Overview of the Workshop on Event Extraction and Understanding: Challenges and Applications

Ali Hürriyetoğlu[1], Surendrabikram Thapa[2], Hristo Tanev[3], Laxmi Thapa[4], Surabhi Adhikari[5]
[1]Wageningen Food Safety Research, Netherlands, [2]Virginia Tech, USA,
[3]European Commission, Joint Research Centre, Italy
[4]O.P. Jindal Global University, India, [5]Columbia University, USA
[1]ali.hurriyetoglu@wur.nl, [2]sbt@vt.edu, [3]hristo.tanev@ec.europa.eu

Abstract

This paper presents an overview of the 9th Workshop on Event Extraction and Understanding: Challenges and Applications (EEUCA 2026), held in conjunction with ACL 2026. Formerly known as CASE, the workshop continues its mission of bringing together researchers from natural language processing, machine learning, computational social science, and related disciplines to advance research on event extraction and understanding. This year's edition particularly emphasized the growing influence of large language models (LLMs), multimodal learning, and weakly supervised methodologies in event extraction research. The workshop featured six regular research papers covering topics such as low-resource event extraction, reflective multi-agent architectures, symbolic auditing of procedural events, geopolitical event extraction, and generative event extraction strategies. In addition, EEUCA 2026 hosted two shared tasks focusing on toxicity detection in gaming communities and multimodal vaccine-critical meme analysis, attracting broad international participation and encouraging research on socially impactful applications of AI. The workshop highlights current advances, emerging challenges, and future directions in multilingual, multimodal, and socially aware event extraction systems.

1 Introduction

The increasing availability of large-scale digital data has transformed the study of events across social, political, economic, and public health domains (Chen et al., 2024; Thapa et al., 2025a). News articles, social media posts, online discussions, multimodal content, and user-generated media continuously document evolving real-world events, creating new opportunities for automated event extraction and understanding (Liu et al., 2020; Shah, 2016; Thapa et al., 2025b; Hey et al., 2025). At the same time, modern societal challenges such as misinformation, political polarization, online extremism, public health crises, and humanitarian emergencies have increased the need for reliable AI systems capable of identifying, structuring, and interpreting complex event information from heterogeneous data sources (Hürriyetoğlu et al., 2025).

Recent advances in large language models (LLMs), multimodal transformers, and instruction-tuned generative systems have significantly reshaped the event extraction landscape (Meng et al., 2024; Liu et al., 2025b; Chen et al., 2024). Contemporary models are increasingly capable of performing event detection, argument extraction, temporal reasoning, and cross-document event understanding with reduced supervision and improved generalization. At the same time, these systems introduce new research questions related to hallucination, interpretability, schema flexibility, multilingual robustness, and ethical deployment. In parallel, multimodal learning has expanded the scope of event analysis beyond traditional text-based pipelines by enabling systems to jointly reason over textual, visual, and contextual information in domains such as misinformation analysis, social media discourse, and crisis monitoring (Liu et al., 2025a; Li et al., 2024; Suwannahong et al., 2026; Ma et al., 2025).

Against this backdrop, the 9th Workshop on Event Extraction and Understanding: Challenges and Applications (EEUCA 2026), formerly known as CASE, continues to provide an interdisciplinary venue for researchers working at the intersection of NLP, machine learning, computational social science, and AI-driven event analytics. Building on previous editions of CASE, now EEUCA, this year's workshop places particular emphasis on multilingual event extraction, multimodal reasoning, generative AI, low-resource settings, and societally relevant applications of event understanding systems (Hürriyetoğlu et al., 2025, 2024, 2023).

EEUCA 2026 featured six regular research papers spanning a broad range of topics, including

Proceedings of the 9th Workshop on Event Extraction and Understanding: Challenges and Applications, pages 1–7

weakly supervised Vietnamese event extraction, benchmark creation for Nepali event extraction, reflective multi-agent event extraction architectures, symbolic auditing of procedural event datasets, generative event extraction training strategies, and LLM-driven geopolitical event extraction pipelines. Collectively, these papers reflect the increasing diversity of methodologies and applications emerging in the event extraction community.

