PhD thesis in LITIS Lab, Rouen Normandy

Robust and Frugal Self-Supervised Federated Learning for Out-of-distribution Data in Time Series

Type: 3-year doctoral thesis (PhD) in LITIS Lab (Rouen)

Start date: Last quarter of 2026

Level: Master's degree (M2) or engineering (or equivalent) in Applied Mathematics, Data Science, Computer Science, or a related field

Funding: Full doctoral contract according to current French regulations (gross salary as per university scale)

Keywords: Self-Supervised Learning, Federated Learning, Time Series, Robustness, Frugal Feature Extraction, Anomaly / Out-of-distribution Detection

Scientific Context

Federated learning (FL) enables a set of distributed clients (e.g., edge devices, wearable sensors, hospitals, or industrial sites) to collaboratively train a model without sharing their raw data, thereby preserving privacy [1]. In most realistic deployments, however, the data that clients collect is unlabeled: annotation is costly, protocols are inconsistent, and models are updated frequently. This positions self-supervised learning (SSL) as the natural learning paradigm for each client, and federated SSL as a principled approach to collaboratively learning representations from unlabeled, decentralized data. The target application is time series: accelerometer and gyroscope traces, body and physiological sensors, energy and weather measurements, or industrial monitoring streams. Such streams are pervasively contaminated by anomalies and out-of-distribution (OOD) samples: sensor faults, dropouts, rare events, novel operating regimes, and unseen classes. In self-supervised federated training, these contaminants corrupt the learned representation and the aggregated global model, since a single client streaming faulty signals can degrade every participant [2,3]. In this context, the primary objective of this PhD thesis is robustness: keeping the representation sound despite contamination, across non-IID clients whose distributions also differ from one another [4]. A robustly trained self-supervised model also yields, as a by-product, a label-free score (its reconstruction or pretext error) that flags anomalous and OOD inputs.

A recent result shows that the robustness of distributed self-supervised objectives is paradigm-dependent [7]: both theory and experiments indicate that masked-reconstruction objectives are intrinsically more robust to corrupted, heterogeneous client data than contrastive ones. This grounds the choice of SSL objective and its aggregation, keeping representations robust to contamination across non-IID clients. Robustness, however, often comes at a price. Reconstruction-based detection is known to depend non-monotonically on encoder capacity: a too large encoder reconstructs anomalies as faithfully as normal inputs and loses discrimination [14,15], while a too small encoder reconstructs neither well, raising a noisy baseline error in which genuine anomalies and OOD inputs are hard to separate from the encoder's own limitations [16]. Frugality

and detection thus pull against each other, and characterizing this trade-off is one of the central questions of the PhD thesis, namely, how small an encoder can be before its self-supervised score can no longer separate genuine anomalies and OOD inputs from its own limitations. Determining where on this capacity curve a frugal federated encoder must operate is, concretely, the first technical question the candidate will study.

The frugality in FL is defined specifically as the computational and memory cost of feature extraction on the client. This axis has become acute because the strongest representations for time series increasingly come from large self-supervised models (foundation models) whose parameter count, memory footprint, and inference cost are far beyond what wearables, sensors, and embedded devices can host [5]. Rather than compressing large encoders, the thesis asks how small a robustly-trained encoder can be before its score stops separating anomalies and OOD inputs from its own underfit. Compactness is pursued through small architectures, parameter-efficient designs, pruning, and low-rank or partial updates [9,10,11], with foundation-model encoders serving only as an oracle upper bound for comparison.

Research Objectives

The objective of this thesis is to propose efficient robustly-trained federated self-supervised encoders that produce a score attributable to an anomaly, an OOD input, or its own underfit, and to act on that attribution. Robustness is the property optimized, frugal client-side feature extraction is the binding constraint, and anomaly/OOD detection is a measured consequence.

Three sub-problems are recognized for this thesis: (i) Score disentanglement: an elevated score may signal an anomaly, an OOD input, or merely an underfitting frugal encoder; separating these from unlabeled signals alone is the core question. (ii) Relative normality: in a federated setting, what counts as anomalous is defined relative to each client's distribution, so an input may be globally OOD yet locally normal; the thesis must specify and evaluate whether detection is judged at the local or global reference, not assume a single global normal. (iii) Honest evaluation without labels: Metrics, such as area under the ROC (AUROC) and area under the precision-recall (AUPR) curves, need labelled anomalies the premise says clients lack, most approaches rely on held-out benchmarks complemented by synthetically contaminated protocols where ground truth is exact.

Methodology and Work Plan

The three years are organized around the three sub-problems above. Year 1 builds a reproducible benchmark for federated self-supervised time-series learning under controlled contamination (injected anomalies and held-out OOD classes) across non-IID clients, reproduces the masked-versus-contrastive robustness comparison [7], and characterizes the capacity curve of reconstruction-based scores from a large reference encoder down to a minimal one [14,15,16]. Year 2 develops the core contribution: a label-free policy that attributes an elevated self-supervised score to a genuine anomaly, an OOD input, or encoder-underfit, and defines normality at a stated local or global reference; it is validated against the decoupled combination of a fixed frugal encoder, an off-the-shelf federated SSL method, and a separate anomaly detector. Year 3 consolidates the framework, adds analysis where tractable, and is devoted to a journal submission, the manuscript writing, and the PhD defense.

