PrivateNLP 2021

The Third Workshop on Privacy in Natural Language Processing

Proceedings of the Workshop

June 11, 2021

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 978-1-954085-43-5

Introduction

The PrivateNLP workshop aims to bring 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.

This year, the workshop accepted 7 papers and one non-archival paper. These accepted papers cover federated learning, text perturbation mechanisms, privacy preserving language models, and secure multiparty computation.

We have 2 invited speakers: Travis Breaux (Carnegie Mellon University) and Adam Dziedzic (Vector Institute and The University of Toronto).

We would like to thank the Program Committee members who kindly reviewed the submissions, as well as the invited speakers, and the workshop co-organizers, Oluwaseyi Feyisetan (Amazon, USA), Sepideh Ghanavati (University of Maine, USA), Shervin Malmasi (Amazon, USA), and Patricia Thaine (University of Toronto, Canada).

Organizing Committee

Oluwaseyi Feyisetan, Amazon (USA)
Sepideh Ghanavati, University of Maine (USA)
Shervin Malmasi, Amazon (USA)
Patricia Thaine, University of Toronto (Canada)

Program Committee

Aleksei Triastcyn (Ecole Polytechnique Federale de Lausanne)
Andreas Nautsch (EURECOM)
Asma Eidhah Aloufi (Rochester Institute of Technology)
Balazs Pejo (Budapest University of Technology and Economics)
Benjamin Zi Hao Zhao (University of New South Wales)
Briland Hitaj (SRI International)
Christian Weinert (Technische Universitat Darmstadt)
Congzheng Song (Apple)
Dinusha Vatsalan (Data61-CSIRO)
Eleftheria Makri (Saxion University)
Fang Liu (University of Notre Dame)
Gerald Penn (University of Toronto)
Isar Nejadgholi (National Research Council Canada)
Jamie Hayes (University College London)
Liwei Song (Princeton)
Mohamed Abdalla (University of Toronto)
Natasha Fernandes (Macquarie University)
Sai Teja Peddinti (Google)
Shomir Wilson (Pennsylvania State University)
Tom Diethe (Amazon)
Travis Breaux (Carnegie Mellon University)
Xavier Ferrer (King's College London)

Invited Speakers

Adam Dziedzic (Vector Institute and The University of Toronto)
Travis Breaux (Carnegie Mellon University)

Table of Contents

Conference Program

Understanding Unintended Memorization in Language Models Under Federated Learning
Om Dipakbhai Thakkar, Swaroop Ramaswamy, Rajiv Mathews and Francoise Beaufays

On a Utilitarian Approach to Privacy Preserving Text Generation
Zekun Xu, Abhinav Aggarwal, Oluwaseyi Feyisetan and Nathanael Teissier

Learning and Evaluating a Differentially Private Pre-trained Language Model
Shlomo Hoory, Amir Feder, Avichai Tendler, Alon Cohen, Sofia Erell, Itay Laish, Hootan Nakhost, Uri Stemmer, Ayelet Benjamini, Avinatan Hassidim and Yossi Matias

An Investigation towards Differentially Private Sequence Tagging in a Federated Framework
Abhik Jana and Chris Biemann

A Privacy-Preserving Approach to Extraction of Personal Information through Automatic Annotation and Federated Learning
Rajitha Hathurusinghe, Isar Nejadgholi and Miodrag Bolic

Using Confidential Data for Domain Adaptation of Neural Machine Translation
Sohyung Kim, Arianna Bisazza and Fatih Turkmen

Private Text Classification with Convolutional Neural Networks
Samuel Adams, David Melanson and Martine De Cock