Abstract

This dissertation aims to explore the application of personalised federated learning for the detection of pneumonia in chest X-ray images. The primary goal of this project is to investigate the challenges posed by non-IID data across numerous distributed datasets, and to explore techniques to mitigate the effects of non-IID data.

Making use of publicly available datasets, a federated learning model for the detection of pneumonia in chest X-ray images is created. Techniques inspired by the state-of-the-art are then implemented, to mitigate the effects of non-IID data on the performance of the models. The findings demonstrate the importance of personalisation techniques in improving the performance of federated learning models, particularly in the context of medical image classification.