

The Expressed Public: Political Opinion on Social Media in the Age of Large Language Models

Abstract:

Social media is an appealing source of data for studying public opinion: it is large, public, and immediate. But how good is it, really, as a measure of what the public thinks? This dissertation takes that question seriously and arrives at a precise answer. Social media does not capture the opinions people hold so much as the statements people choose to make public, and the two are not necessarily the same thing. What ends up in the data depends not only on what people think about an issue, but also on whether, how, and why they choose to say it.

The dissertation develops tools that use large language models to infer who is posting and what stance they take on an issue, with a stated level of confidence rather than just a guess. It then uses those tools to test the quality of the data: whether the people who post can stand in for members of the public with whom they share an opinion, and how the answer changes across platforms and events. Participation on social media platforms turns out to be selective and event-driven, so that platforms from Bluesky to Truth Social behave as distinct publics. The work reframes social media not as a flawed substitute for surveys, but as a measure of political expression with its own strengths and its own conditions for use.

Bio:

Mao Li is a Ph.D. candidate at the University of Michigan, advised by Dr. Fred Conrad. His research focuses on computational social science, survey methodology, large language models, and public opinion. His work examines how emerging AI systems and social media data can be used to study human attitudes, communication, and survey response processes.