MSc Project:
Unsupervised Anomaly Detection for Medical Imaging

Abstract: Unsupervised anomaly detection methods aim to detect irregular and rare data instances that deviate from the normal, expected distribution. Although much work has been done on the use of auto-encoders (AE) to detect anomalies [1], it has been shown that learning the healthy distribution is stil cumbersome, with AEs being able to reconstruct some types of anomalies even better than samples from the trained distribution [2, 3]. Strategies to constrain the latent manifold of AE include adversarial training [4,5], probabilistic modeling [6,7], or deformable auto-encoders [8]. However, synthesizing healthy reconstructions from pathological input scans still remains a challenging task.

The objective of this project is to develop novel unsupervised anomaly detection methods to detect and segment pathology on challenging medical imaging datasets.

Requirements:

- Prior experience and good understanding in machine learning and statistics.
- Very good programming skills in Python (and PyTorch).
- Interest in medical imaging.

Contact:

Cosmin Bercea (cosmin.bercea@tum.de)
Prof. Julia Schnabel (julia.schnabel@helmholtz-muenchen.de)

References

[1] Ruff, Lukas, et al. "A unifying review of deep and shallow anomaly detection." Proceedings of the IEEE 109.5 (2021): 756-795.

[2] P. Perera, R. Nallapati, and B. Xiang. Ocgan: One-class novelty detection using gans with constrained latent representations. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 28982906, 2019.

[3] Schirrmeister, Robin, et al. "Understanding anomaly detection with deep invertible networks through hierarchies of distributions and features." Advances in Neural Information Processing Systems 33 (2020): 21038-21049.

[4] S. Pidhorskyi, R. Almohsen, and G. Doretto. Generative probabilistic novelty detection with adversarial autoencoders. Advances of Neural Information Processing Systems (NeurIPS), 2018.

[5] Daniel, Tal, and Aviv Tamar. "Soft-IntroVAE: Analyzing and improving the introspective variational autoencoder." Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2021.

[6] Abati, Davide, et al. "Latent space autoregression for novelty detection." Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 2019.

[7] W. H. L. Pinaya, P.-D. Tudosiu, R. Gray, G. Rees, P. Nachev, S. Ourselin, and M. J. Cardoso. Unsupervised brain anomaly detection and segmentation with transformers. In Medical Imaging with Deep Learning, pages 596617. PMLR, 2021.

[8] Bercea, Cosmin I., Daniel Rueckert, and Julia A. Schnabel. "What do we learn? Debunking the Myth of Unsupervised Outlier Detection." arXiv preprint arXiv:2206.03698 (2022).