

NLP4PI 2026

The Fifth Workshop on NLP for Positive Impact

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 979-8-89176-495-8

Preface

Welcome to the Fifth Workshop on NLP for Positive Impact (NLP4PI 2026), held on October 28, 2026 in Budapest, Hungary, co-located with EMNLP 2026.

The NLP4PI workshop brings together researchers and practitioners who use natural language processing to benefit society, and who study the ethical and societal implications of language technologies. Every submission to the workshop was asked to discuss the societal impact of the work.

This year the workshop received 76 direct submissions and 12 commitments of papers reviewed through ACL Rolling Review. This volume contains the 45 archival papers accepted to the workshop. Their topics include health and mental-health support, education, accessibility, low-resource and endangered languages, humanitarian response, content moderation, and fairness and bias in language models. A number of further accepted papers are presented at the workshop as non-archival contributions.

We thank all authors for their submissions and our program committee for their careful reviews. We are grateful to our steering committee for their guidance and to Google for sponsoring the workshop. We also thank the EMNLP 2026 workshop chairs for their support.

The NLP4PI 2026 Organizers

Organizing Committee

Organizing Committee

Katherine Atwell, Northeastern University
Angana Borah, University of Michigan
Daryna Dementieva, Technical University of Munich
Elisa Kreiss, University of California
Neema Kotonya, Dataminr
Jiarui Liu, Carnegie Mellon University
Liz Olson, Dataminr
Ruyuan Wan, Pennsylvania State University
Jieyu Zhao, University of Southern California

Steering Committee

Rada Mihalcea, University of Michigan
Joel Tetreault, Dataminr
Zhijing Jin, University of Toronto

Program Committee

Program Committee

Ashutosh Ahuja, Humayun Akhtar, Aashish Anantha Ramakrishnan, Antonios Anastasopoulos

Nikolay Babakov, Longju Bai, Yann Billeter, Chaithanya Reddy Bogadi, Lorena Bregaj

Divya Chaudhary, Shou-Kai Cheng

Debarati Das, Bonnie J Dorr

Tobias Eder

Advitya Gemawat, Iacopo Ghinassi, Nishant Gupta

Zhongyuan Han, Souha Ben Hassine, Zeyu He, Daniel Hershcovich, Chia-Chien Hung

Muhammad Okky Ibrohim

Labiba Jahan

Antonia Karamolegkou, Sonal Khosla, Sai Tulasi Kolapudi, Alok Kumar, Anamika Kumari

Ryan Luo Li, Zhengyang Li, Jindřich Libovický, Shijia Liu, Gabriel Loiseau

Hannah Murray, Benedetta Muscato, Arianna Muti

Eman Nabeel, Poli Nemkova, Vera Neplenbroek

Shu Okabe, Emily Öhman

Sangameshwar Patil, Sanchita Porwal, Irina Proskurina, Michal Ptaszynski

Shihao Ran, Shivam Rana, Viswanathan Ranganathan, Naquee Rizwan, Daniel Ruffinelli

