Uncovering Latent Arguments in Social Media Messaging by Employing LLMs-in-the-Loop Strategy

Tunazzina Islam, Dan Goldwasser

Department of Computer Science,

Purdue University, West Lafayette, IN 47907, USA

NAACL 2025

Date: April 29-May 4, 2025

PURDUE UNIVERSITY

Public Opinion

- Responsive governance
- Policy alignment with public interests
- Societal harmony
- Continuous policy refinement *(Glynn & Huge, 2008; Price, 1988)*

Distributed Landscape of Social Media

Distributed Landscape of Social Media

Users **generate** and **consume** a variety of content.

Analyzing Public Opinion

- Automatically analyzing public opinion on social media platforms.

Analyzing Public Opinion

- Automatically analyzing public opinion on social media platforms.
- **Argument Mining**.
    - automatically extracts the reasons, claims, and talking points/arguments.
    - shedding light on how and why specific opinions are formed.

Argument Mining - Previous Works

- Topic Modeling.
  - Shallow Themes.
- Manual and qualitative coding *(Hagen et al., 2022; Nguyen et al., 2021; Del Valle et al., 2020)*.

Argument Mining - Previous Works

- Topic Modeling.
  - Shallow Themes.
- Manual and qualitative coding *(Hagen et al., 2022; Nguyen et al., 2021; Del Valle et al., 2020)*.
- Theme Discovery *(Islam & Goldwasser, 2025; Pacheco et al., 2023; Islam et al.,2023b; Islam & Goldwasser, 2022; Pacheco et al., 2022a)*.
  - Can not to recognize conflicting arguments under same theme, i.e.,

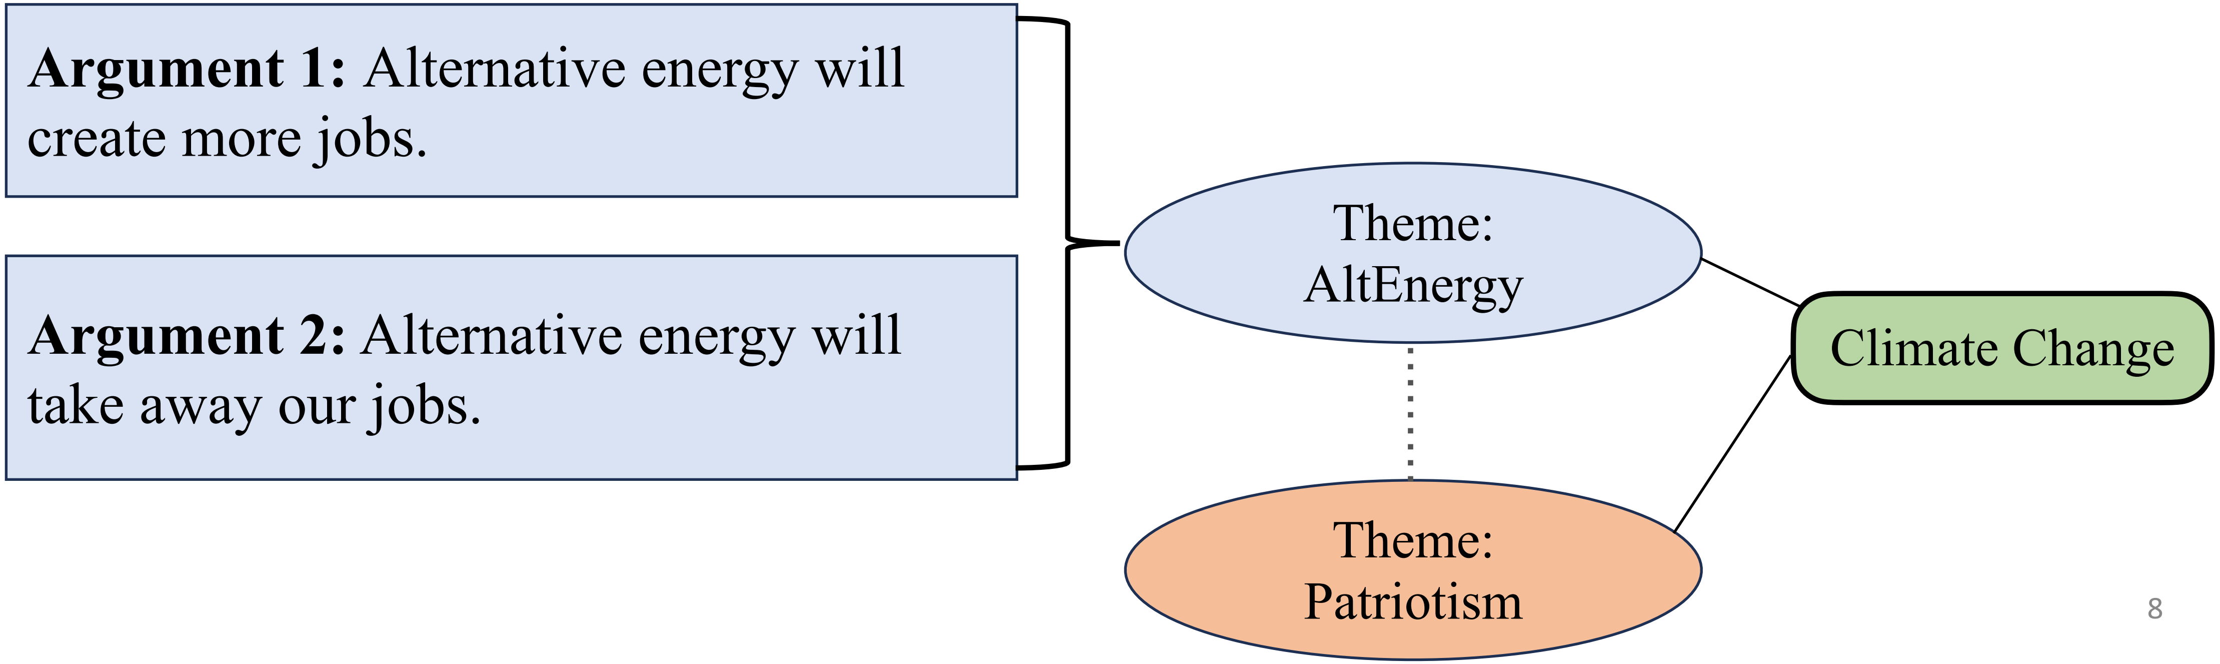

Argument Mining - Previous Works

- Human-in-loop *(Pacheco et al. 2022b;a)*.
  - Costly scalability.
  - Time consuming.

Machine-in-loop Approach

- Human-in-loop *(Pacheco et al. 2022b;a)*.
  - Costly scalability.
  - Time consuming.

- **Machine-in-the-Loop: LLMs-in-the-Loop.**
  - LLMs possess **extensive domain insights**.
  - **Reasoning** capabilities.
  - **Accelerate** the process of refinement.

Sketch of *LLMs-in-the-Loop* Approach

- **Themes are pre-defined**.
- Theme-specific **clustering**.
- **Summarizing** sub-clusters.
  - Zero-shot multi-document summarization using GPT-4 on top-k instances.
- **Generating** and **refining** arguments.
  - Implement a redundancy check to identify and merge similar arguments.

