Beyond Intensity Thresholding:
Towards Unsupervised Anomaly Screening

Abstract

The paramount benefit of unsupervised anomaly detection is the ability to detect irregular and rare data instances even in the absence of training labels. Even though much work has been done on using auto-encoders (AE) for anomaly detection, there are still two critical challenges to overcome. First, learning compact and detailed representations of healthy distribution is cumbersome. Recent work shows that AEs can reconstruct some types of anomalies even better than samples from the training distribution. Second, most unsupervised approaches are tailored to detect hyper-intense lesions on FLAIR brain MR scans. Recent works show that simple thresholding techniques achieve superior white lesion segmentation results. Moreover, we found that state-of-the-art (SOTA) AEs fail to detect several classes of abnormalities that are not hyperintense, including brain atrophy, edema, or dural thickening. In this work, we propose *reversed AEs (RA)* to generate accurate pseudo-healthy reconstructions and localize pathologies in brain MRIs. We extensively evaluate our method on brain T1w MR scans containing multiple global and local pathologies and artifacts and showed superior detection and localization performance.

Keywords: Unsupervised Anomaly Screening.

Figure 1: *RA* produces pseudo-healthy reconstructions of different pathological brain scans. This allows the detection and localization of multiple anomalies.