory cuhz 112,000 people have died in the year 2011...

- (+) Objective ranking for quality
- (-) May just reveal which arguments are bad

Which argument is more convincing? Analyzing and predicting convincingness of Web arguments using bidirectional LSTM

Persuasion ⟶ Social Media ⟶ Influence

Tan et al. (2016)

Goal: *Personal* persuasion

Habernal and Gurevych (2016)

Goal: *Objectively* ranking arguments

Das et al., (2016)

Goal: Analyze intent in social networks
Data: Manually generated and Twitter
Method: Crowdsourcing and LDA

Hyundai cars just suck. Mine broke down right after their guarantee period.

- (+/-) Measure persuasion by change in sentiment
- (-) Controlled, artificial experiments

Persuasion ⟶ Social Media ⟶ Influence

Tan et al. (2016)

Goal: *Personal* persuasion

Habernal and Gurevych (2016)

Goal: *Objectively* ranking arguments

Das et al., (2016)

Goal: Analyze intent in social networks
Data: Manually generated and Twitter
Method: Crowdsourcing and LDA

Hyundai cars just suck. Mine broke down right after their guarantee period.

- (+/-) Measure persuasion by change in sentiment
- (-) Controlled, artificial experiments

Das et al., (2016)

Goal: Analyze intent in social networks (*global* influence)

Rosenthal and McKeown (2017)

Goal: Predict *personal* influence
Data: LiveJournal, Wikipedia Talk, Twitter, CreateDebate
Method: Cascaded supervised system
Features: Persuasion, Argument, Sentiment, Dialog, Agreement

- (-) Evaluation assumes at least one influencer
- (+) Domain adaptation

Detecting Influencers In Multiple Online Genres

Persuasion ⟶ Social Media ⟶ Influence

Das et al., (2016)

Goal: Analyze intent in social networks (*global* influence)

Rosenthal and McKeown (2017)

Goal: Predict *personal* influence
Data: LiveJournal, Wikipedia Talk, Twitter, CreateDebate
Method: Cascaded supervised system
Features: Persuasion, Argument, Sentiment, Dialog, Agreement

- (-) Evaluation assumes at least one influencer
- (+) Domain adaptation

Detecting Influencers In Multiple Online Genres

Das et al., (2016)

Goal: Analyze intent in social networks (*global* influence)

Rosenthal and McKeown (2017)

Goal: Predict *personal* influence

Walker et al. (2012)

Goal: Predict stance
Data: CreateDebate
Method: MaxCut, Logistic Regression
Features: Sentiment, Argumentation

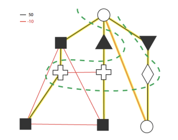

- (+) Naturally-labeled data, (+) proxy for persuasion
- (+) Model social interaction, (-) limited set of topics

Stance Classification using Dialogic Properties of Persuasion

Persuasion ⟶ Social Media ⟶ Stance

Das et al., (2016)

Goal: Analyze intent in social networks (*global* influence)

Rosenthal and McKeown (2017)

Goal: Predict *personal* influence

Walker et al. (2012)

Goal: Predict stance
Data: CreateDebate
Method: MaxCut, Logistic Regression
Features: Sentiment, Argumentation

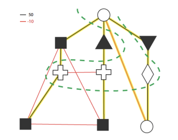

- (+) Naturally-labeled data, (+) proxy for persuasion
- (+) Model social interaction, (-) limited set of topics

Stance Classification using Dialogic Properties of Persuasion

Persuasion ⟶ Social Media

1. *What makes an argument more persuasive than a logical sequence of reasons*?
    - Social Interaction
        - Walker et al. (2012) - graph partitions
        - Das et al. (2016) - neighbor content similarity
        - Tan et al. (2016) - word overlap
        - Rosenthal and McKeown (2017) - dialog patterns
    - Emotional Content
2. How are persuasive arguments structured?

Persuasion ⟶ Social Media

1. *What makes an argument more persuasive than a logical sequence of reasons*?
    - Social Interaction
    - Emotional Content
        - Das et al. (2016) - emotion and logic depending on topic
        - Habernal and Gurevych (2016) - negative often less convincing
        - Tan et al. (2016) - presence of sentiment
        - Rosenthal and McKeown (2017) - sentiment for attempts to persuade
        - Walker et al. (2012) - sentiment for stance
2. How are persuasive arguments structured?

