Fighting Words: Using Natural Language Processing to Detect Partisan Polarization in Text

Masha Krupenkin

July 18, 2018

Explosion of partisan media sources

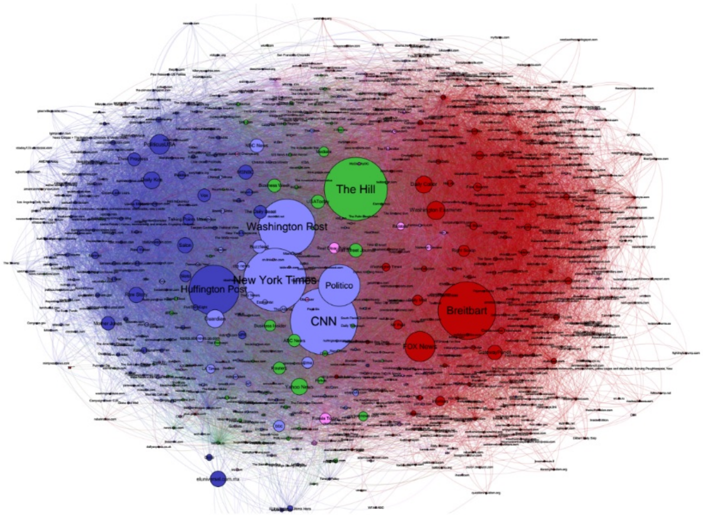

(Benkler et al 2017)

How much of this media is agreeing with one's supporters, and how much is cutting down the opposing party?

Americans more negative than ever...

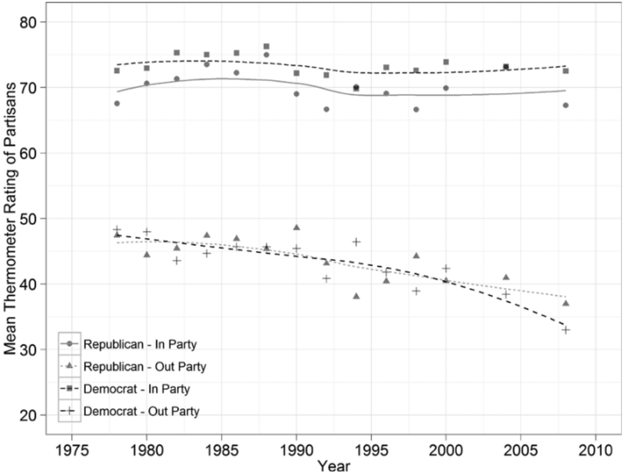

(Iyengar, Sood, and Lelkes 2012)

Measuring a text's sentiment towards the Democratic and Republican parties

Targeted sentiment analysis

- Determine a sentence's sentiment towards a specific figure or object
- Used in marketing literature to code product reviews (Yi et al 2003)
- Prior work in political science looking at non-targeted sentiment:
  - General sentiment of a sentence (Young and Soroka 2012)
  - Emotional context of text - "indignant disagreement" (Messing et al 2017)

Why use targeted sentiment analysis?

"Trump criticized the failing Obama admininistration"

- Sentence is negative towards Obama, but **positive** towards Trump.
- Simple dictionary method would code this sentence as negative towards both Trump and Obama (50 % correct)

Theory

Components of a targeted sentiment model

"Bill Clinton reduced the federal budget deficit."

1. Which sentiment words are relevant to the target (**Bill Clinton**)?
2. Which negations are relevant to the target?
3. Which words do the relevant negations negate?

1. Which sentiment words are relevant to the target?

Bill Clinton reduced the federal budget **deficit**.

Bill Clinton reduced the federal budget **deficit** and the Bush tax cuts dug us deeper into debt.

2. Which negations are relevant to the target?

Bill Clinton **reduced** the federal budget deficit.

Bill Clinton **reduced** the federal budget deficit, but the early 2000 recession decreased employment.

3. Which words do the relevant negations negate?

Bill Clinton **reduced** <u>the federal budget deficit</u>.

Bill Clinton **reduced** <u>the federal budget deficit</u>, but he lied under oath.

Final Result

Bill Clinton **reduced** the federal budget **deficit**.

Clinton Sentiment **+1**

Fitting the Model

Overview

- Supoervised machine learning approach
- Use a separate SuperLearner classifier for each of the three aforementioned steps
  1. DV: sentiment word relevant to target?
  2. DV: negation word relevant to target?
  3. DV: word relevant to negation?
- Combine predicted results to estimate sentiment of a sentence

Fitting the Model I

1. From an archive of 5000+ congressional candiate websites from 2002 - 2016, randomly sample 1000 sentences that contain a partisan keyword:
   - Trump, Clinton, Obama, McCain, Kerry, Bush, Gore, Republican or Democrat
2. Identify possible sentiment words + negations using a dictionary
   - Modified LSD and Topics Dictionary for sentiment (Young and Soroka 2012)
   - +/- Effect words dictionary for negations (Choi et al 2014)
3. Hand-code sentiment and negation words as relevant or irrelevant to the target to create training set.

