Who Supports Bernie? Analyzing identity and ideological variations of Bernie Supporters on Twitter

Alexander R. Caceres-Wright, Grant Bunn, Stef M. Shuster, Kenneth Joseph

Introduction

Our work introduces:

1) The concept of *explicit* stance detection
2) A novel data consisting of over 1000 labelled tweets to be used as a benchmark in this task

And ongoing work is focused on:

1) Using Large Language Models such as Meta's Llama-2 to accomplish this task
2) Creating a model to identify tweets related to, but not containing, keyword(s)
3) Identifying conspiracy theories in tweets

Percent of gold label tweets annotated by Llama text generation

Candidate	Percent Annotated	Candidate	Percent Annotated
Bennet	65.97	Patrick	52.92
Biden	77.08	Steyer	70.71
Bloomberg	76.07	Sanders	68.26
Buttigieg	75.09	Trump	53.82
Delaney	83.08	Warren	68.25
Gabbard	72.73	Yang	75.08
Klobuchar	73.84		

Krippendorff's Alpha on Human Annotations

Candidate	Mentions Agreement	Full Agreement	Agreement w/o Neutral
Bennet	0.80	0.58	0.76
Biden	-0.01	0.54	0.61
Buttigieg	0.61	0.55	0.66
Klobuchar	1.00	0.45	0.63
Patrick	0.96	0.62	0.90
Sanders	0.92	0.42	0.59
Steyer	1.00	0.57	0.56
Trump	0.55	0.48	0.64
Yang	0.56	0.50	0.66
Gabbard	1.00	0.71	0.80
Bloomberg	1.00	0.52	0.58
Delaney	1.00	0.48	0.70
Warren	0.79	0.51	0.86

Data Analysis

We made use of several different models and approaches for our task:

1) Meta's Llama-2 Model:
   a) Generate free response text
   b) A pipeline
2) Deberta Polistance:
   a) A pipeline

Results

We found the following:

1) Off-the-shelf Large Language Models, such as Meta's Llama-2 struggle with identification in text
2) Fine tuned models tend to perform better, especially when making use of a pipeline, which allows for specific formats for the outputs
3) All three models we tested struggled with the "neutral" label the most

