# Toward a digital twin of U.S. Congress

**Hayden Helm**[1], **Tianyi Chen**[2], **Harvey McGuinness**[3], **Paige Lee**[4], **Brandon Duderstadt**[4], **Carey E. Priebe**[2]

[1]Helivan Research [2]Johns Hopkins University Department of Applied Mathematics and Statistics

[3]Johns Hopkins University Paul H. Nitze School of Advanced International Studies [4]Nomic AI

## 1. Summary


- Introduce a daily-updated dataset that contains every Tweet from every U.S. congressperson during their respective terms.
- Show large language model(LLM) + congressperson-specific tweets can generate Tweets that are largely indistinguishable from actual Tweets.
- Show generated Tweets can predict roll-call vote and quantify the likelihood of congresspersons crossing party lines.


## 2. Digital Twin of US Congress

Definition of Digital twin [1]: A) relevant and dynamic data, B) physical to virtual feedback, C) valuable inference capabilities, and D) virtual to physical feedback.

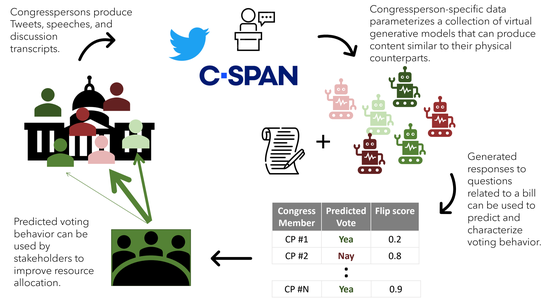


## 3. Nomic Congressional Twitter dataset

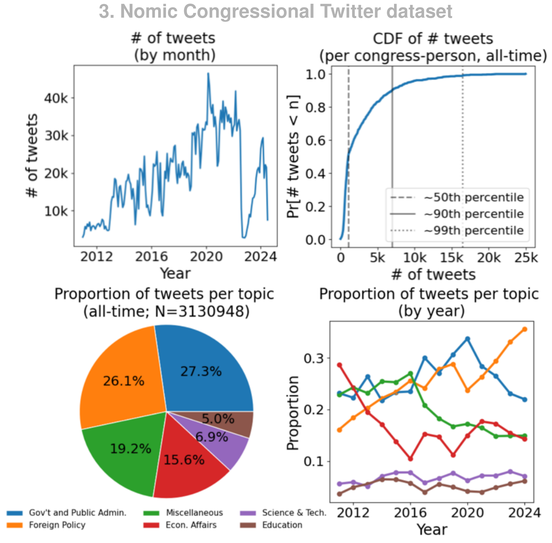


Characteristics of the Nomic Congressional Twitter dataset from October 10th, 2024.The dataset is updated daily and available at https://atlas.nomic.ai/data/hivemind/.

## 4. Physical to Virtual: generated tweets are indistinguishable from real tweets

| | | |
|---|---|---|
| Generic system prompt with no augmentation (**-SP -RAG**) | You are a helpful assistant. | Complete the following Tweet: {start of real Tweet}. Respond with the full Tweet. |
| Specific system prompt with no augmentation (**+SP -RAG**) | You are U.S. congressperson {name}. | Complete the following Tweet: {start of real Tweet}. Respond with the full Tweet. |
| Specific system prompt with augmentation (**+SP +RAG**) | You are U.S. congressperson {name}. | Complete the following Tweet: {start of real Tweet}. Here is an example Tweet potentially related to the to-be-completed Tweet: "{retrieved Tweet}". Respond with the full Tweet. |
| Specific system prompt with augmentation (**+SP +RAG Generated**) | You are U.S. congressperson {name}. | "Write a Tweet that addresses the following question: {question}. Here is an example Tweet potentially related to the to-be-completed Tweet: "{retrieved Tweet}". Start your response with "Here is the full Tweet:" |

Example detectability analysis: Morgan Griffith

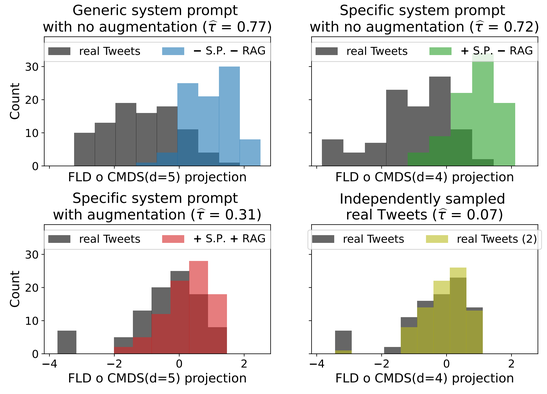


$\hat{\tau}$ is the empirical classification accuracy using Fisher linear discriminant. Smaller $\hat{\tau}$ means distribution are more indistinguishable[2]. Figure shows the inclusion of previously written Tweets via RAG decreases detectability significantly.

