

Learning what's normal: Deep-learning human brain anatomy for unsupervised anomaly delineation

Baur C¹, Navab N¹, Zimmer C², Albarqouni S¹ & Wiestler B²

¹Computer-Aided Medical Procedures (CAMP), TU Munich, Germany

²Department of Diagnostic and Interventional Neuroradiology, TUM University Hospital, Munich, Germany

Introduction

Deep-learning is a disruptive technology that will transform the work of radiologists in the near future. Supervised learning, requiring well labeled data for the computer to learn from, already today supports radiologists (e.g. MS lesion detection) in practice.

Truly unsupervised image analysis (i.e. without human input) is an exciting next step for modern (neuro)radiology. To this end, reliable detection of pathological findings is a pivotal first step.

Purpose

To detect pathologies, defining what is “normal” is essential. Combining several advances in deep-learning in terms of latent space mappings (auto-encoders) and generative reconstruction (generative adversarial networks), we aimed to teach a computer to reconstruct normal anatomy from a given MR scan (**Figure 1**).




Figure 1: Overview of the pipeline.

Methods

We leveraged recent unsupervised deep-learning techniques, i.e. auto-encoders (AE) and adversarial training (GAN), to learn to compress and faithfully reconstruct images of healthy brain anatomy.

As a training dataset, we used 106 MRI scans (3D-FLAIR) of healthy participants.

We hypothesized that bottleneck selection is crucial for reconstruction accuracy in different types of pathologies, and therefore experimented with several combinations of AE and GAN.

When feeding the trained model with an image containing pathologies, it reconstructs a healthy version of the image, such that anomalies can be detected by subtracting the reconstruction from the input. To test this pipeline, we analyzed images of patients with Multiple Sclerosis (typically many small lesions) and Glioblastoma (large lesions with mass effect), diseases the network has never seen during training.

Results

Dice scores for both diseases (indicating overlap between manual ground-truth segmentation and the segmentation generated by the software) were in the range of state-of-the-art unsupervised segmentation strategies: 0.69 for Multiple Sclerosis and 0.67 for Glioblastoma (examples shown in **Figure 2**).

As we expected, different levels of compression and bottleneck significantly influenced the segmentation scores: For space-occupying lesions (like Glioblastoma), dense AE yielded generally higher dice scores, as the reconstruction of these networks yielded much higher reconstruction errors in large lesions. On the other hand, spatial AE performed better in Multiple Sclerosis patients, as these networks tended to yield “sharper”, more detailed images, thereby clearly revealing even small inflammatory lesions in the subtraction maps.

Conclusion & Outlook

The network architecture we present is capable of faithfully reconstructing “normal” brain anatomy and thereby detecting pathologies the network has never seen before.

This approach has great potential to advance the field of machine-driven, unsupervised image analysis. Most importantly, our approach does not require any labeled data, and thereby is able to analyze even large retrospective cohorts such as those found in PACS archives in hospitals.

Bottleneck and compression strategy selection is crucial for segmentation quality. We are currently evaluating several strategies to enable automatic, individualized choice of these parameters.

Following detection, pathology classification is the next major step towards automated reporting of images. We are therefore now developing classifiers based on the segmentation output.







Figure 2: Example of a Glioblastoma (top) and Multiple Sclerosis segmentation (bottom). Green delineates a “true positive” segmentation, orange is “false positive” and red is “false negative”.