In addition to regular papers, the workshop hosted two shared tasks designed to encourage research on socially relevant and multimodal problems. The first shared task focused on understanding toxic behavioral intent in gaming chat logs using the GameTox dataset, while the second addressed multimodal identification of vaccine-critical content on social media using the VaxMeme dataset. These shared tasks attracted broad international participation and highlighted the growing importance of multimodal reasoning, domain adaptation, and robust classification under severe class imbalance and noisy real-world conditions.

This overview paper summarizes the accepted papers, shared tasks, and broader research themes represented at EEUCA 2026. We further discuss emerging trends in event extraction research, including the role of LLMs, multimodal systems, weak supervision, and reflective reasoning architectures, while outlining future directions for building more reliable, interpretable, and socially responsible event understanding systems.

2 Accepted Papers

This year, 6 regular papers were accepted. Below, we provide brief descriptions of accepted papers:

Hieu et al. (2026) proposed a weakly supervised framework for Vietnamese event extraction that addresses the scarcity of annotated data by constructing a large-scale silver corpus from unlabeled news articles. Their approach first pseudo-labels Vietnamese news data using an existing BKEE event extraction model, then applies a cross-document n-ary relation filtering strategy that retains event structures consistently observed across multiple articles discussing the same topic, thereby reducing noisy annotations. The authors further improve diversity through schema-based data augmentation, where event schemas containing triggers and arguments are diversified and instantiated using LLMs to generate structurally valid synthetic training examples. Built on top of the FourIE joint event extraction architecture, their framework expands the training data to over 46,000 silver-labeled sentences and achieves consistent improvements on the Vietnamese BKEE benchmark, particularly for entity mention detection and event argument extraction.

Maharjan et al. (2026) introduced NepEE, a manually annotated benchmark dataset for low-resource Nepali event extraction, focusing on trigger phrase identification and event type classification in morphologically complex Devanagari text. The dataset contains 10,226 Nepali sentences annotated by five native speakers through a rigorous three-phase annotation protocol involving pilot annotation, instruction refinement, and conflict resolution, resulting in high inter-annotator agreement scores of Fleiss' κ = 0.812 for trigger identification and $\kappa = 0.855$ for event classification. The authors defined eight event categories, including political, economic, health, disaster, and education events, and developed detailed guidelines to capture complex Nepali trigger structures such as nominalized triggers and compound verb constructions. They further benchmarked a broad range of approaches, including classical machine learning models, multilingual and Indic-specific Transformer encoders, and instruction-tuned LLMs under zero-shot and few-shot prompting settings. Experimental results showed that Indic-specialized Transformer models achieved the strongest performance for event classification, while generative LLMs struggled with exact trigger span extraction due to Nepali's morphological complexity and ambiguity in trigger boundaries.

Rim et al. (2026) presented a symbolic audit framework for procedural event annotations by introducing Entity Qualia Structure (EQS), a lexical-semantic representation grounded in Generative Lexicon theory to distinguish semantically central entity state changes from incidental ones in procedural text. Using lexical resources such as the Brandeis Semantic Ontology, CoreLex, and WordNet, the authors categorized entities into coarse sortal types including natural, artifactual, and instrument classes, and applied this framework to the OpenPI food-domain procedural dataset. Their analysis revealed that only 51.1% of OpenPI transformation annotations corresponded to actual

food entities, while 30.2% tracked incidental instrument-related state changes such as bowls, knives, or ovens, highlighting substantial annotation noise in procedural state tracking datasets. The study further compared EQS-based filtering with prior human and LLM-based cleanup methods, demonstrating that the symbolic approach uniquely identified 15.6% of problematic annotations missed by both human re-annotation and LLM salience scoring approaches. Additionally, the authors analyzed the AGENTIVE quale feature and showed that most agentive-positive annotations involved instruments rather than food entities, emphasizing that procedural state interpretation requires compositional reasoning between entity qualia and verb semantics.

Ravikumar and Batista-Navarro (2026) conducted a systematic study of training strategies for generative event extraction, focusing on how event detection (ED) and event argument extraction (EAE) should be coordinated when fine-tuning LLMs. The authors proposed a taxonomy of seven training strategies spanning three paradigms: disjoint training, fully shared training, and hybrid parameter-sharing approaches, and evaluated them across ACE2005 and RichERE using multiple instruction-tuned LLMs ranging from 3B to 12B parameters. Their framework formulated ED, EAE, and joint event extraction as conditional text generation tasks using structured JSON outputs, enabling direct comparison between pipeline and joint generative approaches under consistent settings. Experimental results demonstrated that training strategy substantially affects extraction performance, with the strongest overall results achieved by an "ED Backward Transfer" approach that initializes event detection adapters from pretrained event argument extraction adapters. In contrast, fully joint modeling approaches that generated complete event structures in a single pass consistently underperformed, particularly for trigger classification. The study further showed that event detection benefits from cross-task transfer and partial parameter sharing, whereas argument extraction performs best with dedicated task-specific adapter capacity, highlighting the importance of carefully balancing parameter sharing and specialization in generative event extraction systems.

Tanev et al. (2026) proposed MAREA, a reflective multi-agent architecture for Semi-Open Event Extraction (SOEE) that combines fixed event schema fields with dynamically generated event attributes inferred through self-reflective reasoning using LLMs. Unlike traditional closed-schema event extraction systems, their SOEE framework preserves a core set of standardized fields such as event type, date, and location while allowing the system to iteratively expand templates with context-specific attributes generated at runtime. The proposed architecture consists of three layers: an expert layer that generates initial event templates and answers follow-up questions, a reflective layer that formulates questions to uncover missing or implicit event information, and a coordination layer that manages interactions among agents. The reflective component employs multiple strategies, including BERT-based question mapping, prompt-driven question generation, and keyword-based reasoning, to discover new event attributes beyond the predefined schema. Evaluated on health-related news articles using LLaMA-3.1-70B-Instruct, MAREA achieved strong extraction performance on core event fields such as event type, actors, disease, and mitigation measures, while also generating additional semantically relevant attributes that improved template completeness and contextual richness. The study demonstrates how reflective multi-agent reasoning can support flexible, extensible, and semantically richer event extraction beyond rigid fixed-schema approaches.

Dell'Orto and Kommandeur (2026) introduced GENOME, a continuously updated geopolitical event extraction pipeline and dataset designed to capture both conflictual and cooperative international interactions using LLMs and the PLOVER ontology. Addressing limitations in existing resources such as POLECAT (Halterman et al., 2023), the authors proposed a two-stage extraction and classification framework that leverages GPT-based one-shot prompting with enforced structured outputs to extract events from large-scale English-language newswire data. GENOME extends the traditional Actor–Recipient event representation by introducing a novel Third Party role, enabling richer multi-entity geopolitical representations that better capture complex international interactions and contextual participants. The pipeline further incorporates entity normalization and embedding-based clustering for resolving geopolitical actors,

as well as a multi-criteria deduplication module that merges duplicate reports of the same event across multiple sources. Evaluated against the POLECAT dataset over a five-month overlap period, GENOME demonstrated strong alignment on conflict-related event types while capturing a substantially more balanced distribution of cooperative and verbal interactions, particularly diplomatic consultations and agreements that were largely absent in POLECAT. The study also showed that GENOME more accurately associates events with their inferred occurrence dates rather than publication dates and provides finer-grained geopolitical entity resolution for organizations such as NATO, IMF, and WTO. Overall, the work highlights the potential of LLM-based structured extraction pipelines for building scalable and temporally grounded geopolitical event databases for international relations research and early-warning applications.

