Recognizing stances, arguments, viewpoints

Ruth Morrison, Julian Chan

Somasundaran and Wiebe (2009), "Recognizing Stances in Online Debates"

Souneil et. al (2011), "Cotrasting Opposing Views of News Articles on Contentious Issues"

Walker et. al (2011), "That's Your Evidence? Classifying Stance in Online Political and Social Debate"

Thomas et. al (2006), "Get Out the Vote: Determining Support or Opposition from Congressional Floor-Debate Transcripts"

Abu-Jbara et. al (2011), "Subgroup Detection in Ideological Discussions"

Overview

- First paper: Two-sided online debates (pro/against).
- Second paper: News articles about contentious issues.
- Different definition of “side”.

Somasundaran and Wiebe (2009): Overview

- Goal: An unsupervised method of detecting stance in two-sided online debates
- Problems
  - Debators will alternate between multiple topics and polarities per post, sometimes per sentence
  - Debators will refer to aspects and features of the debate topic rather than repeating the topic name itself
  - Debators make concessions to the other side
- Approach
  - Learn aspects that correspond to sides, and apply linear programming to compute the side of individual posts

Mining the web for opinions on features

- Connect opinions to target features and topics using a 8000-word subjectivity lexicon , the Stanford dependency parser, and some syntactic rules:

DIRECT OBJECT Rule: dobj(opinion, target) In words: The target is the direct object of the opinion Example: I love$_{opinion1}$ Firefox$_{target1}$ and defended$_{opinion2}$ it$_{target2}$
NOMINAL SUBJECT Rule: nsubj(opinion, target) In words: The target is the subject of the opinion Example: IE$_{target}$ breaks$_{opinion}$ with everything.
ADJECTIVAL MODIFIER Rule: amod(target, opinion) In words: The opinion is an adjectival modifier of the target Example: The annoying$_{opinion}$ popup$_{target}$
PREPOSITIONAL OBJECT Rule: if prep(target1,IN) $\Rightarrow$ pobj(IN, target2) In words: The prepositional object of a known target is also a target of the same opinion Example: The annoying$_{opinion}$ popup$_{target1}$ in IE$_{target2}$ ("popup" and "IE" are targets of "annoying")
RECURSIVE MODIFIERS Rule: if conj(adj2, opinion$_{adj1}$) $\Rightarrow$ amod(target, adj2) In words: If the opinion is an adjective (adj1) and it is conjoined with another adjective (adj2), then the opinion is tied to what adj2 modifies Example: It is a powerful$_{opinion(adj1)}$ and easy$_{opinion(adj2)}$ application$_{target}$ ("powerful" is attached to the target "application" via the adjective "easy")

Associating positive and negative attitudes towards features with topics

- First, positive and negative opinions of the debate topics are mined, and nearby features (within 5 sentences) are also noted
- Conditional probabilities for P (topic+/- | target+/-) are calculated
- Some examples:

	$side_1$ (pro-iPhone)		$side_2$ (pro-blackberry)	
$term^p$	$P(iPhone^+ \mid term^p)$	$P(blackberry^- \mid term^p)$	$P(iPhone^- \mid term^p)$	$P(blackberry^+ \mid term^p)$
$storm^+$	0.227	0.068	0.022	0.613
$storm^-$	0.062	0.843	0.06	0.03
$phone^+$	0.333	0.176	0.137	0.313

Calculating the debate side of a post

- For each instance of a target-polarity pair in the post, values *w* and *u* are calculated for the two sides:

$$w_j = P(topic_1^+|target_i^p) + P(topic_2^-|target_i^p) \quad (2)$$

$$u_j = P(topic_1^-|target_i^p) + P(topic_2^+|target_i^p) \quad (3)$$

- Then linear programming is applied:

$$\sum_{j=1}^{N}(w_j x_j + u_j y_j)$$

$$x_j \in \{0,1\},\ \forall j$$

$$y_j \in \{0,1\},\ \forall j$$

$$x_j + y_j = 1,\ \forall j$$

$$x_j - x_{j-1} = 0,\ j \in \{2..N\}$$

$$y_j - y_{j-1} = 0,\ j \in \{2..N\}$$

Dealing with concessions

- Concessions are identified using the Penn Discourse Treebank list of discourse connectives (from the concession and contra-expectation categories):
- *While iPhone may appeal to younger generations and BB to older...*
- *Vista will close the gap on the interface some but...*
- Opinions found in conceded clauses count towards the opposite side (*w* and *u* are reversed)

