Algorithms and Ideology: Measuring Political Bias in ChatGPT, Claude, Gemini, and DeepSeek

Thompson, Ellis (School: Stuyvesant High School)

Large language models (LLMs) like ChatGPT, Claude, Gemini, and DeepSeek increasingly influence how people access and interpret political information, raising concerns that they may subtly shape public opinion at scale. This study examines whether these models exhibit measurable ideological bias, a tendency to respond more favorably to certain ideologies than others, in their outputs. Using key terms drawn from term frequency analysis of the Democratic and Republican party platforms, each model was prompted with three distinct prompts for each of 20 terms, one framed neutrally, one with liberal themes, and one with conservative themes, generating 240 unique responses. All four models then scored each response on a 1–10 positivity scale; these scores were averaged and compared using Cohen's d to quantify the direction and magnitude of bias. Across all LLMs, responses to prompts containing terms selected from the Democratic platforms received higher average positivity scores than those to prompts with terms selected from the Republican platforms, with ChatGPT and DeepSeek showing the largest effects (d ˜ 0.6–0.7). Bias did persist under neutral prompt conditions, which suggests that ideological tilt is intrinsic to model behavior rather than simply a function of framing. While framing slightly increased positivity, its effect was small for most models, with the exception of DeepSeek. Results were consistent across models, suggesting broad applicability and indicating that political bias is persistent in current LLMs.

Awards Won: