

Report on the 1st Symposium on NLP for Social Good (NSG 2023)

Procheta Sen
University of Liverpool
UK
procheta.sen@liverpool.ac.uk

Tulika Saha
University of Liverpool
UK
tulika.saha@liverpool.ac.uk

Danushka Bollegala
University of Liverpool
UK
danushka.bollegala@liverpool.ac.uk

Abstract

Artificial intelligence (AI), specifically, Natural Language Processing (NLP) is being hailed as a new breeding ground for immense innovation potential. Researchers believe that NLP-based technologies could help to solve societal issues such as equality and inclusion, education, health, and hunger, climate action etc., and many more. Tackling these questions requires a concerted, collaborative effort across all sectors of society. The first Symposium on NLP for Social Good (NSG) was a novel effort that aimed to enable NLP researchers and scholars from inter-disciplinary field who want to think about the societal implications of their work for solving humanitarian and environmental challenges. The objective of the symposium was to support fundamental research and engineering efforts and empower the social sector with tools and resources, while collaborating with partners from all sectors to maximise effect in solving problems within public health, nature & society, accessibility, crisis response etc. In its inception, we invited speakers from academia and industry to provide an overview of some areas from NLP applications such as education, healthcare and legal domains in order to provide a platform to stimulate discussion regarding the current state of NLP in these varied fields.

Date: 8–9 June 2023.

Website: <https://nlp4social.github.io/nlp4socialgood/>.

1 Introduction and Motivation for NSG

Artificial Intelligence (AI) has widespread impact on today’s world. Recent advancement in deep learning and machine learning methods have made AI models more effective in terms of predictive and generative capabilities. As a result of this, a thread of recent research studies have applied AI to solve different societal issues (e.g. Healthcare, Education etc.). The goal of applying AI in

different areas of society is to improve the overall quality of life for common people. There has been a number of recent workshops and conference tracks which were focused on AI for social good. ACL 2021 hosted the first ever workshop on NLP for Positive Impact [Field et al., 2021]. ACL has also started a special track on AI for social Good. HARVARD CRCS group had also organised a AI for Social Good workshop in 2021¹. A recent workshop on AI for social good was held in AAAI 2023².

With this idea in mind, we propose a symposium called “NLP for Social Good” to discuss about the societal and environmental issues where NLP has already shown promising results and discover new areas where there is potential for the application of AI. Additionally, the objective of the Symposium on NSG was also to open up potential opportunities for collaboration across universities and also between academia and the industry.

2 NSG 2023

NSG-2023 (website link ³) was held as an hybrid event from 8th to 9th June, 2023 in the University of Liverpool (UoL), United Kingdom. The symposium was organized by NLP academicians from the Department of Computer Science, UoL. The event hosted two keynote speakers and four invited talks from academia to share their research works related to NSG and their insights in regards to the potential of NSG. There were 177 registered participants for NSG 2023 from across the globe out of which 62.14% were “males”, 35.59% were “females” and 0.03% were from “others”. In addition to the participation from United Kingdom, there were participants from 42 other countries. The next section describes in detail the keynote and invited talks at NSG 2023.

2.1 Invited Talks at NSG 2023

Here, we enlist the six keynote/invited talks delivered in NSG 2023. We are grateful to all the speakers for their insightful talks and the attendees for the meaningful discussion in the Q&A sessions. The video recordings of the invited talks are available on the YouTube channel⁴.

2.1.1 Keynote Talk #1 by Prof. Iryna Gurevych

Speaker Bio: The first keynote speaker, Prof. Iryna Gurevych, is a Professor of Computer Science and Director of the Ubiquitous Knowledge Processing (UKP) Lab at the Technical University (TU) of Darmstadt in Germany. Her main research interests are in machine learning for large-scale language understanding and text semantics. Iryna’s work has received numerous awards like ACL fellow award 2020 and the first-ever Hessian LOEWE Distinguished Chair award (2,5 mil. Euro) in 2021. She is the co-director of the NLP program within ELLIS, a network of excellence in machine learning. She is currently the president of the Association for Computational Linguistics. In 2022, she has been awarded an ERC Advanced Grant.

¹<https://crcs.seas.harvard.edu/event/ai-social-good-workshop-2020>

²<https://aaai-23.aaai.org/ws23/>

³<https://nlp4social.github.io/nlp4socialgood/>

⁴<https://www.youtube.com/@NLPforSocialGood>

Title: NLP meets Psychotherapy: from Estimating Depression Severity to Estimating the Client’s Well-Being

Abstract: Language plays a crucial role in psychotherapy and the related processes. Linguistic expression not only reveals emotional states, but also provides insight into the immediate cognitions of the client. In this talk, she presented two projects related to the use of NLP for estimating the client’s depression severity and well-being. In the first project, they examined emotional coherence between the subjective experience of emotions and emotion expression in therapy and whether this coherence is associated with clients’ well-being. She presented an end-to-end approach where they used emotion predictions from their transformer-based emotion recognition model to study emotional coherence and its diagnostic potential in psychotherapy research. In the second project, her team proposed an efficient semantic pipeline to study depression severity in individuals based on their social media posts. They represented each individual by a set of selected sentences and used them as evidence for predicting the individual’s symptom severity. For that, they explored different aggregation methods to answer one of four Beck Depression Inventory (BDI) options per symptom. She concluded by a short overview of her further ongoing and planned projects in NLP for Mental Health.

2.1.2 Keynote Talk #2 by Prof. Amit Sheth

Speaker Bio: The second keynote speaker, Prof. Amit Sheth, is an Educator, Researcher, and Entrepreneur. He is the founding director of the university-wide AI Institute at the University of South Carolina. He is a Fellow of IEEE, AAI, AAAS, ACM and AAIA. His awards include IEEE CS W. Wallace McDowell and IEEE TVSVC Research Innovation awards. He has (co-)founded four companies, including the first Semantic Search company in 1999 that pioneered technology similar to what is found today in Google Semantic Search and Knowledge Graph, ezDI which developed knowledge-infused clinical NLP/NLU, and Cognovi Labs at the intersection of emotion and AI.

Title: Language Understanding using Neuro-symbolic AI: Why we need varied, comprehensive, and stratified knowledge

Abstract: ‘Data alone is not enough.’ This was the section heading in Pedro Domingos’ 2012 seminal paper. I have been a believer in this for a long time. In our Semantic Search engine, commercialized in 2000, we complemented machine learning classifiers with a comprehensive World-Model™ or knowledge bases (now referred to as knowledge graphs) for improved named entity and relationship extraction and semantic search. It was an early demonstration of the complementary nature of data-driven statistical learning (since replaced by neural networks) and knowledge-supported symbolic AI methods. In this talk, I want to observe three important issues about the Why, What, and How of using knowledge in neuro-symbolic AI systems to advance from NLP to NLU. While the transformer-based models have achieved tremendous success in many NLP tasks, the pure data-driven approach comes up short when we need NLU, where knowledge is key to understanding the language, as required for the explanation, safety, and ensuring adherence to decision-making processes that must be followed (e.g., in clinical diagnosis). Throughout

the talk, I will use examples from the social good domains, specifically mental health and addiction, to demonstrate the need for "understanding" (for example, for safety with explanations) and why/how knowledge-infused learning offers better outcomes compared to data-driven only alternatives."

2.1.3 Invited Talk #3 by Dr. Saptarshi Ghosh

Speaker Bio: The third speaker, Dr. Saptarshi Ghosh, is an Associate Professor at the Department of Computer Science and Engineering, Indian Institute of Technology, Kharagpur. His research interests include Law and AI, Social media analytics, and Algorithmic bias and fairness (on which he presently leads a Max Planck Partner Group at IIT Kharagpur). His works have been awarded at top Law-AI conferences, including the Best Paper Award at JURIX 2019 and the Best Student Paper Award at ICAIL 2021.

Title: NLP for the Legal Domain: Challenges and Opportunities.

Abstract: Legal NLP - the application of NLP over various types of legal text - has become an important application domain for NLP, primarily because of the recent proliferation of publicly available legal data, and the socio-economic benefits of mining legal insights. There are several associated challenges, ranging from the need to deal with complex legal language, to the extreme length of legal documents, to the need for explainability in this domain. This talk discussed some of the challenges and opportunities in applying NLP in the legal domain. The talk also discussed some of the popular research problems being studied by the Legal NLP community, including summarization of long legal documents, complex classification tasks over legal documents, and pretrained language models for the legal domain.

2.1.4 Invited Talk #4 by Dr. Mrinmaya Sachan

Speaker Bio: Dr. Mrinmaya Sachan is an Assistant Professor of Computer Science at ETH Zurich. His research is in the area of Natural language processing and the interface of Machine learning and Education. He has received several awards for his work, including an outstanding paper award at ACL 2015, an IBM PhD fellowship, the Siebel scholarship and the CMU CMLH fellowship.