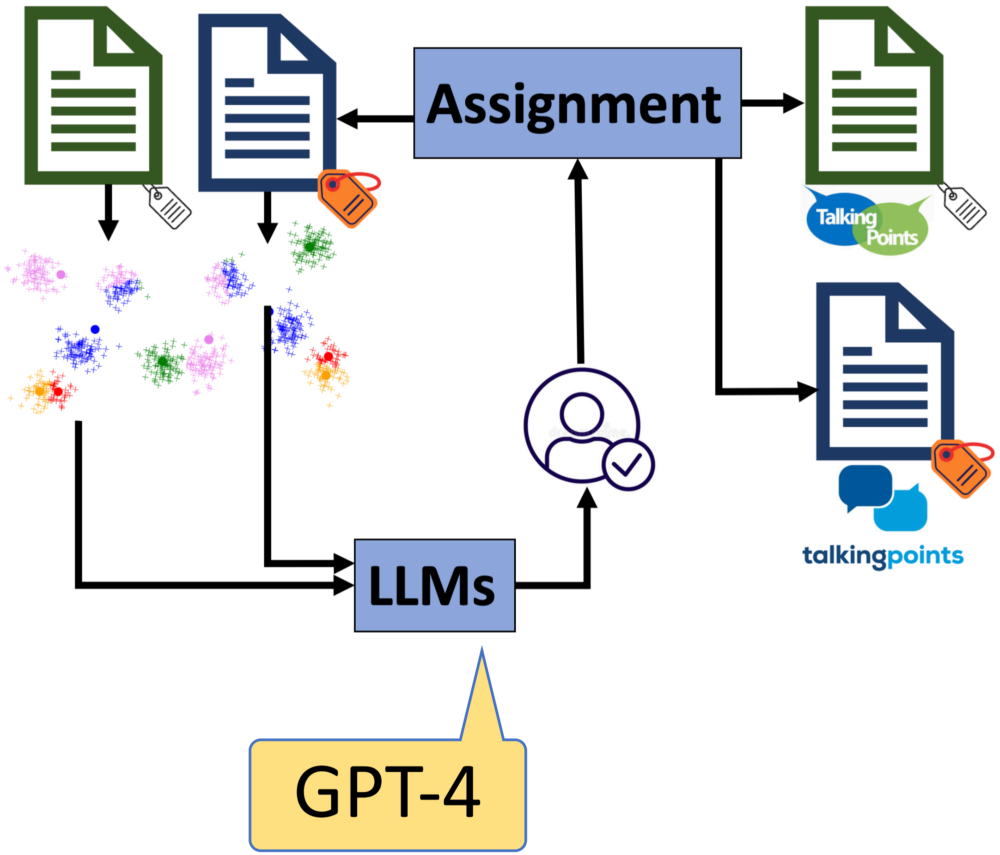

Sketch of *LLMs-in-the-Loop* Approach

- **Themes are pre-defined**.
- Theme-specific **clustering**.
- **Summarizing** sub-clusters.
  - Zero-shot multi-document summarization using GPT-4 on top-k instances.
- **Generating** and **refining** arguments.
  - Implement a redundancy check to identify and merge similar arguments.
- **Human evaluation**.
  - Quality and relevance of the generated arguments.

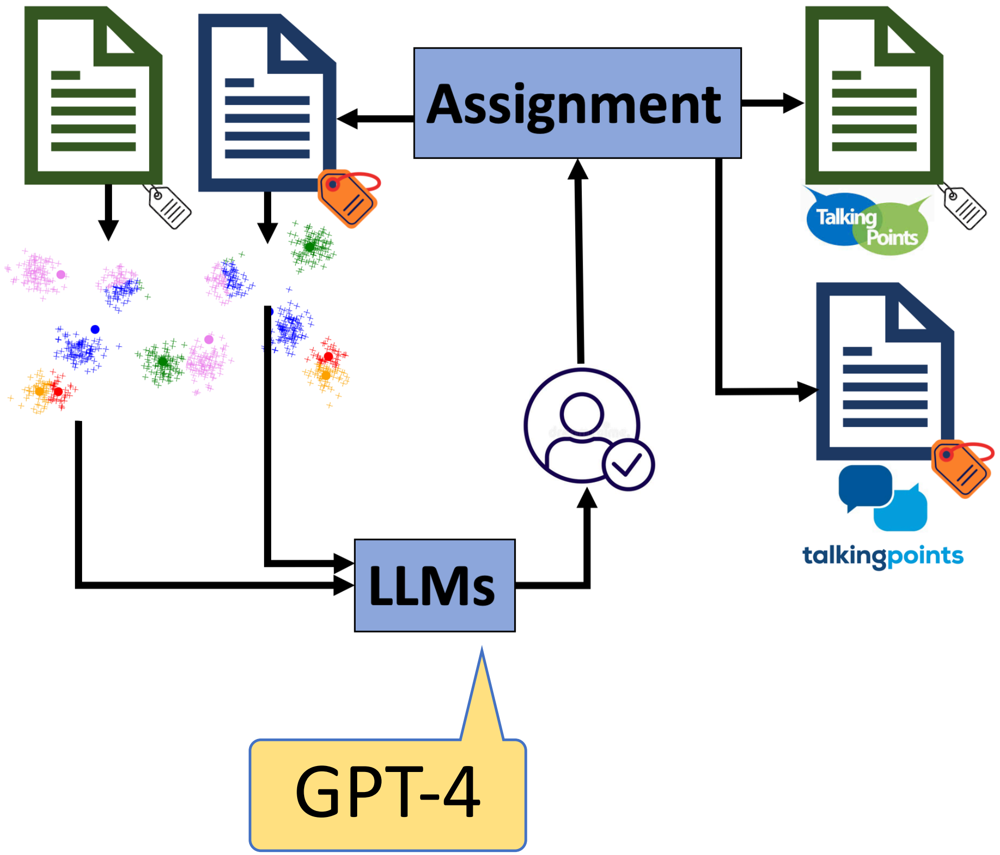

Sketch of *LLMs-in-the-Loop* Approach

- **Themes are pre-defined**.
- Theme-specific **clustering**.
- **Summarizing** sub-clusters.
  - Zero-shot multi-document summarization using GPT-4 on top-k instances.
- **Generating** and **refining** arguments.
  - Implement a redundancy check to identify and merge similar arguments.
- **Human evaluation**.
  - Quality and relevance of the generated arguments.
- **Mapping** instances to arguments.
  - Distance-based approach for mapping.

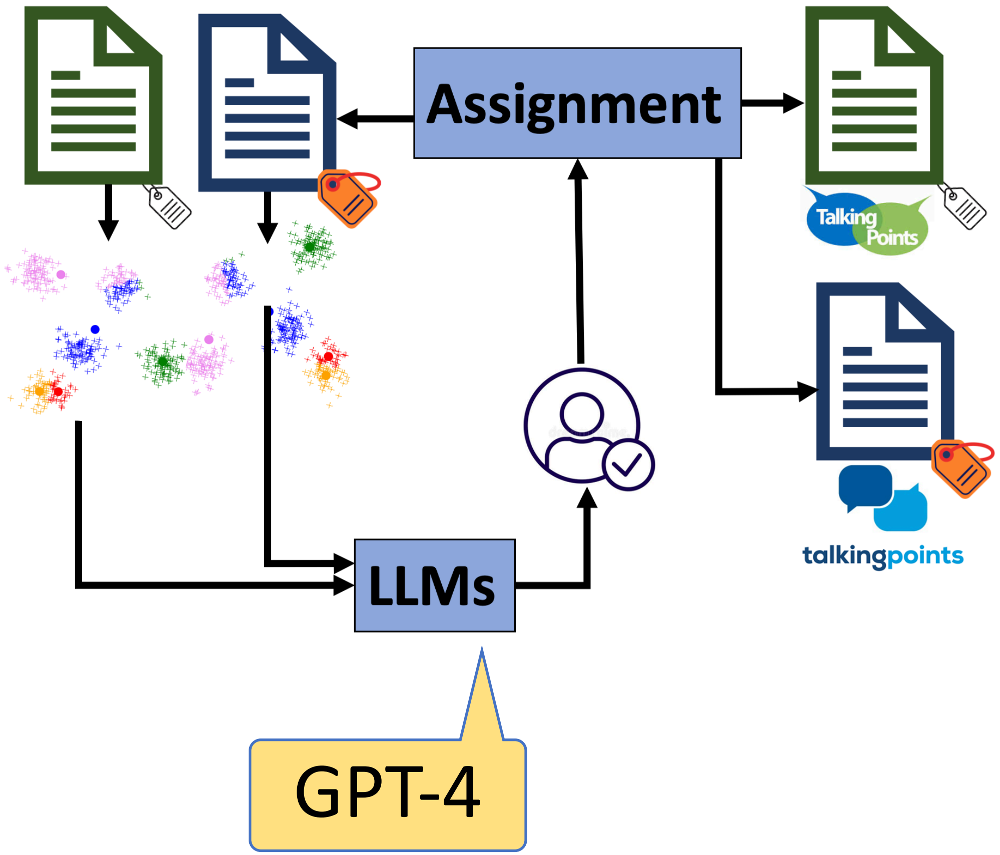

Sketch of *LLMs-in-the-Loop* Approach

- **Themes are pre-defined**.
- Theme-specific **clustering**.
- **Summarizing** sub-clusters.
  - Zero-shot multi-document summarization using GPT-4 on top-k instances.
- **Generating** and **refining** arguments.
  - Implement a redundancy check to identify and merge similar arguments.
- **Human evaluation**.
  - Quality and relevance of the generated arguments.
- **Mapping** instances to arguments.
  - Distance-based approach for mapping.
- **Repeat**: Unassigned ads from iteration1.

