



LREC 2026

**The 3rd Workshop on Natural Language Processing
for Political Sciences (PoliticalNLP 2026)
@ LREC 2026**

Workshop Proceedings

Editors

**Haithem Afli, Houda Bouamor, Wajdi Zaghouani,
Sahar Ghannay, Shehenaz Hossain**

16 May 2026

Proceedings of The 3rd Workshop on Natural Language Processing for Political Sciences
(PoliticalNLP 2026) @ LREC 2026

© ELRA Language Resources Association (ELRA), 2026
These proceedings are licensed under a Creative Commons Attribution-
NonCommercial 4.0 International Licence (CC BY-NC 4.0)

ISBN 978-2-493814-76-0

Preface

PoliticalNLP 2026 is the third edition of the international workshop dedicated to the intersection of Natural Language Processing, political science, legal studies, media analysis, and computational social science. Following the success of PoliticalNLP 2022 at LREC in Marseille and PoliticalNLP 2024 at LREC-COLING in Torino, this third edition takes place on 16 May 2026 at the Palau de Congressos de Palma, Mallorca, Spain, as a half-day hybrid workshop co-located with the Joint International Conference on Language Resources and Evaluation (LREC 2026).

This year's workshop is organised around the theme of *Trust, Transparency & Generative AI in Political Discourse Analysis*. Large Language Models (LLMs) and generative AI systems have rapidly transformed the landscape of computational political analysis. On the one hand, they offer unprecedented capabilities for processing, interpreting, and generating political text at scale; on the other, they introduce pressing concerns regarding bias, cultural representation, explainability, misinformation amplification, and the integrity of democratic discourse.

The workshop received submissions addressing a broad range of topics, reflecting both the depth and diversity of contemporary political NLP research. These include the analysis of political stances, narratives, and media framing in news outlets and social media platforms; the capabilities and limitations of LLMs for cross-lingual political discourse analysis; AI policy, governance, and regulatory frameworks; parliamentary and legislative NLP; hate speech detection and civil discourse modelling; culturally-aware and culturally-adaptive language models; and methodological questions of reproducibility and transparency in computational social science.

The accepted papers reflect the growing maturity and breadth of the field. They draw on a rich variety of corpora, languages, and methodological approaches, and their findings carry direct relevance to questions of political communication, governance, and civic technology. Several contributions engage explicitly with the ethical dimensions of deploying NLP tools in politically sensitive contexts and collectively they signal a community that is increasingly attentive to the real-world stakes of its methods and findings.

We extend our sincere gratitude to all authors who submitted their work to PoliticalNLP 2026. The quality and diversity of this year's submissions were particularly encouraging. We are deeply indebted to the members of the Programme Committee, whose careful and thorough reviews were indispensable to the selection process. We thank the LREC 2026 organising committee for their continued support in hosting this workshop, and ELRA for their commitment to publishing and disseminating open language resources research. Finally, we are grateful to all participants — both in-person in Mallorca and joining remotely — whose engagement and intellectual curiosity make PoliticalNLP a vibrant and welcoming research community.

We hope that these proceedings serve as a valuable reference and a catalyst for continued collaboration across the disciplines of NLP, political science, media studies, legal informatics, and computational social science. We look forward to a stimulating day of discussion.

Haithem Afli
General Chair, PoliticalNLP 2026
ADAPT Centre, Munster Technological University, Ireland

Organising Committee

General Chair: Haithem Afli, ADAPT Centre, Munster Technological University, Ireland

Programme Chairs: – Houda Bouamor, Carnegie Mellon University, Qatar

- Wajdi Zaghouani, Northwestern University in Qatar, Qatar
- Sahar Ghannay, Université Paris-Saclay, CNRS, LISN, France
- Shehenaz Hossain, Munster Technological University, Ireland

Programme Committee: – Bruno Andrade, Munster Technological University, Ireland

- Lenka Dražanová, European University Institute, Italy
- Georgios Stavropoulos, Centre for Research and Technology Hellas, Greece
- Andrea Iana, University of Mannheim, Germany
- Valentin Barrieren, Télécom ParisTech, France
- Mohammed Hasanuzzaman, Munster Technological University, Ireland
- Patrick Paroubek, Université Paris-Saclay, CNRS, LISN, France
- Suman Adhya, Indian Association for the Cultivation of Science, India
- Valentin Barriere, Centro Nacional de Inteligencia Artificial, Chile