## 5. Virtual to Physical: predicting roll-call vote using generated tweets

Construction of data kernel perspective space(DKPS) which are low-dimensional representations of digital congressperson/LLMs as a summary of the relative position of each congressperson with respect to the bill-related questions[3][4].

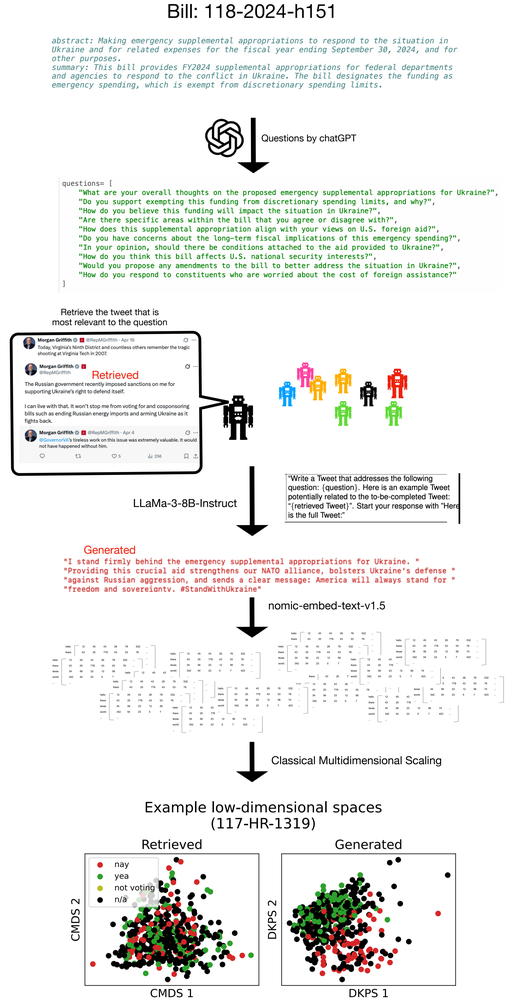


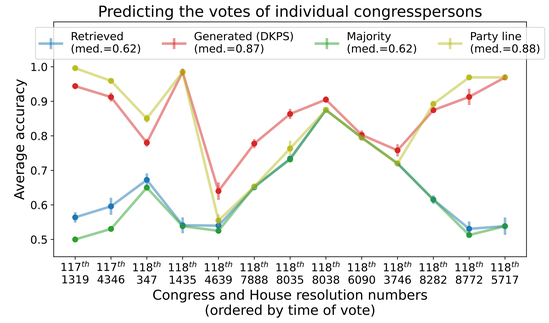


Using DKPS representation obtained by generated tweets shows capability of predicting vote with relatively high accuracy.

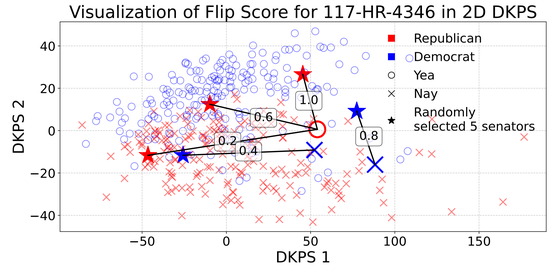


Enlarged $\times$ and $\circ$ is the House member used to calculate senator's flip score $\mathcal{H}(S)$. For a given Senator $S$ define $\mathcal{H}(S) := \{H : H \in \text{House}, H \text{ and } S \text{ in same party}, H \text{ voted across party lines}\}$, then flip score$(S) = \frac{1}{\min_{H \in \mathcal{H}(S)} \|X_H - X_S\|}$. Senators closer to a cross-party line voter in their party are assigned a higher flip score. Flip score can be used by stakeholders to improve resource allocation.

## 6. Conclusion and Limits

We provided evidence that a collection of language models each equipped with a congressperson-specific dataset satisfies the four requirements for a virtual model to be a digital twin for a collection of congresspersons. We plan to involve other sources of data such as campaign speeches etc; model congressperson-to-congressperson or congressperson-to-public interactions.

## References


[1] National Academy of Engineering and National Academies of Sciences, *Foundational Research Gaps and Future Directions for Digital Twins*, The National Academies Press, Washington, DC, 2024.

[2] Hayden Helm, Carey E. Priebe, and Weiwei Yang, "A statistical turing test for generative models," 2023.

[3] Aranyak Acharyya, Michael W. Trosset, Carey E. Priebe, and Hayden S. Helm, "Consistent estimation of generative model representations in the data kernel perspective space," 2024.

[4] Hayden Helm, Aranyak Acharyya, Brandon Duderstadt, Youngser Park, and Carey E. Priebe, "Embedding-based statistical inference on generative models," 2024.