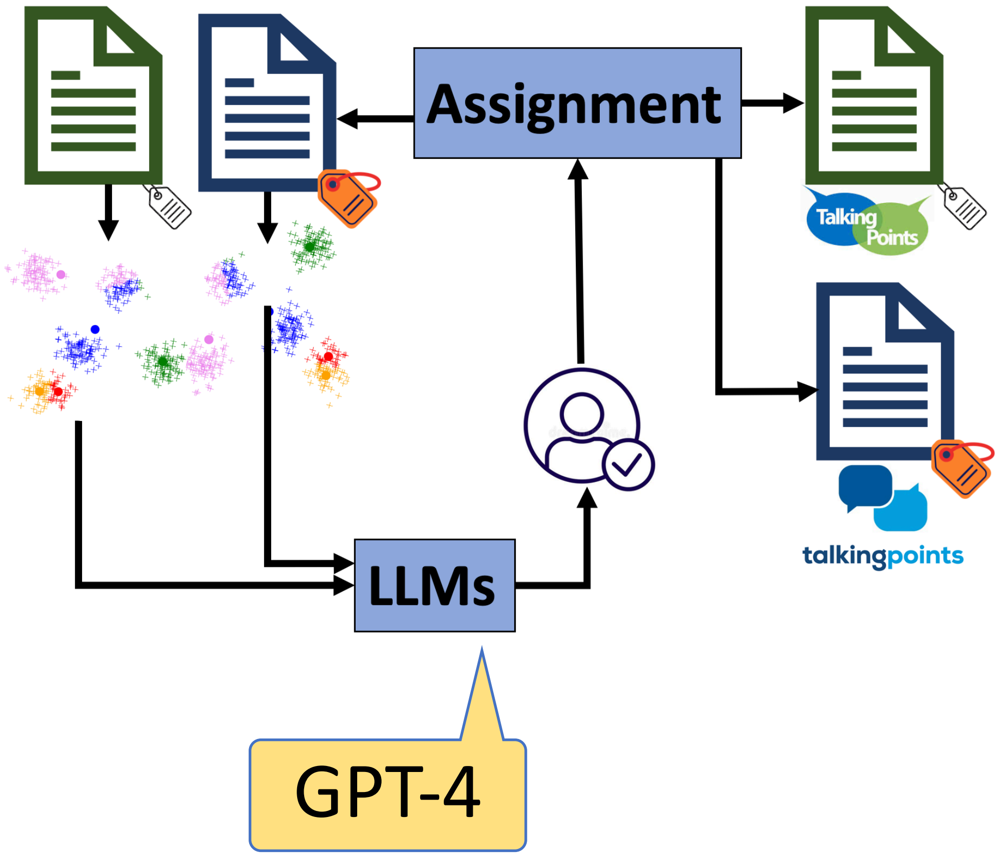

Case Studies

- **Climate** campaigns.
  - **14k ads** *(Islam et al. 2023, Islam and Goldwasser 2025)*, January 2021 to January 2022.
  - Stance (e.g., pro-energy, clean-energy) and theme (e.g., support climate policy).

Case Studies

- **Climate** campaigns.
  - **14k ads** *(Islam et al. 2023, Islam and Goldwasser 2025)*, January 2021 to January 2022.
  - Stance (e.g., pro-energy, clean-energy) and theme (e.g., support climate policy).
- **COVID-19** vaccine campaigns.
  - **9k ads** *(Islam and Goldwasser 2022)*, December 2020 to January 2022.
  - Moral foundation (e.g., care/harm) *(Haidt and Graham, 2007)* and theme (e.g., vaccine equity).

Evaluation

- **Sort** the ads according to their semantic distance to their assigned arguments.
- Compute the **three** quartiles and sample a set of **12 ads** per theme, such that **3 ads** are randomly sampled from each quartile.
- **300** ads in the 1st iteration and another **300** ads from the 2nd iteration from **climate** case study.
- **168** ads in the 1st iteration and another **168** ads from the 2nd iteration from **COVID-19** case study.
- Manually annotate **936 ads** whether the mapping is correct or not.

Results

- **Better performance** in the **lower distance** between ad and argument.
- **Improvement** in performance both in **coverage** and **mapping quality** after **subsequent iterations**.

Case Study	Iter.	# Args	Coverage	$\leq Q_1$	$\leq Q_2$	$\leq Q_3$	All
Climate	1	113	37.38%	76.00%	70.67%	58.67%	57.33%
	2	213	44.40%	88.00%	74.67%	70.67%	64.00%
COVID-19	1	47	36.18%	78.57%	61.90%	61.90%	52.38%
	2	78	40.47%	82.93%	73.81%	64.29%	57.14%

Table: Coverage and mapping quality w.r.t. Human Judgments.

Ablation Study

- Comparable results in terms of coverage:
  - Arguments from the top k instances of a cluster **without summarizing** vs. **with summarizing**.

Case Study	Iter.	Number of covered ads thr < 0.6	thr < 0.5	thr < 0.4	thr < 0.3
Climate	1	13319	10677	5355	1132
	1-w/o sum.	13394	10613	5189	1164
	2	13669	11541	**6360**	**1458**
	2-w/o sum.	**13759**	**11592**	6143	1384
COVID-19	1	7962	6525	3589	850
	1-w/o sum.	8133	6507	3477	787
	2	8197	**6833**	**4015**	**1089**
	2-w/o sum.	**8426**	6767	3710	908

Table: Ablation study (coverage). sum: summary, thr: threshold.

Downstream Task: Stance Prediction

Stance prediction task **improves** when **talking points** are **added** with **text**.

Model	Acc	F1
$Longformer_{text}$	90.13%	89.89%
$Longformer_{tp}$	83.43%	83.44%
$\mathbf{Longformer_{text+tp}}$	**93.30%**	**93.29%**
$RoBERTa_{text}$	93.07%	92.96%
$RoBERTa_{tp}$	82.73%	82.81%
$\mathbf{RoBERTa_{text+tp}}$	**93.65%**	**93.56%**
$llama3_{text}$	92.00%	90.95%
$llama3_{tp}$	81.00%	78.68%
$\mathbf{llama3_{text+tp}}$	**92.50%**	**91.49%**

Table: Contribution of talking point (tp) in stance classifier for climate campaigns dataset.

Argumentative Cohesion Comparison

- **COVID-19** vaccine campaigns.
  - Pearson **correlation** between **arguments** and **moral foundations**.
  - Random **15** arguments.

- **Climate** campaigns.
  - Pearson **correlation** between **arguments** and **stances**.
  - Random **25** arguments.

