“We Demand Justice!”: Towards Social Context Grounding of Political Texts

Rajkumar Pujari Chengfei Wu Dan Goldwasser

Purdue University, West Lafayette, USA

{rpujari,wu1491,dgoldwas}@purdue.edu

PURDUE UNIVERSITY

Department of Computer Science

Abstract

- Understanding political discourse on social media fully by reading only the text is difficult. However, knowledge of the social context information makes it easier.
- We characterize the social context required to understand such ambiguous discourse fully.
- We propose two datasets that require an understanding of social context and benchmark them using large pre-trained language models and several novel structured models.
- We show that structured models, explicitly modeling social context, outperform larger models on both tasks, but still lag significantly behind human performance.

Motivation

- Humans familiar with a politician’s stances and, possessing knowledge about the event, can easily understand the *intended meaning*.
- Our main question is - ***Can an NLP model find the right meaning?***

Dataset Statistics

Target Task Data Statistics		Vague Text Data Statistics	
Unique Tweets	865	Unique Vague Texts	93
Positive Targets	1512	Positive Examples	739
Negative Targets	1085	Negative Examples	2217
Neutral Targets	784	Total Examples	2956
Non-Targets	2509	Hard Test Examples	180
Total Data Examples	5891	Number of Events	9
Number of Events	3		

- We use Amazon Mechanical Turk (AMT), LLMs, and contextual expert annotations to obtain the data.

