

# Findings and Insights from the 8th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text

Ali Hürriyetoğlu<sup>1</sup>, Surendrabikram Thapa<sup>2</sup>, Hristo Tanev<sup>3</sup>, Surabhi Adhikari<sup>4</sup>

<sup>1</sup>Wageningen Food Safety Research, Netherlands, <sup>2</sup>Virginia Tech, USA,

<sup>3</sup>European Commission, Joint Research Centre, Italy, <sup>4</sup>Columbia University, USA

<sup>1</sup>ali.hurriyetoğlu@wur.nl, <sup>2</sup>sbt@vt.edu,

<sup>3</sup>hristo.tanev@ec.europa.eu, <sup>4</sup>surabhi.adhikari@columbia.edu

## Abstract

This paper presents an overview of the 8th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE), held in conjunction with RANLP 2025. The workshop featured a range of contributions, including regular research papers, system descriptions from shared task participants, and an overview paper on shared task outcomes. Continuing its tradition, CASE brings together researchers from computational and social sciences to explore the evolving landscape of event extraction. With the rapid advancement of large language models (LLMs), this year's edition placed particular emphasis on their application to socio-political event extraction. Alongside text-based approaches, the workshop also highlighted the growing interest in multimodal event extraction, addressing complex real-world scenarios across diverse modalities.

## 1 Introduction

In an increasingly interconnected and digitized world, the vast availability of textual and multimodal data related to socio-political, economic, environmental, and humanitarian events presents unprecedented opportunities for data-driven analysis across the social sciences and humanities (Hürriyetoğlu et al., 2024, 2021a; Chen et al., 2023). Governments, international organizations, journalists, and civil society actors increasingly rely on such data to gain timely, granular, and actionable insights into events such as protests, conflicts, public health emergencies, migration patterns, and policy shifts (Shu and Ye, 2023; Hürriyetoğlu et al., 2021c).

Recent years have seen a dramatic shift in the landscape of event extraction due to the rapid advancement of large language models (LLMs) (Thapa et al., 2025c; Hou and Huang, 2025). These

models, capable of understanding, generating, and reasoning over text with minimal supervision, have opened new avenues for tackling long-standing challenges in socio-political event extraction such as low-resource languages, implicit events, cross-document reasoning, and multilingual understanding (Ziems et al., 2024; Anthis et al.; Shen et al., 2023). Techniques such as prompt-based learning, instruction tuning, and alignment have enabled more adaptive and generalizable extraction pipelines, reducing dependence on handcrafted features and rigid annotation schemas (Hou et al., 2024; Kirk et al., 2024; Khan et al., 2025).

Beyond text, the field is increasingly moving toward multimodal event extraction, integrating information from images, videos, and social media metadata (Thapa et al., 2025d). This trend is especially relevant in crisis monitoring, misinformation detection, and humanitarian response, where visual and textual signals must be jointly interpreted. At the same time, emerging agentic AI frameworks that combine LLMs with external tools and structured reasoning offer a promising direction for building systems that can autonomously collect, verify, and contextualize event data in dynamic environments (Raheem and Hossain, 2025; Hughes et al., 2025).

In this context, the 8th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE), held at RANLP 2025, continues to serve as a critical platform for advancing interdisciplinary research at the intersection of computational methods and socio-political analysis. Building on the momentum of previous editions (Hürriyetoğlu et al., 2020, 2021b, 2022, 2023, 2024), CASE 2025 highlights innovative approaches for extracting, representing, and interpreting event information spanning traditional NLP pipelines, LLM-centric methods, and multimodal frameworks.

This edition of the workshop features regular

research papers, system descriptions from shared task participants, and keynote talks from experts across disciplines. It also includes a shared task designed to benchmark the capabilities of current systems on complex event extraction problems in multimodal scenarios. This paper provides a brief overview of the CASE 2025 workshop, outlining its themes, activities, and contributions to the broader research community.