Persuasion

1. What makes an argument more persuasive than a logical sequence of reasons?
2. *How are persuasive arguments structured?*

Social Media	**Persuasive Essays**
Tan et al. (2016)	Peldszus and Stede (2015)
Habernal and Gurevych (2016)	Ghosh et al. (2016)
Das et al. (2016)	Somasundaran et al. (2016)
Rosenthal et al. (2017)	Forbes-Riley et al. (2016)
Walker et al. (2012)	

Peldszus and Stede (2015)

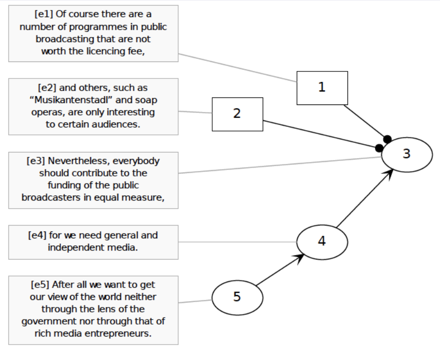

Goal: Argumentation parsing
Data: Manually generated German and (-) translated English essays
Method: Logistic regression, MST

- Claims/premises and support/attack relations
- (+) Joint prediction, (-) but components modeled individually

Joint prediction in MST-style discourse parsing for argumentation mining

Persuasion ⟶ Essays

Peldszus and Stede (2015)

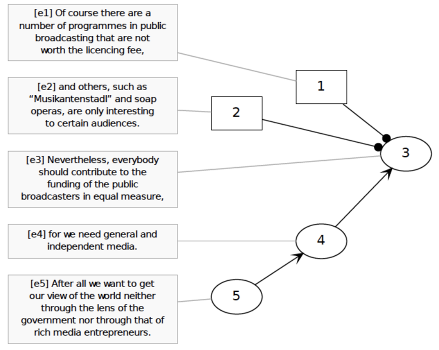

Goal: Argumentation parsing
Data: Manually generated German and (-) translated English essays
Method: Logistic regression, MST

- Claims/premises and support/attack relations
- (+) Joint prediction, (-) but components modeled individually

---

Joint prediction in MST-style discourse parsing for argumentation mining

Persuasion ⟶ Essays

Ghosh et al. (2016)

Goal: Persuasive essay scoring
Data: TOEFL essays
Method: Linear regression
Features: Argumentation

- (+/-) Coarse-grained claims/premises and support/attack relations

Goal: Argumentation parsing
Data: Manually generated German and (-) translated English essays
Method: Logistic regression, MST

- Claims/premises and support/attack relations
- (+) Joint prediction, (-) but components modeled individually

Coarse-grained Argumentation Features for Scoring Persuasive Essays

Persuasion ⟶ Essays

Somasundaran et al. (2016)

Goal: Automatic essay scoring
Data: GRE essays
Methods: Linear Regression
Features: PageRank and graph-based

$$\begin{tikzcd} \text{countries} \arrow[r] \arrow[dr] & \text{values} \\ & \text{culture} \end{tikzcd}$$

- Model (+) *globally* as graphs with each word as a node
- (-) All nodes of the same word are collapsed

Evaluating Argumentative and Narrative Essays using Graphs

Persuasion ⟶ Essays

Somasundaran et al. (2016)

Goal: Automatic essay scoring
Data: GRE essays
Methods: Linear Regression
Features: PageRank and graph-based

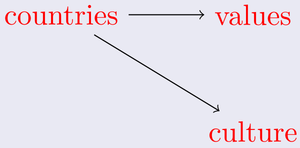

- Model (+) *globally* as graphs with each word as a node
- (-) All nodes of the same word are collapsed

Evaluating Argumentative and Narrative Essays using Graphs

Somasundaran et al. (2016)

Goal: Automatic essay scoring

- Model (+) *globally* as graphs with each word as a node
- (-) All nodes of the same word are collapsed

Forbes-Riley et al. (2016)

Goal: Analyze and predict Penn Discourse Tree bank relations
Data: AP English essays
Methods: Crowdsourcing and pre-trained discourse parser

- Mostly sequential *local* relations
- More Contingency relations, (-) missing Justification and Claim

Extracting PDTB Discourse Relations from Student Essays

Somasundaran et al. (2016)

Goal: Automatic essay scoring

- Model (+) *globally* as graphs with each word as a node
- (-) All nodes of the same word are collapsed

Forbes-Riley et al. (2016)

Goal: Analyze and predict Penn Discourse Tree bank relations
Data: AP English essays
Methods: Crowdsourcing and pre-trained discourse parser

