

NLP4PI 2022

Second Workshop on NLP for Positive Impact

Proceedings of the Workshop

December 7, 2022

©2022 Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
209 N. Eighth Street
Stroudsburg, PA 18360
USA
Tel: +1-570-476-8006
Fax: +1-570-476-0860
acl@aclweb.org

ISBN 978-1-959429-19-7

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 developed with responsible methods. Some top areas that we prioritize in this workshop correspond to the United Nations Sustainable Development Goals, such as applications of NLP to address poverty, healthcare, education, climate change, and so on.

This volume contains the proceedings of the Second Workshop on NLP for Positive Impact held in conjunction with the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP 2022). The workshop received 48 submissions of papers of which 22 were accepted (17 archival and 5 non-archival), for an acceptance rate of 46%. Additionally, 10 Findings of EMNLP papers will be presented at the workshop. We thank all 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 to help applications like physical and mental health, climate change, crisis response, social mobility, education, employment, and culture preservation, 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:

Mike Bailey, Meta

Sam Bowman, New York University & Anthropic AI

Rada Mihalcea, University of Michigan

Preslav Nakov, MBZUAI

Milind Tambe, Harvard University

Panelists:

Luis Chiruzzo, Universidad de la República, Uruguay

Tara Chklovski, Technovation

Dora Demszky, Stanford University

Rada Mihalcea, University of Michigan

We are grateful to all the people who have contributed to this workshop, including speakers, authors, reviewers, and attendees, and we would additionally like to thank the EMNLP workshop chairs and program chairs for making the workshop happen.

We hope that our workshop can encourage future work on NLP for positive social impact and we look forward to welcoming you all to our hybrid workshop!

- Laura Biester, Dora Demszky, Zhijing Jin, Mrinmaya Sachan, Joel Tetreault, Steven Wilson, Lu Xiao, Jieyu Zhao

Organizing Committee

Organizers

Laura Biester, University of Michigan
Dorottya Demszky, Stanford University
Zhijing Jin, Max Planck Institute & ETH
Mrinmaya Sachan, ETH Zurich
Joel Tetreault, Dataminr
Steven Wilson, Oakland University
Lu Xiao, Syracuse University
Jieyu Zhao, University of Maryland

Program Committee

Reviewers

Agostina Calabrese, The University of Edinburgh
Mahdi Abavisani, Dataminr, Inc.
Ameeta Agrawal, Portland State University
Abeer Aldayel, King Saud University
Bashar Alhafni, New York University
Yejin Bang, Hong Kong University of Science and Technology
Delphine Bernhard, Lilpa, Université de Strasbourg
Eleftheria Briakou, University of Maryland
Chris Brockett, Microsoft Research
Agostina Calabrese, The University of Edinburgh
Sky CH-Wang, Columbia University
Serina Chang, Stanford University
Martin Chodorow, Hunter College and the Graduate Center of CUNY
Christos Christodoulopoulos, Amazon Research
Oana Cocarascu, King's College London
Mark Dredze, Johns Hopkins University
Yupe Du, Utrecht University
Pablo Duboue, Textualization Software Ltd.
Steffen Eger, Bielefeld University
Jennifer Foster, Dublin City University
Kathleen C. Fraser, National Research Council Canada
Andrea Galassi, University of Bologna
Mozhdeh Gheini, University of Southern California
Kartik Goyal, Toyota Technological Institute at Chicago
Aylin Gunal, University of Michigan
Ivan Habernal, Technische Universität Darmstadt
Lisa Anne Hendricks, DeepMind
Valentin Hofmann, University of Oxford
Oana Ignat, University of Michigan
Muhammad Imran, Qatar Computing Research Institute
Radu Tudor Ionescu, University of Bucharest
David Jurgens, University of Michigan
Ashkan Kazemi, University of Michigan
Ashiqur KhudaBukhsh, Rochester Institute of Technology
Hyunwoo Kim, Seoul National University
Svetlana Kiritchenko, National Research Council Canada
Ekaterina Kochmar, University of Bath
Hemank Lamba, Dataminr, Inc.
Mirella Lapata, University of Edinburgh
Anne Lauscher, University of Hamburg
Nayeon Lee, Hong Kong University of Science and Technology
Ji-Ung Lee, UKP Lab Technische Universität Darmstadt
Diane Litman, University of Pittsburgh
Robert L Logan IV, Dataminr, Inc.
Li Lucy, University of California, Berkeley
Jakub Macina, ETH Zurich