Title: NLP for Education

Abstract: This talk showed the application of large language models (LLMs) to support education in various ways, such as creating practice problems and scaffolds, providing meaningful feedback to students, having a Socratic dialog with students, etc. - while - customizing the learning experiences of students according to their own individual needs. Finally, I will share some challenges as well as opportunities in this emerging field.

2.1.5 Invited Talk #5 by Dr. Zaiqiao Meng

Speaker Bio: Zaiqiao Meng is currently a Lecturer in the IDA section of the University of Glasgow. Previously, he worked as a Postdoctoral Researcher at the Language Technology Laboratory of the University of Cambridge and the Terrier team of the University of Glasgow. He currently leads a research team consisting of PhDs, masters, and undergraduate students working on Natural Language Processing, Knowledge Extraction, Representation & Reasoning Learning, particularly some biomedical applications.

Title: Probing and Infusing Biomedical Knowledge for Pre-trained Language Models

Abstract: Pre-trained language models (PLMs) have made significant advancements in few- or zero-shot language understanding tasks by pre-training model parameters in a task-agnostic way and transferring knowledge to specific downstream tasks via fine-tuning. The integration of factual knowledge from knowledge graphs (KGs) is crucial for knowledge-intensive tasks such as question-answering and fact-checking. In the biomedical domain, where public training data is limited and noisy, trusted biomedical KGs are essential for accurate inferences. Therefore, measuring the amount of world knowledge stored in PLMs (knowledge probing) and integrating factual knowledge into these models (knowledge infusion) are challenging tasks in the NLP field. In this talk, we will introduce a new biomedical knowledge probing benchmark, MedLAMA, and a novel probing approach called Contrastive Probe, for probing biomedical knowledge of PLMs. We will also present our proposed knowledge infusion approach, Mixture-of-Partitions (MoP), which infuses factual knowledge based on partitioned KGs into PLMs and automatically routes useful knowledge from these adapters to downstream tasks. Finally, we will review ChatGPT's impact on the healthcare domain.

2.1.6 Invited Talk #6 by Inna Wanyin Lin

Speaker Bio: Inna Wanyin Lin is a PhD student in the Paul G. Allen School of Computer Science & Engineering at the University of Washington, working with Prof. Tim Althoff. Her research focuses on natural language processing and human-AI collaboration for health and social applications. She has received Marilyn Fries Endowed Regental Fellowship and the WWW 2021 Best Paper Award. Inna's research has been covered internationally by news outlets including The Wall Street Journal, The Washington Post, NPR, Deutschlandfunk, and infobae.

Title: Human-Centered NLP for Mental Health.

Abstract: The increasing prevalence of mental health issues worldwide has amplified the need for innovative and effective solutions to address the growing demand for mental healthcare. Recent NLP advancements have offered the promising potential to tackle this issue. In this talk, I will delve into the opportunities and risks in developing NLP applications for mental health by discussing two recent projects. I will first talk about HAILEY, an AI-in-the-loop agent that provides just-in-time feedback to help online peer supporters respond more emphatically to those seeking help. Through a randomized controlled trial with real-world peer supporters on Talklife, a large online peer-to-peer support platform, we demonstrate that our Human-AI collaboration

approach significantly increases conversational empathy. Despite promising opportunities, the use of language technologies also presents risks. In this regard, I will discuss our recent project examining gendered mental health stigma in masked language models, which are frequently used in downstream applications related to mental health. We find that masked language models capture societal stigma about gender in mental health, and different models capture dimensions of stigma differently for men and women. Overall, I hope to shed light on the opportunities and challenges in building effective and ethical NLP applications for mental healthcare.

