Layman's Summary: Anomaly detection and segmentation methods using Variational Autoencoders in brain MRI

Introduction: Variational Autoencoders (VAEs) have emerged as a promising technique for unsupervised anomaly detection and segmentation in brain MRI. The principle behind this deep generative modeling is learning a model of healthy brain anatomy by encoding this information into a latent representation and reconstructing it. Anomalies can be detected and localized by discrepancies between the reconstructed data and the original input. This technique relieves the need for pixel-level segmented data and provides the possibility to detect arbitrary anomalies. Challenges have been reported in recent publications regarding the sensitivity towards bright lesions and the limitations of commonly used anomaly scoring methods based on reconstruction error and residual images. These issues are addressed in this work.

Methods: Firstly, a VAE is trained using a simpler dataset, Fashion MNIST. T-shirts are set as normal item and the rest as anomalous items. Anomaly detection methods are tested. Secondly, axial slices from the Medical Out-of-Distribution dataset are used to train another model. Anomaly detection along with anomaly segmentation methods are investigated. Thirdly, axial patches are extracted from these images to train a final network and anomaly segmentation algorithms are analyzed. These three experiments include an evaluation using toy generated images containing anomalies with different characteristics.

Results: The Fashion MNIST experiment showed a higher mean AP for outlier detection in latent space and items similar to t-shirts had higher results than using the reconstruction error. In the case of brain images, outlier detection in latent space and activation map extraction outperformed the commonly used reconstruction error and residual image except for the case of gray anomalies.

Conclusion: This work suggests that dark anomalies can be detected and segmented by changing the background of brain MR images. The investigated techniques for anomaly detection and segmentation based on outlier detection and the extraction of activation maps show potential to replace reconstruction error and residual image computation. An evaluation using real data is discussed as the next possible step.