Anomaly detection with deep learning in an activity detection use case for post-stroke patients

Internship under the supervision of
Carole Frindel, CREATIS - INSA Lyon and Jan Ramon, MAGNET Inria team

Topics: Machine learning / Anomaly detection

Keywords: Auto-encoders, Variational auto-encoders, Bi-GAN, PAC-Bayed analysis

Context:

Monitoring daily physical activity is essential to understanding patient recovery. Post-stroke patients exhibit high levels of sedentary behavior [1], which has been identified as a risk factor for secondary cardiovascular disease and mortality [2]. Wearable sensors, coupled with activity recognition models and machine learning techniques, can identify various mobility-related activities, such as sitting, lying, standing, walking, and stair use, in the clinic and in the home and community. This enables therapists to develop personalized, data-driven programs to advise patients and improve activity levels [3]. Thus, efficiently measuring mobility activities through a wearable activity recognition system is a major quantitative outcome measure for studying new therapeutic interventions for stroke survivors.

One critical task in physical activity monitoring for post-stroke patients is detecting abnormal events such as deterioration of the patient's physical condition or risk of relapse. Generally, abnormal events rarely occur as compared to normal activities. Therefore, to alleviate the waste of labor and time for neurologists, developing intelligent machine learning algorithms for automatic anomaly detection is a pressing need. The goal of a practical anomaly detection system is to timely signal an activity that deviates from normal patterns and identify the time window of the occurring anomaly.

The underlying strategy for most approaches to anomaly detection is to first model normal behavior, and then exploit this knowledge to identify anomalies. The first step, referred to as the training step, involves building a model of normal behavior using available data. The second step in the anomaly detection loop, the test step, introduces the concept of threshold-based anomaly tagging. Based on the range of scores assigned by the model, one can select a threshold rule that drives the anomaly tagging process; e.g., scores above a given threshold are tagged as anomalies, while those below it are tagged as normal (Figure 2).

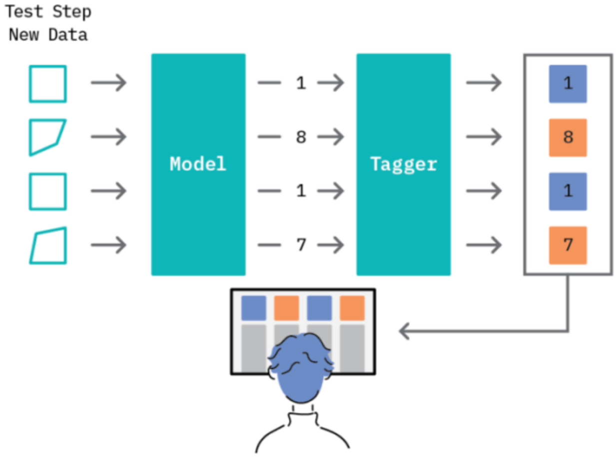

Figure 1: The test step in the anomaly detection loop.

Deep learning anomaly detection approaches exploiting autoencoders (AE) have shown good performances [4]. Autoencoder-based anomaly detection consists in training an autoencoder to reconstruct a set of examples and then to detect as anomalies those inputs that show a sufficiently large reconstruction error. Unfortunately, deep non-linear architectures are able to perform high dimensionality reduction while keeping reconstruction error low. This problem is in part due to the lack of regularity in the latent space. To alleviate the above problem, recently some authors have proposed to exploit Variational autoencoders (VAE) and bidirectional Generative Adversarial Networks (Bi-GAN), which arise as a variant of standard autoencoders designed for generative purposes, both enforcing the organization of the latent space guaranteeing continuity [5].

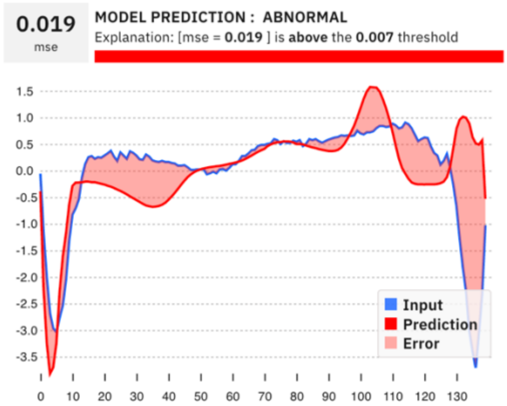

Figure 2: EEG signal tested via a model learned by an AE: in blue the tested signal, in red the signal reconstructed by the AE, in pink the error per signal point then averaged at the top left. This signal is detected as an anomaly because it is above the threshold given in the upper right corner.

Objective: The goal of this internship is to test the three DL architectures (AE, VAE and Bi-GAN) on a database of signals from accelerometer and gyroscope sensors recorded on post-stroke patients for anomaly detection related to activities to be recognized. The idea will be to compare their performances and to use a PAC-Based approach in order to automatically choose the detection thresholds based on empirically calculated generalization bounds [6].

Prerequisites and expected skills:
- Efficient programming level in Python
- Knowledge of machine learning and TensorFlow and/or Pytorch libraries
- Interest in DL mathematical theory would be a plus
- Good collaboration and communication skills (written / oral)

Hosting structure: This internship will be hosted either in Lyon or Lille campus. This work will be done in collaboration with Antoine Boutet and Jan Aalmoes from PRIVATICS INRIA team.
Contact: carole.frindel@insa-lyon.fr, jan.ramon@inria.fr

About CREATIS: CREATIS is a research unit in medical imaging with about 200 people whose research areas are at the crossroads of two major areas: major health issues and the theoretical challenges in signal and image processing. CREATIS responds to these challenges with a transdisciplinary approach involving four research teams belonging to information and communication sciences and technologies, engineering sciences and life sciences.

About Magnet INRIA team: A primary objective of Magnet is in making artificial intelligence more acceptable to society by solving some ethical issues of Machine Learning (ML) and on empowering end users of artificial intelligence. We study graph-based machine learning methods which are the common foundations of the research group and we rely on methods coming from statistical and computational learning theory, graph theory, representation learning, (distributed) optimization and statistics.

References

[1] Morton, Sarah, et al. "Sedentary behavior after stroke: A new target for therapeutic intervention." *International Journal of Stroke* 14.1 (2019) : 9-11.

[2] Bahls, Martin, et al. "Physical activity, sedentary behavior and risk of coronary artery disease, myocardial infarction and ischemic stroke: a two-sample Mendelian randomization study." Clinical Research in Cardiology 110.10 (2021): 1564-1573.

[3] Lin, David J., Seth P. Finklestein, and Steven C. Cramer. "New directions in treatments targeting stroke recovery." Stroke 49.12 (2018): 3107-3114.

[4] Chalapathy, Raghavendra, and Sanjay Chawla. "Deep learning for anomaly detection: A survey." arXiv preprint arXiv:1901.03407 (2019).

[5] Angiulli, Fabrizio, Fabio Fassetti, and Luca Ferragina. “Latent Out: an unsupervised deep anomaly detection approach exploiting latent space distribution." Machine Learning (2022): 1-27.

[6] Chérief-Abdellatif, Badr-Eddine, et al. "On PAC-Bayesian reconstruction guarantees for VAEs." *International Conference on Artificial Intelligence and Statistics*. PMLR, 2022.