3 Shared Task on Understanding Toxic Behavioral Intent in Gaming Chat Logs

The shared task on Understanding Toxic Behavioral Intent in Gaming Chat Logs[1] addressed the challenge of fine-grained toxicity detection in online gaming communities using the GameTox dataset, a large-scale corpus of approximately 53,000 annotated chat utterances collected from the multiplayer game *World of Tanks* (Naseem et al., 2025; Thapa et al., 2026c). Participants were required to classify each utterance into one of six intent categories: *Non-toxic*, *Insults and Flaming*, *Other Offensive Texts*, *Hate and Harassment*, *Threats*, and *Extremism*, reflecting the diverse and highly imbalanced nature of toxic communication in gaming environments. The task highlighted several key challenges specific to gaming chat, including short and noisy utterances, multilingual and code-switched communication, gaming-specific slang, and severe long-tail class imbalance where high-risk categories such as threats and extremism were extremely rare. A total of 102 participants registered for the competition, with 35 teams submitting systems that explored a broad range of approaches, including domain-adaptive pretraining, multilingual transfer learning, supervised contrastive learning, token-attribution guided architectures, ensemble methods, and LLM-based synthetic data augmentation for minority classes. Systems were evaluated using macro-averaged F1-score to emphasize balanced performance across all toxicity categories, and the best-performing system achieved a Macro F1-score of 0.7041. Overall, the shared task provided a comprehensive benchmark for studying toxicity detection in gaming communities and highlighted the importance of domain adaptation, rare-class modeling, and robust multilingual learning for developing safer and healthier online gaming environments.

4 Shared Task on Multimodal Identification of Vaccine Critical Content on Social Media

The shared task on Multimodal Identification of Vaccine Critical Content on Social Media focused on detecting vaccine stance in social media memes[2] using the VaxMeme dataset, a large-scale multimodal collection of over 10,000 vaccination-related memes containing both images and associated textual content (Naseem et al., 2023; Thapa et al., 2026b,a). Participants were tasked with classifying each meme into one of three categories: *Vaccine-critical*, *Neutral*, or *Pro-vaccine*, requiring systems to jointly reason over visual cues, embedded OCR text, sarcasm, humor, and multimodal context. The task highlighted the challenges of multimodal public health misinformation analysis, where stance is often conveyed implicitly through image-text interactions, cultural references, and visual metaphors rather than explicit textual claims alone. A total of 77 participants registered for the competition, with 25 teams submitting systems that explored a wide range of approaches, including transformer-based multimodal architectures, vision-language models, cross-modal attention mechanisms, ensemble strategies, OCR-enhanced pipelines, and instruction-tuned LLMs. Systems were evaluated using macro-averaged F1-score to ensure balanced performance across stance categories despite moderate class imbalance, and the best-performing system achieved a Macro F1-score of 0.8494. Overall, the shared task provided a benchmark for multimodal vaccine stance detection and offered insights into the strengths and limitations of current multimodal AI systems for analyzing vaccine-related discourse, misinformation, and public health narratives on social media platforms.

[1] https://www.codabench.org/competitions/12083/

[2] https://www.codabench.org/competitions/12085/

5 Future Direction

The rapid evolution of LLMs, multimodal AI systems, and agentic reasoning frameworks continues to redefine the future of event extraction and understanding research. Future editions of EEUCA will further expand beyond traditional text-centric pipelines toward systems capable of integrating information across multiple modalities, languages, and sources while maintaining robustness, interpretability, and scalability.

One important future direction involves multilingual and low-resource event extraction. Despite recent progress, many languages still lack sufficiently large annotated corpora, standardized schemas, and benchmark datasets. Future research must continue to explore weak supervision, synthetic data generation, transfer learning, and cross-lingual adaptation techniques that can improve event extraction capabilities for underrepresented languages and regions. Building multilingual and culturally aware event extraction systems remains essential for ensuring equitable global coverage of socio-political and public health events.