Social Commonsense

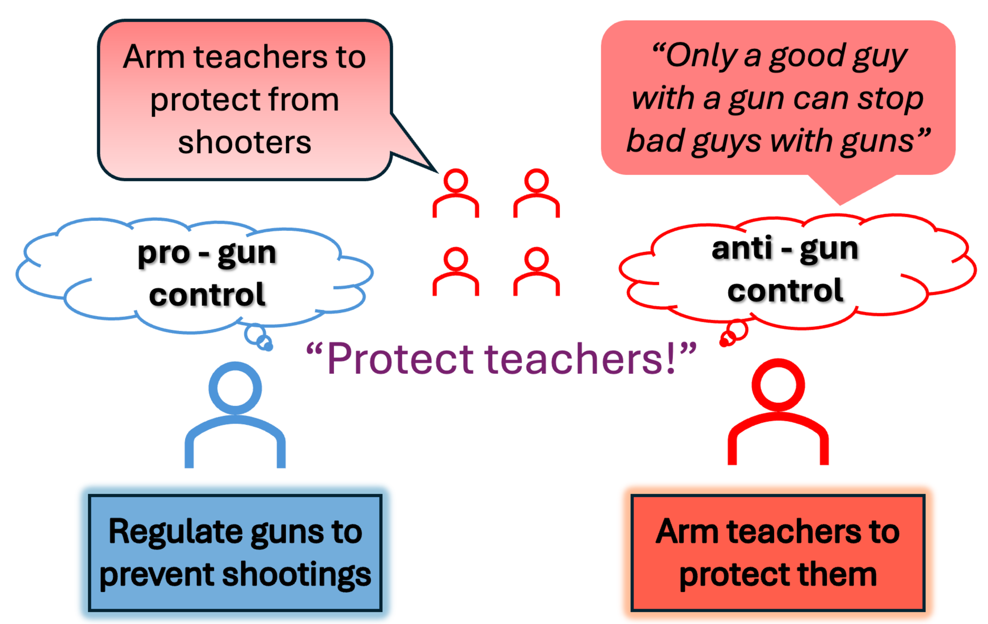

Entity-Sentiment Perspective of an Event

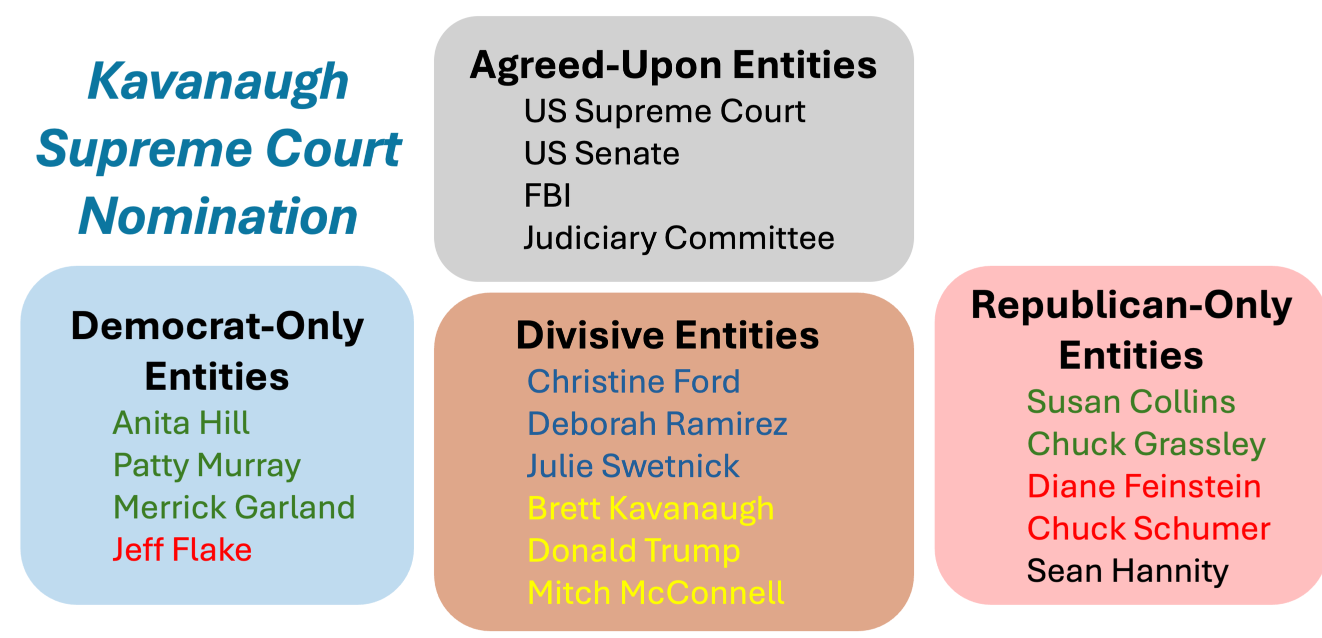

Proposed Datasets

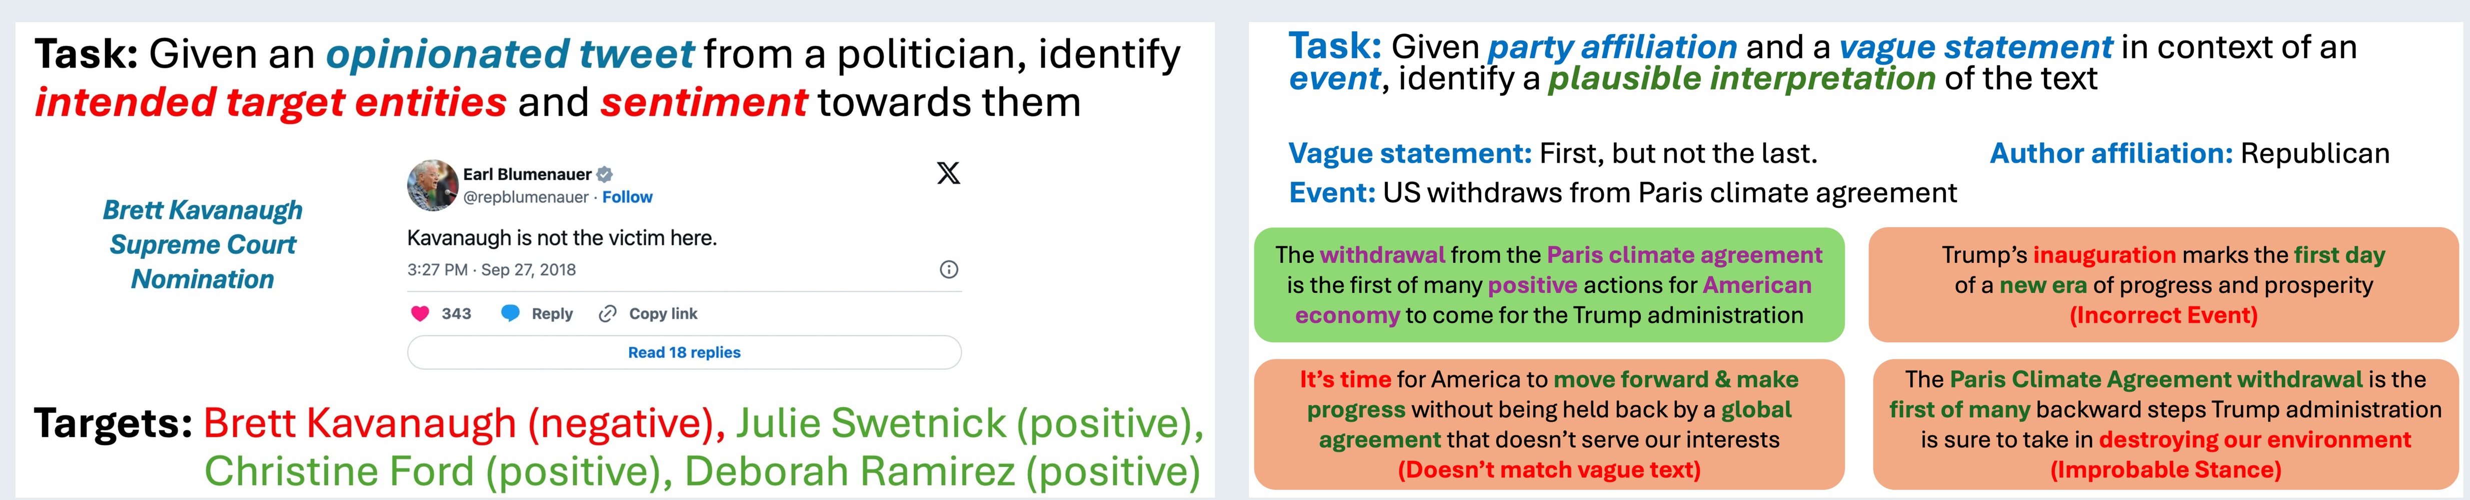

(a) Target-Entity & Sentiment Task (b) Vague Text Disambiguation Task

Target Entity and Sentiment Results

	Model	Target-Entity Task		Target-Sentiment Task	
		Macro-F1	Acc	Macro-F1	Acc
No-Context Baselines	BERT-large (B-l)	68.83	70.56	58.95	58.37
	RoBERTa-base (R-b)	65.14	66.40	61.36	60.65
Text-Context Baselines	B-l + twitter-bio	69.34	71.66	60.13	59.86
	R-b + twitter-bio	64.79	66.30	59.94	59.46
	B-l + wiki	60.33	61.05	53.9	53.32
	R-b + wiki	68.62	70.27	58.07	58.28
LLMs	GPT-3 zero-shot	69.77	73.78	54.18	56.80
	GPT-3 few-shot	66.45	67.03	55.00	57.15
Contextualized Embeddings	B-l + PAR-embs	60.25	60.56	55.89	55.80
	R-b + PAR-embs	67.67	69.18	55.51	55.40
	B-l + DCF-embs	68.32	69.97	61.22	60.75
	R-b + DCF-embs	**73.56**	**75.82**	62.90	63.03
Discourse Contextualized Models	B-l + DCF	71.17	72.94	**65.34**	**65.31**
	R-b + DCF	70.39	72.15	63.37	63.23

Vague Text Results

	Model	Vague Text Disambiguation Task	
		Macro-F1	Acc
No-Context Baselines	BERT-large (B-l)	50.28	53.75
	BERT-base (B-b)	54.53	56.08
Text-Context Baselines	B-l + wiki	66.87	76.03
	B-b + wiki	64.36	75.69
LLMs	GPT-3 zero-shot	62.58	63.50
	GPT-3 few-shot	61.86	62.04
Contextualized Embeddings	B-l + PAR-embs	65.53	73.79
	B-b + PAR-embs	65.49	72.59
	B-l + DCF-embs	67.55	73.10
	B-b + DCF-embs	**71.71**	**78.45**
