

NLP4PI 2024

Third Workshop on NLP for Positive Impact

Proceedings of the Workshop

November 15, 2024

©2024 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 979-8-89176-175-9

Introduction

The widespread and indispensable use of language-oriented AI systems presents new opportunities to have a positive social impact. Much existing work on NLP for social good focuses on detecting or preventing harm, such as classifying hate speech, mitigating bias, or identifying signs of depression. However, NLP research also offers the potential for positive proactive applications that can improve user and public well-being or foster constructive conversations. Nevertheless, “positive impact” remains difficult to define, and well-intentioned NLP technology can raise concerns about ethics and privacy.

This volume contains the proceedings of the Third Workshop on NLP for Positive Impact held in conjunction with the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024). The workshop received 48 submissions of technical papers of which 31 were accepted (25 archival and 6 non-archival), for an acceptance rate of 65%. Non-archival papers are included in the schedule and presented during the workshop, but are not included in the proceedings, whereas archival papers are included. We thank Program Committee members for providing high quality reviews in assembling these proceedings. These papers cover diverse aspects of NLP for positive impact, including developing NLP technology for applications like healthcare, criminal law, education, social media analyses, and consumer privacy as well discussing challenges and ethical implications of using NLP in these areas.

In addition to technical papers, this workshop also features invited keynote speakers and panelists to facilitate discussion and enhance knowledge of NLP for positive impact.

Keynote speakers:

Anjalie Field, Johns Hopkins University
Stephen Mayhew, Duolingo
Cordelia Moore, NGOs Advisor on Digital Violence
Louis-Philippe Morency, Carnegie Mellon University
Veronica Perez-Rosa, University of Michigan
Mrinmaya Sachan, ETH Zürich
Yulia Tsvetkov, University of Washington

Panelists:

Anjalie Field, Johns Hopkins University
Stephen Mayhew, Duolingo
Cordelia Moore, NGOs Advisor on Digital Violence
Jieyu Zhao, University of Southern California

We are grateful to all the people who have contributed to this workshop, including speakers, authors, reviewers, and attendees.

We hope that our workshop can encourage future work on pro-social NLP and we look forward to welcoming you all to our workshop!

- Daryna, Oana, Zhijing, Rada, Giorgio, Joel, Steven, and Jieyu

Organizing Committee

Program Chairs

Daryna Dementieva, Technical University of Munich

Oana Ignat, Santa Clara University

Zhijing Jin, Max Planck Institute and ETH Zürich

Rada Mihalcea, University of Michigan

Giorgio Piatti, ETH Zürich

Joel Tetreault, Dataminr

Steven Wilson, University of Michigan-Flint

Jieyu Zhao, University of Southern California

Program Committee

Reviewers

Udita Patel, Amazon
Tornike Tsereteli, Universität Mannheim
John E. Ortega, Northeastern University
Eda Okur, Intel Labs
Hao Chen, Zhejiang University
Shihao Ran, University of Houston
Resmi Ramachandranpillai, Institute for Experiential AI
Nikolay Babakov, University of Santiago de Compostela
Theodora Chaspari, University of Colorado at Boulder
Shubham Shukla, Nordstrom
Pratinav Seth, Arya.ai
Pablo Duboue, Textualization Software Ltd.
Chan Young Park, University of Washington
Ehsanul Haque Nirjhar, Texas A&M University - College Station
Hannah Rashkin, Google
Sarah Ebling, University of Zurich
Ekaterina Artemova, Toloka AI
Raluca-Andreea Gînga, University of Bucharest
Daniel Spokoyny, Carnegie Mellon University
Christos Christodoulopoulos, Amazon
Jiawen Wang, Ludwig-Maximilians-Universität München
Tobias Eder, Technische Universität München
Antonios Anastasopoulos, Athena Research Center
Jakub Macina, ETHZ - ETH Zurich
Divya Chaudhary, Northeastern University
Ke Zhang, Dataminr, inc
Akhila Yerukola, Carnegie Mellon University
Di Lu, Dataminr
Qinlan Shen, Oracle
Chia-Chien Hung, NEC Laboratories Europe
Kevin Zhu, Algoverse AI Research
Xiaofan Zheng, Xi'an Jiaotong University
Mascha Kurpicz-Briki, BFH - Bern University of Applied Sciences
Minnan Luo, Xi'an Jiaotong University
Yann Billeter, ZHAW - Zürcher Hochschule für Angewandte Wissenschaften
Sangameshwar Patil, Indian Institute of Technology, Madras
Zhongyuan Han, Foshan University
Robert L. Logan IV, University of California, Irvine
Emily Öhman, Waseda University
Abeer Aldayel, King Saud University
James R. Foulds, University of Maryland, Baltimore County
Jing Ma, University of Zurich
Labiba Jahan, Southern Methodist University
Iacopo Ghinassi, Queen Mary University of London
Andrea Galassi, University of Bologna
Yangxinyu Xie, University of Pennsylvania

