Research internship

Insider threat detection: alerts with a controlled error rate

Nicolas Lachiche, Gabriel Frey

ICube laboratory (UMR 7357, University of Strasbourg, CNRS), SDC team

1. Context

Insider threats come from people who hold legitimate access: employees, contractors, compromised accounts [1]. There is no intrusion to detect, only abnormal use of valid privileges. The dominant approach learns the usual behaviour of each user, or of their peer group, and flags deviations, for instance with Isolation Forest [3]. But in many datasets, malicious activity is extremely rare: a false-positive rate that looks low still produces dozens of false alerts per day. And the alert threshold is usually set by hand, or after the fact.

Conformal anomaly detection turns any anomaly score into a p-value and controls the false discovery rate (FDR) without distributional assumptions [5]. The method has limits, however:

- Base rate: the guarantee bounds false alerts, not precision.
- Exchangeability: it is broken by behavioural drift, and by an adversary who slowly gets the system used to them.
- Contamination: an insider who is already active could enter the calibration data and shield themselves [6].
- Marginal guarantees: an overall rate can hide over-flagged groups (administrators, atypical working hours, . . .).
- Atypical $\neq$ malicious: the ranking is that of the raw score. The benefit lies in the stability of the threshold over time and across sites.

2. Objectives

The internship will seek to establish **when a conformal guarantee remains reliable for insider detection, and how to restore it when it does not**. It will aim to:

- Build a reproducible pipeline: user-day profiles, individual, peer-group and long-term baselines, detectors (Isolation Forest and recent variants), explanations (DIFFI [4]).
- Measure, at a fixed alert budget and with a temporal split, the FDR actually achieved and the detection power on datasets such as CERT, ChimeraLog [2] and the LANL authentication data, using the *nonconform* library.
- Stress-test the guarantee: drift, contaminated calibration, slow insider, per-group calibration under data scarcity, fairness across groups.
- Evaluate recent remedies: weighted or sliding calibration, active cleaning of the calibration set, p-values that exploit a few labelled anomalies.

Studying the limits of the methods and of the data, both synthetic and real, will be an integral part of the work.

3. Work plan

1. **Literature review and scoping**: insider threats, conformal inference, evaluation pitfalls in security.
2. **Data and first models**: CERT, ChimeraLog and/or LANL datasets. Unsupervised baseline detectors. Conformal calibration.
3. **Experiments**: stress protocols, remedies, explanations, comparison across datasets.
4. **Writing and defence**: consolidation of results, thesis and defence.

Team and working environment

The internship takes place in the **SDC team (Data Science and Knowledge)** of the **ICube** laboratory (UMR 7357, University of Strasbourg, CNRS). It is co-supervised by Nicolas Lachiche and Gabriel Frey.

- **Location**: ICube, 300 bd Sébastien Brant, CS 10413, 67412 Illkirch Cedex, south of Strasbourg.
- **Remote work**: the internship is mainly on site. Occasional remote work is possible, by agreement with the supervisors.
- **Stipend**: according to the regulations in force at the time of the internship.

Duration and dates

- Duration: **5 to 6 months**.
- Start: **from February 2027**.

Profile

Student in the second year of a Master's in computer science (M2), or in the final year of an engineering school in computer science.

- Solid background in probability and statistics (hypothesis testing, p-values, multiple testing).
- Good command of Python and its scientific ecosystem (NumPy, pandas, scikit-learn).
- Knowledge of machine learning.
- Interest in computer security. No prior expertise is required: the necessary background is covered at the start of the internship.

How to apply

Send an e-mail to both supervisors:

- Nicolas Lachiche: nicolas.lachiche@unistra.fr
- Gabriel Frey: g.frey@unistra.fr

The e-mail must contain:

- A CV.
- Copies of diplomas and academic transcripts.
- A short cover letter.

Applicants from outside the European Union: due to security regulations and the time required to obtain visas and residence permits, applicants must hold, at the time they apply, a residence permit authorising them to carry out an internship in a French laboratory.

References

1. Homoliak, I. et al. (2019). *Insight into Insiders and IT: A Survey of Insider Threat Taxonomies, Analysis, Modeling, and Countermeasures.* ACM Computing Surveys 52(2). arXiv:1805.01612
2. Yu, J. et al. (2026). *Chimera: Harnessing Multi-Agent LLMs for Automatic Insider Threat Simulation.* NDSS 2026. arXiv:2508.07745
3. Liu, F. T., Ting, K. M., Zhou, Z.-H. (2008). *Isolation Forest.* IEEE ICDM.
4. Carletti, M., Terzi, M., Susto, G. A. (2020). *Interpretable Anomaly Detection with DIFFI: Depth-based Isolation Forest Feature Importance.* arXiv:2007.11117
5. Bates, S., Candès, E., Lei, L., Romano, Y., Sesia, M. (2023). *Testing for Outliers with Conformal p-values.* Annals of Statistics 51(1). arXiv:2104.08279
6. Bashari, M., Sesia, M., Romano, Y. (2025). *Robust Conformal Outlier Detection under Contaminated Reference Data.* arXiv:2502.04807