Technische Hochschule Brandenburg
University of Applied Sciences
Fachbereich Informatik und Medien

Hyperparameter Selection for Anomaly Detection with Stacked Autoencoders – a Deep Learning Application

Tobias Meyer

Bachelors Thesis • Studiengang Informatik • Fachbereich Informatik und Medien • 01.10.2016

Concept and Motivation

The aim of the project is an evaluation of the consequences of autoencoder variations on the results of an algorithm for anomaly detection on image data.

Autoencoders are a specific deep learning architecture, that gets trained to reconstruct certain input data from a dimensionality reduced, encoded state [GBC16]. Data for which the autoencoder is not optimized, is not reconstructed correctly. This is used to identify and detect anomalous data.

The choice of the correct autoencoder plays an important role for the success of the algorithm for anomaly detection. In this work, different autoencoder setups are varied and their performance for the task of anomaly detection is evaluated.

Anomaly Detection

The detection of anomalies has numerous applications. These include:

- Industrial defect detection, e.g. in assembly and maintenance
- Fraud detection
- Intruder detection in computer networks
- Medical and biological applications, e.g. for diagnosis of diseases and vermin detection
- Video surveillance in civil or military contexts

Algorithm

The Algorithm (figure 1) is a generic approach for anomaly detection on image data.

After a preprocessing step, training data is used to train an autoencoder. The fully trained autoencoder is used to reconstruct images from a new dataset. Differences between input and reconstruction are evaluated and detected as possible anomalies.

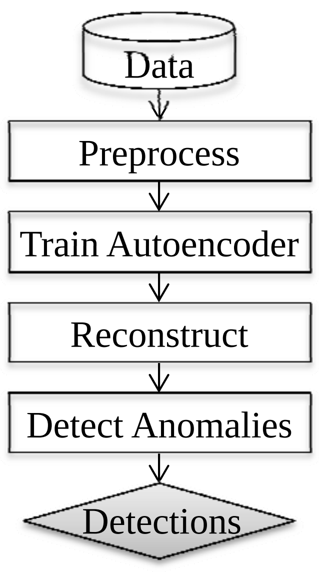

Fig. 1: Main steps of the algorithm for anomaly detection

Improving Algorithm Performance: Better Malaria Detection

Through variation of autoencoder macro architecture (number and size of layers), regularization strength and application of other algorithm variations like denoising [VLL+10] and dropout [SHK+14], that both corrupt data at different autoencoder stages with random noise, algorithm performance can be significantly increased (compare figure 2). For a malaria detection dataset, the amount of correct detections among all detections (precision) and the amount of correct detections among all anomalies (recall) approaches or even exceeds human performance.

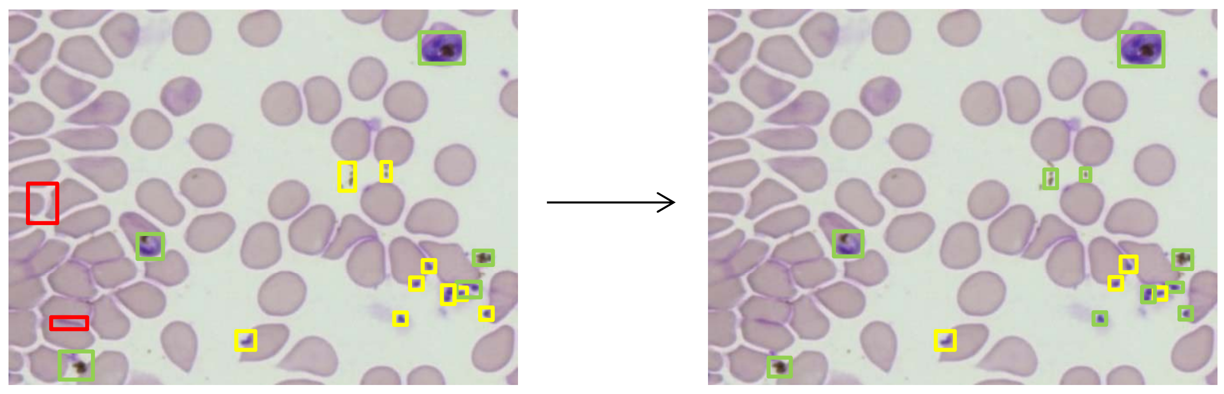

Fig. 2: Increased detection quality of models, caused through usage of better autoencoder variation. Green are correct detections, yellow are missed anomalies, red are false detections.

