Predicting the “Future” – Can old LLMs predict current events?

Teamproject

UNIVERSITY OF MANNHEIM

Maximilian Kreutner

04.09.2024

Background

- **Large Language Models** (LLMs) e.g., GPT can simulate human samples
- When provided with **socioeconomic attributes**, LLMs can model **behavior patterns**
- Example: Using prompts (e.g., "*Act as a male student with specific background*"), LLMs have been shown to approximate how U.S. citizens voted in the 2016 election

Out of One, Many:
Using Language Models to Simulate Human Samples

Lisa P. Argyle[1], Ethan C. Busby[1], Nancy Fulda[2], Joshua Gubler[1], Christopher Rytting[2], and David Wingate[2]

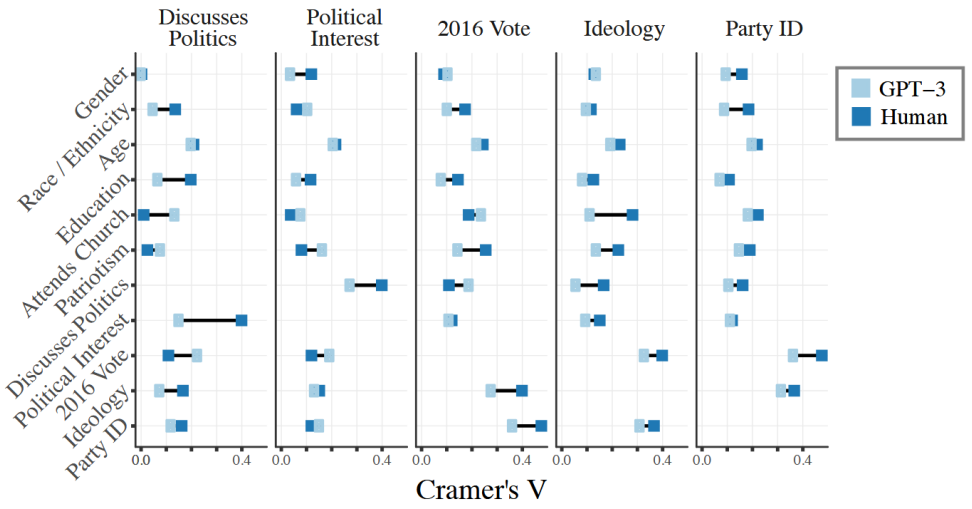

Figure 4: Cramer's V Correlations in ANES vs. GPT-3 Data

The Project

UNIVERSITY OF MANNHEIM

- To evaluate these samples, we need ground truth data
  - **Problem:** Training data of LLMs contains this data
  - **Solution:** Use old LLMs that have a training cutoff before predicted event
- For example: Llama 3 has a pretraining cutoff in December 2023
  - **Not** included in training data: US Election, German Election, etc.
  - How do LLMs predict the "*Future*" based on their old data
  - And how does it compare to ours? Do LLMs predict Trump as president?
- If the prediction is wrong, how do we fix it?
  - Simple add more information to the prompt?
  - Use Retrieval Augmented Generation to add news headlines of the specific month?
  - Other techniques?

General Information

UNIVERSITY OF MANNHEIM

- Responsible person: Maximilian Kreutner
- Email: maximilian.kreutner@uni-mannheim.de
- Language: English
- Duration: 6 Months
- Min/max number of participants: 3-5
- Prerequisites: strong python skills
- Applicable to MMDS: Yes