- Mostly sequential *local* relations
- More Contingency relations, (-) missing Justification and Claim

---

Extracting PDTB Discourse Relations from Student Essays

Persuasion ⟶ Essays

1. What makes an argument more persuasive than a logical sequence of reasons?
2. *How are persuasive arguments structured?*
    - Ghosh et al. (2016) and Peldszus and Stede (2015) use tree structures
    - Somasundaran et al. (2016) study graphs of word interactions
    - Forbes-Riley et al. (2016) analyze local discourse relations

Persuasion

Goals of persuasive argumentation:

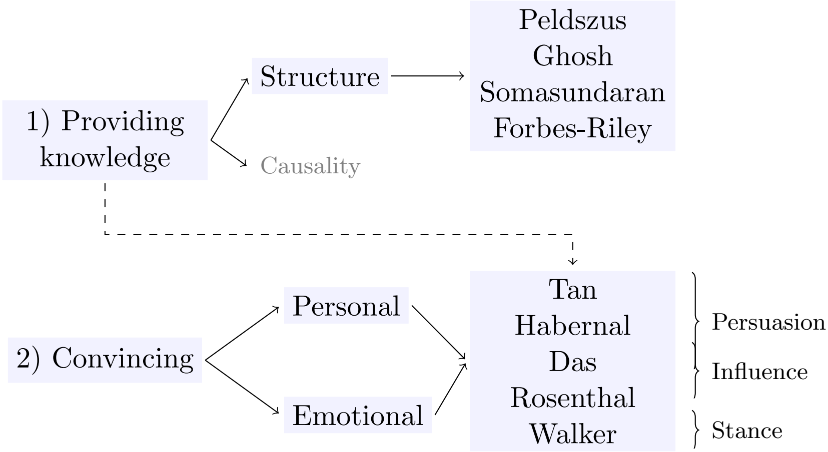

Causal Relations

Causal relations for persuasive argumentation:

1. Mining factual causal relations
2. Modeling causal relations in persuasive argumentation

Goals:

1. How can we better represent and model causal relations?
2. How can we model sequences of reasoning?

Causal Relations

Causal relations for persuasive argumentation:

1. Mining factual causal relations
2. Modeling causal relations in persuasive argumentation

Goals:

1. How can we better represent and model causal relations?
2. How can we model sequences of reasoning?

Causal Relations

1. How can we better represent and model causal relations?
2. How can we model sequences of reasoning?

Contextual	**Distributional**
Ji et al (2016)	Biran and McKeown (2013)
Prasad et al. (2010)	Braud and Denis (2016)
Dunietz et al. (2017)	Sharp et al. (2016)
Riaz and Girju (2014)	
	Rocktaschel et al. (2015)
	Das et al. (2017)

Causal Relations

1. How can we better represent and model causal relations?
2. How can we model sequences of reasoning?

Contextual

Ji et al (2016)

Prasad et al. (2010)

Dunietz et al. (2017)

Riaz and Girju (2014)

Distributional

Biran and McKeown (2013)

Braud and Denis (2016)

Sharp et al. (2016)

Formal Logic

Rocktaschel et al. (2015)

Das et al. (2017)

Causal Relations

1. *How can we better represent and model causal relations?*
2. How can we model sequences of reasoning?

Contextual

Ji et al (2016)

Prasad et al. (2010)

Dunietz et al. (2017)

Riaz and Girju (2014)

Distributional

Biran and McKeown (2013)

Braud and Denis (2016)

Sharp et al. (2016)

Formal Logic

Rocktaschel et al. (2015)

Das et al. (2017)

Causal Relations → Contextual Approaches

Ji et al (2016)

Goal: Predict implicit discourse relations

John was tired. He left early.

Data: Wall Street Journal (PDTB)

Model: LSTM with discourse relation as latent variable

- (+) Discourse-aware language modeling
- (-) Implicit discourse relation detection still very difficult
- (-) No reporting of individual class performance

A Latent Variable Recurrent Neural Network for Discourse Relation Language Models

Ji et al (2016)

Goal: Predict implicit discourse relations

John was tired. He left early.

Data: Wall Street Journal (PDTB)

Model: LSTM with discourse relation as latent variable

- (+) Discourse-aware language modeling
- (-) Implicit discourse relation detection still very difficult
- (-) No reporting of individual class performance

A Latent Variable Recurrent Neural Network for Discourse Relation Language Models

Ji et al (2016)

Goal: Predict implicit discourse relations (still (-) very difficult)
John was tired. He left early.