2.2 Organizers of NSG-2023

- **Dr. Procheta Sen**⁵ is a Lecturer (Assistant Professor) in the Department of Computer Science at University of Liverpool (UOL), United Kingdom. She is a part of the NLP group at UOL. Before joining UOL, she worked with Emine Yilmaz as a postdoctoral researcher in the Web Intelligence Group at University College London (UCL). She finished her PhD. in 2021 from Dublin City University, Ireland. Her recent research interest lies in building intelligent systems fair and explainable. Her work on explaining information retrieval systems from an IR practitioner’s perspective and measuring trust of an information retrieval system got published in SIGIR 2020 and CIKM 2022 respectively. She has given a tutorial on Explainability for IR and NLP models at FIRE 2022, SIGIR 2023. She was one of the organizers of the first symposium on NLP for Social Good at the University of Liverpool. Some of her publications are as follows :
 1. Automated Argument Generation from Legal Facts [Tuvey and Sen, 2023]. arXiv preprint arXiv:2310.05680.
 2. Towards socially responsible ai: Cognitive bias-aware multi-objective learning [Sen and Ganguly, 2020]. In Proceedings of the AAAI Conference on Artificial Intelligence.
- **Dr. Tulika Saha**⁶ (corresponding author) is a Lecturer/Assistant Professor of Computer Science in the University of Liverpool, United Kingdom (UK). Her current research interests include NLP typically Dialogue Systems, Social Good, Social Media Analysis using Deep Learning and Reinforcement Learning etc. She was a postdoctoral research fellow at the National Centre for Text Mining, University of Manchester, UK. Previously she earned her Ph.D. from Indian Institute of Technology Patna, India. Her research articles are published in top-tier conferences such as ACL, ACM SIGIR, EMNLP, NAACL, CoLing etc. and several peer-reviewed journals. She has organized several workshops/symposium namely NLP for Social Good 2023, AICAI 2023⁷, LLMIT 2023 [Saha et al., 2023a] and presented tutorials in ECIR [Saha et al., 2023b], InterSpeech⁸ 2023 etc. Some of her recent publications in the Social Good theme are as follows :
 1. A shoulder to cry on: towards a motivational virtual assistant for assuaging mental agony [Saha et al., 2022b]. In Proceedings of the 2022 conference of the North American chapter of the association for computational linguistics: Human language technologies.

⁵<https://procheta.github.io/sprocheta/>

⁶<https://sahatulika15.github.io/>

⁷<https://aicai2023.site/>

⁸<https://interspeech2023.org/tutorials/>

-
2. Towards motivational and empathetic response generation in online mental health support [Saha et al., 2022a]. In Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval.
 3. Symptoms are known by their companies: towards association guided disease diagnosis assistant [Tiwari et al., 2022]. BMC bioinformatics, 23(1).
 4. Domain infused conversational response generation for tutoring based virtual agent [Jain et al., 2022]. In 2022 International Joint Conference on Neural Networks (IJCNN), IEEE.
 5. T-VAKS: A Tutoring-based Multimodal Dialog System via Knowledge Selection [Jain et al., 2023]. In Proceedings of the 26th European Conference on Artificial Intelligence, 2023.
 6. Can you Summarize my learnings? Towards Multi-modal Perspective-based Educational Dialogue Summarization. In Findings of EMNLP, 2023.

- **Prof. Danushka Bollegala**⁹ is a Professor in the Department of Computer Science at the University of Liverpool and a Amazon Scholar with Amazon Search. He has published over 190 peer-reviewed papers related to NLP/ML and has worked as a Program Chair, Senior Area chair, Senior PCs at NLP/AI conferences. His research interests are lexical and compositional semantics, information extraction, domain adaptation and text summarization. He has received various awards such as the conference best paper awards (GECCO, PRICAI) and IEEE Young Author Award in the past for his research excellence. Some of his publications relevant to the theme are as follows :

1. Unmasking the Mask – Evaluating Social Biases in Masked Language Models [Kaneko and Bollegala, 2022], Proc. of the 36th AAI Conference on Artificial Intelligence, 2022.
2. Sense Embeddings are also Biased – Evaluating Social Biases in Static and Contextualised Sense Embeddings [Zhou et al., 2022], Proc. of the 60th Annual Meeting of the Association for Computational Linguistics, 2022.
3. Debiasing isn't enough! – On the Effectiveness of Debiasing MLMs and their Social Biases in Downstream Tasks [Kaneko et al., 2022a], Proc. of the 29th International Conference on Computational Linguistics (COLING 2022).
4. Gender Bias in Meta-Embedding [Kaneko et al., 2022b], Proc. of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP 2022).
5. The Impact of Debiasing on the Performance of Language Models in Downstream Tasks is Underestimated [Kaneko et al., 2023], Proceedings of the 13th IJCNLP and the 3rd Conference of the ACL, 2023.

3 Concluding Remarks

In NSG 2023, we discussed about the emerging topics of NLP which has relevance to social good such as Healthcare, Legal-NLP and Education-NLP. The discussions were broadly categorised into

⁹<https://danushka.net/>

two areas. One was focused on the challenges and accessibility of the datasets to pursue research in these areas. The other one was about the application of the recent and advanced NLP and Machine Learning techniques to solve the different challenges in these areas. The discussions with the invited speakers were useful for both students and senior researchers attending NSG.

