

ClimateNLP 2026

**The 3rd Workshop on Natural Language Processing Meets
Climate Change**

Proceedings of the Workshop

October 28, 2026

©2026 Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
317 Sidney Baker St. S
Suite 400 - 134
Kerrville, TX 78028
USA
Tel: +1-855-225-1962
acl@aclweb.org

ISBN 979-8-89176-465-1

Introduction

We are delighted to welcome you to ClimateNLP 2026, the third Workshop on Natural Language Processing Meets Climate Change. This year, the workshop is co-located with EMNLP 2026 and takes place on October 28, 2026, in the beautiful city of Budapest, Hungary.

ClimateNLP aims to be the premier publication venue for research at the intersection of Natural Language Processing (NLP) and climate change. The workshop brings together researchers and practitioners to discuss how NLP methods can be incorporated into climate change science and climate change action. This year, the program includes four keynote talks by Chung-Chi Chen (National Institute of Informatics), Ramit Debnath (University of Cambridge), Maria João Sousa (Climate Change AI), and Nelly Bencomo (Durham University). Furthermore, we hold a panel discussion on future directions for ClimateNLP. Four oral paper presentations and two poster sessions round up the day.

We received 37 submissions this year and recruited 38 active Program Committee (PC) members who are experts in NLP, climate change, or both. The vast majority of submissions received two reviews. When making our selections for the program, we carefully considered the reviews and held extensive discussions among the members of the organizing committee. The members of the Program Committee did an excellent job in reviewing the submitted papers, and we thank them for their essential role in selecting the accepted papers and helping produce a high-quality program for the workshop. In line with our goal of discussing and learning about the intersection of NLP and climate change, we aimed to create an inclusive program that accommodates as many favorably rated papers as possible. We accepted 24 papers (acceptance rate 64.9%), 21 of which are included in these proceedings. The accepted papers span a broad range of topics, including climate-related benchmarks and datasets, information extraction from corporate sustainability and ESG disclosures, analysis of climate discourse in news and social media, climate adaptation, and the efficiency and environmental footprint of language models themselves.

We thank our Program Committee members for committing their time to help us select an excellent technical program.

We thank all the authors who submitted to the workshop and all workshop participants for making ClimateNLP 2026 a success and for growing the research area of NLP for climate change with their fine work.

The ClimateNLP 2026 Organizing Committee

Organizing Committee

Program Chairs

Tobias Schimanski, University of Zurich, Switzerland
Aida Usmanova, Leuphana University Lüneburg, Germany
Gaku Morio, Hitachi, Ltd. and Hitotsubashi University, Japan
Jingwei Ni, ETH Zurich, Switzerland
Kalyan Dutia, Climate Policy Radar, UK
Seyed Saeid Vaghefi, University of Zurich, Switzerland
Dominik Stambach, Princeton University, USA
Markus Leippold, University of Zurich, Switzerland