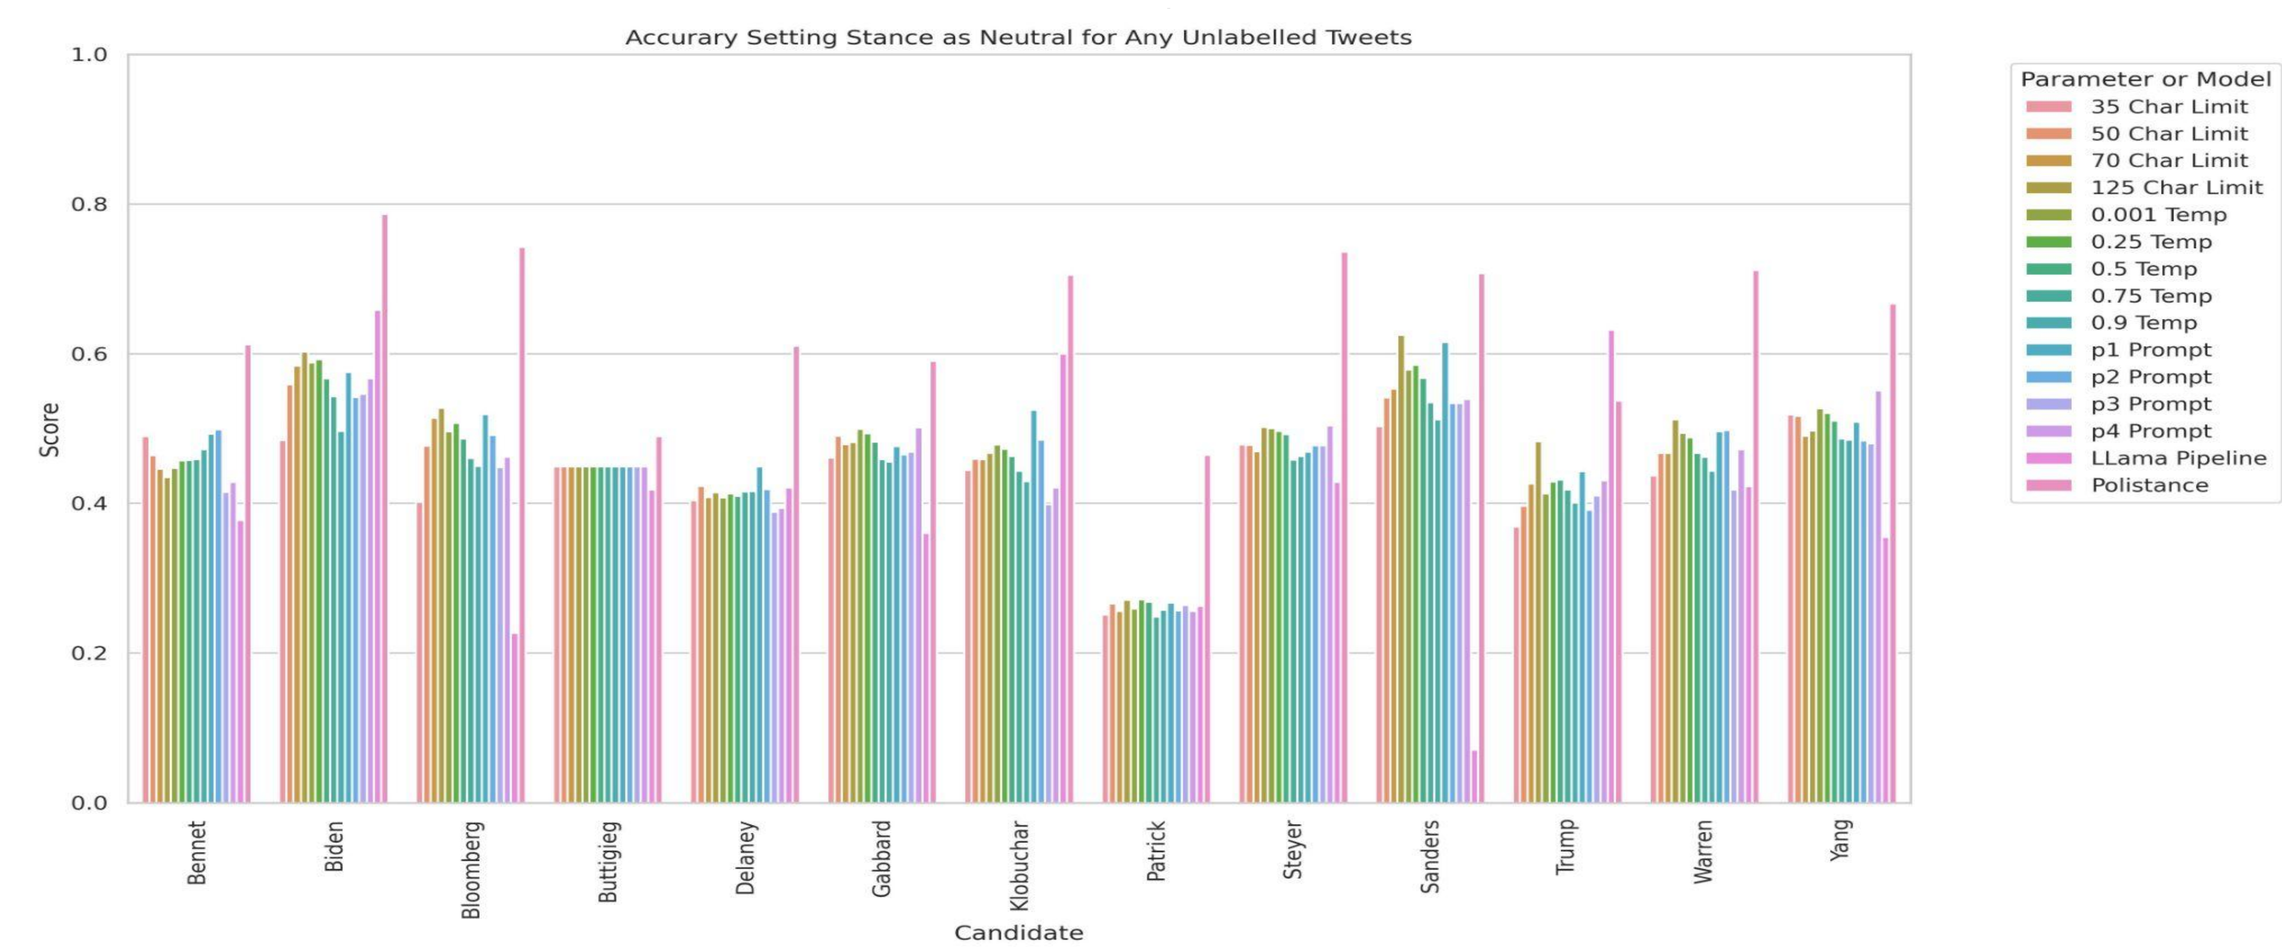

Accuracy setting neutral for unlabelled tweets (top) and removing any unlabelled tweets from the gold label set (bottom)