A successful model can detect 99% of malaria infections while introducing less than 50 false detections per 100 correct detections (recall 0.99 with precision 0.68, compare figure 3, left side, purple plot). Also, 100% detected anomalies are possible!

Denoising can repair bad decisions in the autoencoder selection process to still get a well performing model (compare figure 3, right side).

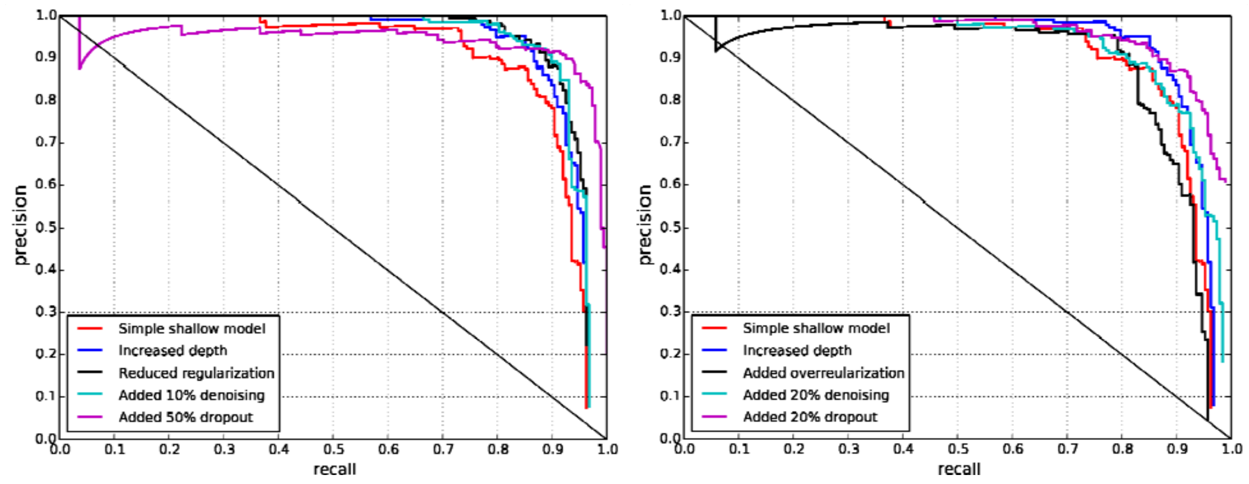

Fig. 3: Precision-recall curves resulting from different autoencoder hyperparameters

Exceeding human performance

The most successful models get detections challenging even human hand annotations. Detections include plasmodia missed during hand annotation, not annotated sample contaminations and cell heaps not correctly distributed during sample layout.

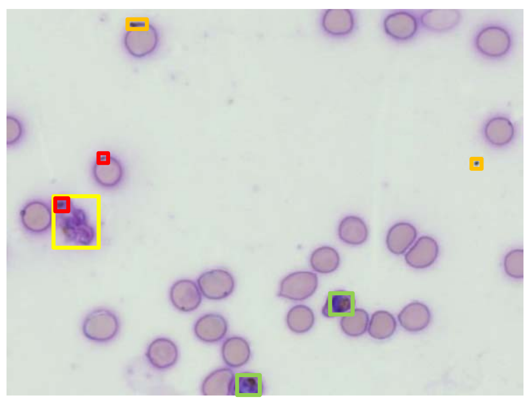

Fig. 4: Sample detections challenging human performance. Green boxes are correct detections, red boxes show wrong detections. Orange boxes are considered false detections, but show plasmodia that were missed in the provided hand annotation. The yellow box shows a missed detection that overlaps with a detection that is too small to be considered correct.

Conclusion

This work shows that the generic algorithm for anomaly detection can be very successful, when autoencoders are chosen carefully.

Additionally, the work includes a recommendation for autoencoder selection and an evaluation of the algorithm on a different dataset.

References

[GBC16] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. Book in preparation for MIT Press, 2016.

[VLL+10] Pascal Vincent et al. *Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion*. JMLR, 11:3371-3408, 2010.

[SHK+14] Nitish Srivastava et al. *Dropout: A simple way to prevent neural networks from overfitting*. Journal of Machine Learning Research, 15:1929-1958, 2014.

Supervision: Dipl.-Inform. Ingo Boersch • Prof. Dr.-Ing. Sven Buchholz • Technische Hochschule Brandenburg • Dr.-Ing. Christian Wojek• Carl Zeiss AG