Discourse Contextualized Models	B-l + DCF	69.94	76.55
	B-b + DCF	70.06	76.55

Human Performance Benchmarking

- Humans familiar with context find the Vague Text Disambiguation task fairly easy. Shows the robustness of the dataset.
- Current NLP models lag significantly behind humans on the *Vague Text Disambiguation* task

Model	Accuracy
BERT-base + wiki	54.89
BERT-base + DCF-embs	63.38
BERT-base + DCF	64.79
Humans	**94.85**

Feature Importance Ablation

- Ablation of different examples classes on the Vague Text task. Shows that no single feature is sufficient for solving the task. A joint understanding of contextual elements is necessary to solve the task.

Data Split	Unseen Party	Unseen Event	Flip Tweet	Flip Event	Flip Party
	Ma-F1	Ma-F1	Acc	Acc	Acc
Random	44.70	29.69	75	75	75
BERT-base +wiki	57.58	39.76	**88.14**	89.77	**87.77**
BERT-base +DCF-embs	61.79	**47.88**	86.10	**93.18**	84.57
BERT-base +DCF	**65.18**	45.65	82.03	89.77	84.04

Conclusion

- We conceptualize and operationalize two holistic social context grounding tasks in English on the US political domain.
- We evaluate existing state-of-the-art models and humans on these tasks and present interesting observations.
- Future work directions include building explicit and interpretable models for *Social Context Grounding*, expanding proposed datasets, designing diverse tasks, and other flavors of social context such as cross-cultural understanding, emergency response, etc.

Resources

https://github.com/pujari-rajkumar/language-in-context