Synthetic 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.

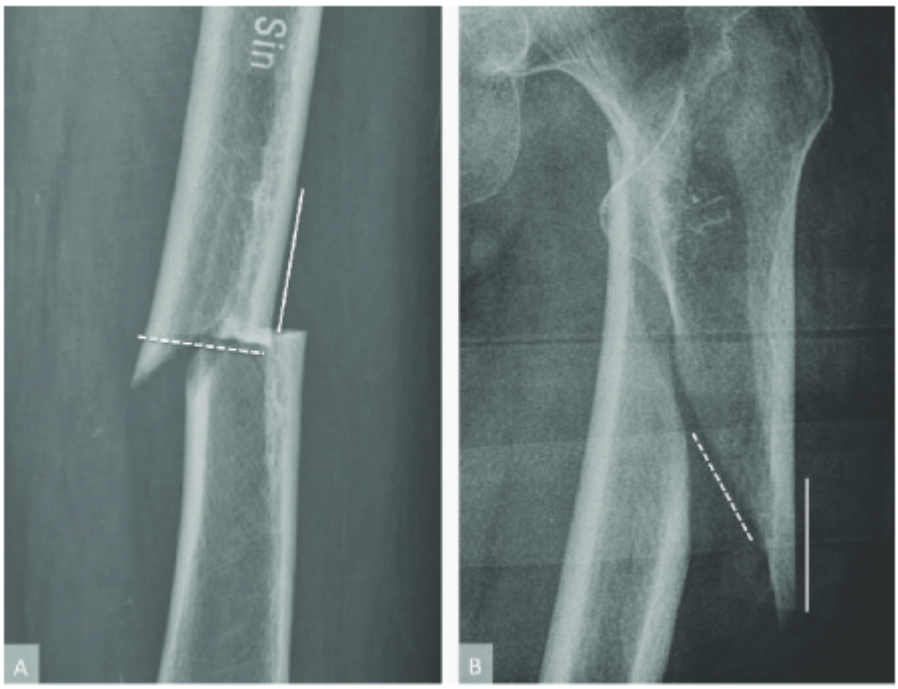

Figure 1. Atypical femur fractures (AFF, image A) and normal femur fractures (NFF, image B).

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. Synthetic images from GANs (generative adversarial networks) and diffusion models can potentially be shared freely, as they do not belong to a specific patient. We have previously had several master theses on

generating synthetic images of brain tumors, and now want to investigate if GANs and diffusion models can also generate synthetic x-ray images of femur fractures.

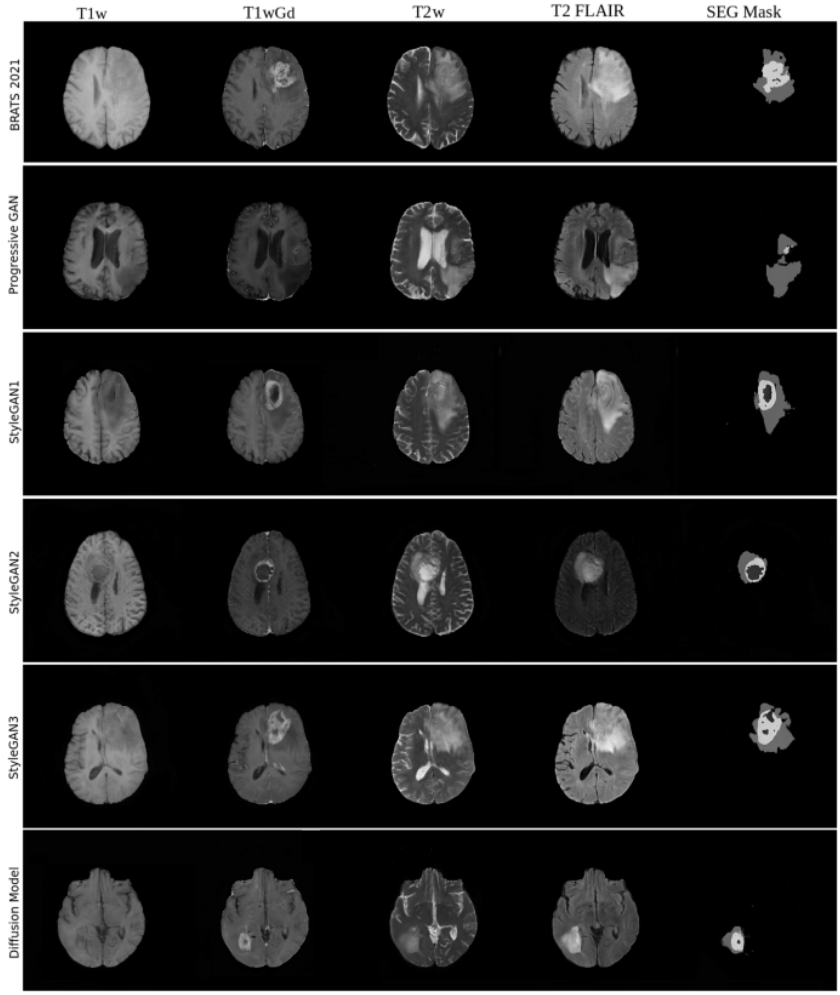

Synthetic brain tumor images generated by GANs and a diffusion model.

Objectives

Train GANs and diffusion models to generate synthetic x-ray images of femur fractures.

Evaluate the synthetic images using different metrics, and use them for training a classifier that classifies AFF/NFF (and evaluate the classifier with real images). A medical doctor can potentially look at the synthetic images for further evaluation.

Should a single generative model be used to generate both AFF and NFF images, or should separate generative models be used for AFF and NFF images?

The images in our dataset have very different size, from 1 to 15 megapixels, how should this be handled during training? GANs and diffusion models normally want all training images to have the same size.

Data

About 4300 X-ray images from some 1200 patients, of which about 20% are AFF. These images originate from 72 Swedish hospitals and vary in image size.

Required background

Machine learning, deep learning, Python programming

Computing resources

The student will have access to very good computing resources (graphics cards) for training GANs and diffusion models. It may also be possible to use the supercomputer Berzelius (752 graphics cards) for trainings.

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