Argumentative Cohesion Comparison: COVID-19

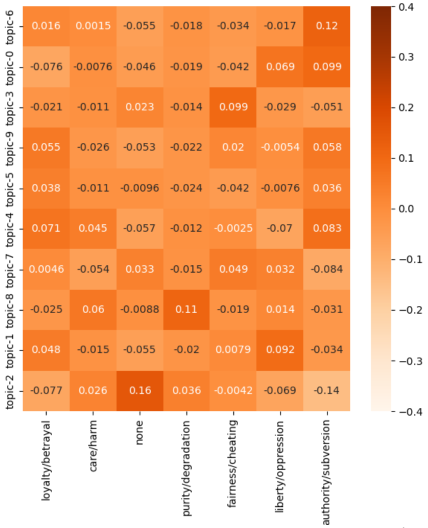

Baseline: 10 LDA Topics

Baseline: 20 LDA Topics

Ours: After 2nd round of iteration

Argumentative Cohesion Comparison: Climate

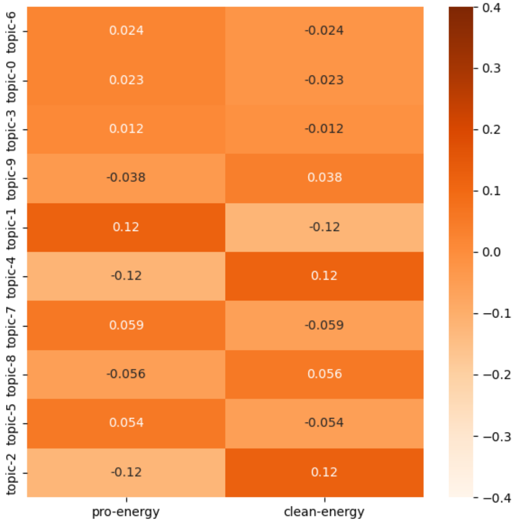

Baseline: 10 LDA Topics

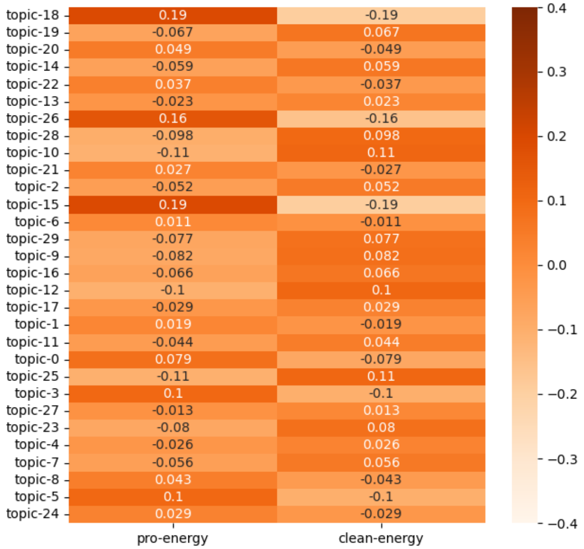

Baseline: 30 LDA Topics

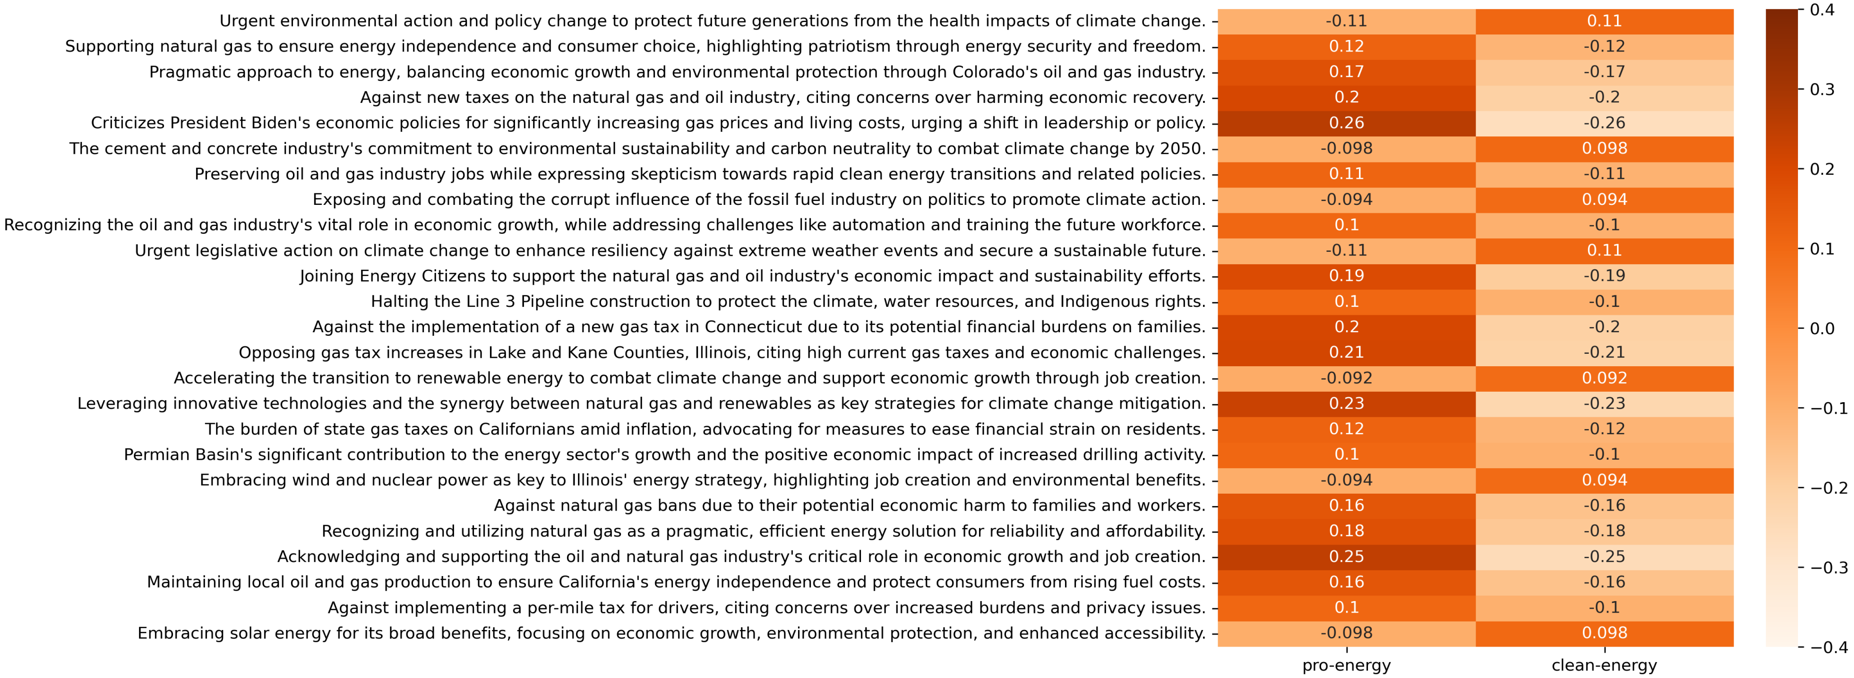

Ours: After 2nd round of iteration

Demographic Targeting

- **Three** age categories.
  i.e.,
  - Young people (ages 13-24)
  - Working-age people (ages 25-54)
  - Older population (age 55+)
- Florida and Texas.

Demographic Targeting

- **Three** age categories. i.e.,
  - Young people (ages 13-24)
  - Working-age people (ages 25-54)
  - Older population (age 55+)
- Florida and Texas.

Age Group	State	Entity	Talking Points
13-24	TX	Children, Parents, reproductive health.	Advocates for the safety of COVID-19 vaccines for children, emphasizing mild side effects and community protection through vaccination.
25-54	FL	Ron DeSantis, Dr. Joseph Ladapo, Surgeon General.	Advocates for building and restoring public trust in the COVID-19 vaccine and the medical community.
55+	TX	seniors, Pfizer, who have passed away or are hospitalized due to Covid, I.	Strongly advocates for COVID-19 vaccination, highlighting its safety, efficacy, and crucial role in preventing severe illness and ending the pandemic.
	FL	Governor Ron DeSantis, seniors, loved one, Johnson & Johnson.	Efforts and challenges in equitable vaccine distribution and access for seniors across various counties.

Arguments Shift Triggered by Key Events

- Event1: **Hurricane Ida**, Date: **August 29, 2021**.

