Popular science summary
Enhancing the resolution of MR images using deep learning.

Malo Gicquel

August 2024

Medical scanners are widely used for patient care and for research, as they take 3D pictures highlighting body properties. With the recent developments in AI, new medical image processing tools have been created and it has become crucial to keep developing them to improve patient's health.

Perhaps the most used scanner type is Magnetic Resonance Imaging (MRI), which highlights tissue properties by measuring their interactions with magnetic fields. A scanner can generate several image types (or modalities) that each highlight different body properties. In this work we care about brain anatomy, displayed by the T1-weighted modality. Each scanner has a resolution, that depends on the strength of the field, measured in Teslas (T). The more teslas, the better the resolution.

Most clinical scanners have a resolution of 1.5T. Wealthy hospitals, specialized clinics and research laboratories often use 3T resolution. A few research centers also have a 7T scanner, a resolution only ever used for research as it is expensive and hard to handle. Yet, 7T images offer a very detailed view of the body, which helps researching and diagnosing diseases.

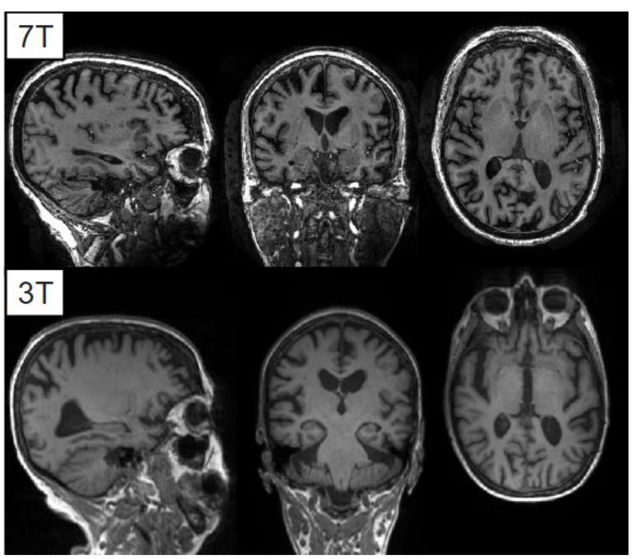

Figure 1: Example of three 2D sections along each dimension of T1-weighted MRI scans of a patient in two resolutions.

In our lab, we have access to thousands of brain T1-weighted scans with a 3T resolution as well a dataset of

172 7T scans, each associated with a 3T scan of the same patient. The idea is to use the dataset to train a deep learning model (an AI) that predicts how a 3T scan would be in 7T resolution. Basically, we create a tool that turns a 3T scan into a 7T scan using the patterns it learned in our dataset. This tool can then be used to enhance the resolution of all our unpaired 3T images.

We investigated the commonly used models and built an advanced AI that learned to generate high resolution images. The techniques we used were inspired by the models used to build AIs performing other image generation tasks, such dall.e or midjourney, that generate an image representing a given sentence. We also used an advanced strategy where we built a second AI that learns to differentiate generated images from real ones, giving a feedback to the first AI. We then assessed our results by visualising generated images and by using mathematical tools that compared the contrast, luminance and structure of the images as well as the pixel to pixel differences. Moreover, we segmented the images, i.e we located all the brain structures, and we compared the segmentations given by the real and generated images.

The models are a good proof of concept and we obtained promising results. The results are not yet good enough for any clinical application, but in a near future such techniques could lead to great improvement in radiology and have many applications in research, by giving information otherwise invisible in low resolution images.