Test data

- 4 debates (Firefox vs. IE, PC vs. Mac, PS3 vs. Wii, Opera vs. Firefox)
- 117 posts of at least 5 sentences
- All posts were automatically gold-labelled by convinceme.net

Baselines

- OpTopic
  - Only considers opinion words directly tied to topic names
- OpPMI
  - PMI: Pointwise Mutual Information
  - Measures of Semantic Relatedness engine searches Google to find "related" topics
  - Opinions on these topics count as opinions on their most closely related debate topic
- Both use the same opinion word lexicon and target word identification algorithms as OpPr

Evaluation

- 17%/20% increase in F-measure and 20%/35% increase in accuracy over baselines
- The addition of concession handling also helped a little

	OpTopic	OpPMI	OpPr	OpPr + Disc
Firefox Vs Internet explorer (62 posts)				
Acc	33.87	53.23	64.52	66.13
Prec	67.74	60.0	64.52	66.13
Rec	33.87	53.23	64.52	66.13
F1	45.16	56.41	64.52	66.13
Windows vs. Mac (15 posts)				
Acc	13.33	46.67	66.67	66.67
Prec	40.0	53.85	66.67	66.67
Rec	13.33	46.67	66.67	66.67
F1	20.0	50.00	66.67	66.67
SonyPs3 vs. Wii (36 posts)				
Acc	33.33	33.33	56.25	61.11
Prec	80.0	46.15	56.25	68.75
Rec	33.33	33.33	50.0	61.11
F1	47.06	38.71	52.94	64.71
Opera vs. Firefox (4 posts)				
Acc	25.0	50.0	75.0	100.0
Prec	33.33	100	75.0	100.0
Rec	25.0	50	75.0	100.0
F1	28.57	66.67	75.0	100.0

Error analysis

- False lexicon hits from words with both subjective and objective meanings
- Target identification errors
- "Pragmatic" opinions that require real-world knowledge (e.g. cost)

Critique

- Very successful over baselines
- Smallish test set compared to the other paper using convinceme.net data
- Not entirely "unsupervised" due to opinion lexicon and discourse lexicon
- Domain issues - only focused on product debates
- Noted that the target identification rules were a source of errors (and no data on that)
- Does not work with >2-sided debates
- Issues with concession handling order

Contrasting Opposing Views of News Articles on Contentious Issues

- Souneil Park, KyungSoon Lee, Junehwa Song

- Done on Korean articles

- Goal:

To give the reader a balanced understanding of the contentious issues by showing the positions of each disputant.

News articles != online debates

- Previous paper focused on identifying stance using positive/negative features. ("I like the iPhone because the camera takes great cat pictures.")
- This works well for online debates about products
- BUT, news articles are an entirely different beast…

News Articles on Contentious Issues

- Unlike debate posts or product reviews, news articles on contentious issues tend to:
  - Span over different topics
  - Not take a position explicitly (“Fair & Balanced”)
    - Include carefully selected facts to cast negative/positive light on government.
  - Have no clear positive/negative distinction.
    - Example 1: Contention over referendum on the Sejong project:
      - Opponents: “The president is a jerk!”
      - President’s office: “We are not considering holding a referendum. Learn how to read.”

She said, he said, they said

- **Solution:** frame the problem based on disputes between different groups (aka. “disputants”).

Benefits of an Opponent-based Frame

- Does not require the documents to discuss common topics nor the opposing arguments to be positive vs. negative.

- Focuses on quotes to identify disputants. Quotes are in abundant supply and easy to identify.

- Aligns with how people perceive contentious issues.

Extracting Disputants

- Many disputants appear as the subject of quotes in the news article set.
- Subjects of direct and indirect quotes are extracted.
- Uses the Korean Named Entity Recognizer and simple anaphora resolution.

Disputant Partitioning

- Identify 2 key opponents, each representing one side, and uses them as a pivot for partitioning other disputants.
- The other disputants are divided according to their relation with the key opponents
- Ex. North Korea and South Korea are the key opponents; other disputants (politicians, experts, US, China) mostly speak about the key opponents.
- It is effective to analyze where the disputants stand regarding their attitude toward the key opponents.