Program Committee

Reviewers

Shafiuddin Rehan Ahmed, Ekant Muljibhai Amin

Nan Bai, Jetske Bonenkamp, Marco Bronzini

Henry Franks

Nupoor Gandhi, Dario Garigliotti, Francesca Grasso

Qi He, Angel Hsu

Aditya Jain, Eshaan Jain

Nikhil Avinash Kharwandikar, Sai Koneru, Mogana Kumaran, Murathan Kurfali

Kai-Robin Lange, Valentin Lemaire, Yucheng Lu

Diego Manya

Yusuke Ohtsubo

Srijith Ravikumar, Tim Repke, Jonas Rieger, Harri Rowlands

Anna Steinberg Schulten, Anne J Sietsma, Manfred Stede

Tommy Tran, Mark Tyrrell

Tejesh Vaish

Azmine Touseh Wasi, Hendrik Weichel

Lexing Xie

Zukang Yang

Shorouq Zahra, Ji Zhang

Table of Contents

<i>How Loud Rumbles Hit Newsstands: A Data Analysis of Coverage and Spatial Bias in German News about Landslides Around the World</i>	
Brielen Madureira, Andreas Niekler, Marc Keuschnigg and Mariana Madruga de Brito	1
<i>CA-ClimateNews: A Human-Annotated Multilingual Benchmark for Climate-Relevance Classification of Central Asian News</i>	
Adilet Metinov, Zarina Uvalieva, Bektemir Kumarbai Uulu, Nurtilek Alibekov and Tynchtykbek Tashbaltaev	20
<i>Climatextract: A Systematic Prompt Evaluation For Climate Information Extraction</i>	
Anna Steinberg Schulten, Malte Schierholz, Lisa Reichenbach, Frauke Kreuter and Alexander Fraser	29
<i>Who Warmed the Archives? LLMs Overestimate Historical Warmth</i>	
Claudiu Creanga and Liviu P Dinu	41
<i>Improving Efficiency of Ternary Large Language Models through Pruning</i>	
Pietro Tropeano, Maria Maistro, Tuukka Ruotsalo, Christina Lioma and Nikolaos Aletras	51
<i>Does Tropical Adaptation Help Plant Language Models Generalize? A Split-Sensitivity Evaluation</i>	
Gabriela Morales Soto and Helio Pedrini	85
<i>Extracting information on public water shortages using news media and LLMs: a pilot study in Sweden</i>	
Jeanne Fernandez, Shorouq Zahra, Mariana Madruga de Brito and Johanna Mård	99
<i>Ontology-Guided Query Expansion for Evidence Retrieval in ESRS E1 Climate Disclosures</i>	
Tatsuya Ishisaka	116
<i>The Climate Contracting Knowledge Graph: Construction and Evaluation</i>	
Georgia Ray and Felix Cohen	124
<i>Beyond Dictionaries: Machine-Learned Normalization of Emission Units in Information Extraction</i>	
Lisa Reichenbach, Simon Oehler, Malte Schierholz and Anna Steinberg Schulten	132
<i>FoodClimNER: A Manually Annotated Corpus for Food Practice Entities Related to Climate Change</i>	
Mathieu Lăi-king, Arnaud Ferré, Louise Deléger, Yann Le Bodo and Estelle Chaix	143
<i>Fine-Tuned Encoders versus Zero-Shot Prompting for Detecting Threat and Solution Framing in German Climate News</i>	
Raven Adam, Marie Lisa Kogler and David Maier	158
<i>Retrieving Climate Change Adaptation Concepts from World Heritage Management Plans in the Netherlands</i>	
Vedant Jumle, Orhan Agaoglu, Laurens Mol, Edgar Salas Gironés and Nan Bai	170
<i>ESG-CID-Plus: A Stratified Benchmark for Disclosure Content Index Table Extraction from Corporate Sustainability Reports</i>	
Shafiuddin Rehan Ahmed, Reza Abdi, Tianyu Luan, Teja Kanchinadam, Christopher Clarke and Eugene Siow	183
<i>Evidence-Grounded Auditing of Identification Assumptions in Climate-Policy Causal Evaluations</i>	
Yonghong Zhang, Yong Xie, Isabel M. Parra and Ricardo Correia	199

<i>The Energy Footprint of LLM-Based Environmental Analysis: LLMs and Domain Products</i>	
Alicia Bao, Angel Hsu, Diego Manya, Ji Zhang and Jiamian He	220
<i>Can Multimodal Large Language Models Identify Thematic Maps in Scientific Articles?</i>	
Tomas Turek, Vincent Nguyen, Dharendra Singh, Sarvnaz Karimi and Xiuzhen Zhang	231
<i>AquaCrisis: A Multilingual Benchmark for Classifying Water-Utility Public Communication</i>	
Vishal Kumar and Mourad Oussalah	246
<i>Climate and Migration on German Political Party YouTube: A Corpus with Stance, Propaganda, and Divisiveness Annotation</i>	
Zaher Alkaei, Manfred Stede and Christoph Maximilian Abels	282
<i>Can Vision-Language Models Assess Environmental Reporting Compliance? A Benchmark for Requirement-Level Disclosure Alignment</i>	
Vaishali Aggarwal, Ishan Verma, Tanay Kumar Gupta and Tirthankar Dasgupta	298
<i>Proportionate Responses: The Case for Small Language Models as Sustainable Expert Climate Policy Annotators</i>	
Jetske Bonenkamp, Robbert Biesbroek and Ioannis N. Athanasiadis	315

Program

Wednesday, October 28, 2026

09:00 - 09:05	<i>Opening Remarks</i>
09:05 - 09:30	<i>Keynote Speech 1: Beyond 2030: Turning Sustainability Commitments into Verifiable and Achievable Outcomes (Chung-Chi Chen)</i>
09:30 - 09:45	<i>Oral Presentation 1</i> <i>Improving Efficiency of Ternary Large Language Models through Pruning</i> Pietro Tropeano, Maria Maistro, Tuukka Ruotsalo, Christina Lioma and Nikolaos Aletras
09:45 - 10:30	<i>Poster Session 1</i>
10:30 - 11:00	<i>Coffee Break</i>
11:00 - 11:05	<i>Session 2 Introduction</i>
11:05 - 11:30	<i>Keynote Speech 2: Quantifying Climate Narratives, Public Sentiment, and Policy Reception via Computational Social Science (Ramit Debnath)</i>
11:30 - 11:45	<i>Oral Presentation 2</i> <i>Ontology-Guided Query Expansion for Evidence Retrieval in ESRS E1 Climate Disclosures</i> Tatsuya Ishisaka
11:45 - 12:30	<i>Panel Discussion - Future Directions for ClimateNLP with Chung-Chi Chen, Ramit Debnath, Maria João Sousa</i>
12:30 - 14:00	<i>Lunch Break</i>
14:00 - 14:05	<i>Session 3 Introduction</i>
14:05 - 14:30	<i>Keynote Speech 3: Public Interest Technology and Climate Change AI (Maria João Sousa)</i>
14:30 - 14:45	<i>Oral Presentation 3</i>

Wednesday, October 28, 2026 (continued)

Retrieving Climate Change Adaptation Concepts from World Heritage Management Plans in the Netherlands

Vedant Jumle, Orhan Agaoglu, Laurens Mol, Edgar Salas Gironés and Nan Bai

14:45 - 15:30 *Poster Session 2*

15:30 - 16:00 *Coffee Break*

16:00 - 16:05 *Session 4 Introduction*

16:05 - 16:30 *Keynote Speech 4: Beyond Accuracy: Bridging the Reasoning Gap in AI-Assisted Climate Law (Nelly Bencomo)*

16:30 - 16:45 *Oral Presentation 4*

Proportionate Responses: The Case for Small Language Models as Sustainable Expert Climate Policy Annotators

Jetske Bonenkamp, Robbert Biesbroek and Ioannis N. Athanasiadis

16:45 - 17:00 *Closing Remarks by Markus Leippold*