

NLP for Social Good: A Survey and Outlook of Challenges, Opportunities, and Responsible Deployment

Antonia Karamolegkou¹,

Angana Borah², Eunjung Cho³, Sagnik Ray Choudhury⁴, Martina Galletti^{5,6}, Pranav Gupta⁹,
Oana Ignat¹⁰, Priyanka Kargupta¹¹, Neema Kotonya¹², Hemank Lamba¹², Sun-Joo Lee¹³, Arushi Mangla¹⁰,
Ishani Mondal⁸, Fatima Zahra Moudakir, Deniz Nazar¹⁴, Poli Nemkova^{4,13}, Dina Pisarevskaya¹⁵,
Naquee Rizwan¹⁶, Nazanin Sabri¹⁷, Keenan Samway, Dominik Stammbach¹⁸, Anna Steinberg^{19,20},
David Tomás²¹, Steven R Wilson²², Bowen Yi^{2,23}, Jessica Zhu⁸, Arkaitz Zubiaga¹⁵, Anders Søgaard¹,
Alexander Fraser^{20,27}, Zhijing Jin^{3,24,25,26}, Rada Mihalcea², Joel R. Tetreault¹², Daryna Dementieva^{20,27}

¹University of Copenhagen ²University of Michigan-Ann Arbor ³ETH Zurich ⁴ University of North Texas ⁵Sony Computer Science Laboratories - Paris

⁶University of Rome “La Sapienza” ⁸University of Maryland, College Park ⁹Lowe’s ¹⁰Santa Clara University

¹¹University of Illinois Urbana-Champaign ¹²Dataminr ¹³ United Nations Development Programme (UNDP) ¹⁴University of Washington

¹⁵Queen Mary University of London ¹⁶ITI Kharagpur ¹⁷University of California San Diego ¹⁸ Princeton University ¹⁹LMU Munich

²⁰Munich Center for Machine Learning (MCML) ²¹University of Alicante ²²University of Michigan-Flint ²³University of Southern California

²⁴Max Planck Institute for Intelligent Systems, Tübingen ²⁵Vector Institute ²⁶University of Toronto ²⁷Technical University of Munich

Accepted to EACL 2026 (Main Conference), Rabat, Morocco.

Correspondence: antka@di.ku.dk, daryna.dementieva@tum.de

Abstract

Natural language processing (NLP) now shapes many aspects of our world, yet its potential for positive social impact is underexplored. This paper surveys work in “NLP for Social Good” (NLP4SG) across nine domains relevant to global development and risk agendas, summarizing principal tasks and challenges. We analyze ACL Anthology trends, finding that inclusion and AI harms attract the most research, while domains such as poverty, peacebuilding, and environmental protection remain underexplored. Guided by our review, we outline opportunities for responsible and equitable NLP and conclude with a call for cross-disciplinary partnerships and human-centered approaches to ensure that future NLP technologies advance the public good.¹

1 Introduction

“Understanding the problem is half the solution.”

— Charles Kettering

To fully realize the potential of NLP, it is essential to look beyond technical achievements and reframe tasks around pressing societal needs. We draw on insights from the United Nations Sustainable Development Goals² (SDGs) and the 2025 Global Risks (GRs) Economic Report³ to provide a foundation for an interdisciplinary recontextualization of NLP, encouraging reflection on how language technologies intersect with today’s most pressing




Figure 1: Mapping NLP applications for Social Good (NLP4SG) with global goals and risks.

challenges. We selected these two agendas as, from a social good perspective, UN SDGs offer a global framework for fostering peace and prosperity for people and the planet. However, while highly influential, these goals were established in 2015—prior to the rapid advancements in artificial intelligence. To contextualize them within today’s technological landscape, we also draw on insights from the 2025 GR Report, which highlights both the transformative potential and the emerging global risks associated with technology and information processing. The resulting mapping of NLP application domains to SDGs and GRs is shown in Figure 1. Thus, our research goal in this work is threefold: **RQ1**—what NLP-based solutions already support positive social impact, **RQ2**—what challenges arise in developing them, and **RQ3**—what promising directions remain overlooked?

¹Full paper: <https://arxiv.org/pdf/2505.22327.pdf>

²<https://sdgs.un.org/goals>

³<https://www.weforum.org/publications/global-risks-report-2025/digest>




Figure 2: Number of ACL Anthology papers per domain showing publication volume and growth trajectories over time (2019–2025).

2 Emerging Challenges and Opportunities:

Despite barriers such as *data scarcity and representational bias, misaligned evaluation metrics, and persistent safety, privacy, and ethical concerns*, alongside ongoing *infrastructural gaps*, we identify several promising directions for progress.

1. *Multimodal and multilingual learning* can help systems better reflect real-world diversity;
2. *Human–AI collaboration* enables more adaptive and interpretable NLP pipelines;
3. *Participatory design and evaluation* ensure that systems are co-developed with affected communities;
4. *Retrieval-augmented and policy-aware methods* provide tools for verifiable, context-sensitive applications;
5. *Explainability and AI literacy* foster critical engagement and equitable access.

3 Call to Action:

To advance NLP4SG, we call on the community to:

- develop joint benchmarks featuring **multilingual, culturally diverse**, and socially grounded data;
- collaborate closely with **domain experts**, such as educators, health practitioners, and civil society organizations, to co-design evaluation frameworks that reflect end-user needs;

- pursue **human-centered** methodologies instead of one-size-fits-all solutions. Progress depends on pluralistic, context-aware roadmaps that align with both local realities and global development goals.

- while modern LLMs offer significant potential, it is crucial to ensure their **affordability** and **accessibility** so they serve the public good rather than exacerbate existing inequalities.

NLP has the tools to move beyond abstract benchmarking and toward socially responsive technologies designed with—and for—impacted communities. Realizing this vision requires not just technical innovation but also sustained interdisciplinary collaboration, inclusive practices, and a commitment to long-term global equity. We hope our findings can help earlier career researchers find their research niche and that more advanced researchers will have a fresh overview of the field to foster NLP4SG applications with a more interdisciplinary paradigm.