## 2 Accepted Papers

This year, 4 regular papers were accepted. Below, we provide brief descriptions of accepted papers:

- [Nadeem et al. \(2025\)](#) investigate political bias in large language models (LLMs) with a focus on multilingual contexts, particularly across Pakistani languages. Building on the Political Compass Test (PCT), they develop a framework that extracts hidden layer activations from decoder-based models such as Mistral and DeepSeek to identify ideological leanings along economic and social axes. The authors introduce Steering Vector Ensembles (SVE), a representation-level debiasing method that aggregates layer-specific vectors derived from contrastive prompts, enabling inference-time mitigation without fine-tuning. Their experiments show that LLMs encode systematic political bias in internal representations, but SVE effectively reduces this bias, especially in socially framed prompts while preserving fluency and coherence.
- [Thapa et al. \(2025a\)](#) present an extensive survey on socio-political event extraction (SPE), analyzing how advances in natural language processing, machine learning, and LLMs are reshaping the field. The paper systematically reviews datasets, annotation frameworks, extraction methods, and evaluation strategies, highlighting both the progress and persistent challenges in capturing complex, real-world events. The authors emphasize the importance of multilingual and low-resource settings, given the global nature of socio-political events, and point to issues of reproducibility, bias, and ethical concerns in applying SPE systems at scale. They also propose future directions, including leveraging multimodal data, improving temporal and causal reasoning, and aligning event extraction systems with policy and humanitarian needs.
- [De Longueville \(2025\)](#) provides a reflective commentary on the rise of LLMs and their implications for NLP, particularly in the domain of automated socio-political event extraction. They argue that while conversational AI like ChatGPT represents both a revolution and an epiphenomenon for NLP, its significance should be contextualized within decades of technological progress, notably the advent of the Transformer architecture. The paper highlights LLMs' unprecedented zero-shot capabilities and versatility but cautions against overreliance, noting limitations such as high computational cost, hallucinations, sycophancy, and the opacity of LLM-as-a-service deployments. [De Longueville \(2025\)](#) emphasizes that despite the hype, core NLP practices such as precision/recall evaluation, gold-standard datasets, and error analysis remain essential. Ultimately, they conclude that LLMs reshape the landscape of NLP without rendering it obsolete, instead calling for a balanced integration of LLMs with established methodologies and domain-specific knowledge systems.
- [Boyd and Mitkov \(2025\)](#) present a comparative evaluation of rule-based machine translation (RBMT), neural machine translation (NMT), and LLMs for French–English translation using the Europarl corpus. The study employs BLEU and METEOR scores with bootstrap statistical testing, finding NMT, particularly Marian NMT, consistently outperforms LLMs and RBMT in both precision and semantic accuracy, while LLMs trained explicitly for translation (e.g., T5) surpass those with only emergent translation abilities (e.g., LLaMA). RBMT lags far behind in performance, though each approach shows domain-dependent strengths, with NMT best for high-precision needs and LLMs offering versatility for broader, creative applications.

## 3 Shared Task on Multimodal Content Analysis on Marginalized Sociopolitical Movements

This year's shared task explored multimodal socio-political discourse by focusing on memes, which are an increasingly popular medium for expressing

opinion, humor, and hate online. With the growing role of social media in shaping public perception, this task aims to evaluate systems’ abilities to interpret stance, hate, and humor in text-embedded images. The shared task used the PrideMM dataset proposed by [Shah et al. \(2024\)](#). The dataset employs a rigorous annotation schema ([Bhandari et al., 2023](#)), consistent with that used in our previous shared tasks ([Thapa et al., 2023, 2024](#)) on multimodal content moderation. The task is divided into four subtasks:

- **Subtask A on Hate Speech Detection:** Binary classification of memes as containing Hate Speech or No Hate Speech, using both text and image modalities.
- **Subtask B on Targets of Hate Speech Detection:** Identifies the target of hateful content as Individual, Community, Organization, or Undirected.
- **Subtask C on Topical Stance Detection:** Determines whether the meme Supports, Opposes, or is Neutral toward a marginalized movement.
- **Subtask D on Intended Humor Detection:** Binary classification of memes based on the presence or absence of Intended Humor.