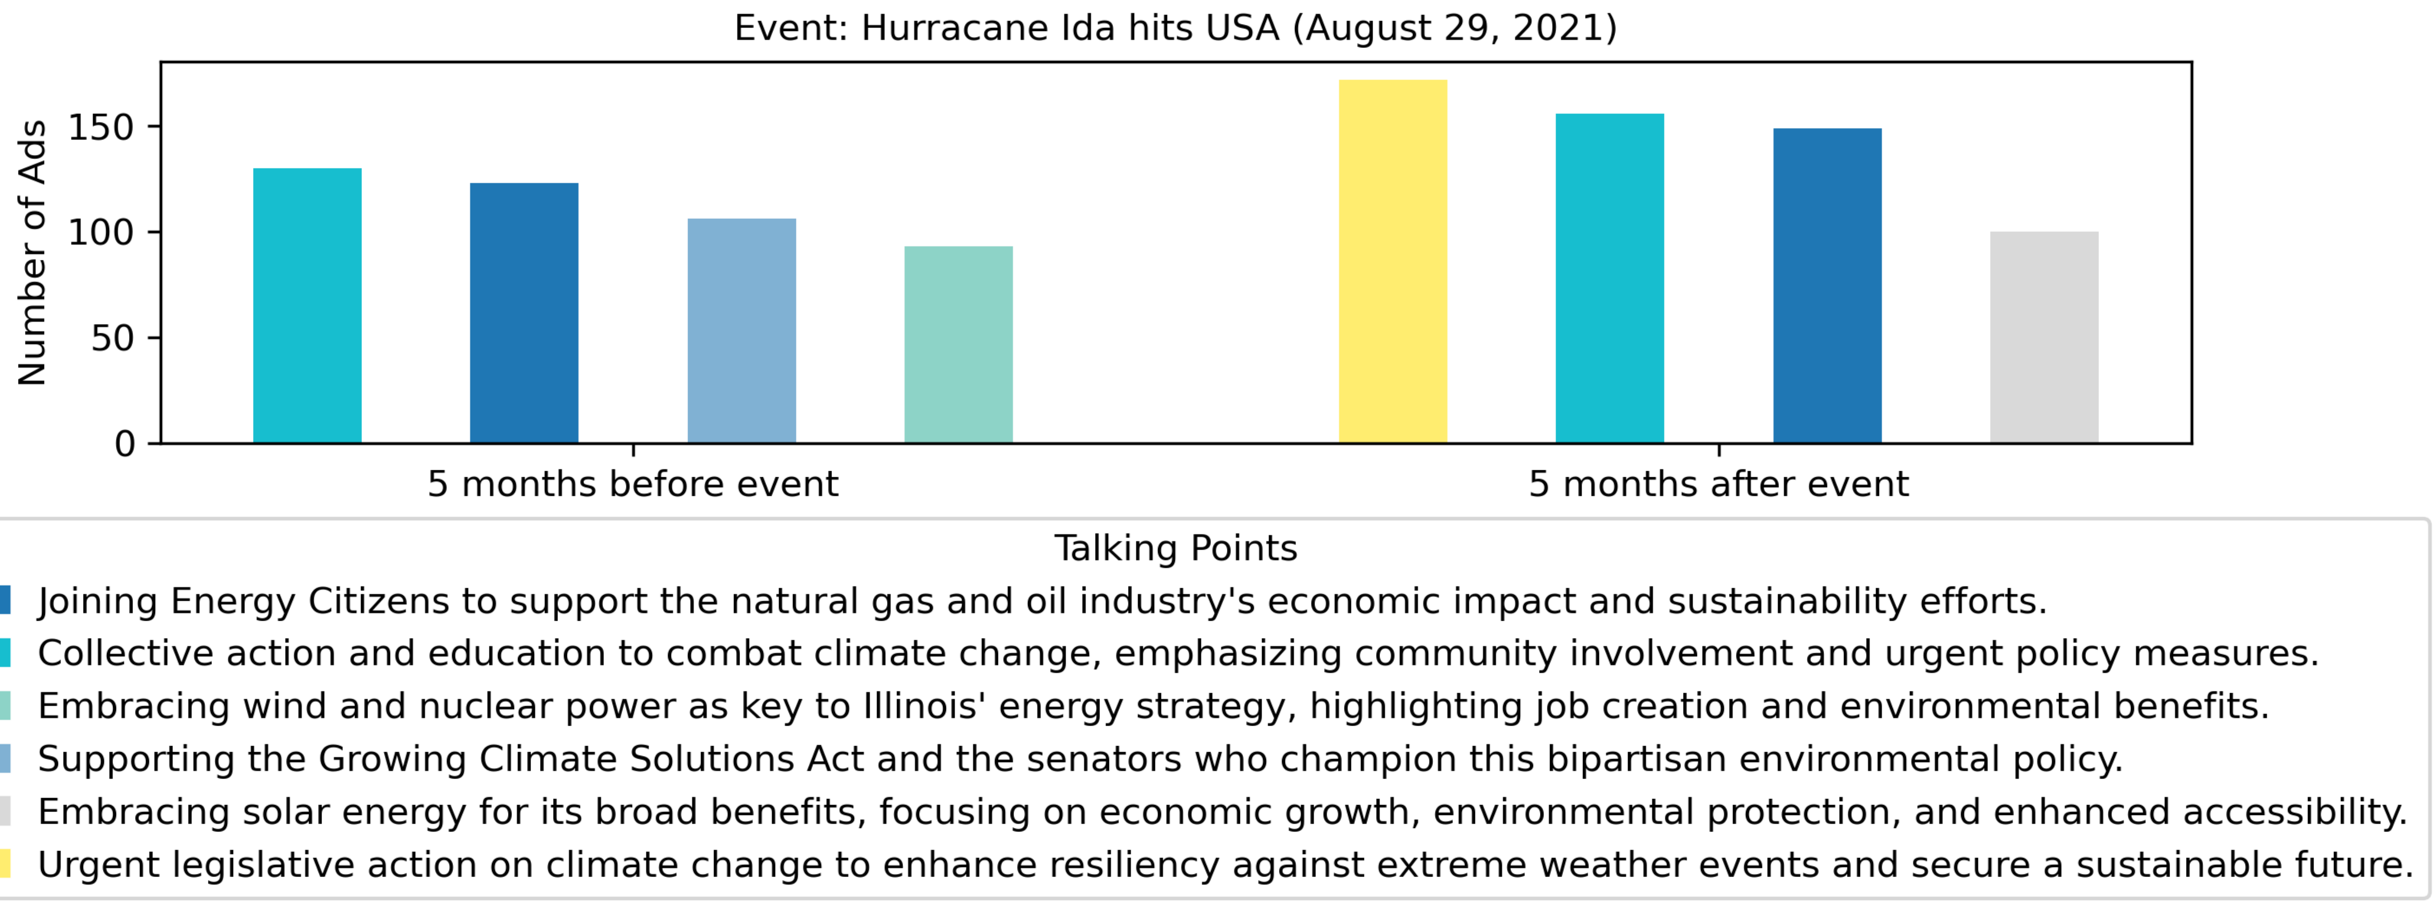

Arguments Shift Triggered by Key Events

- Event1: **Hurricane Ida**, Date: **August 29, 2021**.