Another major research direction concerns multimodal event understanding. Increasingly, important real-world events are communicated not only through text but also through images, videos, memes, and multimodal social media content. The success of this year's shared tasks demonstrates both the promise and the difficulty of multimodal reasoning in socially sensitive settings such as misinformation detection and online toxicity analysis. Future work should focus on more robust cross-modal alignment, multimodal temporal reasoning, sarcasm and implicit intent detection, and multimodal explainability techniques capable of identifying how visual and textual signals jointly contribute to model predictions.

The workshop also highlights growing interest in reflective and semi-open event extraction architectures powered by LLMs. Future systems may increasingly move beyond rigid fixed-schema extraction toward adaptive frameworks capable of dynamically discovering new event attributes, reasoning over incomplete information, and interacting with external knowledge sources. Agentic AI systems combining retrieval, reasoning, verification, and self-reflection may play an important role in improving event completeness, factual grounding, and temporal consistency.

Another critical challenge concerns reliability, fairness, and evaluation. While LLM-based systems have demonstrated impressive generative capabilities, they remain susceptible to hallucination, bias, instability, and poor calibration in high-stakes applications. Future research should therefore prioritize more rigorous evaluation methodologies, uncertainty estimation, bias auditing, and interpretable reasoning frameworks for event extraction systems. Human-in-the-loop evaluation, symbolic validation, and hybrid neuro-symbolic approaches may become increasingly important for ensuring trustworthy event extraction pipelines.

Future research could also explore tighter integration of event extraction with related NLP tasks highlighted in this workshop edition, such as abusive language and hate speech detection, misinformation analysis, stance detection, question answering, and sentiment analysis. Future EEUCA shared tasks will continue to focus on realistic and societally relevant challenges involving multimodal analysis, multilingual event understanding, misinformation, public health communication, and emerging online harms. We also aim to further support participation from early-career researchers and underrepresented communities through mentorship opportunities, collaborative initiatives, and accessible benchmark resources.

6 Conclusion

The 9th Workshop on Event Extraction and Understanding: Challenges and Applications (EEUCA 2026) reflects the continued growth and diversification of the event extraction research community. This year's workshop showcased advances spanning low-resource event extraction, multimodal reasoning, generative architectures, symbolic auditing, geopolitical event analysis, and socially grounded shared tasks involving gaming toxicity and vaccine-related misinformation. The accepted contributions demonstrate how modern event extraction research is increasingly shaped by LLMs, multimodal AI, and interdisciplinary applications that extend far beyond traditional information extraction settings.

The workshop further highlighted both the opportunities and the challenges introduced by rapidly evolving AI technologies. While LLMs and multimodal systems have substantially expanded the capabilities of event extraction pipelines, important questions remain regarding reliability, interpretability, fairness, multilingual inclusivity, and responsible deployment. Through its combination of

regular papers, shared tasks, and interdisciplinary collaboration, EEUCA continues to provide a platform for advancing research on robust, scalable, and socially responsible event understanding systems.

Broader Impact

EEUCA 2026 contributes to the broader advancement of socially impactful AI research by promoting event extraction technologies that support applications in public health monitoring, misinformation analysis, online safety, humanitarian response, and socio-political understanding. The workshop encourages interdisciplinary collaboration between NLP researchers, computational social scientists, and domain experts working on real-world societal challenges.

This year's shared tasks particularly emphasized socially relevant applications involving online toxicity in gaming communities and vaccine-related misinformation on social media, both of which have direct implications for digital well-being, public discourse, and public health communication. By fostering research on multilingual, multimodal, and low-resource event understanding systems, the workshop also supports the development of more inclusive AI technologies capable of addressing global and culturally diverse contexts.

