

Persuasive Influence Detection: the Role of Argument Sequencing

Christopher Hidey Kathleen McKeown

Columbia University

February 1, 2018



Research questions:

- Can we leverage sequential aspects of social media posts for persuasiveness prediction?
- How do humans perform at identifying personalized persuasion?

- 1 Introduction
- 2 Data
- 3 Methods
- 4 Results
- 5 Error Analysis

- Tan et al. (2016); Wei et al. (2016) - predicting persuasiveness/ranking comments in Change my View
- Rosenthal and McKeown (2017) - identifying influencers in social media
- Stab and Gurevych (2016) - objectively ranking convincingness of arguments

Data

OP CMV: Patriotism is the belief that being born on one side of a line makes you better...



I would define patriotism quite simply as supporting one's country, but not *necessarily* disparaging others...

R Someone who assists another country that is in worse shape instead of assisting their own can still be a patriot, but also recognize significant need in other nations and decide to assist them as well



OP This is true, but, I think, supporting the common good is also more important than supporting your country



R Yes, but the two are often one the same, especially when you live in a country as large as the U.S. most acts which serve the common good generally support your country.



OP I see. They're not mutually exclusive so I think I had the wrong definition: Δ

Overall: 5296 positive, 16685 negative

Research questions:

- Can we leverage sequential aspects of social media posts for persuasiveness prediction?
- How do humans perform at identifying personalized persuasion?

Deep Learning for predicting persuasion:

$$y = \sigma (MLP(\mathbf{h}) + \beta^T \phi)$$

- \mathbf{h} - learned document representation (our contribution)
- ϕ - additional document features (Tan et al., 2016)

Interplay - intersection of words between OP and response

Patriotism	is	the	belief	that	...
I	would	define	patriotism	as	...

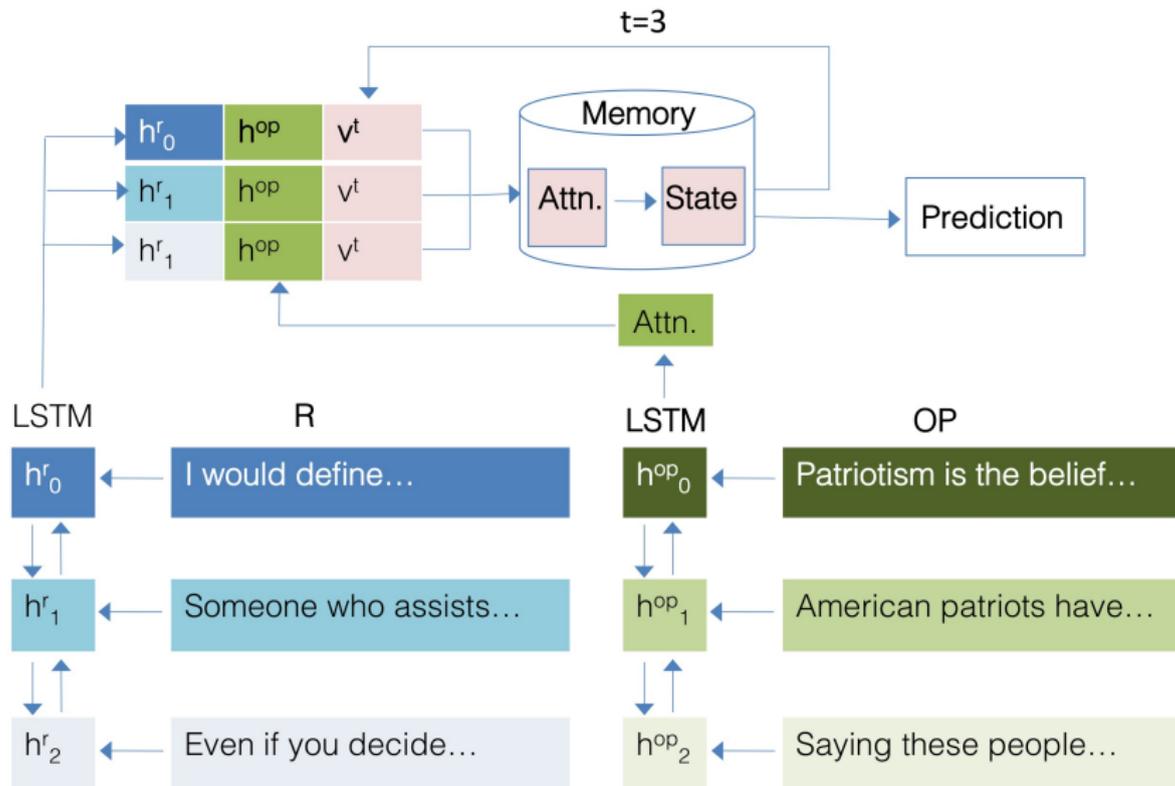
- MLP, β - learned weights

Methods

h - the **initial** document representation



\mathbf{h} - the **complete** document representation



Results

Influence - predict whether a post is persuasive or not

	<i>Model</i>	<i>Acc.</i>	<i>AUC</i>	<i>F1</i>
<i>R</i>	Bag of Words	61.9	72.8	50.3
	MLP	68.8	73.2	50.3
	LSTM	75.1	75.5	53.0
<i>OP</i>	Interplay	72.7	76.7	54.6
	LSTM+Memory	74.3	77.3	55.4
	LSTM+Memory+Interplay	81.0	82.1	60.7

Table 1: Results of Influence Prediction Task

Interplay - intersection of words between original post and response

MLP - sentence vectors from word, frame, discourse embeddings

LSTM - over sentence vectors from word, frame, discourse embeddings

Research questions:

- Can we leverage sequential aspects of social media posts for persuasiveness prediction?
- **How do humans perform at identifying personalized persuasion?**

Crowdsourcing Experiment

- 200 original posts paired with positive and negative arguments
- 3 questions: Would the original poster find the first/second argument convincing? How would the OP rank the arguments?
- Required to provide a justification of 20 words for each of the 3 questions

Yes, because I feel that the argument A coincides more with the original post, referring to the motivation that some inhabitants have to vote

Model	Pairwise	Influence
Annotators	54.84	57.14
Model	71.99	63.00

Table 2: Human Performance

Error Analysis

		Human		Model	
Category	%	P	I	P	I
Government	29	76.3	55.1	64.4	58.5
Sociology	23	71.7	53.3	80.4	68.5
Morality	11	72.7	63.6	77.3	68.2
Economics	9	50.0	50.0	72.2	58.3
Politics	8	62.5	56.3	68.8	62.5
Science	6	66.6	66.6	66.6	62.5
Culture	5.5	54.5	45.5	54.5	63.6

Table 3: Error Analysis on Categorized Data

Humans do better at tasks requiring world knowledge (government)

Error Analysis

OP Weinberg was wrong when he said that “for good people to do evil things, that takes religion”

R I think that someone isn't a good person if they have an ideology I disagree with
I think nationalists are bad, fascists are bad and so on
(**Gold:** Negative **Predicted:** Negative)

OP Countries should have a “no confidence” vote in elections if they want to increase turnout, while achieving a better understanding of the public's perception of the political climate

R The US state of Nevada has had a choice called “none of these candidates” since 1975
(**Gold:** Positive **Predicted:** Negative)

Conclusions

- Modeling argument sequencing and context helps
- Humans are poor judges of personalized persuasiveness
- Future models would benefit from world knowledge and reasoning

Questions

?

Thanks to DARPA-DEFT and all the annotators!