Julia Mendelsohn, University of Michigan
Sewon Min, University of Washington
Negar Mokhberian, University of Southern California
Roser Morante, UNED
Aur lie N v ol, Universit  Paris Saclay, CNRS, LISN
Eda Okur, Intel Labs
Ji Ho Park, Google
Chan Young Park, Carnegie Mellon University
Carla Parra Escart n, RWS Language Weaver
Thierry Poibeau, LATTICE (CNRS & ENS/PSL)
Hannah Rashkin, Google Research
Navid Rekasaz, Johannes Kepler University
Laura Rimell, DeepMind
Bj rn Ross, University of Edinburgh
Mrinmaya Sachan, ETH Zurich
Danae S nchez Villegas, University of Sheffield
Maarten Sap, Carnegie Mellon University
Sofia Serrano, University of Washington
Naomi Shapiro, University of Washington
Qinlan Shen, Oracle
Shubham Shukla, Independent Researcher
Kevin Stowe, Educational Testing Services
Sara Stymne, Uppsala University
Swabha Swayamdipta, University of Southern California
Paolo Torroni, Universit  di Bologna
Josep Valls-Vargas, Adobe
Lucy Vanderwende, University of Washington
Zijian Wang, AWS AI Labs
Ke Zhang, Dataminr, inc

Keynote Talk: Fighting the Global Social Media Infodemic: from Fake News to Harmful Content

Preslav Nakov

Mohamed bin Zayed University of Artificial Intelligence

Abstract: The COVID-19 pandemic has brought us the first global social media infodemic. While fighting this infodemic is typically thought of in terms of factuality, the problem is much broader as malicious content includes not only “fake news”, rumors, and conspiracy theories, but also hate speech, racism, xenophobia, panic, and mistrust in authorities, among others. Thus, we argue for the need for a holistic approach combining the perspectives of journalists, fact-checkers, policymakers, social media platforms, and society as a whole.

We further argue for the need to analyze entire news outlets, which can be done in advance; then, we can fact-check the news before it was even written: by checking how trustworthy the outlet that has published it is (which is what journalists actually do). We will show how this can be automated by looking at variety of information sources.

The infodemic is often described using terms such as “fake news”, which mislead people to focus exclusively on factuality, and to ignore the other half of the problem: the potential malicious intent. We aim to bridge this gap by focusing on the detection of specific propaganda techniques in text, e.g., appeal to emotions, fear, prejudices, logical fallacies, etc. This is the target of the ongoing SemEval-2023 task 3, which focuses on multilingual aspects of the problem, covering English, French, German, Italian, Polish, and Russian. We further present extensions of this work to the automatic analysis of various types of harmful memes: from propaganda to harmfulness and harm’s target identification to role-labeling in terms of who is portrayed as hero/villain/victim, and generating natural text explanations.

