

Identifying Fine-grained Forms of Populism in Political Discourse: A Case Study on Donald Trump's Presidential Campaigns



Ilias Chalkidis ^α Stephanie Brandl ^α Paris Aslanidis ^β

(α) University of Copenhagen Denmark (β) National Centre for Social Research (EKKE) Greece



1. Introduction


Large Language Models (LLMs) have demonstrated remarkable capabilities across a wide range of instruction-following tasks, *yet their grasp of nuanced social science concepts remains underexplored.*

We examine whether LLMs can identify and classify fine-grained forms of populism, a complex and contested concept in both academic and media debates.

Working Definition - A discursive approach
(Laclau, 2005; Stavrakakis, 2024; Aslanidis, 2024)

Populism leverages the principle of **popular sovereignty** to emphasize the idea of the common “people” and its opposition to **various elite groups** that game the system to serve their own ends. — Two core elements:

- (i) **people-centrism**: positive sentiment towards “the people”
- (ii) **anti-elitism**: negative sentiment towards “elites”



2. Datasets / Examples

We release 3 new datasets:

- (a) **TRUMP 2016**: 70 annotated presidential campaign speeches (~15K sentences) by Donald Trump during the Republican Party primaries and the 2016 presidential campaign.
- (b) **TRUMP CHRONOS**: 713 non-annotated speeches (~656k sentences) by Donald Trump, ranging over a decade from June 2015 to April 2025.
- (c) **EU-OOD**: 6 speeches by 5 selected European (🇩🇪 🇫🇷 🇪🇸 🇬🇧 🇮🇹) leaders who recently ran for office and are widely considered to have relied on populist discourse during their campaigns.

And if you look, Ford is building massive plants there instead of Michigan.	○
She's totally bought and paid for by Wall Street, the special interests, the lobbyists , 100 percent.	●
Once again, we will have a government of, by, and for the people .	●
Most importantly, we are going to take power back from the political class in Washington and return that power to you, the American people .	● ●

○ Neutral ● People-centrism ● Anti-elitism

Note: Populist invocations appear relatively sparse in political discourse, even in the case of a characteristically populist politician (93% neutral, 6% anti-elitism, 3% people-centrism)