Bonnie J Dorr, University of Florida
Aoife Cahill, Dataminr
Sofia Serrano, Lafayette College
Hemank Lamba, Dataminr Inc
Aylin Ece Gunal, University of Michigan - Ann Arbor
Anthony Sicilia, Northeastern University
Ines Rehbein, Universität Mannheim
Neema Kotonya, Dataminr
Ashwini Kamaraj, University of Zurich
Christopher Klamm, Universität Mannheim
Manaar Alam, New York University, Abu Dhabi
Longfei Zuo, Ludwig-Maximilians-Universität München
Jimin Mun, CMU, Carnegie Mellon University
Julia Mendelsohn, University of Michigan - Ann Arbor
Daniel Hershovich, University of Copenhagen
Ryan Luo Li, Algoverse
Ana Sabina Uban, Universitatea Bucuresti
Rodolfo Zevallos, Universitat Pompeu Fabra
Lucas Rosenblatt, New York University
Lama Moukheiber, Massachusetts Institute of Technology
Karina H Halevy, Carnegie Mellon University
Ziyu Yao, George Mason University
Alok Kumar, Tata Consultancy Services Limited, India
Johnny Wei, University of Southern California
Nazanin Sabri, University of California, San Diego
Azmine Toushik Wasi, Shahjalal University of Science and Technology
Mai ElSherief, Northeastern University
Dipesh Singla, IEEE
Philipp Seeberger, Technische Hochschule Nürnberg Georg Simon Ohm
Mert Inan, Northeastern University
Prabin Bhandari, George Mason University
Elisa Kreiss, University of California, Los Angeles
Rong Li, University of Zurich
Mira Moukheiber, Massachusetts Institute of Technology

Table of Contents

<i>What is the social benefit of hate speech detection research? A Systematic Review</i> Sidney Gig-Jan Wong	1
<i>Multilingual Fact-Checking using LLMs</i> Aryan Singhal, Thomas Law, Coby Kassner, Ayushman Gupta, Evan Duan, Aviral Damle and Ryan Luo Li	13
<i>Transferring Fairness using Multi-Task Learning with Limited Demographic Information</i> Carlos Alejandro Aguirre and Mark Dredze	32
<i>Selecting Shots for Demographic Fairness in Few-Shot Learning with Large Language Models</i> Carlos Alejandro Aguirre, Kuleen Sasse, Isabel Alyssa Cachola and Mark Dredze	50
<i>Covert Bias: The Severity of Social Views’ Unalignment in Language Models Towards Implicit and Explicit Opinion</i> Abeer Aldayel, Areej Alokaili and Rehab Alahmadi	68
<i>PG-Story: Taxonomy, Dataset, and Evaluation for Ensuring Child-Safe Content for Story Generation</i> Alicia Y. Tsai, Shereen Oraby, Anjali Narayan-Chen, Alessandra Cervone, Spandana Gella, Apurv Verma, Tagyoung Chung, Jing Huang and Nanyun Peng	78
<i>Towards Explainable Multi-Label Text Classification: A Multi-Task Rationalisation Framework for Identifying Indicators of Forced Labour</i> Erick Mendez Guzman, Viktor Schlegel and Riza Batista-Navarro	98
<i>All Models are Wrong, But Some are Deadly: Inconsistencies in Emotion Detection in Suicide-related Tweets</i> Annika Marie Schoene, Resmi Ramachandranpillai, Tomo Lazovich and Ricardo A. Baeza-Yates 113	
<i>Efficient Aspect-Based Summarization of Climate Change Reports with Small Language Models</i> Iacopo Ghinassi, Leonardo Catalano and Tommaso Colella	123
<i>An NLP Case Study on Predicting the Before and After of the Ukraine–Russia and Hamas–Israel Con- flicts</i> Jordan Miner and John E. Ortega	140
<i>Exploring the Jungle of Bias: Political Bias Attribution in Language Models via Dependency Analysis</i> David F. Jenny, Yann Billeter, Bernhard Schölkopf and Zhijing Jin	152
<i>AgriLLM: Harnessing Transformers for Framer Queries</i> Krish Didwania, Pratinav Seth, Aditya Kasliwal and Amit Agarwal	179
<i>SciTechBaitRO: ClickBait Detection for Romanian Science and Technology News</i> Raluca-Andreea Gînga and Ana Sabina Uban	188
<i>Investigating Ableism in LLMs through Multi-turn Conversation</i> Guojun Wu and Sarah Ebling	202
<i>Eliciting Uncertainty in Chain-of-Thought to Mitigate Bias against Forecasting Harmful User Beha- viors</i> Anthony Sicilia and Malihe Alikhani	211

