PrivateNLP 2024

The Fifth Workshop on Privacy in Natural Language Processing

Proceedings of the Workshop

August 15, 2024

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 979-8-89176-139-1

Introduction

Welcome to the Fifth Workshop on Privacy in Natural Language Processing. Co-located with ACL 2024 in Bangkok, Thailand, the workshop is scheduled for August 15, 2024. To facilitate the participation of the global NLP community, we continue running the workshop in a hybrid format.

Privacy-preserving language data processing has become essential in the age of Large Language Models (LLMs) where access to vast amounts of data can provide gains over tuned algorithms. A large proportion of user-contributed data comes from natural language e.g., text transcriptions from voice assistants. It is therefore important to curate NLP datasets while preserving the privacy of the users whose data is collected, and train ML models that only retain non-identifying user data. The workshop brings together practitioners and researchers from academia and industry to discuss the challenges and approaches to designing, building, verifying, and testing privacy preserving systems in the context of Natural Language Processing.

Our agenda features a keynote speech, hybrid talk sessions both for long and short papers, and a poster session. This year we received 29 submissions. We accepted 23 submissions after a thorough peer-review. Five accepted submissions preferred the non-archival option and thus are not included in this proceedings. Moreover, our poster session features additional four ACL-Findings papers.

We would like to deeply thank to all the authors, committee members, keynote speaker, and participants to help us make this research community grow both in quantity and quality.

Workshop Chairs

Organizing Committee

Program Chairs

Ivan Habernal, Ruhr-University Bochum, Germany
Sepideh Ghanavati, University of Maine, United States
Abhilasha Ravichander, Allen Institute for AI, United States
Vijayanta Jain, University of Maine, United States
Patricia Thaine, Private AI, Canada
Timour Igamberdiev, Technical University of Darmstadt, Germany
Niloofar Mireshghallah, University of Washington, United States
Oluwaseyi Feyisetan, Amazon, United States

Program Committee

Program Committee

Andrea Atzeni, Polytechnic Institute of Turin
Asma Aloufi, Taif University
Eleftheria Makri, Leiden University
Erion Cano, Universität Paderborn
Eugenio Martínez-Cámara, Universidad de Jaén
Gergely Acs, Technical University of Budapest
Isar Nejadgholi, National Research Council Canada
Jaydeep Borkar, Northeastern University
Kambiz Ghazinour, SUNY Canton
Lizhen Qu, Monash University
Mattia Salnitri, Polytechnic Institute of Milan
Mousumi Akter, Technische Universität Dortmund
Natasha Fernandes, Macquarie University
Pengwei Li, Meta
Peter Story, Clark University
Pierre Lison, Norwegian Computing Center
Rocky Slavin, University of Texas at San Antonio
Ruyu Zhou, University of Notre Dame
Sai Peddinti, Google
Sebastian Ochs, Technische Universität Darmstadt
Shomir Wilson, Pennsylvania State University
Timour Igamberdiev, Technische Universität Darmstadt
Travis Breaux, Carnegie Mellon University

Table of Contents