Bio: Preslav Nakov is Professor at Mohamed bin Zayed University of Artificial Intelligence. Previously, he was Principal Scientist at the Qatar Computing Research Institute (QCRI), HBKU, where he led the Tanbih mega-project, developed in collaboration with MIT, which aims to limit the impact of fake news, propaganda and media bias by making users aware of what they are reading, thus promoting media literacy and critical thinking. He received his PhD degree in Computer Science from the University of California at Berkeley, supported by a Fulbright grant. He is Chair-Elect of the Association for Computational Linguistics (ACL), Secretary of ACL SIGSLAV, and Secretary of the Truth and Trust Online board of trustees. Formerly, he was PC chair of ACL 2022, and President of ACL SIGLEX. He is also member of the editorial board of several journals including Computational Linguistics, TACL, ACM TOIS, IEEE TASL, IEEE TAC, CS&L, NLE, AI Communications, and Frontiers in AI. He authored a Morgan & Claypool book on Semantic Relations between Nominals, two books on computer algorithms, and 250+ research papers. He received a Best Paper Award at ACM WebSci’2022, a Best Long Paper Award at CIKM’2020, a Best Demo Paper Award (Honorable Mention) at ACL’2020, a Best Task Paper Award (Honorable Mention) at SemEval’2020, a Best Poster Award at SocInfo’2019, and the Young Researcher Award at RANLP’2011. He was also the first to receive the Bulgarian President’s John Atanasoff award, named after the inventor of the first automatic electronic digital computer. Dr. Nakov’s research was featured by over 100 news outlets, including Forbes, Boston Globe, Aljazeera, DefenseOne, Business Insider, MIT Technology Review, Science Daily, Popular Science, Fast Company, The Register, WIRED, and Engadget, among others.

Keynote Talk: The Role of Social Networks in Economic Mobility

Mike Bailey
Meta

Abstract: Social capital—the strength of an individual’s social network and community—has been identified as a potential determinant of outcomes ranging from education to health. We use data on 21 billion friendships in the US to measure and analyze different types of social capital including connectedness between different types of people, social cohesion, and civic engagement. We demonstrate the importance of distinguishing these forms of social capital by analyzing their associations with economic mobility across areas. The share of high-SES friends among individuals with low SES—which we term economic connectedness—is among the strongest predictors of upward income mobility identified to date. In a different paper we use social network data in India to show the importance of social networks to labor migrants and find that increasing social connectedness across space may have considerable economic gains, improving average wages by 3% (24% for the bottom wage-quartile) in a migration model.

Bio: Mike Bailey is a senior social scientist at Meta on the Computational Social Science team. His work focuses on the role of social networks on economic opportunity including migration, health, education, and social capital and his work has been published in top scientific journals such as *Nature* and the *Journal of Political Economy* and covered by outlets such as *The Economist* and *The New York Times*. He is a co-creator of the Social Capital Atlas dataset and the Social Connectedness Index which are publicly available datasets measuring social connectedness. Previously at Facebook Mike founded and led several research science teams including the Economics Research team, the Feed Science team, and the Society Research team. He holds a PhD in Economics from Stanford and a BS in Math and Economics from Utah State and is originally from Utah.

Keynote Talk: AI for social impact: Results from deployments for public health and conservation

Milind Tambe

Harvard University and Google Research

Abstract: With the maturing of AI and multiagent systems research, we have a tremendous opportunity to direct these advances towards addressing complex societal problems. I will focus on domains of public health and conservation, and address one key cross-cutting challenge: how to effectively deploy our limited intervention resources in these problem domains. I will present results from work around the globe in using AI for challenges in public health such as Maternal and Child care interventions, HIV prevention, and in conservation such as endangered wildlife protection. Achieving social impact in these domains often requires methodological advances. To that end, I will highlight key research advances in multiagent reasoning and learning, in particular in, restless multiarmed bandits, influence maximization in social networks, computational game theory and decision-focused learning. In pushing this research agenda, our ultimate goal is to facilitate local communities and non-profits to directly benefit from advances in AI tools and techniques.

Bio: Milind Tambe is Gordon McKay Professor of Computer Science and Director of Center for Research in Computation and Society at Harvard University; concurrently, he is also Principal Scientist and Director AI for Social Good at Google Research. Prof. Tambe's work focuses on advancing AI and multiagent systems for public health, conservation & public safety, with a track record of building pioneering AI systems for social impact. He is recipient of the IJCAI John McCarthy Award, AAMAS ACM Autonomous Agents Research Award, AAAI Robert S. Engelmore Memorial Lecture Award, and he is a fellow of AAAI and ACM. He is also a recipient of the INFORMS Wagner prize for excellence in Operations Research practice and Rist Prize from MORS (Military Operations Research Society). For his work on AI and public safety, he has received Columbus Fellowship Foundation Homeland security award and commendations and certificates of appreciation from the US Coast Guard, the Federal Air Marshals Service and airport police at the city of Los Angeles.