Table of Contents

<i>From News Streams to Narrative Intelligence Briefs: LLM-Assisted Political Discourse Analysis in the Hungarian 2026 Pre-Election Context</i>	
Ekaterina Loginova, Maksim Ermakov and Stephan Khramov	1
<i>Decoding News Narratives: A Critical Analysis of Large Language Models in Framing Detection</i>	
Valeria Pastorino, Jasivan Alex Sivakumar and Nafise Sadat Moosavi	17
<i>From Cairo to Cape Town: How African Twitter Shapes the Global Palestine-Israel Narrative</i>	
Mahmoud Fawzi, Houda Elmimouni and Walid Magdy	29
<i>Multimodal Analysis of State-Funded News Coverage of the Israel–Hamas War on YouTube Shorts</i>	
Daniel Miehling and Sandra Kübler	35
<i>Characterization of User Engagement in Electronic News Media: A Case Study for India</i>	
Abhishek Kumar, Abhik Jana and Debi Prosad Dogra	48
<i>Analyzing Political Stances on Twitter/X in the Lead-up to the 2024 U.S. Election</i>	
Hazem Ibrahim, Farhan Kamrul Khan, Yasir Zaki and Talal Rahwan	58
<i>Giving Voice to the Constitution: Low-Resource Text-to-Speech for Quechua and Spanish Using a Bilingual Legal Corpus</i>	
John E. Ortega, Rodolfo Joel Zevallos and Fabrício Carraro	64
<i>Analysing Calls to Order in German Parliamentary Debates</i>	
Nina Smirnova, Daniel Dan and Philipp Mayr	70
<i>Operationalising the "Right to Be Forgotten" in LLMs: A Lightweight Sequential Unlearning Framework for Privacy-Aligned Deployment in Politically Sensitive Environments</i>	
Esen Kurt and Haithem Afli	77
<i>Exploring Two Decades of Parliamentary Speeches on the Use of Narratives</i>	
Matti Wiegmann, Jürgen Neyer and Benno Stein	87
<i>Supercharging Agenda Setting Research: The ParlaCAP Dataset of 28 European Parliaments and a Scalable Multilingual LLM-Based Classification</i>	
Taja Kuzman Pungeršek, Peter Rupnik, Daniela Širinić and Nikola Ljubešić	94
<i>A Vocabulary Analysis of News Articles in Relation to the Political Orientation of Their Source and Their Thematic</i>	
Laurène Cave and Gaël Lejeune	111
<i>Posts Talk Policy, Stories Don't: Policy-Issue Detection on Instagram with Fine-Tuned Transformers and Prompted LLMs</i>	
Michael Achmann-Denkler, Mario Haim and Christian Wolff	120
<i>Mobilize, Inform, Interact: Classifying Political Calls-to-Action Types on Instagram</i>	
Michael Achmann-Denkler, Clara Helmig, Jakob Fehle, Mario Haim and Christian Wolff	131

<i>Hate Speech and Hate Crime: A Cross-Disciplinary Analysis of Xenophobia in Greece</i> Maria Gavriilidou, Vasiliki Georgiadou, Lamprini Rori and Maria Pontiki	139
<i>Navigating Global AI Regulation: A Multi-Jurisdictional Retrieval-Augmented Generation System</i> Courtney Ford, Ojas Rane and Susan Leavy	149
<i>Beyond Sentiment: Comparing Traditional NLP and LLM-Based Multi-Dimensional Analysis for Political News Evaluation</i> Maryam Fooladi and Federico Bottino	159
<i>Evaluating the Abilities of LLMs and SpeechLMs in Discovering Implicit Contents of Italian Political Speeches</i> Lorenzo Gregori, Walter Paci and Alessandro Panunzi	165
<i>Reproducibility under Threat: Proposing a Framework for Reliable LLM-Research in Psychology and Computational Social Science</i> Kevin Dirk Kiy, Alexander Porshnev, Dounia Lakhzoum, Dermot Lynott, Diarmuid O'Donoghue and Manokamna Singh	171
<i>How Far Did They Go? The Persuasive Tactics of Covert LLM Agents in a Discontinued Field Experiment</i> Kokil Jaidka and Saifuddin Ahmed	180
<i>An Examination of the Party Leanings of Large Language Models</i> Lars Bungum, Charles Huang and Jari Bakken	195
<i>Attitude Identification through Parameter-Efficient Fine-Tuning</i> Mariia Anisimova, Gabriella Lapesa and Sarka Zikanova	204
<i>Do LLMs Transfer Political Framing across Languages? A Cross-Lingual Analysis of LLM-Generated Discourse</i> Nooredeen Awwad, Ebtihal Ismail Enfes and Kolawole John Adebayo	210
<i>When Neutral Turns Negative: Cross-Domain Failure Modes in Hinglish Political Sentiment Analysis</i> Rahul Chennuru and Kolawole John Adebayo	219
<i>Identifying Political Bias in Arabic News Articles</i> Saoussen Chaabane, Omar Trigui and Maher Jaoua	228
<i>PReSS: An Automated Black-Box Framework for Evaluating Political Stance Stability in LLMs</i> Shariar Kabir, Yue Dong and Kevin Esterling	234
<i>Large Language Models Unpack Complex Political Opinions through Target-Stance Extraction</i> Ozgur Togay, Javier Garcia-Bernardo, Florian Kunneman and Anastasia Giachanou ..	248
<i>Can Large Language Models Facilitate Qualitative Political Narrative Analysis?</i> Luke Stephens, Clare Llewellyn, Lauren Rogers, Constantine Kyritsopoulos, Arman Prangere, Feiteng Long, Peyton Snyder and Laura Cram	261
<i>Cultural Adaptation in Large Language Models for Political Discourse</i> Wajdi Zaghouani	272

