Synthetic data generation: formal methods for privacy analysis

PhD thesis proposal

Supervisors: Tristan Allard (Univ. Rennes, IRISA), Barbara Fila (INSA Rennes, IRISA)

Contact: tristan.allard@irisa.fr and barbara.fila@irisa.fr

Keywords: privacy, formal methods, synthetic data generation, differential privacy, membership inference attacks, risk analysis.

1 Context and goal

Health data, social networks, electricity consumption... Vast quantities of personal data are collected today by private companies or public organizations. Various legal[1], monetary[2], or visibility incentives push data holders to envision sharing versions of the collected datasets that provide both *statistical utility and privacy guarantees*. Indeed, sharing data at large, *e.g.,* as *open data*, without jeopardizing privacy, is expected to bring strong benefits (strengthening, *e.g.,* scientific studies, innovation, public policies).

Synthetic data generation is a promising approach. First, synthetic data generation algorithms aim at generating datasets that are as close as possible to the original datasets. Either *synthetically generated data* or the *generative models* trained over the original data could be shared for supporting elaborate data analysis. Second, substantial progress has been made during the last decade about the privacy guarantees of synthetic data generation algorithms. For example, there exist today synthetic data generation algorithms that satisfy variants of *differential privacy*, one of the most prominent family of privacy models [2].

However security is a constant race between the attackers and the defenders. A large number of attacks exists and keeps growing [5]. As a result, because of the complex environment in which synthetic data generation takes place (e.g., utility needs, diversity of information sources, diversity of data generation algorithms), analyzing the risks remains hazardous even when strong privacy-preserving techniques are used.

The main goal of this PhD thesis is to design a formal method based approach allowing data holders to analyze the risks related to their synthetic data publication practices.

The main tasks of the PhD student will be to:

- Study the state-of-the-art about attacks on synthetic data generation algorithms (e.g., membership inference attacks [4, 6]) and about relevant formal methods (*e.g.,* attack tree based risk analysis models [3]). We will focus on tabular data and time series.
- Model the full synthetic data generation environment. Most especially, this includes capturing the attackers' capabilities (e.g., goals [5], background knowledge, computational resources, sequences of steps), the relationships between attackers, the sources of auxiliary information, and the data sharing practices.
- Design efficient algorithms for finding the attacks that illustrate privacy risks, implement them, and evaluate their performance.

[1] See for example the EU Directive 2019/1024 on open data and the re-use of public sector information.

[2] For example, through data marketplaces such as Innodata (https://innodata.com/ai-data-marketplace/) or Defined.ai (https://www.defined.ai/).

In addition to the core tasks of the project, the successful candidate will also contribute to the organisation of competitions where the privacy guarantees of synthetic data generation algorithms are challenged[3] [1].

2 Supervision and environment

This PhD offer is funded by the PEPR Cybersecurity IPoP project 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 Barbara FILA (PhD, HDR), associate professor at INSA Rennes, expert in formal methods for risk assessment.

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.

3 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 cybersecurity.
- Skills in machine learning and/or formal methods will be appreciated.

4 General information

Laboratory, Team IRISA institute (UMR 6074), SPICY team[4].

Supervisors Tristan Allard, Barbara Fila.

Start As soon as the position is filled.

Duration 36 months.

Location Rennes, France.

Funding PEPR Cybersecurity IPoP project[5].

[3] See for example the SNAKE$_1$ challenge: https://snake-challenge.github.io/.
[4] https://www-spicy.irisa.fr/
[5] https://files.inria.fr/ipop/

5 Contact and application

To apply, please send the following documents to both `tristan.allard@irisa.fr` and `barbara.fila@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.

Informal inquiries are welcome.

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] Damien Desfontaines and Balázs Pejó. Sok: Differential privacies. *Proceedings on Privacy Enhancing Technologies*, 2020(2):288–313, 2020.

[3] Barbara Kordy (Fila), Ludovic Piètre-Cambacédès, and Patrick Schweitzer. Dag-based attack and defense modeling: Don't miss the forest for the attack trees. *Comput. Sci. Rev.*, 13-14:1–38, 2014.

[4] Hongsheng Hu, Zoran A. Salcic, Lichao Sun, Gillian Dobbie, P. Yu, and Xuyun Zhang. Membership inference attacks on machine learning: A survey. *ACM Computing Surveys (CSUR)*, 54:1 – 37, 2021.

[5] Ahmed Salem, Giovanni Cherubin, David Evans, Boris Köpf, Andrew Paverd, Anshuman Suri, Shruti Tople, and Santiago Zanella-Béguelin. Sok: Let the privacy games begin! a unified treatment of data inference privacy in machine learning. In *Proceedings of the 2023 IEEE Symposium on Security and Privacy (S&P '23)*, pages 327–345, 2023.

[6] Antonin Voyez, Tristan Allard, Gildas Avoine, Pierre Cauchois, Élisa Fromont, and Matthieu Simonin. Membership inference attacks on aggregated time series with linear programming. In *Proceedings of the 19th International Conference on Security and Cryptography (SECRYPT '22)*, 2022.