Nazanin Sabri, Nutan Sahoo, Sofia Serrano, Pratinav Seth, Mohit Singhal

Yashdeep Thorat

Jiawen Wang, Zhizhi Wang, Steven R Wilson

Yangxinyu Xie

Zukang Yang

Fanqi Zeng, Zixuan Zhu, Longfei Zuo

Table of Contents

<i>A Quantitative Gap Audit of LLM-era NLP for Education: Validity, Deployment, and Multilingual Equity</i>	
Farzana Zahid, Md. Rakibul Islam, Mayesha Kader and Israt Jahan Anika	1
<i>A Safety-Engineered Consumer Mental-Health Assistant: A Deployment Case Report on Retrieval-First Grounding, Release Gates, and the Non-Clinical Boundary</i>	
Robert Sneiderman	17
<i>Accessible Translation from Standard Arabic to Moroccan Sign Language: A Rule-Based System for Low-Resource Sign Languages</i>	
Abdelhadi Soudi, Kristof Van Laerhoven and Mohammed Hannani	22
<i>Agentic LLM for Information Access among Refugees: A Field Report from a Resource-Constrained Organization</i>	
Nari Yoo and Youngah Kim	30
<i>Aligned for Whom? A Brazilian Research Agenda for Culturally Grounded AI Safety Evaluation</i>	
Livy Real and Roberta Viola	49
<i>An Edge AI-Based Offline System for Educational Technology in Resource-Constrained Regions</i>	
Irin Jose, Nickolas Kenyeres and Kailey Cozart-Quesada	57
<i>Before You Trust the Count: Auditing Multilingual Sustainable Development Goal News</i>	
Sergey Volkov	67
<i>Beyond Helpfulness: A Teaching-over-Solving Diagnostic for Measuring Educational Impact in LLM Tutors</i>	
Junyi Yao, Zihao Zheng and Baichuan Li	77
<i>Beyond the Flag: Clinical Framing Closes the Moderation Gap in Suicide Risk Measurement</i>	
Shreyas Krishnan, Gun Ahn and Jungjin Kim	83
<i>BodyCam-VQA: Enhanced Body-Worn Camera Video Captioning via Multimodal Reasoning and Probe Question Generation</i>	
Karish Gupta, Matthew Alex, Alex Li, Yang Wu, Yun-Wei Chu, Kashif Munir, Xiaotian Zhou, Zhengping Ji and Xiaozhong Liu	103
<i>Computational Preservation of Endangered Pamiri Languages through NLP: Insights from Bartangi and a Roadmap for Digital Inclusion</i>	
Warda Tariq	116
<i>Conjoint Auditing: A Framework for Measuring the Social Impact of LLM Deployment</i>	
Katrina M. Jones and Lennart Kasserra	125
<i>Correction Is Not Response: Reframing the Unverifiable Majority of Health Infodemic Discourse as Communicative Acts</i>	
Minsu Ko, Jongyoung Lee and Woonho Choi	140
<i>Detecting Coordinated Inauthentic Behavior on Social Media with Memory-Guided Causal Inference</i>	
Yi Han, Wen Ding, Xuanjing Chen and Zhiguo Tao	151
<i>Evaluating Bias in LLM-Synthetic Data across Majority and Minority Groups on General and Stigmatized Topics: The Case of Multicultural and Monocultural Adolescents in South Korea</i>	
Jiyun Beak, Myokyung Han and Lanu Kim	161

<i>Evaluating Therapist-in-the-Loop NLP in Comprehension Rehabilitation for Children with SEND</i>	
Martina Galletti, Caterina Marchesi and Eleonora Pasqua	175
<i>From Signals to Barriers: An LLM-Assisted Bow-Tie Framework for NGO Civic-Space Risk Monitoring</i>	
Noemí Orué-Guerrero and Bernat Coma-Puig	188
<i>Generating Curriculum Grounded Nepali Questions Based On Bloom’s Taxonomy</i>	
Prabigya Acharya	198
<i>Grounded, Compute-Efficient LLM Policy Agents for Energy-Poverty Equity in Physically-Constrained Peer-to-Peer Energy Markets</i>	
Kunal Jadhav and Siddhesh More	211
<i>Hate Speech as a Conflict-Period Marker: Evidence from Arabic Telegram in Sudan</i>	
Tara Pai and Poli Nemkova	220
<i>How Persona and Message Type in Conversation Initiation Influence Learners’ Experience with Educational Chatbots</i>	
Jingchao Fang, Qingxiaoyang Zhu and Hao-Chuan Wang	232
<i>Human Values in Stories: Narrative-informed Directions for Value Discovery</i>	
Alexandra Barancová, Giovanni Sileno and Sennay Ghebreab	249
<i>LLM Mental Health Response Manipulation</i>	
Cristina Maria Popescu, Anamaria Hodivoianu, Ana Sabina Uban and Sergiu Nisioi	256
<i>MATCHA: LLM-Based Chart Summarization for Screen Reader Users</i>	
Lauren Claudia Kahn, Jaime Soto, Om Arya, Audrey Michal and Steven R Wilson	269
<i>Measuring Public Perceptions of the Social Impacts of Artificial Intelligence in YouTube Discourse</i>	
Aarav Gaur	283
<i>Measuring the Impact of LLM-Generated Scientific Summaries for Neurodivergent Readers</i>	
Abraham Ur Rahman, Skyler Koba, Audrey Michal and Steven R Wilson	298
<i>Meridian: A Bivariate Perplexity Architecture for Trustworthy and Equitable AI Text Detection</i>	
Pranil Raichura	312
<i>NSIDDx: A Design Framework for Neuro-Symbolic, Practitioner-First Differential Diagnosis in Low-Resource Settings</i>	
Aarav Singh	326
<i>On the Neutrality of Model Confidence in German and English Gender-Bias Benchmarks</i>	
Shalaka Satheesh, Nika Jacob and Katharina Beckh	339
<i>Preference Optimization Improves Definition Generation: Semantic Change in EU Institutional Sustainable Finance Disclosures</i>	
Mariia Fedorova, Andrey Kutuzov and Gustavo Cadenas Delascio	349
<i>Reading Between the Signs: Predicting Future Suicidal Ideation from Adolescent Social Media Texts</i>	
Paul Blum, Enrico Liscio, Ruixuan Zhang, Caroline Figueroa and Pradeep K. Murukannaiah	360
<i>Right-Wing Rock or Just Rock? A Computational Linguistic Analysis of Frei.Wild</i>	
Carlotta Schneeberger and Kevin Tang	375
<i>Roleplaying with Structure: Synthetic Therapist-Client Conversation Generation from Questionnaires</i>	
Doan Nam Long Vu, Rui Tan, Svenja Jule Francke, Lena Moench, Daniel Woiwod, Florian Thomas-Odenthal, Sanna Stroth, Tilo Kircher, Christiane Hermann, Udo Dannlowski, Hamidreza Jamalabadi, Simone Balloccu and Shaoxiong Ji	395