Selecting Key Opponents

- Find the "players" and "player haters", the loudmouths.
- Search for disputants who frequently criticize, and are also criticized by other disputants.
- Map out who the disputant criticizes and who criticizes him/her.
- A sentence is considered to express the disputant's criticism to another disputant if:
    1. The sentence is a quote
    2. The disputant is the subject of the quote
    3. Another disputant appears in the quote.
    4. A negative lexicon appears in the sentence

HITS algorithm

- Effective algorithm for identifying the two key opponents
- Each disputant is modeled as a node
- A link is made from a criticizing disputant to a criticized disputant.
- Each node has two scores:
  - Authority score
    - Value of IN links.
    - Increases it is pointed by many nodes with high Hub score.
    - Initially set to number of quotes in which the disputant appears but is NOT the subject
  - Hub score
    - Value of OUT links.
    - Increases if it points to many nodes with high Authority score
    - Initially set to number of quotes in which the disputant is the subject

HITS algorithm cont.

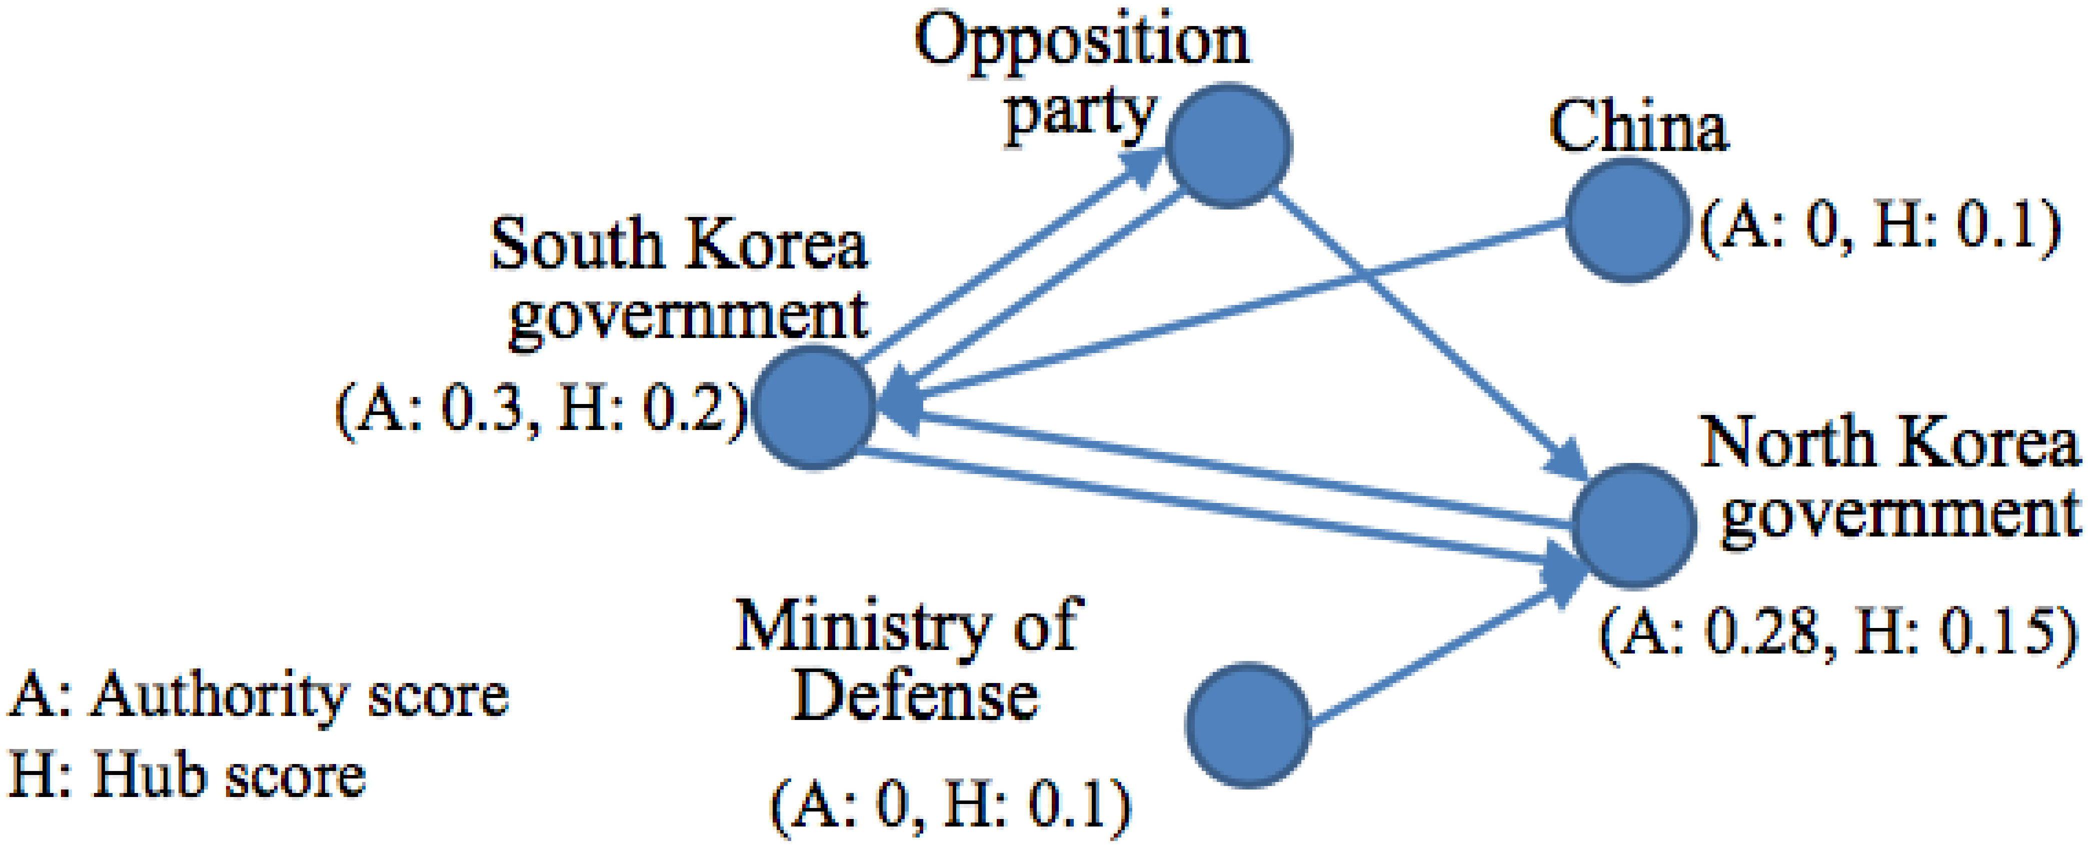

Partitioning minor disputants:

- Positive Quote Rate: Given a key opponent A, and a minor disputant B, the feature measures the ratio of positive quotes between them.
  - The sentence is a direct or indirect quote
  - The 2 disputants appear in the sentence, one is the subject
  - A positive lexicon appears in the sentence.
  - The number of such sentences is divided by the number of all quotes
- Negative Quote Rate: opposite of PQR.
- Frequency of Standing Together (ex. "South Korea and US both criticize North Korea for...")
