Explainability and privacy for synthetic time series generation

PhD thesis proposal

Supervisors: Tristan Allard (Univ. Rennes, IRISA), Romaric Gaudel (Univ. Rennes, IRISA)

Contact: `tristan.allard@irisa.fr` and `romaric.gaudel@irisa.fr`

Keywords: privacy, synthetic time series generation, explainability.

1 Context and goal

Financial transactions, water, gas, or electricity consumption, biomedical signals... A vast amount of personal data is today generated in the form of timestamped sequences of data called *time series* hereafter. These time series are collected and stored by companies or public organizations in order to support a large variety of usages (*e.g.,* fraud detection in financial flows, epidemiology, smartgrid management). They carry detailed information about individual behaviors or health status. As a result, for obvious privacy reasons (*e.g.,* large-scale re-identifications [5]), they are today mostly secluded within the systems that collect them, obliterating the benefits expected from large scale time series sharing.

Generative models are promising solutions. Given an input set of time series, they generate a set of *synthetic* time series that is different from, but statistically close to, the input training set [3]. When protected by sound privacy-preserving mechanisms (*e.g.,* differentially private perturbation [4]), they carry the promise to enable organizations to share (synthetic) time series at a large scale without jeopardizing privacy guarantees. However, utility of synthetic time series is both complex and hard to achieve, especially when strong privacy guarantees, *e.g.,* differential privacy, are met. First no generative model consistently outperforms the others on all the datasets or on all the utility metrics. Second, within a set of synthetic time series, some time series might exhibit punctual anomalies. As a result, time series generative models need to be able to *explain* their outputs both globally and locally in order to ensure their validity, to understand the sources of errors, and eventually to allow reliable usages.

While there exists a rich litterature studying *explainability techniques* for classifiers (*e.g.,* [2]) the issue of explaining generative models has largely been ignored until very recently. The need to provide differential privacy guarantees further complicates the issue because the privacy guarantees must cover the explainability algorithms in addition to the generative model, and they require to inject possibly large random perturbations at training time, introducing additional variance in the results.

The goal of this PhD thesis is to design, implement, and thoroughly evaluate explainability techniques for synthetic time-series generation algorithms with differentially private guarantees.

The main tasks of the PhD student will be to:

- Study the state-of-the-art work about privacy-preserving synthetic time-series generation algorithms, time-series explainability, and privacy-preserving explainability techniques for classifiers.
- Design differentially private explainability techniques for privacy-preserving synthetic time-series generation algorithms and thoroughly demonstrate and evaluate their privacy and utility guarantees.

- Contribute to the organisation of competitions where the privacy guarantees of synthetic time series generation algorithms are challenged[1] [1].

2 Profile of the candidate

- The candidate must have obtained, or be about to obtain, a master degree in computer science or in a related field.
- The candidate must be curious, autonomous, and rigorous.
- The candidate must be able to communicate in English (oral and written). The knowledge of the French language is not required.
- The candidate must have a strong interest in machine learning.
- Skills in cybersecurity, especially in privacy, will be appreciated.

3 Supervision and environment

This PhD offer is funded by the *Chaire CPDDF* (Fondation Univ. Rennes) and proposed by the Security and Privacy team (SPICY) from the IRISA institute in Rennes, France. The work will be supervised jointly by Tristan Allard (PhD, HDR) associate professor at the University of Rennes, expert in privacy in data intensive systems, and Romaric Gaudel (PhD, HDR), expert in machine learning and explainability.

The Chaire CPDDF involves six industrial partners (Apixit, Chambre Interdépartemental des Notaires, Crédit Mutuel Arkea, Enedis, Sogescot, Veolia) and one public organization (Région Bretagne). Regular meetings will be organized with the partners in order to favor collaborations on the topic.

The successful candidate will be working at IRISA – the largest French research laboratory in the field of computer science and information technologies (more than 850 people). IRISA provides an exciting environment where French and international researchers perform cutting edge scientific activities in all domains of computer science.

Rennes is located in the West part of France in the beautiful region of Brittany. From Rennes, you can reach the sea side in about 45 minutes by car and Paris center in about 90 minutes by train. Rennes is a nice and vibrant student-friendly city. It is often ranked as one of the best student cities in France. Rennes is known and appreciated for its academic excellence, especially in the field of cybersecurity, its professional landmark, the quality of its student life, the affordability of its housing offer, its rich cultural life, and much more.

4 Application

To apply, please send the following documents to both `tristan.allard@irisa.fr` and `romaric.gaudel@irisa.fr`:

- Your detailed CV.
- A short letter explaining your motivation for working on this project.
- The grade transcript of all university-level courses taken.
- Your master thesis.
- Names and contact information for two professional references.

The deadline for applying is September 14, however the applications will be evaluated as soon as they are received.

[1] See for example the SNAKE$_1$ challenge: `https://snake-challenge.github.io/`.

5 General information

Laboratory, Team IRISA institute (UMR 6074), SPICY team[2].

Supervisors Tristan Allard, Romaric Gaudel.

Start Autumn 2025.

Duration 36 months.

Location Rennes, France.

Funding Chaire CPDDF (Fondation Univ Rennes)[3].

References

[1] Tristan Allard, Louis Béziaud, and Sébastien Gambs. Snake challenge: Sanitization algorithms under attack. *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23)*, 2023.

[2] Romaric Gaudel, Luis Galárraga, Julien Delaunay, Laurence Rozé, and Vaishnavi Bhargava. s-lime: Reconciling locality and fidelity in linear explanations. In *International Symposium on Intelligent Data Analysis*, 2022.

[3] Zinan Lin, Alankar Jain, Chen Wang, Giulia C. Fanti, and Vyas Sekar. Using gans for sharing networked time series data: Challenges, initial promise, and open questions. *Proceedings of the ACM Internet Measurement Conference*, 2019.

[4] Yulian Mao, Qingqing Ye, Qi Wang, and Haibo Hu. Differential privacy for time series: A survey. *IEEE Data Eng. Bull.*, 48:67–92, 2024.

[5] Antonin Voyez, Tristan Allard, Gildas Avoine, Pierre Cauchois, Elisa Fromont, and Matthieu Simonin. The Privacy Cost of Fine-Grained Electrical Consumption Data. *Scientific Reports*, 15(17391):1–8, May 2025.

[2] Website of the SPICY team.
[3] Short presentation of the chaire CPDDF (in French).