This shared task advances multimodal understanding of contentious online discourse and offers new benchmarks for evaluating models in complex socio-political contexts. [Thapa et al. \(2025b\)](#) provide a detailed overview of the shared task, including participant methods, the task timeline, and a discussion of the key findings. A total of 89 participants took part in the shared task, reflecting strong engagement from the research community. We accepted 13 shared task description papers.

## 4 Future Direction

As the field of socio-political event extraction continues to evolve, future iterations of the CASE workshop aim to further broaden the scope and impact of the research community. One key direction is the continued exploration of multilingual and multimodal event extraction, recognizing the global and diverse nature of socio-political discourse. Understanding how events manifest across languages and modalities like text, image, video, and audio remains a critical challenge, particularly

in crisis monitoring and cross-cultural analysis. We aim to encourage contributions that advance the robustness, adaptability, and inclusivity of event extraction systems in these contexts.

Another important focus will be the integration of alignment techniques in LLM-based event extraction pipelines. As large language models become increasingly central to the field, understanding how to align them with domain-specific ontologies, human feedback, and real-world utility will be crucial. We are particularly interested in prompting strategies, instruction tuning, fine-grained evaluation frameworks, and the use of LLMs within agentic systems that can reason, validate, and act based on extracted event information.

In future editions, we also plan to organize more innovative and task-oriented shared tasks that reflect real-world complexities such as low-resource event extraction, multi-hop event reasoning, and cross-modal fusion. These shared tasks will continue to serve as benchmarks while also driving the development of practical solutions deployable in high-stakes environments. To support and grow the community, we are also looking to introduce mentorship opportunities for early-career researchers and students, especially those from underrepresented regions or working with low-resource languages. We plan to host dedicated mentorship sessions, community-building events, and tutorials to promote inclusion, collaboration, and knowledge transfer across domains.

## 5 Conclusion

The 8th edition of the CASE workshop highlighted significant progress in socio-political event extraction, with contributions spanning multimodal analysis, large language model applications, and fine-grained stance and hate speech detection. This year’s shared task emphasized the growing importance of understanding text-embedded images, reflecting the need to address evolving forms of online discourse. The workshop brought together researchers from diverse disciplines to tackle real-world challenges such as misinformation, polarization, and marginalization through computational methods. Looking ahead, CASE aims to remain an inclusive and interdisciplinary platform that fosters collaboration, supports innovative shared tasks, and promotes research that meaningfully contributes to the understanding of complex socio-political phenomena.

## Acknowledgments

Funding for part of this research has been provided by the European Union's Horizon Europe research and innovation programme EFRA [grant number 101093026] and ECO-Ready [grant number 101084201].

## Broader Impact

The CASE workshop has a far-reaching impact by promoting interdisciplinary research at the intersection of computational methods and socio-political analysis, encouraging the development of tools and models that can interpret complex societal discourse at scale. By addressing real-world challenges such as hate speech, misinformation, political polarization, and public sentiment, the workshop supports the creation of socially responsible technologies that inform policy, empower marginalized voices, and enhance crisis response. Through shared tasks and diverse participation, CASE promotes equitable access to research opportunities and drives forward the responsible use of NLP for social good.

## References

Jacy Reese Anthis, Ryan Liu, Sean M Richardson, Austin C Kozlowski, Bernard Koch, Erik Brynjolfsson, James Evans, and Michael S Bernstein. Position: Llm social simulations are a promising research method. In *Forty-second International Conference on Machine Learning Position Paper Track*.