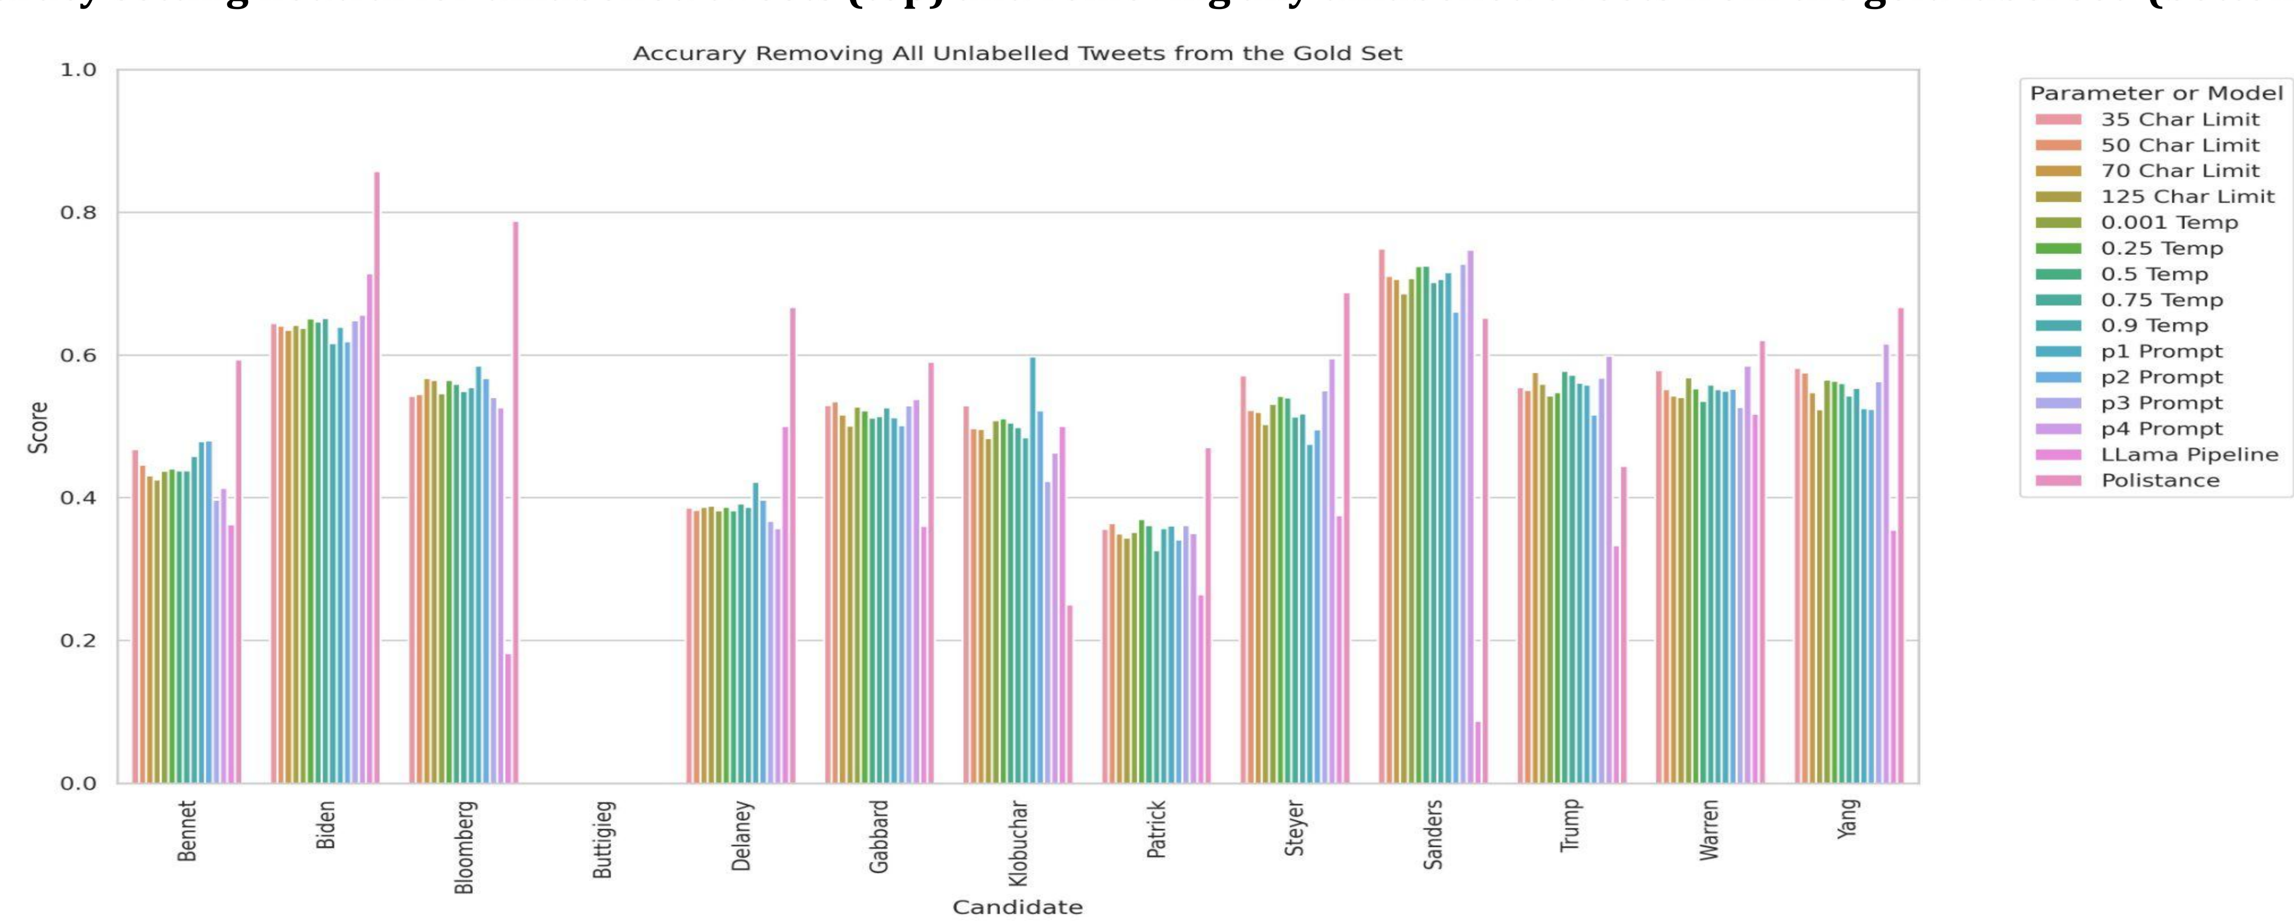

Past Work

Past work focused on:

1) Creating a mixed method approach to identify networks of accounts which interact with one another, or specific tweets
2) Present a computational method to separate groups of users who share support for a given candidate
3) Applied these methods to a set of account which at some point signaled support for Bernie Sanders, and demonstrated several distinct ideologies within these accounts

Percent of tweets sent by users in each of 5 distinct groups of Bernie supporters

