

Facebook plays a central role in political communication; millions of people view political stories in their newsfeeds and engage in lively political discussions through posts and comments daily. These experiences can impact both individuals' beliefs as well as polarization and prejudice on a global scale. I develop computational methods to analyze language in political discussions, and highlight the broader societal implications of these linguistic choices. My work uncovers how content producers *frame* political issues and, in doing so, how they can perpetuate biases against marginalized groups. By synthesizing political communication, psychology, linguistics, and natural language processing, my research is deeply interdisciplinary and pushes the frontiers of both computational and social science.

Framing, highlighting aspects of an issue to promote a particular perspective, can impact public opinion, collective action, and policy. However, studies of framing are mostly limited to small-scale work on mass media, where issues are framed by political and journalistic elites. Nowadays, ordinary people on social media play a major role in generating and propagating frames, so it is crucial to develop large-scale methods to understand framing among the public. My research involves building scalable and theoretically-grounded NLP technology to study framing across a wide range of issues and contexts. I plan to investigate how people's social identities, ego-networks, and news consumption influence framing. Finally, I will combine NLP, causal inference and information diffusion to understand how frames evolve and spread within social networks and how online framing impacts offline outcomes. Support from the Facebook Fellowship would enable me to dedicate time and resources to realize these goals, and I relish the opportunity to learn from Facebook's computational social scientists.

Two projects demonstrate how I apply computational methods to understand framing in sociopolitical discourses and conversely incorporate social science theory into computational models. First, I drew from political communication to model the framing of immigration on Twitter. I created a dataset of tweets labeled for multiple framing typologies, upon which I developed supervised frame detection models. By analyzing 2.6M tweets, I compared the framing of immigration across regions and political ideologies, and leveraged Twitter's engagement behaviors to demonstrate how framing affects audience responses. In earlier work, I took a person-centric approach to framing by introducing a set of neural embedding-based techniques for computationally analyzing dehumanizing language, which I applied in a case study of LGBTQ representation in the media over three decades.

My research has far-reaching implications for Facebook, where users' exposure to political information can have immense consequences. For example, we could leverage a better understanding of political discourse to improve users' experiences by enabling them to discover content with a diversity of perspectives, which can mitigate cross-partisan animosity and lay foundations for more cooperative and open-minded conversations. Furthermore, the ability to detect subtle prejudices within frames can make Facebook safer for members of marginalized groups. Ultimately, my research can make online spaces more equitable and welcoming to everybody, thus embodying Facebook's mission to "give people the power to build community and bring the world closer together".