Prasad et al. (2010)

Goal: Identify alternative discourse markers
GM appears to be stepping up the pace of its factory consolidation to get in shape for the 1990s. **One reason is** mounting competition.
Data: Wall Street Journal (PDTB)
Model: Paraphrases

- (+) Provides lexical signal, (+/-) open class of markers
- (-) Limited to intra-sentence relations

Realization of Discourse Relations by Other Means: Alternative Lexicalizations

Causal Relations → Contextual Approaches

Ji et al (2016)

Goal: Predict implicit discourse relations (still (-) very difficult)
John was tired. He left early.

Prasad et al. (2010)

Goal: Identify alternative discourse markers
GM appears to be stepping up the pace of its factory consolidation to get in shape for the 1990s. **One reason is** mounting competition.
Data: Wall Street Journal (PDTB)
Model: Paraphrases

- (+) Provides lexical signal, (+/-) open class of markers
- (-) Limited to intra-sentence relations

Realization of Discourse Relations by Other Means: Alternative Lexicalizations

Causal Relations → Contextual Approaches

Ji et al (2016)

Goal: Predict implicit discourse relations (still (-) very difficult)

Prasad et al. (2010)

Goal: Identify alternative discourse markers

Dunietz et al. (2017)

Goal: Predict causality and cause/effect spans

For market discipline **to** work, banks cannot expect to be bailed out.

Data: New York Times, Wall Street Journal, Dodd-Frank hearings

Model: Cascaded supervised system

Features: Lexical, Syntactic, Semantic

- (+) Contiguous and non-contiguous, but (-) no temporal
- (-) Closed class at prediction, (-) per-relation classifier

Causal Relations → Contextual Approaches

Ji et al (2016)

Goal: Predict implicit discourse relations (still (-) very difficult)

Prasad et al. (2010)

Goal: Identify alternative discourse markers

Dunietz et al. (2017)

Goal: Predict causality and cause/effect spans

For market discipline **to** work, banks cannot expect to be bailed out.

Data: New York Times, Wall Street Journal, Dodd-Frank hearings

Model: Cascaded supervised system

Features: Lexical, Syntactic, Semantic

- (+) Contiguous and non-contiguous, but (-) no temporal
- (-) Closed class at prediction, (-) per-relation classifier

Automatically Tagging Constructions of Causation and Their Slot-Fillers

Causal Relations → Contextual Approaches

Dunietz et al. (2017)

For market discipline **to** work, banks cannot expect to be bailed out.

- Lexical grounding

Riaz and Girju (2014)

Goal: Predict causality

At least 1,833 people **died** in the **hurricane**.

Data: FrameNet, WordNet, and GigaWord

Model: Semi-supervised ILP

- (+) Non-contiguous, (+) open class
- (+/-) Requires real-world definition of causality
- (-) Missing other causal constructions

In-depth Exploitation of Noun and Verb Semantics to Identify Causation in Verb-Noun Pairs

Causal Relations → Contextual Approaches

Dunietz et al. (2017)

For market discipline **to** work, banks cannot expect to be bailed out.

- Lexical grounding

Riaz and Girju (2014)

Goal: Predict causality

At least 1,833 people **died** in the **hurricane**.

Data: FrameNet, WordNet, and GigaWord

Model: Semi-supervised ILP

- (+) Non-contiguous, (+) open class
- (+/-) Requires real-world definition of causality
- (-) Missing other causal constructions

In-depth Exploitation of Noun and Verb Semantics to Identify Causation in Verb-Noun Pairs

Causal Relations → Contextual Approaches

1. *How can we better represent and model causal relations*?
   - Dunietz et al. (2017)- expand to constructions like "**so ... that**"
   - Prasad et al. (2010)- alternative lexicalizations, "**The reason is**"
   - Riaz and Girju (2014)- verb-noun pairs such as "**died/hurricane**"
   - Ji et al. (2016)- implicit discourse relations as latent variables
2. How can we model sequences of reasoning?

Causal Relations

1. *How can we better represent and model causal relations?*
2. How can we model sequences of reasoning?