- Frequency of Division: opposite of FST.

Partitioning minor disputants (cont.)

classify a to b's side if,

$$(PQR_{ab} - NQR_{ab}) > (PQR_{ac} - NQR_{ac}) \text{ or } ((FST_{ab} > FD_{ab}) \text{ and } (FST_{ac} = 0));$$

classify a to c's side if,

$$(PQR_{ac} - NQR_{ac}) > (PQR_{ab} - NQR_{ab}) \text{ or } ((FST_{ac} > FD_{ac}) \text{ and } (FST_{ab} = 0));$$

classify a to *other*, otherwise.

Article Classification

- News articles are classified by analyzing which side is importantly covered. There are 3 categories - one of the two sides, or "other".
- First considers from which side the article's quotes came
- Then considers the similarity of the rest of the article's text to the arguments of each side.

Evaluation – Disputant Partitioning

- 70% accuracy on average
- False positives were mostly the disputants who appear only a few times both in the article set and the news search results
- Recall is slightly lower than precision. Some disputants were omitted in the disputant extraction stage.

Evaluation - Article Classification

- Baselines:
  - **Similarity-based clustering (Sim)** - tf-idf of unigram and bigrams as features. K-means clustering algorithm.
  - **Quote-based classification (QbC)** - still does disputant extraction and disputant partition, but classification of news articles is done merely based on quote (if > 70% are from on side, or to the "other" category otherwise)

- F-measures: 0.68 (DrC), 0.59 (QbC), 0.48 (Sim)

Critique

- Clever insight of how people actually perceive contentious issues.
- Clever use of HITS algorithm to identify key opponents.
- Having only 2 key opponents could potentially be too simplistic. What if there are 3 groups?