Master Internship Proposal contact edwige.cyffers@ens-lyon.fr
Dauphine Lamsade 2026

Performative Privacy

Motivation

Machine learning relies on the use of large amounts of data to train models. However, this implies privacy risks, as models are prone to memorizing parts of the training set and can leak sensitive information. Differential privacy helps mitigate these risks but tends to limit the model's utility, giving rise to a privacy-utility trade-off. A widely shared opinion is that this trade-off is misleading, as it presents privacy solely as a constraint that limits the maximum achievable utility at a given time, whereas one could argue that people may oppose data collection in the future in case of data leakage. For instance, users may opt out or perturb their input depending on the privacy level. In the long term, enforcing privacy would thus be the best way to maximize utility by maximizing data collection.

The goal of the internship, potentially followed by a PhD, is to mathematically encode such a setting and to prove under which conditions enforcing privacy leads to a better equilibrium.

Context

Differential Privacy [1] mathematically quantifies the worst-case information leakage about a single entity in the dataset through its influence on the algorithm's outputs, thus providing protection against all attacks. More precisely, an algorithm $\mathcal{A}$ is differentially private if, for all pairs of datasets $D \sim D'$ differing by a single participant, and every subset $\mathcal{S} \subset Z$, the following inequality holds:

$$\mathbb{P}(\mathcal{A}(D) \in \mathcal{S}) \leq \exp(\varepsilon)\, \mathbb{P}(\mathcal{A}(D') \in \mathcal{S}) + \delta.$$

DP is used in deployment (e.g., Apple statistics on emoji, Google next-word prediction, LinkedIn audience measurement, the US Census, Wikimedia pageviews) and enables private statistics and private machine learning [2].

Performative Learning [4] addresses distribution changes due to model deployment, aiming to minimize

$$\mathrm{PR}(\theta) = \mathop{\mathbb{E}}_{Z \sim \mathcal{D}(\theta)} \ell(Z; \theta).$$

Unlike classical machine learning where $\mathcal{D}$ is fixed, here it is parameterized by the same parameter θ used in the loss function. Finding a performatively optimal solution requires accounting for the model's performative effect on the distribution. Typical examples include loan applications or hiring, where applicants may modify their features to increase their probability of being classified in the positive class, and the optimal classifier should thus take into account the extent to which a given feature may be modified. This problem has also been studied in terms of regret minimization [3], to ensure that the whole trajectory minimizes the risk.

Performative privacy [5] was used to describe active resistance to data collection due to data misuse, for instance people wearing large hoodies to escape public surveillance, but has not been studied yet from a mathematical point of view. It should, however, be possible to model this discontent as a form of performative feedback in case of data leakage. For instance, users may opt out or perturb their input depending on the privacy level, leading to sparser data collection or the introduction of harmful modifications in the collected data.

For the candidate

The ideal candidate is finishing a Master's degree in Machine Learning, Applied Mathematics, or a related field, and is curious about the social impacts of machine learning. The candidate should have a solid background in basic probability and optimization. Prior knowledge of differential privacy is a plus, but not required. The balance between theoretical and empirical work may depend on the candidate. The project will likely start with the study of simple private statistics computation and simple feedback, and then gradually move to more complex machine learning models.

The internship will take place in Paris intra-muros, at Université Paris Dauphine or PariSantéCampus, within the LAMSADE team, with regular meetings. The internship offers the standard "gratification", around 700 euros per month.

Starting date The expected starting date is between mid-April and June 2026 (flexible).

Extension into a PhD is possible and shall be decided upon mutual agreement. A dedicated funding is available.

To apply send an email to `edwige.cyffers@ens-lyon.fr` with:

1. A short description of your motivation: what you understand from the project, why you believe you are a good fit (maximum 500 words)
2. A CV
3. Your transcript (first year of the Master's and available grades for the second year)

References

[1] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. "Calibrating Noise to Sensitivity in Private Data Analysis". In: *Theory of Cryptography*. Berlin, Germany: Springer Berlin Heidelberg, 2006.

[2] Cynthia Dwork and Aaron Roth. "The Algorithmic Foundations of Differential Privacy". In: *Foundations and Trends® in Theoretical Computer Science* 9.3-4 (2013).

[3] Meena Jagadeesan, Tijana Zrnic, and Celestine Mendler-Dünner. "Regret Minimization with Performative Feedback". In: *ICML*. 2022.

[4] Juan Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. "Performative Prediction". In: vol. 119. Proceedings of Machine Learning Research. PMLR, 2020.

[5] Scott Skinner-Thompson. "Performative Privacy". In: *UC Davis Law Review* 50.4 (2017). NYU School of Law, Public Law Research Paper No. 17-10; U of Colorado Law Legal Studies Research Paper No. 17-4.