Datasets and Baselines

Experiments will use standard time-series sources where each subject or device is a non-IID client — Human Activity Recognition (UCI-HAR, WISDM, HHAR), Sleep Stage Classification (single-channel EEG), and Machine Fault Diagnosis, following the AdaTime protocol [8] — together with dedicated time-series anomaly benchmarks (e.g., SMD, SMAP/MSL, PSM) [17] for detection, complemented by synthetically-contaminated streams where ground truth is exact. Self-supervised pretraining uses the unlabeled signals; small labelled subsets serve only for linear-probe, few-shot, and detection (AUROC/AUPR) evaluation. Orchestration relies on standard FL frameworks (e.g., Flower, FedML), with FEMNIST and CIFAR-10/100 under Dirichlet partitioning as auxiliary non-IID testbeds. Baselines include FedAvg [1] and heterogeneity-oriented optimizers [4], the distributed-SSL robustness reference [7], profile-mapping aggregation [6], federated time-series representation learning [5], and standard reconstruction- and forecasting-based self-supervised anomaly detectors.

Candidate Profile

We are looking for a highly motivated candidate holding (or about to obtain) a Master's degree or engineering diploma in applied mathematics or computer science, with a solid background in machine learning and strong programming skills in Python (PyTorch or equivalent). Prior exposure to time-series modeling, federated or distributed learning, self-supervised learning, or anomaly / out-of-distribution detection is appreciated but not required. A good level of English and an interest in both theoretical analysis and rigorous experimentation are strongly recommended.

Applications should include a detailed CV, a cover letter, academic transcripts from the Master's or engineering's years, and contact details of one or two references.

Supervisors

Fannia Pacheco (fannia.pacheco@univ-rouen.fr)
Paul Honeine (paul.honeine@univ-rouen.fr)
Maxime Berar (maxime.berar@univ-rouen.fr)

Thesis Context

The thesis will be funded by an ANR project (funding granted). The candidate will benefit from a friendly and collaborative environment alongside other interns and PhD students. In addition, there will be opportunities for professional training and international travel.

The thesis will take place at the LITIS laboratory, which has 160 members, including PhD students, and six research groups covering a broad range of fundamental and applied research in AI and computer science. The intern will join the "Apprentissage" team (Machine Learning group), consisting of 22 faculty members from Rouen University and INSA Rouen and approximately 20 PhD students and postdocs. The team focuses on machine learning and pattern recognition for interpreting diverse data such as signals, images, and text.

References

[1] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Agüera y Arcas, “Communication-Efficient Learning of Deep Networks from Decentralized Data,” AISTATS, 2017.

[2] P. Kairouz et al., “Advances and Open Problems in Federated Learning,” Foundations and Trends in Machine Learning, 2021.

[3] H. Zhu, J. Xu, S. Liu, and Y. Jin, “Federated Learning on Non-IID Data: A Survey,” Neurocomputing, 2021.

[4] Q. Li et al., “Federated Learning on Non-IID Data Silos: An Experimental Study (NIID-Bench),” IEEE ICDE, 2022.

[5] S. Chen, G. Long, M. Blumenstein, and J. Jiang, “FeDaL: Federated Dataset Learning for General Time Series Foundation Models,” ICLR, 2026 (arXiv:2508.04045).

[6] M. Li, D. Fenoglio, M. Gjoreski, and M. Langheinrich, “Federated Learning with Profile Mapping under Distribution Shifts and Drifts (FEROMA),” ICLR, 2026.

[7] X. Chen, N. Yang, S. Wang, and D. Yuan, “Understanding the Robustness of Distributed Self-Supervised Learning Frameworks against Non-IID Data,” ICLR, 2026.

[8] M. Ragab, E. Eldele, et al., “AdaTime: A Benchmarking Suite for Domain Adaptation on Time Series Data,” ACM TKDD, 2023.

[9] T. Wu, C. Song, and P. Zeng, “Efficient Federated Learning on Resource-Constrained Edge Devices Based on Model Pruning,” Complex & Intelligent Systems, 2023.

[10] C. Wu, F. Wu, L. Lyu, Y. Huang, and X. Xie, “Communication-Efficient Federated Learning via Knowledge Distillation,” Nature Communications, 2022.

[11] X. Wu, X. Liu, J. Niu, H. Wang, S. Tang, et al., “Decoupling General and Personalized Knowledge in Federated Learning via Additive and Low-Rank Decomposition,” ACM Multimedia, 2024.

[12] H. Liu, C. Guo, Y. Ren, J. Guan, and S. Zhou, “FedOpenMatch: Towards Semi-Supervised Federated Learning in Open-Set Environments,” ICLR, 2026. (Out of scope — future work.)

[13] Z. Niu, H. Dong, and A. K. Qin, “Bridging the Generalization Gap of Heterogeneous Federated Clients Using Generative Models (FedVTC),” ICLR, 2026. (Out of scope — future work.)

[14] D. Gong, L. Liu, V. Le, B. Saha, M. R. Mansour, S. Venkatesh, and A. van den Hengel, “Memorizing Normality to Detect Anomaly: Memory-Augmented Deep Autoencoder for Unsupervised Anomaly Detection (MemAE),” ICCV, 2019.

[15] F. Angiulli, F. Fassetti, and L. Ferragina, “On the Behaviour of Reconstruction-Based Anomaly Detectors and the Over-Generalization of Autoencoders,” 2023.

[16] D. Mantegazza, C. Redondo, F. Espada, L. M. Gambardella, A. Giusti, and J. Guzzi, "Sensing Anomalies as Potential Hazards: Datasets and Benchmarks," *Towards Autonomous Robotic Systems (TAROS 2022)*, LNCS, pp. 205–219, 2022. (arXiv:2110.14706)

[17] J. Xu, H. Wu, J. Wang, and M. Long, “Anomaly Transformer: Time Series Anomaly Detection with Association Discrepancy,” ICLR, 2022.