Why Crosslingual Alignment Fails for Better Crosslingual Transfer

- **Supervisor**: Yihong Liu
- **Examiner**: Prof. Hinrich Schütze
- **Open**: MSc
- **Summary**: Crosslingual alignment refers to the degree of similarity among representations of similar meanings across languages. Therefore, it is often believed that such alignment can benefit better crosslingual transfer. There have been quite a few methods to improve the crosslingual alignment given a pretrained model, typically with a contrastive learning objective. Among them, using parallel corpora to construct the contrast pairs is the mainstream [1, 2]. Additionally, better alignment can be achieved using monolingual corpora, through data augmentation like dropout [3, 4] or transliterations [5, 6]. Nevertheless, explicit alignments do not always contribute to better crosslingual transfer performance in downstream tasks [7, 8]. Unfortunately, the previous studies only offer empirical findings but do no investigate the reasons behind those discoveries. In this project, we want to explore both empirically and theoretically. We will revisit different alignment methods and observe their effect on the target mPLM in a controlled setting. We will analyze the capability of the model and how the representation space varies after alignment. This will finally contribute to a better understanding of crosslingual alignment and its effect on the model and representation space.

- **Prerequisites**: enthusiasm, good mathematical background and programing background (preferably python), good knowledge of NLP, a good command of DL framework (preferably PyTorch)

[1] https://arxiv.org/pdf/2007.07834
[2] https://arxiv.org/pdf/2007.15960
[3] https://arxiv.org/pdf/2212.08378
[4] https://arxiv.org/pdf/2105.11741
[5] https://arxiv.org/pdf/2401.06620
[6] https://arxiv.org/pdf/2406.19759
[7] https://aclanthology.org/2020.emnlp-main.362.pdf
[8] https://aclanthology.org/2023.findings-acl.189.pdf