At the same time, the workshop recognizes the ethical challenges associated with automated event extraction and large-scale content analysis. Event extraction systems may inherit societal biases, misinterpret context, or produce harmful outputs when deployed without appropriate safeguards. Multimodal and LLM-based systems are additionally vulnerable to hallucination, misinformation propagation, privacy concerns, and unfair treatment of marginalized groups. Accordingly, EEUCA promotes responsible AI practices, transparent evaluation methodologies, and research that prioritizes fairness, accountability, and human-centered deployment considerations in event extraction technologies.

References

Ruirui Chen, Chengwei Qin, Weifeng Jiang, and Dongkyu Choi. 2024. Is a large language model a good annotator for event extraction? In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pages 17772–17780.

Alessandro Dell'Orto and Jesse Kommandeur. 2026. Genome: A new geopolitical event methodology and dataset using large language models. In *Proceedings of the 9th Workshop on Event Extraction and Understanding: Challenges and Applications (EEUCA)*.

Andrew Halterman, Benjamin E Bagozzi, Andreas Beger, Phil Schrodt, and Grace Scraborough. 2023. Plover and polecat: A new political event ontology and dataset. In *International Studies Association Conference Paper*.

Spencer Phillips Hey, Julie Walsh, and Eni Mustafaraj. 2025. Comparing human and llm ethical analyses: A case study in computational social science research. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 8, pages 1245–1254.

Pham Xuan Hieu, Tuan Vu Minh, Mai-Vu Tran, and Hoang-Quynh Le. 2026. Constructing a silver corpus for weakly supervised vietnamese event extraction using cross-document n-ary relation filtering. In *Proceedings of the 9th Workshop on Event Extraction and Understanding: Challenges and Applications (EEUCA)*.

Ali Hürriyetoğlu, 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 Hürriyetoğlu, Surendrabikram Thapa, Hristo Tanev, and Surabhi Adhikari. 2025. Findings and insights from the 8th workshop on challenges and applications of automated extraction of socio-political events from text. In *Proceedings of the 8th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Texts*, pages 1–5.

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.

Haoxuan Li, Zhengmao Yang, Yunshan Ma, Yi Bin, Yang Yang, and Tat-Seng Chua. 2024. Mm-forecast: a multimodal approach to temporal event forecasting with large language models. In *Proceedings of the 32nd ACM international conference on multimedia*, pages 2776–2785.

Shaoyu Liu, Jianing Li, Guanghui Zhao, Yunjian Zhang, Xin Meng, Fei Richard Yu, Xiangyang Ji, and Ming Li. 2025a. Eventgpt: Event stream understanding with multimodal large language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 29139–29149.

Tianpeng Liu, Feng Xue, Jian Sun, and Xiao Sun. 2020. A survey of event analysis and mining from social multimedia. *Multimedia Tools and Applications*, 79(45):33431–33448.

Wenxuan Liu, Zixuan Li, Long Bai, Yuxin Zuo, Daozhu Xu, Xiaolong Jin, Jiafeng Guo, and Xueqi Cheng. 2025b. Towards event extraction with massive types: Llm-based collaborative annotation and partitioning extraction. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 34365–34387.

Junxiao Ma, Jingjing Wang, Jiamin Luo, Peiying Yu, and Guodong Zhou. 2025. Sherlock: Towards multi-scene video abnormal event extraction and localization via a global-local spatial-sensitive llm. In *Proceedings of the ACM on Web Conference 2025*, pages 4004–4013.

Sujal Maharjan, Astha Shrestha, Lakshmojee Koduru, Sweta Poudel, Shuvam Shiwakoti, Rabin Thapa, Kritesh Rauniyar, and Surendrabikram Thapa. 2026. Benchmarking models for low-resource nepali event extraction with trigger phrase identification and event classification. In *Proceedings of the 9th Workshop on Event Extraction and Understanding: Challenges and Applications (EEUCA)*.