<i>Inferring Mental Burnout Discourse Across Reddit Communities</i> Nazanin Sabri, Anh C. Pham, Ishita Kakkar and Mai ElSherief	224
<i>Decoding Ableism in Large Language Models: An Intersectional Approach</i> Rong Li, Ashwini Kamaraj, Jing Ma and Sarah Ebling	232
<i>Explainable Identification of Hate Speech towards Islam using Graph Neural Networks</i> Azmine Toushik Wasi	250
<i>From Text to Maps: LLM-Driven Extraction and Geotagging of Epidemiological Data</i> Karlyn K. Harrod, Prabin Bhandari and Antonios Anastasopoulos	258
<i>Crafting Tomorrow’s Headlines: Neural News Generation and Detection in English, Turkish, Hungarian, and Persian</i> Cem Üyük, Danica Rovó, Shaghayeghkolli Shaghayeghkolli, Rabia Varol, Georg Groh and Daryna Dementieva	271
<i>Reference-Based Metrics Are Biased Against Blind and Low-Vision Users’ Image Description Preferences</i> Rhea Kapur and Elisa Kreiss	308
<i>MultiClimate: Multimodal Stance Detection on Climate Change Videos</i> Jiawen Wang, Longfei Zuo, Siyao Peng and Barbara Plank	315
<i>AAVENUE: Detecting LLM Biases on NLU Tasks in AAVE via a Novel Benchmark</i> Abhay Gupta, Ece Yurtseven, Philip Meng and Kevin Zhu	327
<i>DiversityMedQA: A Benchmark for Assessing Demographic Biases in Medical Diagnosis using Large Language Models</i> Rajat Rawat, Hudson McBride, Dhiyaan Chakkresh Nirmal, Rajarshi Ghosh, Jong Moon, Dhruv Karthik Alamuri and Kevin Zhu	334
<i>Improving Industrial Safety by Auto-Generating Case-specific Preventive Recommendations</i> Sangameshwar Patil, Sumit Koundanya, Shubham Kumbhar and Alok Kumar	349

Program

Friday, November 15, 2024

- 09:00 - 09:05 *Opening Remark*
- 09:05 - 09:30 *Opening Talk by Rada Mihalcea*
- 09:30 - 09:55 *Talk by Prof Yulia Tsvetkov (UW) & 5 min Q&A*
- 10:00 - 10:25 *Talk by Prof Anjalie Field (JHU) & 5 min Q&A*
- 10:30 - 11:00 *NGO Lightning Talk*
- 11:00 - 12:00 *Poster Session (In-Person and Virtual)*

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

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

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

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

CLIMB: A Benchmark of Clinical Bias in Large Language Models
Yubo Zhang, Shudi Hou, Mingyu Derek Ma, Wei Wang, Muhao Chen and Jieyu Zhao

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

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

Friday, November 15, 2024 (continued)

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

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

Efficient Aspect-Based Summarization of Climate Change Reports with Small Language Models

Iacopo Ghinassi, Leonardo Catalano and Tommaso Colella

Uchaguzi-2022: A Dataset of Citizen Reports on the 2022 Kenyan Election

Roberto Mondini, Neema Kotonya, Robert L. Logan IV, Elizabeth M Olson, Angela Oduor Lungati, Daniel Duke Odongo, Tim Ombasa, Hemank Lamba, Aoife Cahill, Joel R. Tetreault and Alejandro Jaimes