Keynote Talk: Recentering NLP Around ALL People

Rada Mihalcea

University of Michigan

Abstract: The field of NLP has come a long way, with many exciting achievements along several research directions, including language generation, large language models, machine translation, and more. However, while most of the NLP technologies built today are branded as one size fits all, the reality is that they are one size fits the majority, with many languages and many minorities left 'on the side'. In this talk, I will highlight some of the drawbacks associated with this strategy of building 'generic' NLP technologies, and make suggestions for ways to move towards NLP for ALL.

Bio: Rada Mihalcea is the Janice M. Jenkins Collegiate Professor of Computer Science and Engineering at the University of Michigan and the Director of the Michigan Artificial Intelligence Lab. Her research interests are in computational linguistics, with a focus on lexical semantics, computational social sciences, and multimodal language processing. She serves or has served on the editorial boards of the Journals of Computational Linguistics, Language Resources and Evaluations, Natural Language Engineering, Journal of Artificial Intelligence Research, IEEE Transactions on Affective Computing, and Transactions of the Association for Computational Linguistics. She was a program co-chair for Empirical Methods in Natural Language Processing 2009 and Association for Computational Linguistics (ACL) 2011, and a general chair for North American ACL 2015 and *SEM 2019. She directs multiple diversity and mentorship initiatives, including Girls Encoded and the ACL Year-Round Mentorship program. She currently serves as ACL Past President. She is the recipient of a Presidential Early Career Award for Scientists and Engineers awarded by President Obama (2009), and was named an ACM Fellow (2019) and an AAAI Fellow (2021). In 2013, she was made an honorary citizen of her hometown of Cluj-Napoca, Romania.

Keynote Talk: What’s the deal with AI safety?

Sam Bowman

New York University

Abstract: Over the last few years, a research community has been forming to study questions about the potential negative impacts of future AI systems with broadly human-level capabilities. This community was initially largely separate from academic ML, with deeper roots in philosophy departments and industry labs. This has started to change, though, with AI safety researchers increasingly focusing on questions about progress in large language models, and with safety-related motivations increasingly steering investments in NLP at large labs like OpenAI and DeepMind. This talk presents the basic goals and projects of the AI safety research community, with a focus on large language models and connections to NLP and on connections to concerns about present-day deployed language technology.

Bio: Sam Bowman is a newly-tenured associate professor at NYU and, during a 2022–2023 sabbatical year, a member of technical staff at Anthropic. At NYU, he is a member of the Center for Data Science, the Department of Linguistics, and the Courant Institute’s Department of Computer Science. His research focuses primarily on developing techniques and datasets for use in controlling and evaluating large language models, and additionally on applications of machine learning to scientific questions in linguistic syntax and semantics. He is the senior organizer behind the GLUE and SuperGLUE benchmark competitions and his work has been funded by the US NSF (including through a CAREER award), Google, Apple, Samsung, Schmidt Futures, and Open Philanthropy, among others.