Contextual

Ji et al (2016)

Prasad et al. (2010)

Dunietz et al. (2017)

Riaz and Girju (2014)

Distributional

Biran and McKeown (2013)

Braud and Denis (2016)

Sharp et al. (2016)

Formal Logic

Rocktaschel et al. (2015)

Das et al. (2017)

Biran and McKeown (2013)

Goal: Distributed representations for implicit discourse

Method: Calculate weighted word-pairs for each explicit connective

- (-) Unable to score unseen word pairs
- (+/-) Simple pre-processing, (-) no evaluation

Aggregated Word Pair Features for Implicit Discourse Relation Disambiguation

Biran and McKeown (2013)

Goal: Distributed representations for implicit discourse

Method: Calculate weighted word-pairs for each explicit connective

- (-) Unable to score unseen word pairs
- (+/-) Simple pre-processing, (-) no evaluation

Aggregated Word Pair Features for Implicit Discourse Relation Disambiguation

Causal Relations → Distributional Approaches

Goal: Distributed representations for implicit discourse
TF-IDF and PMI-IDF, with IDF over connectives

Biran and McKeown (2013)

- (-) Requires lots of training data, unable to score unseen word pair
- (+/-) Simple pre-processing, (-) no evaluation

Braud and Denis (2016)

Method: Each word is a weighted d-dimensional vector

- (+) Evaluation of pre-processing
- (-) Expanding to additional markers increases sparsity

Learning Connective-based Word Representations for Implicit Discourse Relation Identification

Causal Relations → Distributional Approaches

Goal: Distributed representations for implicit discourse
TF-IDF and PMI-IDF, with IDF over connectives

Biran and McKeown (2013)

- (-) Requires lots of training data, unable to score unseen word pair
- (+/-) Simple pre-processing, (-) no evaluation

Braud and Denis (2016)

Method: Each word is a weighted d-dimensional vector

- (+) Evaluation of pre-processing
- (-) Expanding to additional markers increases sparsity

Learning Connective-based Word Representations for Implicit Discourse Relation Identification

Causal Relations → Distributional Approaches

Biran and McKeown (2013), Braud and Denis (2016)

Goal: Distributed representations for implicit discourse

Sharp et al. (2016)

Goal: Distributed representations for causality
Method: skip-gram, word-context pairs are from causes and effects

- (-) Simple pre-processing, (+/-) some evaluation of span selection
- (+) Both intrinsic and extrinsic evaluation

Creating Causal Embeddings for Question Answering with Minimal Supervision

Causal Relations

1. How can we better represent and model causal relations?
2. *How can we model sequences of reasoning?*

Contextual

Ji et al (2016)

Prasad et al. (2010)

Dunietz et al. (2017)

Riaz and Girju (2014)

Distributional

Biran and McKeown (2013)

Braud and Denis (2016)

Sharp et al. (2016)

Formal Logic

Rocktaschel et al. (2015)

Das et al. (2017)

Rocktaschel et al. (2015)

Goal: Perform inductive reasoning on a knowledge base
Data: New York Times (train) and Freebase (train/test)
Methods: Matrix factorization and probabilistic logic rules

$$r_s(x, y) \implies r_t(x, y)$$

$$[\mathcal{A} \implies \mathcal{B}] = [\mathcal{A}]\,([\mathcal{B}] - 1) + 1$$

Injecting Logical Background Knowledge into Embeddings for Relation Extraction

Causal Relations ⟶ Distributional ⟶ Formal Logic

Goal: Perform inductive reasoning on a knowledge base

Rocktaschel et al. (2015)

Methods: Matrix factorization and probabilistic logic rules

Das et al. (2017)

Data: Freebase

Methods: RNN over paths in a knowledge base

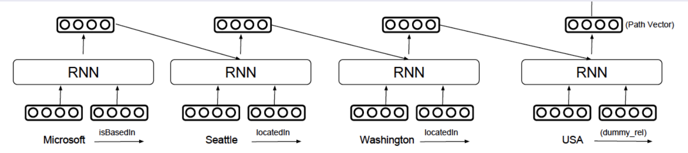

Chains of Reasoning over Entities, Relations, and Text using Recurrent Neural Networks

Goal: Perform inductive reasoning on a knowledge base

Rocktaschel et al. (2015)

Methods: Matrix factorization and probabilistic logic rules

Das et al. (2017)

Methods: RNN over paths in a knowledge base

- (+) Open set of relations
- (-) Difficult to model confounding variables and other complex interactions

1. *How can we better represent and model causal relations?*
    - Biran and McKeown (2013) - word pairs for explicit connectives
    - Braud and Denis (2016) - word co-occurrence vectors
    - Sharp et al. (2016) - skip-gram for cause/effect word pairs
2. How can we model sequences of reasoning?