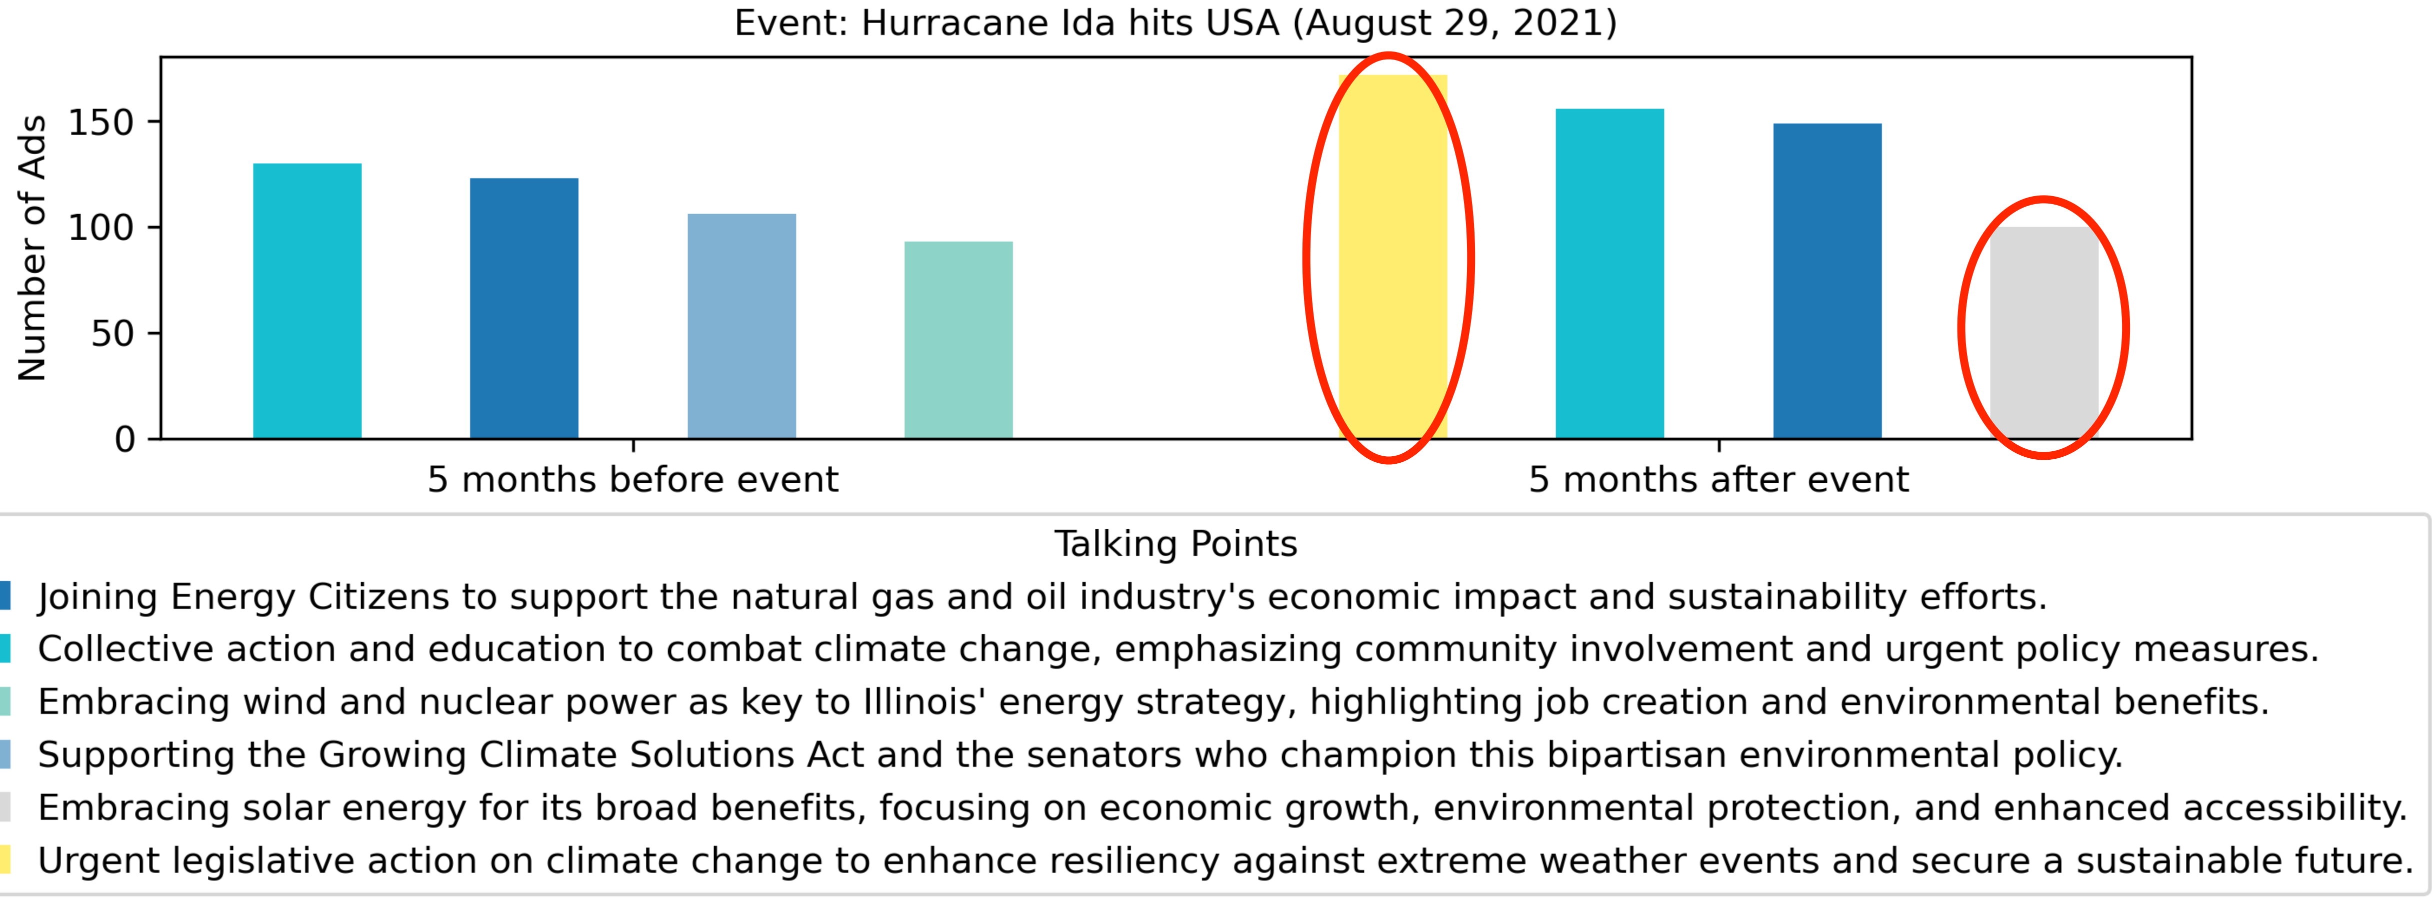

Arguments Shift Triggered by Key Events

- Event2: **Federal COVID-19 vaccine mandate**, Date: **September 09, 2021**.