One of the core contributions of this Symposium is its ability to foster interdisciplinary engagement. NSG transcends traditional academic boundaries, attracting researchers from various fields such as computer science, linguistics, social sciences, public policy, and more. This diversity of perspectives enriches the academic discourse and promotes innovative approaches to solving real-world problems and creates a thriving ecosystem of knowledge exchange and research partnerships. Through this Symposium on NSG, we emphasize the importance of addressing global challenges, such as healthcare, education, environmental sustainability, and social justice etc. This not only widens the scope of NLP research but also encourages researchers to tackle critical issues, resulting in research that has a meaningful impact on society.

The video recordings of NSG-2023 is available on the Symposium on NLP for Social Good (NSG) YouTube channel¹⁰. We plan to organise future editions of NSG, incorporating all the feedback and suggestions obtained from NSG-2023.

Acknowledgements

The organizers express their gratitude to all the keynote/invited speakers and the attendees for their invaluable contributions to the symposium. The organizers would also like to extend heartfelt thanks to the Dept. of CS, University of Liverpool for providing funding to organize NSG-2023.

References

- Anjalie Field, Shrimai Prabhumoye, Maarten Sap, Zhijing Jin, Jieyu Zhao, and Chris Brockett. Proceedings of the 1st workshop on nlp for positive impact. In *Proceedings of the 1st Workshop on NLP for Positive Impact*, 2021.
- Raghav Jain, Tulika Saha, Souhitya Chakraborty, and Sriparna Saha. Domain infused conversational response generation for tutoring based virtual agent. In *2022 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2022.
- Raghav Jain, Tulika Saha, and Sriparna Saha. T-vaks: A tutoring-based multimodal dialog system via knowledge selection. 2023.
- Masahiro Kaneko and Danushka Bollegala. Unmasking the mask–evaluating social biases in masked language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11954–11962, 2022.
- Masahiro Kaneko, Danushka Bollegala, and Naoaki Okazaki. Debiasing isn’t enough!–on the effectiveness of debiasing mlms and their social biases in downstream tasks. *arXiv preprint arXiv:2210.02938*, 2022a.
- Masahiro Kaneko, Danushka Bollegala, and Naoaki Okazaki. Gender bias in meta-embeddings. *arXiv preprint arXiv:2205.09867*, 2022b.

¹⁰<https://www.youtube.com/@NLPforSocialGood>

-
- Masahiro Kaneko, Danushka Bollegala, and Naoaki Okazaki. The impact of debiasing on the performance of language models in downstream tasks is underestimated. *arXiv preprint arXiv:2309.09092*, 2023.
- Tulika Saha, Vaibhav Gakhreja, Anindya Sundar Das, Souhitya Chakraborty, and Sriparna Saha. Towards motivational and empathetic response generation in online mental health support. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*, pages 2650–2656, 2022a.
- Tulika Saha, Saichethan Reddy, Anindya Das, Sriparna Saha, and Pushpak Bhattacharyya. A shoulder to cry on: towards a motivational virtual assistant for assuaging mental agony. In *Proceedings of the 2022 conference of the North American chapter of the association for computational linguistics: Human language technologies*, pages 2436–2449, 2022b.
- Tulika Saha, Debasis Ganguly, Sriparna Saha, and Prasenjit Mitra. Large language models’ interpretability and trustworthiness (llmit). 2023a.
- Tulika Saha, Abhishek Tiwari, and Sriparna Saha. Trends and overview: The potential of conversational agents in digital health. In *European Conference on Information Retrieval*, pages 349–356. Springer, 2023b.
- Procheta Sen and Debasis Ganguly. Towards socially responsible ai: Cognitive bias-aware multi-objective learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(03):2685–2692, Apr. 2020.
- Abhishek Tiwari, Tulika Saha, Sriparna Saha, Pushpak Bhattacharyya, Shemim Begum, Minakshi Dhar, and Sarbajeet Tiwari. Symptoms are known by their companies: towards association guided disease diagnosis assistant. *BMC bioinformatics*, 23(1):1–23, 2022.
- Oscar Tuvey and Procheta Sen. Automated argument generation from legal facts, 2023.
- Yi Zhou, Masahiro Kaneko, and Danushka Bollegala. Sense embeddings are also biased—evaluating social biases in static and contextualised sense embeddings. *arXiv preprint arXiv:2203.07523*, 2022.