Causal Relations → Distributional Approaches

1. How can we better represent and model causal relations?
2. *How can we model sequences of reasoning*?
    - Rocktaschel et al. (2015) - matrix factorization with injected logic
    - Das et al. (2017) - RNNs over paths in knowledge graph

Causal Relations

Goals of persuasive argumentation:

1) Providing knowledge → Structure
1) Providing knowledge → Causal relations

Causal relations → Contextual
Causal relations → Distributional

Contextual:

Ji	PDTB
Prasad	Alt. lex.
Dunietz	construction
Riaz	verb-noun

Distributional:

Biran Braud	PDTB
Sharp	causal
Rocktaschel Das	logic

2) Convincing → Personal
2) Convincing → Emotional

Generation

Natural language generation for persuasive argumentation:

1. Content-framed
2. Context-driven
3. Goal-oriented
4. Globally coherent

1. How can we customize generation to emphasize persuasion?
2. How can we generate goal-oriented and globally coherent arguments?

Generation

Natural language generation for persuasive argumentation:

1. Content-framed
2. Context-driven
3. Goal-oriented
4. Globally coherent

1. How can we customize generation to emphasize persuasion?
2. How can we generate goal-oriented and globally coherent arguments?

Generation

1. How can we customize generation to emphasize persuasion?
2. How can we generate goal-oriented and globally coherent arguments?

Persuasion

Ding and Pan (2016)

Bilu and Slonim (2016)

Li et al. (2016)

Dodge et al. (2016)

Other

Andreas and Klein (2016)

Hu et al. (2017)

Chen et al. (2009)

Kiddon et al. (2016)

Generation

1. How can we customize generation to emphasize persuasion?
2. How can we generate goal-oriented and globally coherent arguments?

Framing

Ding and Pan (2016)

Bilu and Slonim (2016)

Goal-oriented

Li et al. (2016)

Dodge et al. (2016)

Context-driven

Andreas and Klein (2016)

Hu et al. (2017)

Coherent

Chen et al. (2009)

Kiddon et al. (2016)

Generation

1. *How can we customize generation to emphasize persuasion?*
2. How can we generate goal-oriented and globally coherent arguments?

Framing

Ding and Pan (2016)

Bilu and Slonim (2016)

Goal-oriented

Li et al. (2016)

Dodge et al. (2016)

Context-driven

Andreas and Klein (2016)

Hu et al. (2017)

Coherent

Chen et al. (2009)

Kiddon et al. (2016)

Framing Content

Bilu and Slonim (2016)

Goal: Generate valid claims (template-based)

Data: idebate

Banning violent video games **is a violation of free speech**

Censoring internet content **is a violation of free speech**

Method: Logistic regression

Features: similarity, relevance, fluency

- (+) Parameter sharing across topics
- (+/-) Text-to-text generation, (-) closed set

Claim Synthesis via Predicate Recycling

Framing Content

Bilu and Slonim (2016)

Goal: Generate valid claims (template-based)

Data: idebate

Banning violent video games **is a violation of free speech**

Censoring internet content **is a violation of free speech**

Method: Logistic regression

Features: similarity, relevance, fluency

- (+) Parameter sharing across topics
- (+/-) Text-to-text generation, (-) closed set

Claim Synthesis via Predicate Recycling

Framing Content

Bilu and Slonim (2016)

Goal: Generate valid claims

Ding and Pan (2016)

Goal: Determine effects of personality on persuasion
Data: Personality tests
Method: Metric Pairwise Constrained K-Means
Features: Big5, Schwartz

- (-) Domain-specific
- (-) No control for how personality affects generation decisions

Personalized Emphasis Framing for Persuasive Message Generation

Framing Content

Bilu and Slonim (2016)

Goal: Generate valid claims

Ding and Pan (2016)

Goal: Determine effects of personality on persuasion
Data: Personality tests
Method: Metric Pairwise Constrained K-Means
Features: Big5, Schwartz

- (-) Domain-specific
- (-) No control for how personality affects generation decisions

Personalized Emphasis Framing for Persuasive Message Generation

Generation

1. *How can we customize generation to emphasize persuasion?*
2. How can we generate goal-oriented and globally coherent arguments?