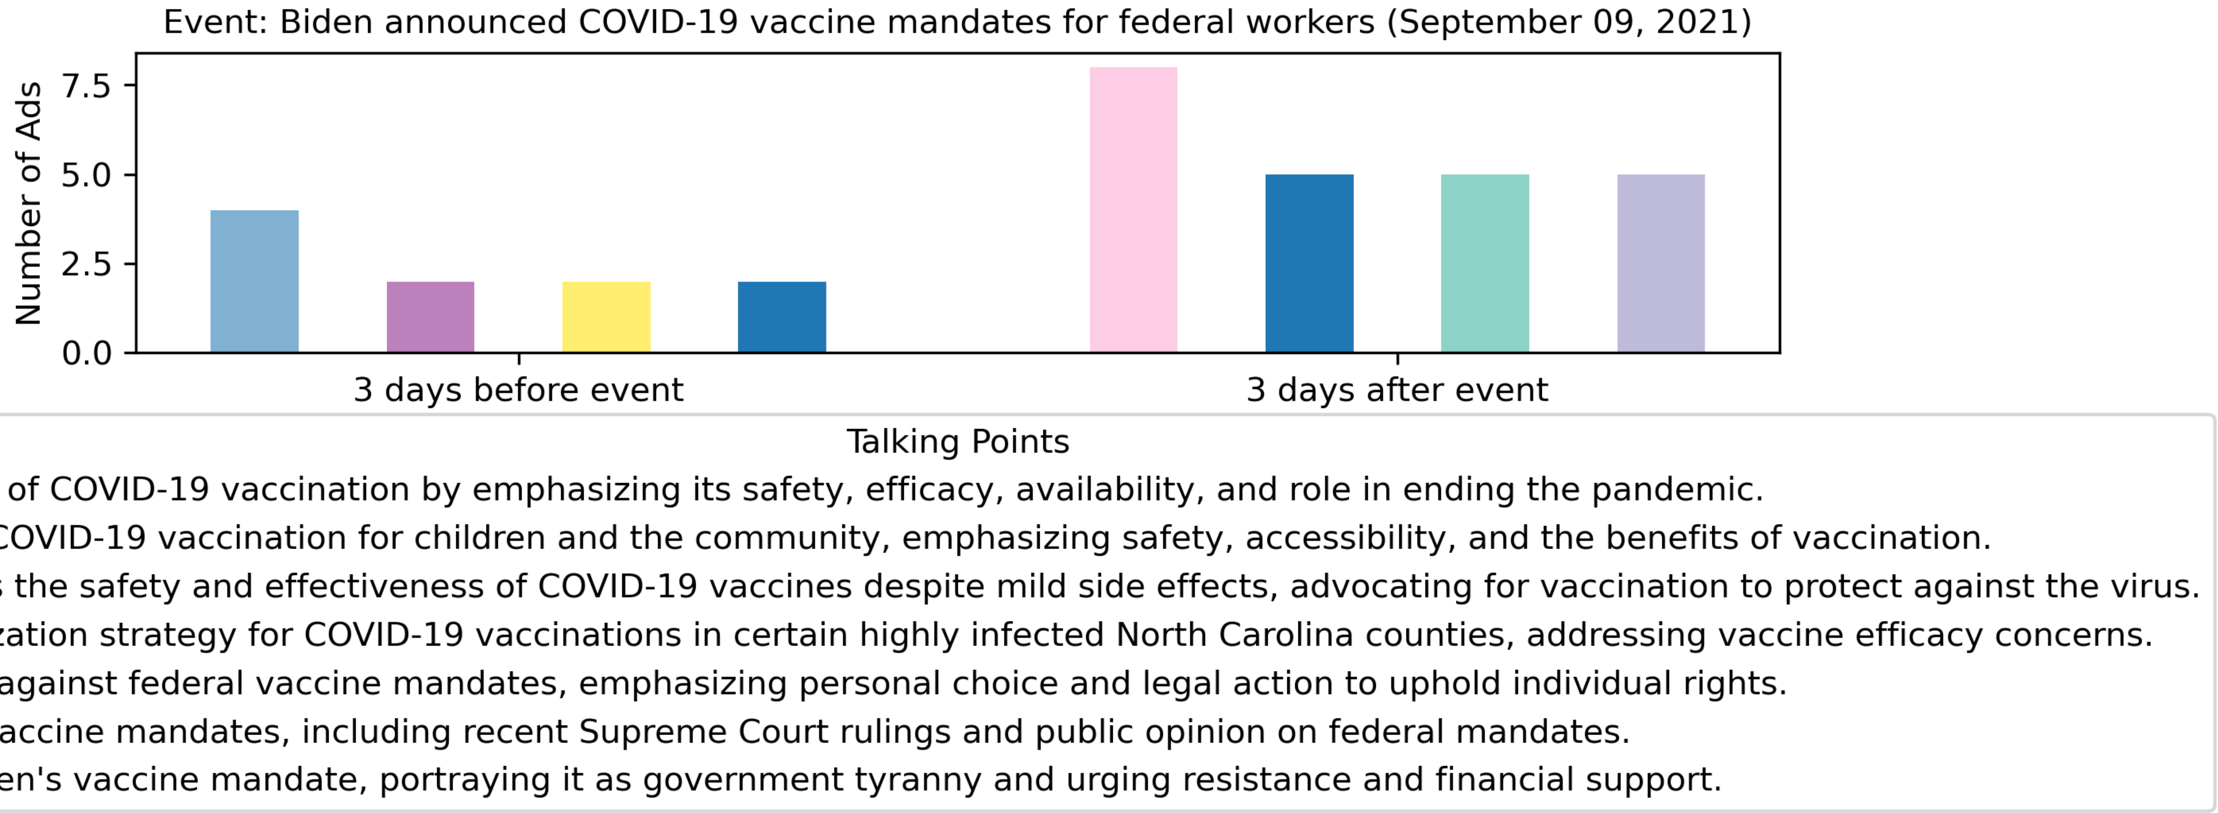

Arguments Shift Triggered by Key Events

- Event2: **Federal COVID-19 vaccine mandate**, Date: **September 09, 2021**.