Table of Contents

<i>A unified framework for cross-domain and cross-task learning of mental health conditions</i> Huikai Chua, Andrew Caines and Helen Yannakoudakis	1
<i>Critical Perspectives: A Benchmark Revealing Pitfalls in PerspectiveAPI</i> Lucas Rosenblatt, Lorena Piedras and Julia Wilkins	15
<i>Securely Capturing People’s Interactions with Voice Assistants at Home: A Bespoke Tool for Ethical Data Collection</i> Angus Addlesee	25
<i>Leveraging World Knowledge in Implicit Hate Speech Detection</i> Jessica Lin	31
<i>A Dataset of Sustainable Diet Arguments on Twitter</i> Marcus Hansen and Daniel Hershcovich	40
<i>Impacts of Low Socio-economic Status on Educational Outcomes: A Narrative Based Analysis</i> Motti Kelbessa, Ilyas Jamil and Labiba Jahan	59
<i>Enhancing Crisis-Related Tweet Classification with Entity-Masked Language Modeling and Multi-Task Learning</i> Philipp Seeberger and Korbinian Riedhammer	70
<i>Misinformation Detection in the Wild: News Source Classification as a Proxy for Non-article Texts</i> Matyas Bohacek	79
<i>Modelling Persuasion through Misuse of Rhetorical Appeals</i> Amalie Pauli, Leon Derczynski and Ira Assent	89
<i>Breaking through Inequality of Information Acquisition among Social Classes: A Modest Effort on Measuring Fun</i> Chenghao Xiao, Baicheng Sun, Jindi Wang, Mingyue Liu and Jiayi Feng	101
<i>Using NLP to Support English Teaching in Rural Schools</i> Luis Chiruzzo, Laura Musto, Santiago Gongora, Brian Carpenter, Juan Filevich and Aiala Rosa 113	
<i>Am I Answering My Job Interview Questions Right?”: A NLP Approach to Predict Degree of Explanation in Job Interview Responses</i> Raghu Verrap, Ehsanul Nirjhar, Ani Nenkova and Theodora Chaspari	122
<i>Identifying Condensing Language: A Tale of Two Distinct Phenomena?</i> Carla Perez Almendros and Steven Schockaert	130
<i>BELA: Bot for English Language Acquisition</i> Muskan Mahajan	142
<i>Applicability of Pretrained Language Models: Automatic Screening for Children’s Language Development Level</i> Byoung-doo Oh, Yoon-koung Lee and Yu-seop Kim	149
<i>Transformers-Based Approach for a Sustainability Term-Based Sentiment Analysis (STBSA)</i> Blaise Sandwidi and Suneer Pallitharammal Mukkolakal	157

Hate-CLIPper: Multimodal Hateful Meme Classification based on Cross-modal Interaction of CLIP Features

Gokul Karthik Kumar and Karthik Nandakumar 171

Program

Wednesday, December 7, 2022

13:15 - 14:00 *Opening Remarks and Introduction to NLP for Social Good*

14:00 - 14:30 *Invited Talk by Preslav Nakov*

14:30 - 14:45 *Preslav Nakov Live Q&A*

14:45 - 15:00 *Break 1*

15:00 - 16:00 *Physical Poster Session*

Securely Capturing People's Interactions with Voice Assistants at Home: A Bespoke Tool for Ethical Data Collection

Angus Addelee

Towards Countering Essentialism through Social Bias Reasoning

Emily Allaway, Nina Taneja, Sarah-jane Leslie and Maarten Sap

Enhancing Crisis-Related Tweet Classification with Entity-Masked Language Modeling and Multi-Task Learning

Philipp Seeberger and Korbinian Riedhammer

Misinformation Detection in the Wild: News Source Classification as a Proxy for Non-article Texts

Matyas Bohacek

Breaking through Inequality of Information Acquisition among Social Classes: A Modest Effort on Measuring Fun

Chenghao Xiao, Baicheng Sun, Jindi Wang, Mingyue Liu and Jiayi Feng

Hate-CLIPper: Multimodal Hateful Meme Classification based on Cross-modal Interaction of CLIP Features

Gokul Karthik Kumar and Karthik Nandakumar

Participatory Translations of Oshiwambo: Towards Culture Preservation with Language Technology

Wilhelmina Nekoto, Julia Kreutzer, Jenalea Rajab, Millicent Ochieng and Jade Abbott

Challenges and Opportunities in Information Manipulation Detection: An Examination of Wartime Russian Media

Yulia Tsvetkov, Anjalie Field, Julia Mendelsohn and Chan Young Park

Wednesday, December 7, 2022 (continued)

Gender Bias in Meta-Embeddings

Naoaki Okazaki, Danushka Bollegala and Masahiro Kaneko

Fair NLP Models with Differentially Private Text Encoders

Aurélien Bellet, Mikaela Keller, Pascal Denis and Gaurav Maheshwari

Mitigating Covertly Unsafe Text within Natural Language Systems

William Yang Wang, Kathleen McKeown, Bruce Bimber, Desmond Patton, John Judge, Emily Allaway, Melanie Subbiah, Sharon Levy, Anisha Kabir and Alex Mei