Framing

Ding and Pan (2016)

Bilu and Slonim (2016)

Goal-oriented

Li et al. (2016)

Dodge et al. (2016)

Context-driven

Andreas and Klein (2016)

Hu et al. (2017)

Coherent

Chen et al. (2009)

Kiddon et al. (2016)

Considering Context

Andreas and Klein (2016)

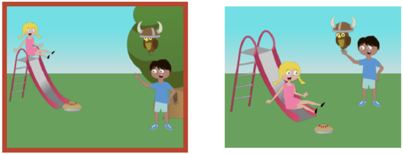

(a) target (b) distractor

the owl is sitting in the tree

Goal: Generate reference text
Data: Abstract Scenes Dataset
Method: Referent ranker, text generator

- (+) Contextual, social interaction
- (+) Agnostic to input representation
- (-) Sampling instead of joint modeling

Reasoning about Pragmatics with Neural Listeners and Speakers

Considering Context

Andreas and Klein (2016)

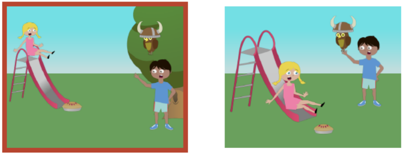

(a) target (b) distractor

the owl is sitting in the tree

Goal: Generate reference text
Data: Abstract Scenes Dataset
Method: Referent ranker, text generator

- (+) Contextual, social interaction
- (+) Agnostic to input representation
- (-) Sampling instead of joint modeling

Reasoning about Pragmatics with Neural Listeners and Speakers

Considering Context

Andreas and Klein (2016)

Goal: Generate reference text

Hu et al. (2017)

Goal: Generate controllable text
Data: IMDB, Stanford Sentiment Treebank-2, TimeBank
Method: Variational Auto-Encoder
the film is strictly routine !
the film is full of imagination .

- (+) Semi-supervised, requires little labeled data
- (-) Unclear how to extend to multi-dimensional attributes with complex interactions

Controllable Text Generation

Considering Context

Andreas and Klein (2016)

Goal: Generate reference text

Hu et al. (2017)

Goal: Generate controllable text
Data: IMDB, Stanford Sentiment Treebank-2, TimeBank
Method: Variational Auto-Encoder
the film is strictly routine !
the film is full of imagination .

- (+) Semi-supervised, requires little labeled data
- (-) Unclear how to extend to multi-dimensional attributes with complex interactions

Generation

1. How can we customize generation to emphasize persuasion?
2. *How can we generate goal-oriented and globally coherent arguments?*

Framing

Ding and Pan (2016)

Bilu and Slonim (2016)

Goal-oriented

Li et al. (2016)

Dodge et al. (2016)

Context-driven

Andreas and Klein (2016)

Hu et al. (2017)

Coherent

Chen et al. (2009)

Kiddon et al. (2016)

Accomplishing Goals

Li et al. (2016)

Goal: Generate dialog for maximizing the length of the conversation
Data: OpenSubtitles
Method: Deep reinforcement learning

A: Where are you going?
B: I'm going to the restroom.
A: See you later.
B: See you later.
A: See you later.
...

- (+) Models both agents in dialog simultaneously
- (-) Preventing loops may contrast with other goals

Deep Reinforcement Learning for Dialogue Generation

Accomplishing Goals

Li et al. (2016)

Goal: Generate dialog for maximizing the length of the conversation
Data: OpenSubtitles
Method: Deep reinforcement learning

A: Where are you going?
B: I'm going to the restroom.
A: See you later.
B: See you later.
A: See you later.
...

- (+) Models both agents in dialog simultaneously
- (-) Preventing loops may contrast with other goals

Deep Reinforcement Learning for Dialogue Generation

Accomplishing Goals

Li et al. (2016)

Goal: Generate dialog likely to result in continued dialog

Dodge et al. (2016)

Goal: Generate dialog for question answering
Data: Online Movie Database, Reddit movies sub-reddit
Method: Memory network

A: I liked Tombstone and The Net. I'm looking for a Fantasy film.
B: Jumanji
A: Who directed that?
B: Joe Johnston
A: I like Tim Burton movies more...

- (+) Ability to store and query factual information
- (-) No shared representation between memory elements

Evaluating Prerequisite Qualities for Learning End-to-End Dialog Systems

Accomplishing Goals

Li et al. (2016)

Goal: Generate dialog likely to result in continued dialog

Dodge et al. (2016)

Goal: Generate dialog for question answering
Data: Online Movie Database, Reddit movies sub-reddit
Method: Memory network

A: I liked Tombstone and The Net. I'm looking for a Fantasy film.
B: Jumanji
A: Who directed that?
B: Joe Johnston
A: I like Tim Burton movies more...

