HiWi for NLP Wanted

PrivaLingo: Privacy-Preserving Machine Translation

Motivation

The **PrivaLingo** project aims to design a privacy-preserving neural machine translation (NMT) system by incorporating various state-of-the-art privatization methods into the training regime of neural network models. We particularly work in the paradigm of differential privacy (DP), which provides formal privacy guarantees for data analysis algorithms. We are currently designing a modular and open-source framework for private NMT, using the JAX and Flax libraries. The framework will include various NMT architectures (LSTMs, transformer-based models) and datasets, trained with the DP-SGD algorithm.

Task Description

- Prepare a modular framework for private NMT models using the JAX/Flax libraries
- Train a variety of NMT models with and without DP on different datasets (e.g. WMT 2016 English-German), comparing performance using various evaluations (e.g. BLEU, ROUGE)
- Compile a literature review of work on private NMT (paper summaries, including descriptions of methodology and results)

Requirements

- Proficiency in Python programming, particularly with ML/DL libraries
- Basic knowledge of machine learning and deep learning for NLP
- Proficiency in English

Application

- Curriculum Vitae
- Transcript
- Starting date: As early as possible
- Preferably 40-80 hrs / month

Contact

Dr. Ivan Habernal

Timour Igamberdiev

timour.igamberdiev@tu-darmstadt.de