<i>SEAL: SocioEconomic Alignment in LLMs</i>	
Elisa Bassignana, Lorena Calvo-Bartolomé, Esra Dönmez, Marlene Lutz, Arianna Muti, Vera Neplenbroek and Donya Rooein	418
<i>SOLIAS: A Reporting Standard for Social-Impact Claims in NLP</i>	
Srimonti Dutta and Akshata Kishore Moharir	435
<i>StageWell: A Process-Aligned Chinese Corpus for Positive-Psychology Support Dialogue</i>	
Yuxiong Wang, Ziwei Lin, Bo Wang, Yu Zhang and Shiguang Ni	452
<i>Temporal Epistemic Graph of Thought: Uncertainty over Latent Reasoning Structure for Selective Ab- stention in Clinical Prediction</i>	
Govind Arun, Sadhvik Bathini and Balu Bhasuran	482
<i>The AI Therapist: ChatGPT Therapy Users at the Boundary of Human Therapy, Disillusionment, and AI Companionship on Reddit</i>	
Nari Yoo and Pengyu Chen	493
<i>Towards Automated Situation Awareness: A RAG-Based Framework for Peacebuilding Reports</i>	
Poli Nemkova, Suleyman Olcay Polat, Rafid Ishrak Jahan, Sagnik Ray Choudhury, Sun-Joo Lee, Shouryadip Sarkar and Mark V. Albert	511
<i>UkrMeme: Multimodal Explainable Persuasion Detection in Ukrainian Memes</i>	
Nazar Andrushko, Vera Schmitt and Veronika Solopova	530
<i>Uncertainty Quantification to Improve the Reliability of LLM-Based Hate Speech Detection</i>	
Nehir Idil Yuksel, Daryna Dementieva and Alexander Fraser	544
<i>When Cultures Move: Measuring and Improving Multicultural Text-to-Video Generation</i>	
Shuowei Li, Yuming Zhao, Parth Bhalerao and Oana Ignat	575
<i>Where Do Predictions Go? Class-Level Shifts and Carbon Cost in Stance Detection</i>	
Nityaa Kalra, Jisk J Attema and Dimitar Shterionov	592
<i>Who Gets Annotated? Impact-Aware Active Learning for Robust Humanitarian Text Classification</i>	
Divyajot Singh	605
<i>Whose Good, Whose Place? The Moral Geography of Agentic AI for Social Good</i>	
Poli Nemkova, Haeshitha Indukuri and Jaedon Charles	615