- (+) Ability to store and query factual information
- (-) No shared representation between memory elements

Evaluating Prerequisite Qualities for Learning End-to-End Dialog Systems

Generation

1. How can we customize generation to emphasize persuasion?
2. *How can we generate goal-oriented and globally coherent arguments?*

Framing

Ding and Pan (2016)

Bilu and Slonim (2016)

Goal-oriented

Li et al. (2016)

Dodge et al. (2016)

Context-driven

Andreas and Klein (2016)

Hu et al. (2017)

Coherent

Chen et al. (2009)

Kiddon et al. (2016)

Chen et al. (2009)

Goal: Model topic transitions
Data: Wikipedia
Method: Generalized Mallows Model

- (+) Works well for domain-specific modeling
- (-) Bag-of-words generation

Global Models of Document Structure Using Latent Permutations

Modeling Coherence

Chen et al. (2009)

Goal: Model topic transitions
Data: Wikipedia
Method: Generalized Mallows Model

- (+) Works well for domain-specific modeling
- (-) Bag-of-words generation

Global Models of Document Structure Using Latent Permutations

Modeling Coherence

Chen et al. (2009)

Goal: Improve topic transitions by global constraints on ordering

Kiddon et al. (2016)

Goal: Generate text from an agenda
Data: Recipes, Hotel dialogs
Method: Neural LM with soft checklist

Sift **flour**, measure, and sift with **baking powder** and **salt**. Fold in stiffly beaten **egg whites**.

- Able to balance long-term goals with short-term word generation

Globally Coherent Text Generation with Neural Checklist Models

Generation

1. *How can we customize generation to emphasize persuasion?*
    - *Framing*
        - Bilu and Slonim (2016) - template-based generation of claims
        - Ding and Pan (2016) - emphasis of attributes based on personality
    - Context
2. How can we generate goal-oriented and globally coherent arguments?
    - Goals
    - Coherence

Generation

1. *How can we customize generation to emphasize persuasion?*
   - Framing
   - *Context*
     - Andreas and Klein (2016) - pragmatic reasoning for descriptions
     - Hu et al. (2017) - text generation conditioned on attributes
2. How can we generate goal-oriented and globally coherent arguments?
   - Goals
   - Coherence

Generation

1. How can we customize generation to emphasize persuasion?
    - Framing
    - Context
2. *How can we generate goal-oriented and globally coherent arguments?*
    - *Goals*
        - Li et al. (2016) - maximizing conversation length for dialogue
        - Dodge et al. (2016) - question answering for dialogue
    - Coherence

Generation

1. How can we customize generation to emphasize persuasion?
    - Framing
    - Context
2. *How can we generate goal-oriented and globally coherent arguments?*
    - Goals
    - *Coherence*
        - Chen et al. (2009) - topic modeling and ordering
        - Kiddon et al. (2016) - agenda-driven generation

Generation

Goals of persuasive argumentation:

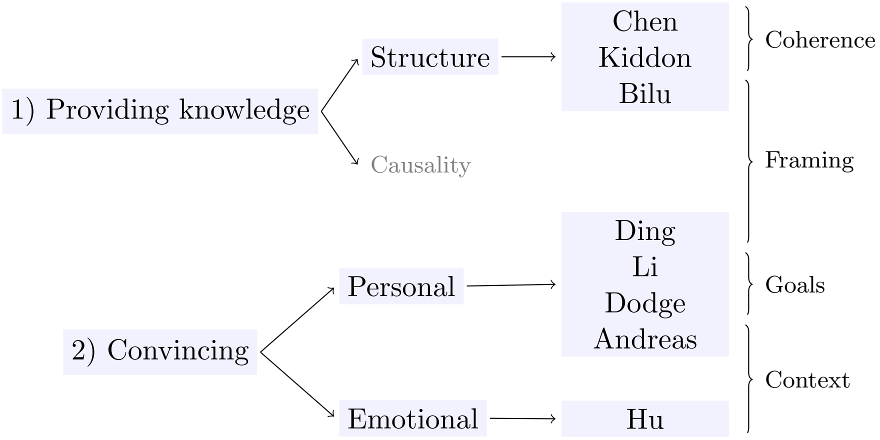

Conclusion

Goals of persuasive argumentation:

1) Providing knowledge

Structure

Trees/Graphs
Coherence
Framing

Causality

Contextual
Distributional
Formal Logic

2) Convincing

Personal

Social Interaction
Framing
Pragmatics

Emotional

Sentiment
Topic/Context