Logical Fallacy Detection

Bernhard Schoelkopf, Rada Mihalcea, Mrinmaya Sachan, Zhiheng Lyu, Yiwen Ding, Xiaoyu Shen, Tejas Vaidhya, Abhinav Lalwani and Zhijing Jin

15:00 - 16:00 *Virtual Poster Session 1*

A Dataset of Sustainable Diet Arguments on Twitter

Marcus Hansen and Daniel Hershcovich

ClimaBench: A Benchmark Dataset For Climate Change Text Understanding in English

Tanmay Laud, Daniel Spokoyny, Tom Corringham and Taylor Berg-kirkpatrick

Modelling Persuasion through Misuse of Rhetorical Appeals

Amalie Pauli, Leon Derczynski and Ira Assent

Identifying Condescending Language: A Tale of Two Distinct Phenomena?

Carla Perez Almendros and Steven Schockaert

Beyond Model Interpretability: On the Faithfulness and Adversarial Robustness of Contrastive Textual Explanations

Mariette Awad and Julia El Zini

HARALD: Augmenting Hate Speech Data Sets with Real Data

Dan Vilenchik and Tal Ilan

Wednesday, December 7, 2022 (continued)

- 16:00 - 17:00 *Lightning Talk Session*
- 17:00 - 17:30 *Break 2*
- 17:30 - 18:00 *Invited Talk by Mike Bailey*
- 18:00 - 18:30 *Invited Talk by Milind Tambe*
- 18:30 - 18:45 *Mike Bailey & Milind Tambe Live Q&A*
- 18:45 - 19:15 *Invited Talk by Rada Mihalcea*
- 19:15 - 19:45 *Invited Talk by Sam Bowman*
- 19:45 - 20:00 *Rada Mihalcea and Sam Bowman Live Q&A*
- 20:00 - 20:15 *Break 3*
- 20:15 - 21:00 *Panel*
- 21:00 - 21:20 *Interactive Session*
- 21:20 - 21:30 *Closing Remarks and Best Paper Awards*
- 21:30 - 22:30 *Virtual Poster Session 2*

A unified framework for cross-domain and cross-task learning of mental health conditions

Huikai Chua, Andrew Caines and Helen Yannakoudakis

Critical Perspectives: A Benchmark Revealing Pitfalls in PerspectiveAPI

Lucas Rosenblatt, Lorena Piedras and Julia Wilkins

Wednesday, December 7, 2022 (continued)

Impacts of Low Socio-economic Status on Educational Outcomes: A Narrative Based Analysis

Motti Kelbessa, Ilyas Jamil and Labiba Jahan

Using NLP to Support English Teaching in Rural Schools

Luis Chiruzzo, Laura Musto, Santiago Gongora, Brian Carpenter, Juan Filevich and Aiala Rosa

Am I Answering My Job Interview Questions Right?": A NLP Approach to Predict Degree of Explanation in Job Interview Responses

Raghu Verrap, Ehsanul Nirjhar, Ani Nenkova and Theodora Chaspari

Generate Me a Bedtime Story: Leveraging Natural Language Processing for Early Vocabulary Enhancement

Trevor Hall, Maria Valentini, Eliana Colunga and Katharina Kann

BELA: Bot for English Language Acquisition

Muskan Mahajan

Transformers-Based Approach for a Sustainability Term-Based Sentiment Analysis (STBSA)

Blaise Sandwidi and Suneer Pallitharammal Mukkolakal

Understanding COVID-19 Vaccine Campaign on Facebook using Minimal Supervision

Tunazzina Islam and Dan Goldwasser

Conditional Supervised Contrastive Learning for Fair Text Classification

Yuan Tian, Han Zhao, Kai-Wei Chang, Yaodong Yu, William Shand and Jianfeng Chi

Handling and Presenting Harmful Text in NLP Research

Leon Derczynski, Bertie Vidgen, Abeba Birhane and Hannah Kirk

Don't Just Clean It, Proxy Clean It: Mitigating Bias by Proxy in Pre-Trained Models

Qinlan Shen, Michael Wick, Ari Kobren and Swetasudha Panda