

StereACuLT 2026

**The 1st Workshop on Stereotypes Across Cultures in
Language Technologies**

Proceedings of the Workshop

July 3, 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 979-8-89176-408-8

Message from the Organizers

This volume contains the proceedings of the 1st Workshop on Stereotypes Across Cultures in Language Technologies (StereACuLT 2026), held on July 3, 2026, in San Diego, California, USA, co-located with ACL 2026.

StereACuLT was organized to advance research on stereotypes in language technologies with an explicit focus on cultural context. While stereotypes and biases in NLP systems have been studied extensively, much of the existing work has centered on English and a limited set of social and cultural settings. The goal of the workshop is to broaden this conversation by bringing together researchers working on culturally grounded definitions of stereotypes, multilingual and cross-cultural evaluation, mitigation methods, data practices, and the societal implications of deploying language technologies across diverse communities.

The workshop program includes invited and contributed presentations that examine stereotypes across languages, regions, and cultural settings, as well as methods for analyzing, evaluating, and mitigating such harms in modern language technologies. We hope the workshop helps build a shared foundation for more transparent, rigorous, and culturally informed research in this area.

We would like to thank all authors for submitting their work to the workshop, the program committee and additional reviewers for their time and thoughtful feedback, and our invited speaker, Diyi Yang, for supporting the workshop. We are also grateful to ACL 2026 for hosting the workshop.

Weicheng Ma
Soroush Vosoughi
Nabeel Gillani
Rolando Coto-Solano

Program Committee

Program Chairs

Rolando Coto-Solano, Dartmouth College
Nabeel Gillani, Northeastern University
Weicheng Ma, Oakland University
Soroush Vosoughi, Dartmouth College

Area Chairs

Weicheng Ma, Oakland University
Soroush Vosoughi, Dartmouth College

Reviewers

Andrew Aurand

Nischal Reddy Chandra

John J. Guerrero

Farnoosh Hashemi

Shiyu Ji, Yaning Jia

Emran Kanaan

Tian Lan

Zhongyu Ouyang

Fagun Patel

Peijun Qing

Sivakumar Selvaraj

Ivory Yang

Chunhui Zhang, Hefan Zhang

Table of Contents

| | |
|--|-----|
| <i>CrowS-Pairs-NL: A Benchmark to Evaluate Dutch Stereotype Bias in LLMs</i> Jens van der Weide, Dong Nguyen, Marianne Schaaphok and Roos M. Bakker | 1 |
| <i>Lost in Translation: Cross-Cultural Bias in LLM-Assisted Medical Symptom Interpretation</i> Yuting Tian, Salar Khaleghzadegan, Benjamin Huh, Yash Raj and Gena Heng | 13 |
| <i>Exploratory As-Analyzed No-Detection of Culturally-Marked Predicate-Triggered PII Amplification in a Synthetic-English RAG Probe: A Predicate-Resource-Confounded Audit</i> Yanhang Li, Zhichao Fan and Zexin Zhuang | 20 |
| <i>Controlling Cross-Lingual Answer Distributions in Language Models: Enabling Transfer of Factual Preferences</i> Lukas Ellinger, Alexander Manev and Georg Groh | 35 |
| <i>Counterfactual Auditing of Cross-Cultural Variation in LLM-Generated Medical Advice</i> Hyunwoo Yoo and Gail Rosen | 50 |
| <i>Stereotyped by Silence: How LLMs Erase Northeast Indian Languages Through Omission and Orthographic Corruption</i> Badal Nyalang | 62 |
| <i>Whose Pragmatics? Cultural Grounding as a Bottleneck for Stereotype Detection in Egyptian Arabic Social Media</i> Samar A. Assem | 69 |
| <i>Measuring Semantic Flow Without Direction: A Rhizomatic Protocol for Stereotype Translation in Cross-Cultural Language Technology</i> Gustavo Aviña Cerecer | 79 |
| <i>Signals Are Not States: Neuro-Symbolic Safeguards for Culturally Aware Classroom AI</i> Sina Bagheri Nezhad | 88 |
| <i>AmchiBias: Measuring Stereotypical Bias in Goan Identity Groups with a Minimal Pair Dataset in English and Konkani</i> Michelle Barbosa, Sebastian Padó and Franziska Weeber | 101 |
| <i>Translation Is Not Representation: English-Hub Routing in Cross-Lingual Bias Benchmarks</i> Hak Hyun Kim and Benjamin Huh | 116 |
| <i>IndicSteer: Inference-Time Safety Steering for Indic LLMs</i> Ruhaib Muhammad, Saahas Vijayalakshmi Rajaram and Suriya Priyan Durairaj | 126 |