Third Workshop on NLP for Positive Impact (NLP4PI 2024)

Miami, Florida, USA
15 November 2024

ISBN: 979-8-3313-0851-3

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571

Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2024) by the Association for Computational Linguistics
All rights reserved.

Printed with permission by Curran Associates, Inc. (2025)

For permission requests, please contact the Association for Computational Linguistics
at the address below.

Association for Computational Linguistics
209 N. Eighth Street
Stroudsburg, Pennsylvania 18360

Phone: 1-570-476-8006
Fax: 1-570-476-0860

acl@aclweb.org

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2633
Email: curran@proceedings.com
Web: www.proceedings.com

Table of Contents

What is the social benefit of hate speech detection research? A Systematic Review
Sidney Gig-Jan Wong 1

Multilingual Fact-Checking using LLMs
Aryan Singhal, Thomas Law, Coby Kassner, Ayushman Gupta, Evan Duan, Aviral Damle and Ryan Luo Li 13

Transferring Fairness using Multi-Task Learning with Limited Demographic Information
Carlos Alejandro Aguirre and Mark Dredze 32

Selecting Shots for Demographic Fairness in Few-Shot Learning with Large Language Models
Carlos Alejandro Aguirre, Kuleen Sasse, Isabel Alyssa Cachola and Mark Dredze 50

Covert Bias: The Severity of Social Views' Unalignment in Language Models Towards Implicit and Explicit Opinion
Abeer Aldayel, Areej Alokaili and Rehab Alahmadi 68

PG-Story: Taxonomy, Dataset, and Evaluation for Ensuring Child-Safe Content for Story Generation
Alicia Y. Tsai, Shereen Oraby, Anjali Narayan-Chen, Alessandra Cervone, Spandana Gella, Apurv Verma, Tagyoung Chung, Jing Huang and Nanyun Peng 78

Towards Explainable Multi-Label Text Classification: A Multi-Task Rationalisation Framework for Identifying Indicators of Forced Labour
Erick Mendez Guzman, Viktor Schlegel and Riza Batista-Navarro 98

All Models are Wrong, But Some are Deadly: Inconsistencies in Emotion Detection in Suicide-related Tweets
Annika Marie Schoene, Resmi Ramachandranpillai, Tomo Lazovich and Ricardo A. Baeza-Yates 113

Efficient Aspect-Based Summarization of Climate Change Reports with Small Language Models
Iacopo Ghinassi, Leonardo Catalano and Tommaso Colella 123

An NLP Case Study on Predicting the Before and After of the Ukraine–Russia and Hamas–Israel Conflicts
Jordan Miner and John E. Ortega 140

Exploring the Jungle of Bias: Political Bias Attribution in Language Models via Dependency Analysis
David F. Jenny, Yann Billeter, Bernhard Schölkopf and Zhijing Jin 152

AgriLLM:Harnessing Transformers for Framer Queries
Krish Didwania, Pratinav Seth, Aditya Kasliwal and Amit Agarwal 179

SciTechBaitRO: ClickBait Detection for Romanian Science and Technology News
Raluca-Andreea Gînga and Ana Sabina Uban 188

Investigating Ableism in LLMs through Multi-turn Conversation
Guojun Wu and Sarah Ebling 202

Eliciting Uncertainty in Chain-of-Thought to Mitigate Bias against Forecasting Harmful User Behaviors
Anthony Sicilia and Malihe Alikhani 211

Inferring Mental Burnout Discourse Across Reddit Communities
Nazanin Sabri, Anh C. Pham, Ishita Kakkar and Mai ElSherief 224

Decoding Ableism in Large Language Models: An Intersectional Approach
Rong Li, Ashwini Kamaraj, Jing Ma and Sarah Ebling 232

Explainable Identification of Hate Speech towards Islam using Graph Neural Networks
Azmine Toushik Wasi 250

From Text to Maps: LLM-Driven Extraction and Geotagging of Epidemiological Data
Karlyn K. Harrod, Prabin Bhandari and Antonios Anastasopoulos 258

Crafting Tomorrow's Headlines: Neural News Generation and Detection in English, Turkish, Hungarian, and Persian
Cem Üyük, Danica Rovó, Shaghayeghkolli Shaghayeghkolli, Rabia Varol, Georg Groh and Daryna Dementieva 271

Reference-Based Metrics Are Biased Against Blind and Low-Vision Users' Image Description Preferences
Rhea Kapur and Elisa Kreiss 308

MultiClimate: Multimodal Stance Detection on Climate Change Videos
Jiawen Wang, Longfei Zuo, Siyao Peng and Barbara Plank 315

AAVENUE: Detecting LLM Biases on NLU Tasks in AAVE via a Novel Benchmark
Abhay Gupta, Ece Yurtseven, Philip Meng and Kevin Zhu 327

DiversityMedQA: A Benchmark for Assessing Demographic Biases in Medical Diagnosis using Large Language Models
Rajat Rawat, Hudson McBride, Dhiyaan Chakkresh Nirmal, Rajarshi Ghosh, Jong Moon, Dhruv Karthik Alamuri and Kevin Zhu 334

Improving Industrial Safety by Auto-Generating Case-specific Preventive Recommendations
Sangameshwar Patil, Sumit Koundanya, Shubham Kumbhar and Alok Kumar 349