Aashish Bhandari, Siddhant B Shah, Surendrabikram Thapa, Usman Naseem, and Mehwish Nasim. 2023. Crisishatemm: Multimodal analysis of directed and undirected hate speech in text-embedded images from russia-ukraine conflict. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1993–2002.

William Jock Boyd and Ruslan Mitkov. 2025. Machine translation in the ai era: Comparing previous methods of machine translation with large language models. In *Proceedings of The Eighth Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.

Yan Chen, Kate Sherren, Michael Smit, and Kyung Young Lee. 2023. Using social media images as data in social science research. *New Media & Society*, 25(4):849–871.

Bertrand De Longueville. 2025. Natural language processing vs large language models: this is the end of the world as we know it, and i feel fine. In *Proceedings of The Eighth Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.

Chenyu Hou, Gaoxia Zhu, Juan Zheng, Lishan Zhang, Xiaoshan Huang, Tianlong Zhong, Shan Li, Hanxiang Du, and Chin Lee Ker. 2024. Prompt-based and fine-tuned gpt models for context-dependent and-independent deductive coding in social annotation. In *Proceedings of the 14th learning analytics and knowledge conference*, pages 518–528.

Yuxin Hou and Junming Huang. 2025. Natural language processing for social science research: A comprehensive review. *Chinese Journal of Sociology*, 11(1):121–157.

Laurie Hughes, Yogesh K Dwivedi, Tegwen Malik, Mazen Shawosh, Mousa Ahmed Albashrawi, Il Jeon, Vincent Dutot, Mandanna Appanderanda, Tom Crick, Rahul De', et al. 2025. Ai agents and agentic systems: A multi-expert analysis. *Journal of Computer Information Systems*, pages 1–29.

Ali Hürriyetoglu, Osman Mutlu, Erdem Yörük, Farhana Ferdousi Liza, Ritesh Kumar, and Shyam Ratan. 2021a. Multilingual protest news detection - shared task 1, CASE 2021. In *Proceedings of the 4th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2021)*, pages 79–91, Online. Association for Computational Linguistics.

Ali Hürriyetoglu, Hristo Tanev, Osman Mutlu, Surendrabikram Thapa, Fiona Anting Tan, and Erdem Yörük. 2023. Challenges and applications of automated extraction of socio-political events from text (CASE 2023): Workshop and shared task report. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text*, pages 167–175, Varna, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.

Ali Hürriyetoglu, Hristo Tanev, Vanni Zavarella, Jakub Piskorski, Reyyan Yeniterzi, Osman Mutlu, Deniz Yuret, and Aline Villavicencio. 2021b. Challenges and applications of automated extraction of socio-political events from text (CASE 2021): Workshop and shared task report. In *Proceedings of the 4th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2021)*, pages 1–9, Online. Association for Computational Linguistics.

Ali Hürriyetoglu, Hristo Tanev, Vanni Zavarella, Reyyan Yeniterzi, Osman Mutlu, and Erdem Yörük. 2022. Challenges and applications of automated extraction of socio-political events from text (CASE 2022): Workshop and shared task report. In *Proceedings of the 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, pages 217–222, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Ali Hürriyetoğlu, Surendrabikram Thapa, Gökçe Uludoğan, Somaiyeh Dehghan, and Hristo Tanev. 2024. A concise report of the 7th workshop on challenges and applications of automated extraction of socio-political events from text. In *Proceedings of the 7th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2024)*, pages 248–255.

Ali Hürriyetoğlu, Erdem Yörük, Osman Mutlu, Fırat Duruşan, Çağrı Yoltar, Deniz Yüret, and Burak Gürel. 2021c. Cross-context news corpus for protest event-related knowledge base construction. *Data Intelligence*, 3(2):308–335.

Ali Hürriyetoğlu, Vanni Zavarella, Hristo Tanev, Erdem Yörük, Ali Safaya, and Osman Mutlu. 2020. [Automated extraction of socio-political events from news \(AESOPEN\): Workshop and shared task report](#). In *Proceedings of the Workshop on Automated Extraction of Socio-political Events from News 2020*, pages 1–6, Marseille, France. European Language Resources Association (ELRA).