Zihao Meng, Tao Liu, Heng Zhang, Kai Feng, and Peng Zhao. 2024. Cean: Contrastive event aggregation network with llm-based augmentation for event extraction. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 321–333.

Usman Naseem, Jinman Kim, Matloob Khushi, and Adam G Dunn. 2023. A multimodal framework for the identification of vaccine critical memes on twitter. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, pages 706–714.

Usman Naseem, Shuvam Shiwakoti, Siddhant Bikram Shah, Surendrabikram Thapa, and Qi Zhang. 2025. GameTox: A comprehensive dataset and analysis for enhanced toxicity detection in online gaming communities. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 440–447, Albuquerque, New Mexico. Association for Computational Linguistics.

Rishi Ravikumar and Riza Batista-Navarro. 2026. When tasks share structure: A comparative study of training strategies for generative event extraction. In *Proceedings of the 9th Workshop on Event Extraction and Understanding: Challenges and Applications (EEUCA)*.

Kyeongmin Rim, Marc Verhagen, and James Pustejovsky. 2026. A qualia-based audit of procedural event annotations. In *Proceedings of the 9th Workshop on Event Extraction and Understanding: Challenges and Applications (EEUCA)*.

Rajiv Ratn Shah. 2016. Multimodal analysis of user-generated content in support of social media applications. In *Proceedings of the 2016 ACM on International Conference on Multimedia Retrieval*, pages 423–426.

Phattaradanai Suwannahong, Piphatpong Wannapanon, Aphisara Klayburee, Takafumi Nakanishi, and Ponlawat Chophuk. 2026. Explainable multimodal earthquake event classification using aime attribution and llm-based narrative reasoning. In *2026 18th International Conference on Knowledge and Smart Technology (KST)*, pages 598–603. IEEE.

Hristo Tanev, Michel de Bollivier, and Bertrand De Longueville. 2026. A self-reflective llm-based architecture for semi-open event extraction. In *Proceedings of the 9th Workshop on Event Extraction and Understanding: Challenges and Applications (EEUCA)*.

Laxmi Thapa, Aryaman Jain, Lakshmojee Koduru, Surabhi Adhikari, Junaid Rashid, Jungeun Kim, Surendrabikram Thapa, and Usman Naseem. 2026a. Concept-grounded detection of vaccine misinformation in multimodal content using interpretable vision-language models. In *Companion Proceedings of the ACM on Web Conference 2026*.

Surendrabikram Thapa, Shuvam Shiwakoti, Siddhant Bikram Shah, Surabhi Adhikari, Hariram Veeramani, Mehwish Nasim, and Usman Naseem. 2025a. Large language models (llm) in computational social science: prospects, current state, and challenges. *Social Network Analysis and Mining*, 15(1):4.

Surendrabikram Thapa, Shuvam Shiwakoti, Siddhant Bikram Shah, Kritesh Rauniyar, Laxmi Thapa, Surabhi Adhikari, Kristina T. Johnson, Ali Hürriyetoğlu, Hristo Tanev, and Usman Naseem. 2026b. Multimodal identification of vaccine content stance on social media. In *Proceedings of the 9th Workshop on Event Extraction and Understanding: Challenges and Applications (EEUCA)*.

Surendrabikram Thapa, Shuvam Shiwakoti, Siddhant Bikram Shah, Kritesh Rauniyar, Laxmi Thapa, Surabhi Adhikari, Kristina T. Johnson, Ali Hürriyetoğlu, Hristo Tanev, and Usman Naseem. 2026c. Understanding toxic behavior in gaming communities using ai to promote healthier digital spaces. In *Proceedings of the 9th Workshop on Event Extraction and Understanding: Challenges and Applications (EEUCA)*.

Surendrabikram Thapa, Hariram Veeramani, Liang Hu, Qi Zhang, Wei Wang, and Usman Naseem. 2025b. 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.