PhD Thesis

Deep Learning with Generative Models for Anomaly Detection

Keywords

Deep learning, generative (probabilistic) models, diffusion probabilistic models, normalizing flows, anomaly detection, time series

PhD Thesis Description

The broad interest in deep neural networks has driven recent advances in anomaly detection, also called out-of-distribution or novelty detection. Deep anomaly detection methods fall within three major categories: Deep one-class, variational autoencoders (VAEs) and generative adversarial networks (GANs) [1, 2]. While these methods do not allow an exact evaluation of the probability density of new samples, they also suffer from notorious training instability (mode collapse, posterior collapse, vanishing gradients and non-convergence), as corroborated by many research studies [3]. For these reasons, we aim to go beyond the weaknesses of these methods, by investigating novel classes of generative models with deep learning to address anomaly detection.

The goal of this PhD thesis is to explore novel generative models, such as diffusion probabilistic models (DPM) and normalizing flows (NF). These classes consist of models that can generate, through a deep latent space, a probability distribution for a given dataset from which we can then sample. With solid theoretical foundations and often interconnections with Optimal Transport, several variants of generative models have been proposed based on different definitions of their main components, namely the forward and backward processes and the sampling procedure. Of particular interest are NF and DPM. NF are generative models where both sampling and density evaluation are efficient and exact, and where the latent representation is learned through an invertible transformation, thus providing explainable models [4, 5]. DPM rely on diffusion processes, inspired from nonequilibrium thermodynamics, with their flagship being denoising diffusion probabilistic models [6]. Diffusion models have been demonstrating record-breaking performance in many applications in computer vision, mainly for image synthesis [7, 8] and medical imaging [9].

The proposed PhD research program aims to investigate these recent advances in generative models with deep learning for anomaly detection. Recent studies have explored generative probabilistic models for anomaly detection, mainly in images [10-12] with some attempts in signal processing [13, 14], demonstrating preliminary results on their relevance in anomaly detection and opening the way to new questions [15]. The PhD candidate will investigate such generative probabilistic models in a more in-depth research study, in order to take full advantage of their underlying theory. Moreover, the PhD candidate will go beyond image processing, with a focus on anomaly detection in time series, by considering the specificities of time series. The proposed framework and devised methods will be assessed in a variety of scenarios and real-world time series datasets.

Research Environment

The PhD candidate will conduct her/his research within the Machine Learning group in the LITIS Lab, under the supervision of Prof. Paul Honeine, Dr. Fannia Pacheco and Dr. Maxime Berar. This PhD thesis is within a research project gathering 9 permanent researchers of the LITIS Lab and the PhD candidate will also interact with several PhD students and interns also working on deep anomaly detection with a focus on time series.

Applicant Profile

- Master or Engineering degree, in data science, AI, applied mathematics, or related fields.
- Strong skills in advanced statistics and Machine Learning, including Deep Learning
- Good programming experience in Python

Location

LITIS Lab, Université de Rouen Normandie, Saint Etienne du Rouvray (Rouen, France).

Application

Applicants are invited to send their CV and grade transcripts by email to:
paul.honeine@univ-rouen.fr, fannia.pacheco@univ-rouen.fr, maxime.berar@univ-rouen.fr.

References

[1] L. Ruff *et al.*, "A unifying review of deep and shallow anomaly detection," *Proceedings of the IEEE*, 2021.
[2] G. Pang *et al.*, "Deep learning for anomaly detection: A review," *ACM Computing Surveys*, 2021.
[3] D. Saxena and J. Cao, "Generative adversarial networks (GANs) challenges, solutions, and future directions," *ACM Computing Surveys*, 2021.
[4] I. Kobyzev *et al.*, "Normalizing flows: An introduction and review of Current Methods," *IEEE T PAMI*, 2021.
[5] G. Papamakarios *et al.,* "Normalizing flows for probabilistic modeling and inference," *JMLR*, 2021.
[6] J. Ho *et al.*, "Denoising diffusion probabilistic models," *NeurIPS*, 2020.
[7] L. Yang *et al.*, "Diffusion models: A comprehensive survey of methods and applications," *ACM Computing Surveys*, 2023.
[8] F.-A. Croitoru *et al.,* "Diffusion models in vision: A survey," *IEEE T PAMI*, 2023.
[9] A. Kazerouni *et al.*, "Diffusion models in medical imaging: A comprehensive survey," *Medical Image Analysis*, 2023.
[10] J. Wolleb *et al.*, "Diffusion models for medical anomaly detection." MICCAI, 2022.
[11] W.H. Pinaya *et al.*, "Fast unsupervised brain anomaly detection and segmentation with diffusion models," MICCAI, 2022.
[12] A. Kascenas *et al.*, "The role of noise in denoising models for anomaly detection in medical images." *Medical Image Analysis*, 2023.
[13] Y. Chen *et al.*, "ImDiffusion: Imputed Diffusion Models for Multivariate Time Series Anomaly Detection," in *Proc. of the VLDB Endowment*, 2023.
[14] R. Hu *et al.*, "Unsupervised Anomaly Detection for Multivariate Time Series Using Diffusion Model," *IEEE ICASSP*, 2024.
[15] P. Kirichenko *et al.*, "Why normalizing flows fail to detect out-of-distribution data," *NeurIPS*, 2020.