Fitting the Model II

4. Run dependency parser from Stanford Core NLP on sentences to identify potential model features.
5. Run three separate SuperLearner classifiers to classify:
   - Which sentiment words are relevant to the target?
   - Which negations are relevant to target?
   - Which words do relevant negations negate?

Natural Language Processing (Stanford Core NLP)

Enhanced++ Dependencies:

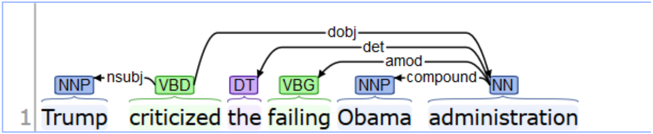

Open IE:

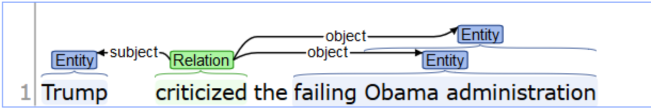

Technical Specifications I

1. DV is binary variable (whether or not a word is relevant to the target)
2. Potential features include dependency relationships between the word and the target, distance between the word and the target,sentiment of the word, whether the target is a subject or object in the sentence, etc
3. Feature selection/screening - drop any variables that fail Pearson correlation test at $p<0.1$

Technical Specifications II

4. Used a SuperLearner with the following learners to fit the model:
   - Elastic Net, Extreme Gradient Boosting, Support Vector Machine, Bayesian Additive Regression Trees, and Random Forest
5. 10-fold cross validation to generate out-of-sample predictions for the full dataset
6. Used classification threshold that maximized F1 score

Results

Superlearner Prediction Accuracy

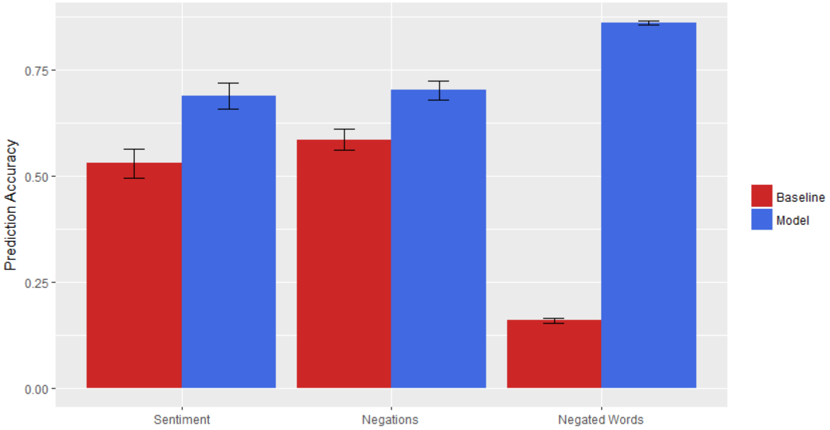

Sentiment by Sentence

- Aggregate all relevant sentiment for each sentence and flip the sentiment for the negated words
- Dictionary MSE: 4.40
- Dependency Model MSE: 1.36
- Dependency Model and dictionary method disagreed 49% of the time. In sentences that included a negated term, the models disagreed 83% of the time!
- Dependency model got the exactly correct sentiment 60 % of the time. Dictionary method got the correct sentiment 46% of the time

Sentence Classification Accuracy (Neg, Neu, Pos)

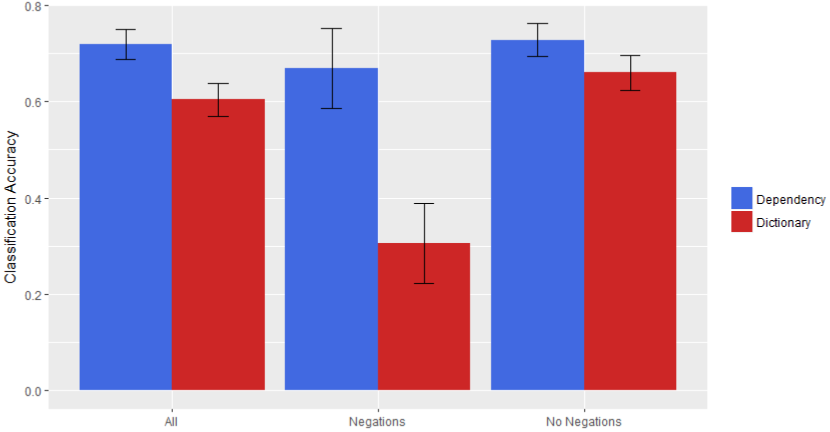

Conclusion and Next Steps

- Model reduced MSE per sentence
- Significant increases in sentence classification accuracy
- Next Steps
  - Qualitative validation
  - Interactions
  - Count coreferences as target ("he", "her", etc.)
  - Expand to adjacent sentences using coreferences
- Application - has partisan negativity increased in congressional campaign websites since 2002?
  - Partisan negativity ("Republicans are out of touch!") vs Candidate negativity ("Sen Smith is out of touch!")