4. Experiments & Results

(a) Main Results

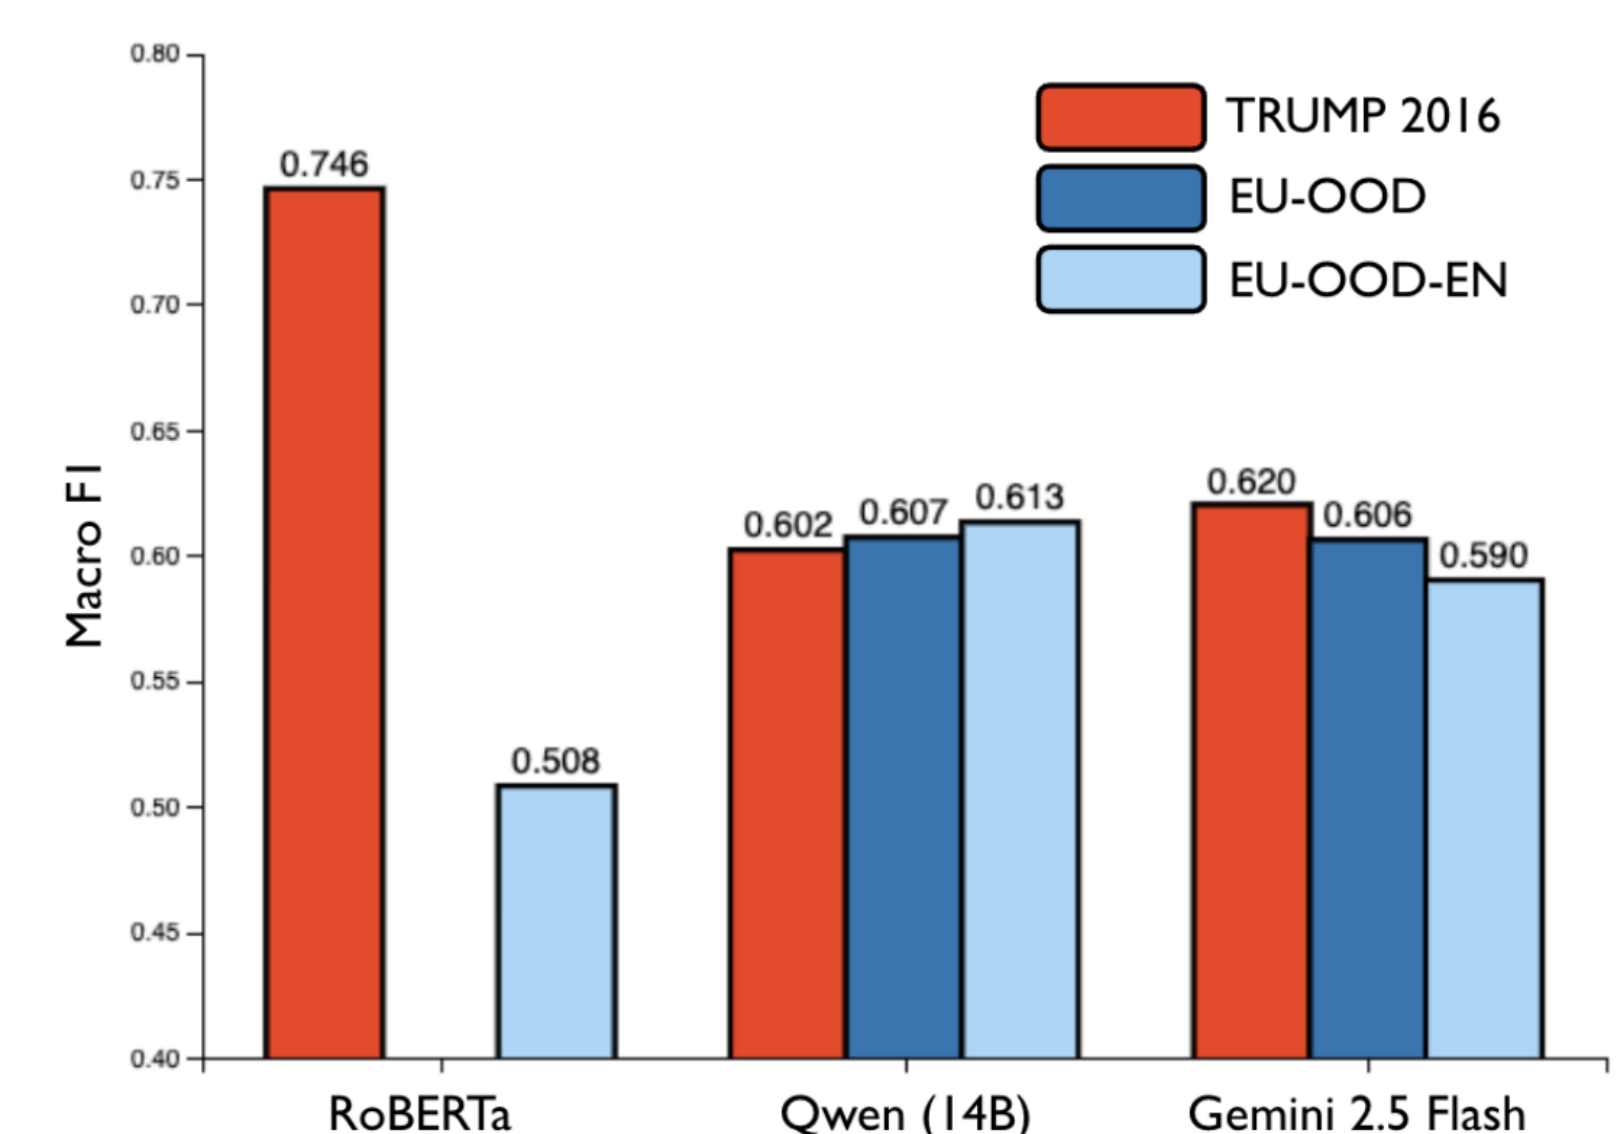
Model Name	Size	Neutral	Anti-Elitism	People-centrism	Macro F1
TFIDF-SVM	30K	.974	.531	.385	.630
ROBERTa	355M	.976	.661	.602	.746
Qwen 3	8B	.864	.368	.286	.506
	14B	.967	.621	.553	.714
GPT 4.1	14B	.916	.491	.340	.582
	32B	.972	.674	.575	.740
Gemini 2.5 Flash	n/a	.875	.458	.292	.542
GPT 4.1	n/a	.850	.461	.201	.503
Gemini 2.5 Flash	n/a	.903	.513	.306	.574

● Fine-tuned ● Proprietary (as a Service)

(c) Meta Analysis

Annotator	Expert 🧑	Students 🧑
Students	0.737	-
ROBERTa	0.800	0.746
Qwen 3 (8B)	0.730	0.714
Qwen 3 (14B)	<u>0.768</u>	0.740
Gemini 2.5 Flash	0.633	0.628*

(d) Out-Of-Distribution



(b) Prompt Tuning Results

Qwen 3 (14B)		Gemini 2.5 Flash	
Mode	Macro F1	K	Mode
Baseline	<u>0.582</u>	0	Base
Context-Aware	0.472		+ Reasoning
Distribution Aware	0.602		+ Dist-Aware
K-Shot (K=8)	0.628	32	Base
K-Shot (K=16)	0.615		+ Reasoning
K-Shot (K=32)	0.583		+ Dist-Aware
K-Shot (K=64)	0.597	128	Base
			+ Reasoning
			+ Dist-Aware

(e) Pol-Sci RQs and Analysis

- RQ1:** Does Donald Trump's populism substantially fluctuate over time?
H1: Donald Trump modulates the intensity of his populist discourse across electoral campaigns to serve strategic objectives.
 → Clear peak of populism volume in 2016 presidential elections.
- RQ2:** Does Trump's populist discourse intensity vary between swing states and non-swing states?
H0: There is no difference in populist discourse intensity between swing and non-swing states.
 → No clear geographical targeting; Contrary signs of more “moderate” tone in 2024 presidential elections in swing states.
- RQ3:** How does Donald Trump's populism spread out within a speech?
H1: Donald Trump bookends his speeches with populist content to maximize efficacy.
 → Book-ending (planning) across speech; Clear backloading of people-centrism.

5. Takeaways

- A fine-tuned RoBERTa classifier vastly outperforms both open-weight and top-notch proprietary LLMs in zero-shot fine-grained populism identification.
- Although LLMs struggle with the ambiguity and contextualization of populism detection, the best-performing models perform on a comparable level to human (student) annotators.

🧑 Resources



📄 Paper

