Robust Deep Learning Anomaly Detection
Application to intrusion detection for electric vehicle charging points

PhD candidate position

Duration: 36 months
Expected starting: October/November 2024
Location: LITIS, INSA Rouen Normandie

Keywords Machine learning, time-series, deep learning, anomaly detection, cybersecurity, electric charging points

Research environnement The research will take place within the LITIS laboratory at INSA Rouen, France, you will work in a research team with topics covering machine learning for structured and unstructured data. You will contribute to the strengthening of these activities through the SHARP project which aims to design, implement and deploy at large scale data-driven Artificial Intelligence algorithms to secure electric vehicle charging points.

This PhD position is part of SHARP project under the supervision of Gilles Gasso and Mokhtar Z. Alaya in close relationship with a private company providing testbed datasets.

Context and scientific objectives The current growth of the Electric Vehicle (EV) market comes with the large scale deployment of EV Charging Points (CP). Such a large deployment may render the charging points subject to impactful cybersecurity risks of which the most prominent attacks that can imperil EV wired charging points include tampering (prevent charging, get sensitive information), energy repudiation (cheat on billing, energy theft), deny of charging, etc [1]. Hence, detecting and mitigating such adverse situations are of high interest.

To address this challenge, the SHARP project aims to propose machine learning algorithms that robustly detect abnormal events on the charging points in an end-to-end manner. Specifically, provided heterogeneous and multi-modal data i.e. time series sampled at different rates and collected on some few charging points (the samples consist of different time series generated by the CP, such as network communication, alive status of CP ("ping"), operational status, transactional data (charging, energy consumption, payment data)...) we aim to elaborate robust anomaly detection algorithms using end-to-end deep learning.

Anomaly detection is generally unsupervised, since abnormal events are rare, of varied nature and hard to annotate. While deep neural networks (DNN) offer flexibility to learn meaningful representation from multi-modal temporal data [2], anomaly detection with deep architecture remains a difficult task.

A representative method is for instance deep one-class classifier [3]. This latter learns a space feature representation alongside the one-class classifier. However, this method suffers issues such as the collapse of the data in the latent space, calling for self-supervised representation learning [4] and has a skewed focus on anomaly detection in images. The pursued goal of this PhD is to jointly learn meaningful representation space of the data by leveraging their multi-modal and temporal aspects and the abnormality detection model such as to address online detection on streaming samples. The robustness of the detection algorithm is crucial to guarantee safe and secure operating CP. For a given testing sample, a deep one-class classifier customarily outputs a one-class score that is thresholded to assess whether the sample corresponds to a malicious operating condition or not. This threshold is to be tuned, and hence to set a tradeoff between maximization of the true abnormal detection rate and minimization of the false alarm rate. Techniques such as conformal inference methods [5] are to be investigated.

Thus, the research work to be conducted can be divided into the following tasks:

- State-of-the-art survey on deep learning anomaly detection on time-series.
- Study new DNN representations to address multi-modal time series with different sampling periods.
- Investigate relevant statistical model (one-class) in the deep latent space for online detection.
- Design and evaluate statistical guarantees of the abnormal detection model by investigating relevant approaches such as conformal inference prediction.

The developped algorithms will be evaluated on real benchmarks issued from penetration testing (pen-testing) targeting different types of attacks.

Outcomes The outcomes will lead to publications in the machine learning and signal processing communities. The thesis will be conducted as a part of the ANR Project SHARP (Machine Learning for Safe Vehicle Charging Points) funded by ANR and held at LITIS. The candidate is expected to develop new machine learning algortihms for abnormal event detection on EV operating charging points and related Python toolbox. On this technical aspect, the PhD candidate will benefit from the expertise of ongoing collaborations with other academic partners on the subject.

Candidate profile Applicants are expected to be graduated in computer science and/or machine learning and/or signal & image processing and/or applied mathematics/statistics, and show an excellent academic profile. Beyond, good programming skills on DNN through Pytorch or Tensorflow are expected.

Application procedure Applicants are requested to submit:

- A resume
- Academic transcripts
- A Cover letter applying for the position.

Applications are to be sent by email to Gilles Gasso (gilles.gasso@insa-rouen.fr) and Mokhtar Z. Alaya (alayaelm@utc.fr).

References

[1] A. Brighente, M. Conti, D. Donadel, R. Poovendran, F. Turrin, and J. Zhou, "Electric vehicles security and privacy: Challenges, solutions, and future needs," *arXiv preprint arXiv:2301.04587*, 2023.

[2] A. Das, W. Kong, R. Sen, and Y. Zhou, "A decoder-only foundation model for time-series forecasting," *arXiv preprint arXiv:2310.10688*, 2023.

[3] L. Ruff, R. Vandermeulen, N. Goernitz, L. Deecke, S. A. Siddiqui, A. Binder, E. Müller, and M. Kloft, "Deep one-class classification," in *International conference on machine learning*. PMLR, 2018, pp. 4393–4402.

[4] K. Sohn, C.-L. Li, J. Yoon, M. Jin, and T. Pfister, "Learning and evaluating representations for deep one-class classification," *arXiv preprint arXiv:2011.02578*, 2020.

[5] S. Bates, E. Candès, L. Lei, Y. Romano, and M. Sesia, "Testing for outliers with conformal p-values," *The Annals of Statistics*, vol. 51, no. 1, pp. 149–178, 2023.