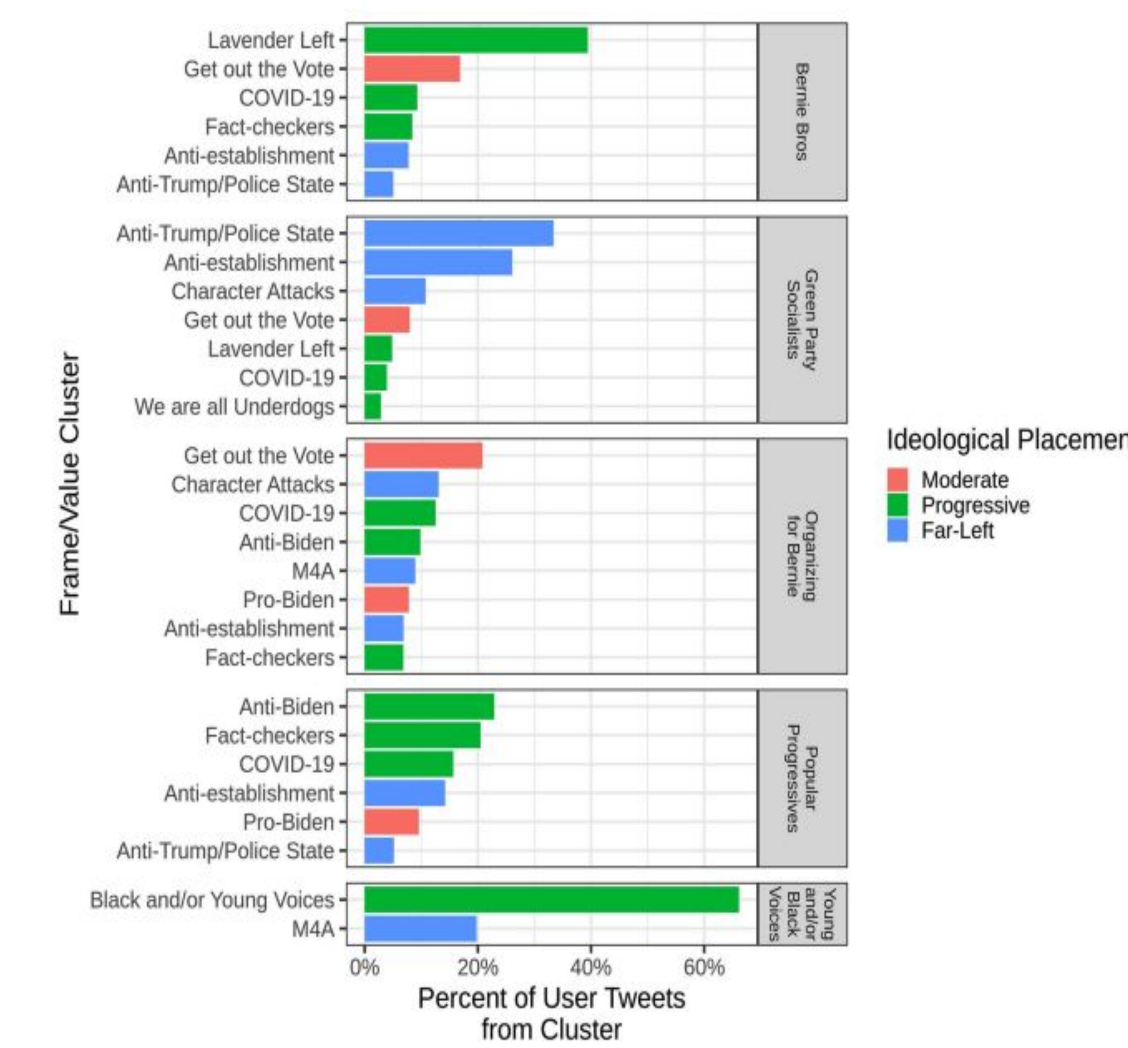

Fig 5. The percentage of tweets (x-axis) sent by ordinary users in each of these five social groups (distinct subplots) that were assigned to each of the fourteen framed values clusters (y-axis) in the who retweets what network. We show frames that make up at least 90% of the tweets in the given cluster of users.

References

1. https://huggingface.co/meta-llama/Llama-2-13b-chat-hf
2. https://huggingface.co/meta-llama/Llama-2-13b-hf
3. https://huggingface.co/mlburnham/deberta-v3-large-polistance-affect-v1.1
4. Can large language models transform computational social science? Ziems et al.,2024
5. Recognizing stances in ideological on-line debates. Somasundran and Wiebe, 2010
6. Who supports Bernie? Analyzing identity and ideological variation on Twitter during the 2020 democratic primaries. Shuster et al., 2024

University at Buffalo The State University of New York

Department of Computer Science and Enginnering
School of Engineering and Applied Science
buffalo.edu