<i>Automated Analysis of Global AI Safety Initiatives: A Taxonomy-Driven LLM Approach</i>	
Takayuki Semitsu, Naoto Kiribuchi and Kengo Zenitani	284

Workshop Program

2026-05-16

- 09:00–09:10 Session OPE:**
Room: Room 2
Chair: Haithem Afli
- 09:00–09:10 *Opening Remarks*
Haithem Afli (General Chair)
- 09:10–10:30 Session S1: Political Narratives, Media Framing and News Discourse**
Room: Room 2
Chair: Wajdi Zaghouani
- 09:10–09:15 *Session 1 Opening – Chair Introduction*
Wajdi Zaghouani
- 09:15–09:27 *From News Streams to Narrative Intelligence Briefs: LLM-Assisted Political Discourse Analysis in the Hungarian 2026 Pre-Election Context*
Ekaterina Loginova, Maksim Ermakov and Stephan Khramov
- 09:27–09:39 *Decoding News Narratives: A Critical Analysis of Large Language Models in Framing Detection*
Valeria Pastorino, Jasivan Alex Sivakumar and Nafise Sadat Moosavi
- 09:39–09:51 *From Cairo to Cape Town: How African Twitter Shapes the Global Palestine-Israel Narrative*
Mahmoud Fawzi, Houda Elmimouni and Walid Magdy
- 09:51–10:03 *Multimodal Analysis of State-Funded News Coverage of the Israel–Hammas War on YouTube Shorts*
Daniel Miebling and Sandra Kübler
- 10:03–10:15 *Characterization of User Engagement in Electronic News Media: A Case Study for India*
Abhishek Kumar, Abhik Jana and Debi Prosad Dogra
- 10:15–10:30 *Session Discussion and Transition to Poster Area*
Wajdi Zaghouani

2026-05-16 (continued)

- 10:30–11:30 Session PS1: Poster Session and Coffee Break**
Room: Poster Area
Chair: Haithem Afli
- 10:30–11:30 *Poster Session and Coffee Break (20 papers)***
- 10:30–11:30 *Analyzing Political Stances on Twitter/X in the Lead-up to the 2024 U.S. Election*
Hazem Ibrahim, Farhan Kamrul Khan, Yasir Zaki and Talal Rahwan
- 10:30–11:30 *Giving Voice to the Constitution: Low-Resource Text-to-Speech for Quechua and Spanish Using a Bilingual Legal Corpus*
John E. Ortega, Rodolfo Joel Zevallos and Fabrício Carraro
- 10:30–11:30 *Analysing Calls to Order in German Parliamentary Debates*
Nina Smirnova, Daniel Dan and Philipp Mayr
- 10:30–11:30 *Operationalising the “Right to Be Forgotten” in LLMs: A Lightweight Sequential Unlearning Framework for Privacy-Aligned Deployment in Politically Sensitive Environments*
Esen Kurt and Haithem Afli
- 10:30–11:30 *Exploring Two Decades of Parliamentary Speeches on the Use of Narratives*
Matti Wiegmann, Jürgen Neyer and Benno Stein
- 10:30–11:30 *Supercharging Agenda Setting Research: The ParlaCAP Dataset of 28 European Parliaments and a Scalable Multilingual LLM-Based Classification*
Taja Kuzman Pungeršek, Peter Rupnik, Daniela Širinić and Nikola Ljubešić
- 10:30–11:30 *A Vocabulary Analysis of News Articles in Relation to the Political Orientation of Their Source and Their Thematic*
Laurène Cave and Gaël Lejeune
- 10:30–11:30 *Posts Talk Policy, Stories Don’t: Policy-Issue Detection on Instagram with Fine-Tuned Transformers and Prompted LLMs*
Michael Achmann-Denkler, Mario Haim and Christian Wolff
- 10:30–11:30 *Mobilize, Inform, Interact: Classifying Political Calls-to-Action Types on Instagram*
Michael Achmann-Denkler, Clara Helmig, Jakob Fehle, Mario Haim and Christian Wolff
- 10:30–11:30 *Hate Speech and Hate Crime: A Cross-Disciplinary Analysis of Xenophobia in Greece*
Maria Gavriilidou, Vasiliki Georgiadou, Lamprini Rori and Maria Pontiki

