Patched Diffusion Models for Unsupervised Anomaly Detection

Abstract

Supervised deep learning techniques for lesion segmentation in brain MRI scans can be challenging due to the large variability of the human brain anatomy and the need for large data sets with pixel-level annotations. An alternative approach is to use unsupervised anomaly detection, which only requires sample-level labels of healthy brain anatomy to create a reference representation. This reference representation can then be compared to unhealthy brains in a pixel-wise manner to identify abnormalities. In order to do this, generative models are needed to create anatomically consistent MRI scans of healthy brains without reproducing the abnormalities present in unhealthy brains. While recent diffusion models have shown promise in this task, accurately reproducing the complex structure of brain MRI scans remains a challenge. In this paper, we propose a method that reformulates the generation task of diffusion models as a patch-based estimation of healthy brain anatomy, using spatial context to guide and improve reconstruction. We evaluate our approach on data of tumors and multiple sclerosis lesions, and demonstrate an improvement of 22.9% in segmentation performance compared to existing baselines.

Keywords: Deep Learning, Unsupervised Anomaly Detection, MRI, Diffusion Models