

AI Politician: Participative democracy powered by chatbots

This project explores the potential for AI chatbots to enhance participative democracy by allowing a politician to engage with a large number of constituents in personalized conversations at scale. By creating a chatbot that emulates a specific politician and is knowledgeable about a key policy issue, we aim to demonstrate how AI could be used to promote civic engagement and democratic participation.

Threat Scenario

Our threat scenario envisions a world where voter disengagement and lack of participation weaken democratic institutions. Many citizens feel disconnected from their political representatives and the democratic process itself. Trust in democratic government is low. All this against a background where malicious actors will employ AI technologies to further the aims of an anti-democratic agenda, for example, by undermining the trust in news media, surveilling or intimidating constituents, or improving autocratic governance through the use of expert systems.

AI chatbots that can engage with voters directly on important issues, educate them, answer questions, and gather their input could be a powerful tool to combat political apathy and alienation. By making interactions with politicians' offices more personalized, informative, and convenient, AI chatbots could draw more citizens into the political process, and ultimately make democracies more competitive with undemocratic governance styles.

Demonstration

We built an enhanced chatbot that roleplays as California Senator Scott Wiener, and engages with the user on the topic of Senate Bill 1047, which is itself about the regulation of artificial intelligence. The bot can engage in open-ended dialogue, build long-term memory of conversations, and reason about policy topics while drawing upon a collection of relevant documents. We also demonstrate logging conversations across multiple users and analyzing aggregate sentiments with LLM assistance (appendix 4), which can also be used maliciously to e.g. analyze the needs and beliefs of constituents of a competitor.

Looking Ahead

As AI language models become more advanced, politician chatbots could become increasingly effective at engaging citizens in the democratic process. Some potential future work includes:

- Additional modalities of interaction including speech and video generation could make the experience more engaging and accessible.
- Chatbot agents improve access to government services by guiding citizens to relevant information and walking them through processes.




Figure 1: A screenshot of our demonstration: A user discusses SB 1047 with the Senator Wiener bot.

Appendix

A1: Technical Implementation Details

This bot is currently hosted on a free tier Hugging Face space. The chat data is collected using the CommitScheduler module of the huggingface_hub package. The application itself was written in Python utilizing the Haystack framework implementation. The basis of this is a RAG implementation that uses the OpenAI modules of the Haystack package and runs gpt-4-turbo as the generating model. We are vector-embedding a number of documents (see A1.1), the content of which can be queried by the model via a function call. The function call then provides the model with the necessary content of a document. To access the application, the model functionality is embedded using the Gradio package in Python. This package provides the necessary features to display a chat interface and change parameters.

A1.1: Building the dataset for Retrieval Augmented Generation

These are documents we gathered for the bot to reference during dialogue:

- [Senate Bill 1047](#)
- [Senator Wiener's press release on SB 1047](#)
- [CA Executive Order N-12-23](#)
- [White House Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence](#)
- [New York Times piece "'The Godfather of A.I.' Leaves Google and Warns of Danger Ahead"](#)
- [Zvi Mowshowitz's analysis of SB 1047 on Don't Worry About the Vase](#)

A1.2: Crafting the system prompt for topicality and style

The system prompt was designed to facilitate productive conversation on the topic of SB 1047. It also contains a style reference collected from Senator Wiener's remarks in [a verbal interview with Ezra Klein](#). The full text of our system prompt can be [found here](#) for reference.

A2 Limitations

As with all LLMs, there are limits on the knowledge the bot possesses. In this hackathon, we focused on providing documents for retrieval augmented generation (RAG) that were relevant to the topic of SB 1047. Some shortcomings in the bots knowledge were revealed, however. For example, it was unaware of Senator Wiener's Chinese name when asked.

Also LLMs are not perfect and may not perform as well as a human politician would in the same conversation. Still, we find that real value is accessible even with today's models.

A3: Discourse on risk of malicious chatbots

AI-driven voter engagement could be used in a manipulative fashion to spread misinformation, target voters with tailored propaganda, or to surveil citizens' political views. Hopefully the deployment of these engaging techniques by actual politicians in a validated way draws engagement away from less altruistic actors' systems.

A4: Automated Analysis of Recorded Conversations

Here are the key takeaways from user conversations that would be useful for Senator Wiener, as a means of demonstrating some of the utility of the project as a means of understanding constituents:

Users had varying levels of familiarity with SB 1047, highlighting the need for clear, accessible information about the bill.

Some users expressed concerns about the bill potentially stifling innovation or being too burdensome for startups and smaller companies. Clarity on how the bill applies to different types of AI models and companies is needed.

Users were interested in understanding the specifics of what the bill would require, such as registration, risk assessments, and compliance with safety standards.

The economic impact of the bill, particularly on California's global competitiveness in the AI industry, was a concern for some users.

One user raised the issue of bad actors not following the rules, emphasizing the need for robust enforcement mechanisms.

Bipartisan support was recognized as crucial for the bill's success, with suggestions to engage lawmakers from both parties and address their concerns.

The influence of tech leaders on public policy was a concern for one user, highlighting the need to balance their perspectives with broader societal interests.

Users appreciated the opportunity to discuss the bill and provide input, underscoring the importance of public engagement in the legislative process.

Overall, the conversations reveal a mix of support, concerns, and desire for more information about SB 1047. Addressing these points through clear communication, stakeholder engagement, and refinement of the bill will be key to its success.

A5: Future Work

This demo focuses on a single politician and a single topic (SB 1047) in order to be a manageable task for a single hackathon, but clearly work could go into fleshing out the system's expertise through retrieval augmented generation (RAG) to match the full range of issues which a single politician could engage on.

A longer project could aim to automatically generate such a chatbot for a given politician by creating an LLM-based agent which does the manual parts of our process: scraping style information from the web and searching for and collecting documents relevant to selected or all topics for RAG.

Related work which we're excited about are systems which can validate the identity of constituents. We note, however, that even those who are not voters in a given politician's district can still be affected by legislation there and/or bring valuable and needed perspective, so we wouldn't want to reject such users even if they could be identified as such.

Finally, advances in foundation models will improve the intelligence, accuracy, and flexibility of the chatbot.

A6: Known Issues

- Multiple users having a conversation with the bot at the same time can get logged as a single interleaved conversation.