2026-05-16 (continued)

- 10:30–11:30 *Navigating Global AI Regulation: A Multi-Jurisdictional Retrieval-Augmented Generation System*
Courtney Ford, Ojas Rane and Susan Leavy
- 10:30–11:30 *Beyond Sentiment: Comparing Traditional NLP and LLM-Based Multi-Dimensional Analysis for Political News Evaluation*
Maryam Fooladi and Federico Bottino
- 10:30–11:30 *Evaluating the Abilities of LLMs and SpeechLMs in Discovering Implicit Contents of Italian Political Speeches*
Lorenzo Gregori, Walter Paci and Alessandro Panunzi
- 10:30–11:30 *Reproducibility under Threat: Proposing a Framework for Reliable LLM-Research in Psychology and Computational Social Science*
Kevin Dirk Kiy, Alexander Porshnev, Dounia Lakhzoum, Dermot Lynott, Diarmuid O'Donoghue and Manokamna Singh
- 10:30–11:30 *How Far Did They Go? The Persuasive Tactics of Covert LLM Agents in a Discontinued Field Experiment*
Kokil Jaidka and Saifuddin Ahmed
- 10:30–11:30 *An Examination of the Party Leanings of Large Language Models*
Lars Bungum, Charles Huang and Jari Bakken
- 10:30–11:30 *Attitude Identification through Parameter-Efficient Fine-Tuning*
Mariia Anisimova, Gabriella Lapesa and Sarka Zikanova
- 10:30–11:30 *Do LLMs Transfer Political Framing across Languages? A Cross-Lingual Analysis of LLM-Generated Discourse*
Nooredeen Awwad, Ebtihal Ismail Enfes and Kolawole John Adebayo
- 10:30–11:30 *When Neutral Turns Negative: Cross-Domain Failure Modes in Hinglish Political Sentiment Analysis*
Rahul Chennuru and Kolawole John Adebayo
- 10:30–11:30 *Identifying Political Bias in Arabic News Articles*
Saoussen Chaabane, Omar Trigui and Maher Jaoua

2026-05-16 (continued)

11:30–13:00	Session S2: LLMs for Political Analysis: Capabilities, Biases and Governance Room: Room 2 Chair: Houda Bouamor
11:30–11:33	<i>Session 2 Opening – Chair Introduction</i> Houda Bouamor
11:33–11:45	<i>PReSS: An Automated Black-Box Framework for Evaluating Political Stance Stability in LLMs</i> Shariar Kabir, Yue Dong and Kevin Esterling
11:45–11:57	<i>Large Language Models Unpack Complex Political Opinions through Target-Stance Extraction</i> Ozgur Togay, Javier Garcia-Bernardo, Florian Kunneman and Anastasia Giachanou
11:57–12:09	<i>Can Large Language Models Facilitate Qualitative Political Narrative Analysis?</i> Luke Stephens, Clare Llewellyn, Lauren Rogers, Constantine Kyritsopoulos, Arman Prangere, Feiteng Long, Peyton Snyder and Laura Cram
12:09–12:21	<i>Cultural Adaptation in Large Language Models for Political Discourse</i> Wajdi Zaghouani
12:21–12:33	<i>Automated Analysis of Global AI Safety Initiatives: A Taxonomy-Driven LLM Approach</i> Takayuki Semitsu, Naoto Kiribuchi and Kengo Zenitani
12:33–13:00	<i>Closing Discussion, Best Paper Announcement and Workshop Closing Remarks</i> Haithem Afli