Ariba Khan, Stephen Casper, and Dylan Hadfield-Menell. 2025. Randomness, not representation: The unreliability of evaluating cultural alignment in llms. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, pages 2151–2165.

Hannah Rose Kirk, Bertie Vidgen, Paul Röttger, and Scott A Hale. 2024. The benefits, risks and bounds of personalizing the alignment of large language models to individuals. *Nature Machine Intelligence*, 6(4):383–392.

Afrozah Nadeem, Mark Dras, and Usman Naseem. 2025. Steering towards fairness: Mitigating political stance bias in llms. In *Proceedings of The Eighth Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.

Tayiba Raheem and Gahangir Hossain. 2025. Agen-tic ai systems: Opportunities, challenges, and trustworthiness. In *2025 IEEE International Conference on Electro Information Technology (eIT)*, pages 618–624. IEEE.

Siddhant Bikram Shah, Shuvam Shiawakoti, Maheep Chaudhary, and Haohan Wang. 2024. [Meme-clip: Leveraging clip representations for multimodal meme classification](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17320–17332, Miami, Florida, USA. Association for Computational Linguistics.

Hong Shen, Tianshi Li, Toby Jia-Jun Li, Joon Sung Park, and Diyi Yang. 2023. Shaping the emerging norms of using large language models in social computing research. In *Companion Publication of the 2023 Conference on Computer Supported Cooperative Work and Social Computing*, pages 569–571.

Xiaoling Shu and Yiwan Ye. 2023. Knowledge discovery: Methods from data mining and machine learning. *Social Science Research*, 110:102817.

Surendrabikram Thapa, Surabhi Adhikari, Hristo Tanev, and Ali Hürriyetoğlu. 2025a. Challenges and applications of automated extraction of socio-political events at the age of large language models. In *Proceedings of The Eighth Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.

Surendrabikram Thapa, Farhan Jafri, Ali Hürriyetoğlu, Francielle Vargas, Roy Ka-Wei Lee, and Usman Naseem. 2023. Multimodal hate speech event detection-shared task 4, case 2023. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text*, pages 151–159.

Surendrabikram Thapa, Kritesh Rauniyar, Farhan Ahmad Jafri, Hariram Veeramani, Raghav Jain, Sandesh Jain, Francielle Vargas, Ali Hürriyetoğlu, and Usman Naseem. 2024. Extended multimodal hate speech event detection during russia-ukraine crisis-shared task at case 2024. In *7th Workshop on Challenges and Applications of Automated Extraction of Socio-Political Events from Text, CASE 2024*, pages 221–228. Association for Computational Linguistics.

Surendrabikram Thapa, Siddhant Bikram Shah, Kritesh Rauniyar, Shuvam Shiawakoti, Surabhi Adhikari, Hariram Veeramani, Kristina T. Johnson, Ali Hürriyetoğlu, Hristo Tanev, and Usman Naseem. 2025b. Multimodal hate, humor, and stance event detection in marginalized sociopolitical movements. In *Proceedings of the 8th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2025)*.

Surendrabikram Thapa, Shuvam Shiawakoti, Siddhant Bikram Shah, Surabhi Adhikari, Hariram Veeramani, Mehwish Nasim, and Usman Naseem. 2025c. Large language models (llm) in computational social science: prospects, current state, and challenges. *Social Network Analysis and Mining*, 15(1):1–30.

Surendrabikram Thapa, Hariram Veeramani, Liang Hu, Qi Zhang, Wei Wang, and Usman Naseem. 2025d. A multimodal prompt-based framework for analyzing code-mixed and low-resource memes. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 19, pages 1913–1923.

Caleb Ziems, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2024. Can large language models transform computational social science? *Computational Linguistics*, 50(1):237–291.