

When Learning-to-Rank Meets Federated Learning

—An Explorative Study Based on RankNet—

Atsuya Kanazawa

In recent years, deep learning techniques have attracted great attention as they yield state-of-the-art performance in many fields, such as image recognition and natural language processing. However, the artificial intelligence (AI) community still faces two major challenges. The first is that in most industries, data exists in the form of isolated islands. The second is the strengthening of data privacy and security. In order to achieve privacy-preserving machine learning, the approach of federated learning has been proposed, which leaves the training data distributed on the mobile devices and learns a shared model by aggregating locally-computed updates. In this thesis, we focus on the problem of privacy in the context of learning-to-rank. In particular, we aim to conduct an explorative study on how to deploy federated learning in the context of learning-to-rank. Specifically, we investigate the effect of varying the number of clients and the number of data owned by each client on the performance of a ranking model. The experimental results show that: both the size of data and the number of clients significantly affects the performance of RankNet when it is trained in a federated manner. Careful examinations of these factors are highly recommended in the development of federated learning-to-rank methods.

(Advisor: Hai-Tao Yu)