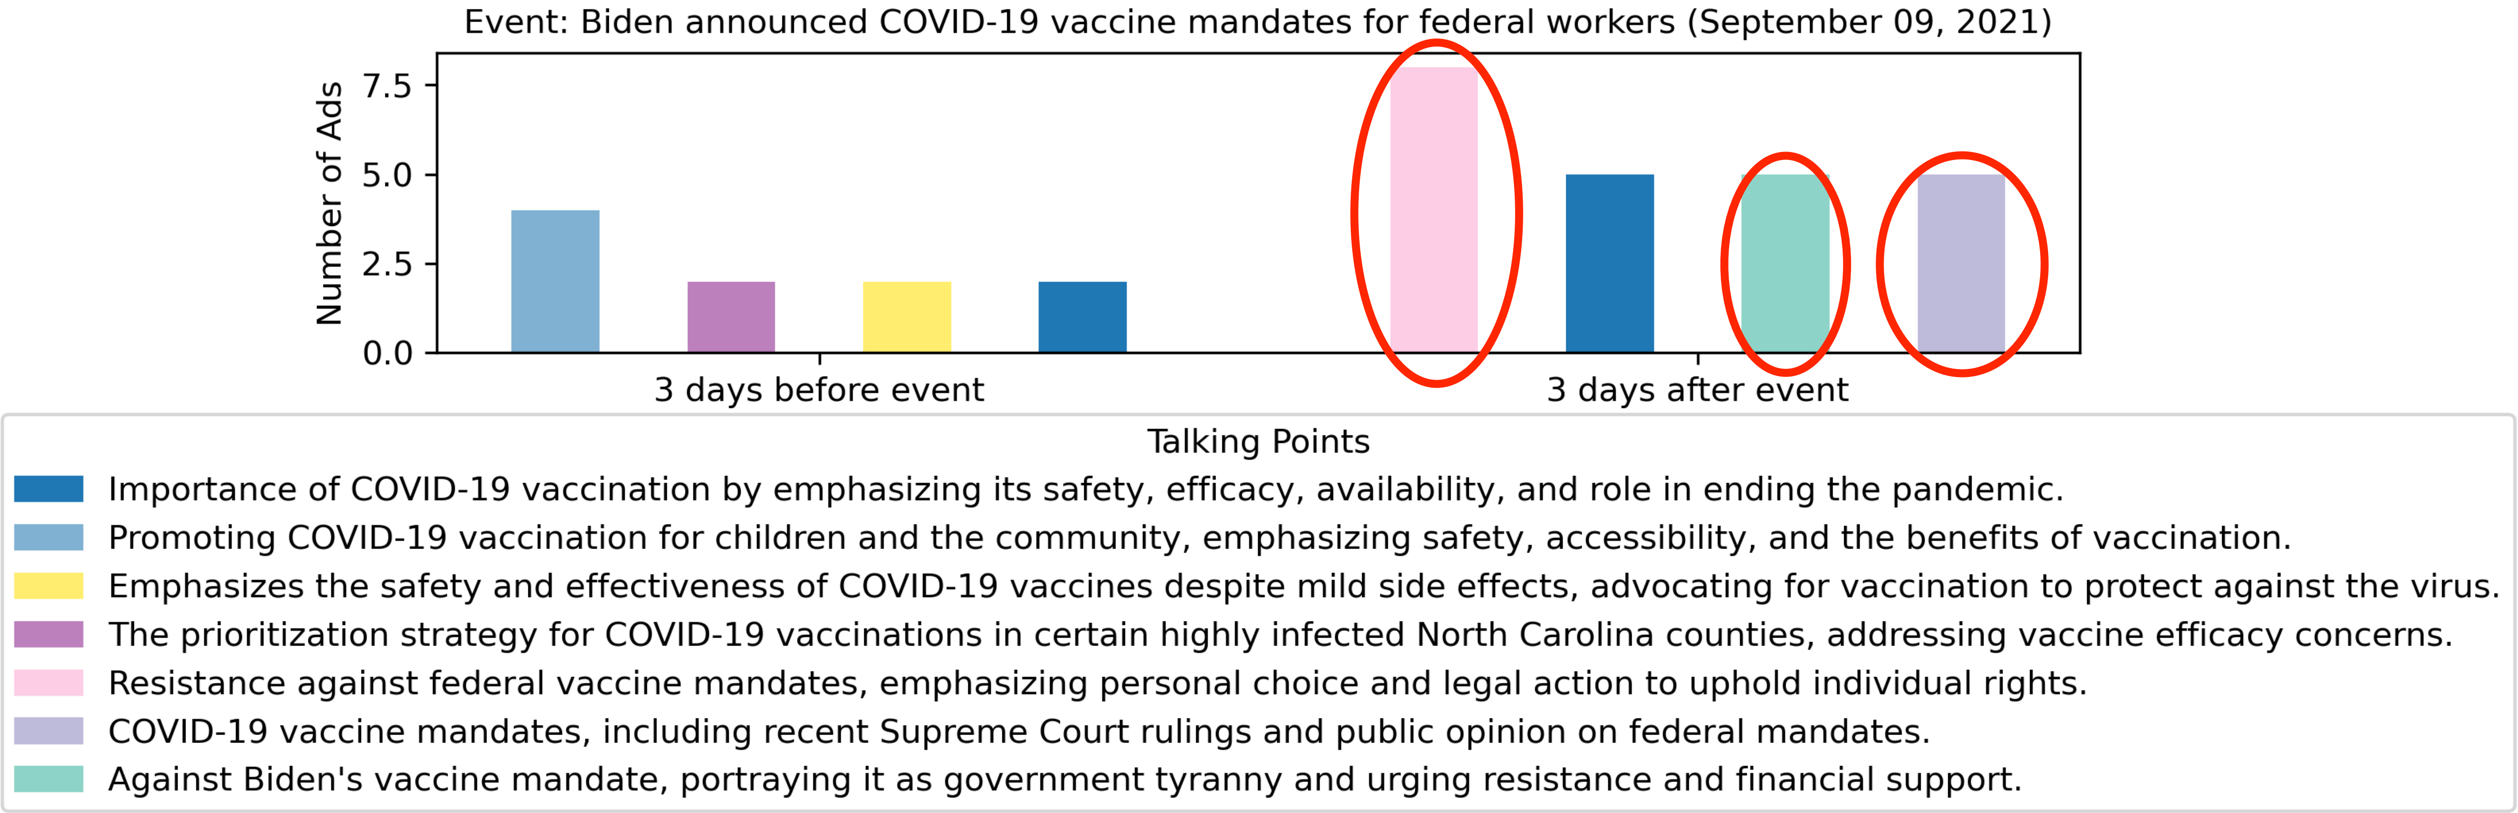

Key Takeaways

- **Iterative LLMs-in-the-Loop** framework for uncovering latent arguments.

Key Takeaways

- **Iterative LLMs-in-the-Loop** framework for uncovering latent arguments.
- Quantitative results:
  - Newly discovered arguments can **cover** a **larger portion** of texts.
  - **Map** texts -> arguments **accurately** w.r.t. human judgment.
  - Arguments are **more strongly correlated** with specific **stances** or **moral foundations** than the LDA topics.

Key Takeaways

- **Iterative LLMs-in-the-Loop** framework for uncovering latent arguments.
- Quantitative results:
  - Newly discovered arguments can **cover** a **larger portion** of texts.
  - **Map** texts -> arguments **accurately** w.r.t. human judgment.
  - Arguments are **more strongly correlated** with specific **stances** or **moral foundations** than the LDA topics.
- Talking point information **improves** the **stance classifier performance**.

Key Takeaways

- **Iterative LLMs-in-the-Loop** framework for uncovering latent arguments.
- Quantitative results:
  - Newly discovered arguments can **cover** a **larger portion** of texts.
  - **Map** texts -> arguments **accurately** w.r.t. human judgment.
  - Arguments are **more strongly correlated** with specific **stances** or **moral foundations** than the LDA topics.
- Talking point information **improves** the **stance classifier performance**.
- Talking points are tailored for **demographic targeting**.
- Talking points **dynamically shift** in response to **real world events**.

References

1. Carroll J Glynn and Michael E Huge. Public opinion. The international encyclopedia of communication, 2008.
2. Vincent Price. On the public aspects of opinion: Linking levels of analysis in public opinion research. Communication research, 15(6):659–679, 1988.
3. J. Haidt and J. Graham, “When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize,” Social Justice Research, vol. 20, no. 1, pp. 98–116, 2007.
4. T. Islam, R. Zhang, and D. Goldwasser, “Analysis of climate campaigns on social media using bayesian model averaging,” in Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society, ser. AIES ’23, Montréal, QC, Canada: Association for Computing Machinery, 2023, pp. 15–25, isbn: 9798400702310.
5. Tunazzina Islam and Dan Goldwasser. Understanding covid-19 vaccine campaign on facebook using minimal supervision. In 2022 IEEE International Conference on Big Data (BigData), pp. 585–595. IEEE, 2022.
6. Tunazzina Islam and Dan Goldwasser. Discovering Latent Themes in Social Media Messaging: A Machine-in-the-Loop Approach Integrating LLMs. arXiv preprint arXiv:2403.10707, 2024.
7. Mar´ıa Leonor Pacheco*, Tunazzina Islam*, Monal Mahajan, Andrey Shor, Ming Yin, Lyle Ungar, and Dan Goldwasser. A holistic framework for analyzing the covid-19 vaccine debate. In Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pp. 5821–5839, 2022a. (* equal contribution)
8. Maria Leonor Pacheco, Tunazzina Islam, Lyle Ungar, Ming Yin, and Dan Goldwasser. Interactively uncovering latent arguments in social media platforms: A case study on the covid-19 vaccine debate. In Proceedings of the Fourth Workshop on Data Science with Human-in-the-Loop (Language Advances), pp. 94–111, 2022b.
9. Maria Leonor Pacheco, Tunazzina Islam, Lyle Ungar, Ming Yin, and Dan Goldwasser. Interactive concept learning for uncovering latent themes in large text collections. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), Findings of the Association for Computational Linguistics: ACL 2023, pp. 5059–5080, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-acl.313. URL https://aclanthology.org/2023.findings-acl.313.

References

10. Loni Hagen, Ashley Fox, Heather O'Leary, DeAndre Dyson, Kimberly Walker, Cecile A Lengacher, and Raquel Hernandez. The role of influential actors in fostering the polarized covid-19 vaccine discourse on twitter: Mixed methods of machine learning and inductive coding. Jmir Infodemiology, 2(1):e34231, 2022.
11. Thu T. Nguyen, Shaniece Criss, Eli K. Michaels, Rebekah I. Cross, Jackson S. Michaels, Pallavi Dwivedi, Dina Huang, Erica Hsu, Krishay Mukhija, Leah H. Nguyen, Isha Yardi, Amani M. Allen, Quynh C. Nguyen, and Gilbert C. Gee. Progress and push-back: How the killings of ahmaud arbery, breonna taylor, and george floyd impacted public discourse on race and racism on twitter. SSM - Population Health, 15:100922, 2021. ISSN 2352-8273. doi: https://doi.org/10.1016/j.ssmph.2021.100922. URL https://www.sciencedirect. com/science/article/pii/S235282732100197X.
12. Marc Esteve Del Valle, Rimmert Sijtsma, Hanne Stegeman, and Rosa Borge. Online deliberation and the public sphere: Developing a coding manual to assess deliberation in twitter political networks. Javnost-The Public, 27(3):211–229, 2020.

THANK YOU ☺

Slide: https://tunazislam.github.io/files/LatentArgumentsLLM.pdf

Questions?

Tunazzina Islam, Ph.D.

Department of Computer Science,

Purdue University, West Lafayette, IN.

Email: **islam32@purdue.edu**

https://tunazislam.github.io/

@Tunaz_Islam

PURDUE UNIVERSITY®