CEHA: A Dataset of Conflict Events in the Horn of Africa

Rui Bai, Di Lu, Shihao Ran, Elizabeth M Olson, Hemank Lamba, Aoife Cahill, Joel R. Tetreault and Alejandro Jaimes

Operationalizing content moderation accuracy in the Digital Services Act

Johnny Wei

An NLP Case Study on Predicting the Before and After of the Ukraine–Russia and Hamas–Israel Conflicts

Jordan Miner and John E. Ortega

Exploring the Jungle of Bias: Political Bias Attribution in Language Models via Dependency Analysis

David F. Jenny, Yann Billeter, Bernhard Schölkopf and Zhijing Jin

AgriLLM: Harnessing Transformers for Framer Queries

Krish Didwania, Pratinav Seth, Aditya Kasliwal and Amit Agarwal

SciTechBaitRO: ClickBait Detection for Romanian Science and Technology News

Raluca-Andreea Gînga and Ana Sabina Uban

Investigating Ableism in LLMs through Multi-turn Conversation

Guojun Wu and Sarah Ebling

Friday, November 15, 2024 (continued)

Eliciting Uncertainty in Chain-of-Thought to Mitigate Bias against Forecasting Harmful User Behaviors

Anthony Sicilia and Malihe Alikhani

WildfireGPT: Tailored Large Language Model for Wildfire Analysis

Yangxinyu Xie, Bowen Jiang, Tanwi Mallick, John K Hutchison, Duane Rudolph Verner, Jordan Branham, M. Ross Alexander, Robert Ross, Yan Feng, Leslie-Anne Levy, Weijie J Su and Camillo Jose Taylor

Inferring Mental Burnout Discourse Across Reddit Communities

Nazanin Sabri, Anh C. Pham, Ishita Kakkar and Mai ElSherief

Decoding Ableism in Large Language Models: An Intersectional Approach

Rong Li, Ashwini Kamaraj, Jing Ma and Sarah Ebling

Explainable Identification of Hate Speech towards Islam using Graph Neural Networks

Azmine Toushik Wasi

From Text to Maps: LLM-Driven Extraction and Geotagging of Epidemiological Data

Karlyn K. Harrod, Prabin Bhandari and Antonios Anastasopoulos

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

Reference-Based Metrics Are Biased Against Blind and Low-Vision Users' Image Description Preferences

Rhea Kapur and Elisa Kreiss

MultiClimate: Multimodal Stance Detection on Climate Change Videos

Jiawen Wang, Longfei Zuo, Siyao Peng and Barbara Plank

AAVENUE: Detecting LLM Biases on NLU Tasks in AAVE via a Novel Benchmark

Abhay Gupta, Ece Yurtseven, Philip Meng and Kevin Zhu

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

Friday, November 15, 2024 (continued)

Improving Industrial Safety by Auto-Generating Case-specific Preventive Recommendations

Sangameshwar Patil, Sumit Koundanya, Shubham Kumbhar and Alok Kumar

From Predictions to Analyses: Explainable Rationales-Augmented Fake News Detection with Large Vision-Language Models

Xiaofan Zheng, Zinan Zeng, Heng Wang, Yuyang Bai, Yuhan Liu and Minnan Luo

- 12:00 - 13:00 *Lunch Break*
- 13:00 - 13:25 *Talk by Prof Mrinmaya Sachan (ETH) & 5 min Q&A*
- 13:30 - 13:55 *Talk by Stephen Mayhew (Duolingo) & 5 min Q&A*
- 14:00 - 14:25 *Talk by Prof Veronica Perez-Rosa (Texas State University) and 5 min Q&A*
- 14:30 - 14:55 *Talk by Prof Louis-Philippe Morency (CMU) & 5 min Q&A*
- 15:00 - 15:30 *Oral Talk Sessions (5 Talks of 5 min each & 5 min Q&A in the end)*
- 15:30 - 15:45 *Coffee Break by EMNLP*
- 15:45 - 16:05 *Special Theme Digital Violence: NGO Talk by Cordelia Moore*
- 16:05 - 17:00 *Panel Conversation*
- 17:00 - 17:45 *Research Brainstorming: NLP for Social Good*
- 17:45 - 18:00 *Best Paper Announcement & Closing*