Federated harmonization of x-ray images of femur fractures

Background

Drugs commonly used in the treatment of osteoporosis (bisphosphonates) inhibit cell function of one specific cell type in bone, leading to increased bone mass and reduced fracture risk. This treatment has been used successfully for decades. Long-term inhibition of these bone-specific cells has recently been shown to cause bone material insufficiency, leading to spontaneous stress fractures in the thigh bone – Atypical Femoral Fractures (AFF). These fractures show features on x-ray images that differentiate them from Normal Femur Fractures (NFF). However, these features are very subtle and can easily be overlooked if not specifically sought for (Figure 1). The detection rate of AFF on clinical plain radiographs is <7%, and reports of drug adverse reactions to the Swedish Drug Agency have an even lower detection rate. While these events are rare compared to fractures that can be prevented, they are of clinical concern and have resulted in decreased use of these medications. As these events are so rare, standard statistical models have failed to identify reliable risk factors that would allow a precision medicine approach to identifying which patients to treat and for how long.

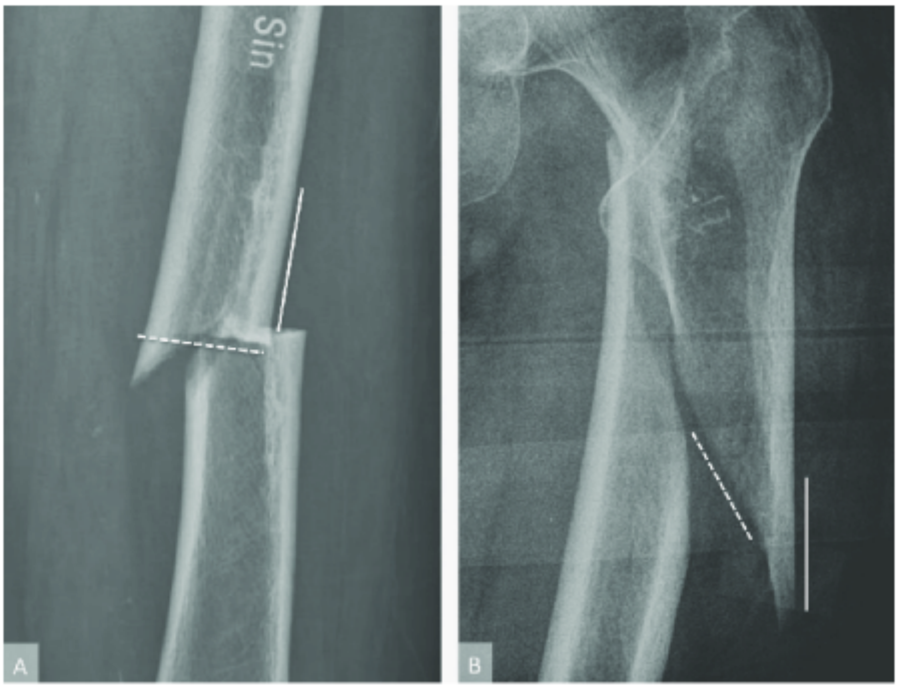

Figure 1. Atypical femur fractures (AFF, image A) and normal femur fractures (NFF, image B).

A general problem in medical imaging is that it is difficult to share data between hospitals or researchers. Another problem is that images from different hospitals have different quality, resolution and appearance due to different imaging equipment. Federated learning can be a solution to using larger datasets; a computer (node) at each hospital joins a federation that together trains a network to classify images as AFF or NFF. No images are sent between the nodes during the training, only updates of the classifier. We have previously had two master theses on federated classification of AFF/NFF, and now want to investigate if the x-ray images can be harmonized without having all images in one computer.

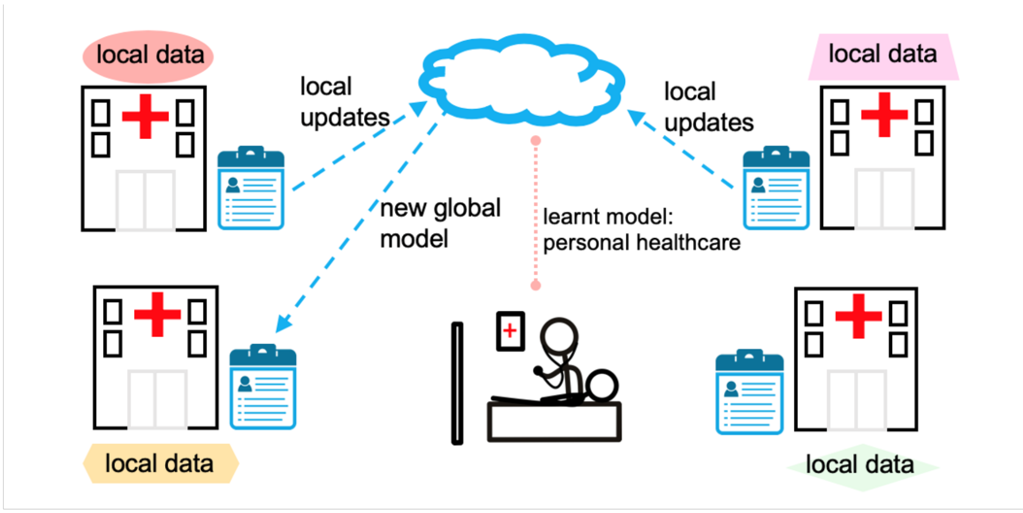

The main idea in federated learning is to not store all data in a single computer, but to instead store for example image data locally at each hospital. Instead of sending medical images and other medical data between the hospitals, the hospitals send updates, or parameters, of deep learning models. This process is then iterated to convergence.

Objectives

Investigate how different the images from the 72 hospitals are, for example using metrics from radiomics.

Implement and evaluate centralized harmonization of the images, i.e. having all images in one computer, to see how it affects AFF/NFF classification performance (using a CNN or a vision transformer) compared to using standard image augmentation during training. The harmonization can for example be done using GANs (generative adversarial networks) or diffusion models.

Implement and evaluate federated harmonization of the images, to see how it affects classification performance. How good is federated harmonization compared to centralized harmonization?

Investigate how big the difference is between centralized and federated harmonization when using different solutions for privacy, such as differential privacy.

Data

About 4300 X-ray images from some 1200 patients, of which about 20% are AFF. These images originate from 72 Swedish hospitals, and the data can be split into 3-6 nodes.

Required background

Machine learning, deep learning, Python programming

Computing resources

The student will have access to very good computing resources (graphics cards) for federated learning. It may also be possible to use the supercomputer Berzelius (752 graphics cards) for simulations.

Contact persons

Anders Eklund, anders.eklund@liu.se, Department of Biomedical Engineering, Department of Computer and Information Science, Center for Medical Image Science and Visualization

Jörg Schilcher, jorg.schilcher@liu.se, Department of Orthopaedic Surgery, University Hospital Linköping