

Data Privacy and Trustworthy Machine Learning

Martin Strobel

National University of Singapore

Reza Shokri

National University of Singapore

Abstract—Privacy risk is a major concern when training machine learning models on sensitive data. Research shows that some constraints trustworthy machine learning poses on training processes can create considerable privacy concerns. We discuss the tradeoffs between data privacy and other goals of trustworthy machine learning (notably fairness, robustness, and explainability).

■ **MACHINE LEARNING ALGORITHMS** try to discover patterns from the training data that can be used to perform particular predictive tasks on new data. The application domains of machine learning touch many aspects of personal life, ranging from image processing used in health care and self-driving cars, natural language processing for personal assistants, to decision support systems in educational and judicial systems.

In critical domains such as medicine, finance, education, and the judicial system, we need to analyze machine learning algorithms with respect to their trustworthiness, beyond their average prediction accuracy. Can these models be trusted to operate on our personal and *private* data? Are these systems *fair*, or do they replicate or potentially increase existing bias in society? Can we rely on their decisions, by being able to *explain* how models arrive at a decision? Can we intervene if they draw conclusions for the wrong reasons? Are the models *robust* to manipulation during deployment and training, or can small changes to the data result in adversarial behavior?

To ensure *trustworthy machine learning*, we need to pose additional constraints on the models we can create. We use specifically designed algorithms to make models privacy-preserving,

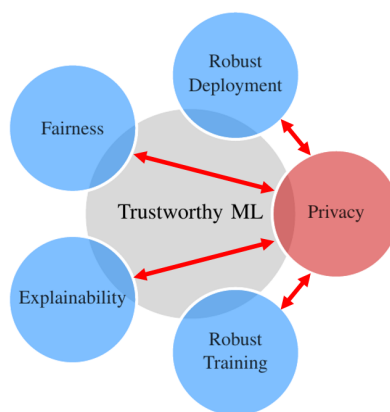


Figure 1. Different pillars of trustworthy machine learning. In this article, we discuss how data privacy interacts and conflicts with the other aspects.

fair, robust, or explainable. These algorithms, however, come with trade-offs. There is a significant body of research on studying the cost of trustworthy ML on the performance of models with respect to their prediction accuracy. What is very concerning is that constraining a model could also pose a risk to the privacy of training data. In this article, we discuss the interaction of data privacy with other pillars of trustworthy




Figure 2. The typical machine learning workflow. Model parameters (learned from private training data) and model predictions can indirectly leak sensitive information about the training data. This leakage can happen even if the entire training process is confidential (e.g., using secure MPC).

machine learning.

Threats to Data Privacy

When considering data privacy, it is important to differentiate between direct and indirect leakage of information. Direct leakage can happen during training or inference time when an untrusted entity can access the data. This threat can be prevented via access control and confidential computing (secure multi-party computation, homomorphic encryption, and trusted hardware). These measures, however, cannot prevent indirect leakage, which is the focus of this article.

An adversary, who can only observe the model and not the training data, can use inference algorithms to reconstruct information about the training data. Obviously, it lies in the nature of useful machine learning that model parameters, and consequentially model predictions, contain *some* information about the training data.

To obtain a precise definition for data privacy, we need to differentiate between general patterns that apply to the entire *population*, which we would want to reveal, and the patterns that apply to specific data of *individual* users, which we want to keep private. Hence, learning anything about individual data records beyond the general patterns should be considered a privacy violation. This is the basis of *differential privacy*, which measures the worst-case sensitivity of an algorithm to changes in individual users' data.

We design a (hypothetical) *indistinguishably*

game, between the algorithm and the adversary, to measure the amount of information leakage of an algorithm about its training data. In the game, we consider two possible worlds. In the first world, a model is trained on a dataset containing the information of a particular individual x . In the second world, the model is trained on the same dataset, except this time x 's data is removed. The adversary interacts with the model, in an unknown world, and tries to infer if the model's training set contains x 's data. The capacity to which the adversary can win this game tells us the extent of individual private information leakage. In the white-box setting, the attacker has complete access to the model's parameters. In the black-box setting, the attacker can only interact with the model's predictions. Shokri et al. [12] designed *membership inference attacks*, which are algorithms to simulate the game and measure the information leakage of models about their training data. Besides, the reports by the US National Institute for Standards and Technology (NIST)¹ and the UK Information Commissioner's Office (ICO)² specifically mention membership inference as a confidentiality violation and potential threat to the training data in AI.

In the black-box setting, membership inference attacks attempt to exploit the signals contained in a model's predictions. Major examples

¹<https://nvlpubs.nist.gov>

²<https://ico.org.uk>

of such a signal are the prediction error and uncertainty. Predictions tend to be more accurate and certain for members of the training set compared to unseen points. Powerful membership inference attacks have been demonstrated for many different scenarios in machine learning [10, 15].

We can use membership inference attacks to quantitatively measure the privacy risks of machine learning algorithms in many diverse scenarios. Throughout this article, we will use membership inference attacks as a tool to measure information leakage of private data. This way, we can study the privacy implications of the other aspects of trustworthy machine learning.

Given a set of points $X \subseteq \mathbb{R}^n$ with membership $m: X \rightarrow \{0, 1\}$ to a training set, we define the average *privacy risk* of a trained model as³

$$\max_{A \in \mathcal{A}} \frac{1}{|X|} \sum_{x \in X} [A(s(x)) = m(x)].$$

Here, $s: X \rightarrow \mathbb{R}^k$ is the signal the attacker can observe for each point after a model is trained. The signal can be the model’s prediction, the model’s loss, an explanation, or similar observable signals about a data point (which differentiates members from non-members). \mathcal{A} is the set of possible attack algorithms the attacker can use. For actual evaluations [3, 11, 12, 13, 14] specific attack algorithms are designed. Usually, X is assumed to be balanced, i.e., the lower bound on the privacy risk is 0.5 obtained by an attacker randomly guessing membership.

Murakonda and Shokri [9] provide a tool called *Privacy Meter*⁴ that can quantify privacy risks of a wide range of machine learning algorithms (e.g., models for classification, regression, computer vision, and natural language processing). The tool relies on state-of-the-art inference techniques and generates extensive reports.

A common attack strategy are *threshold-based attacks* [11, 15]. Here, the attacker directly compares the signal to a predefined threshold τ . For example, it is common for the loss of training set members to be lower than for non-members. In a threshold-based attack using the point’s loss as a signal, each point with a loss lower than the threshold is considered a member, and each

³See Ye et al. [15] for a more comprehensive discussion on how to measure privacy risk in machine learning.

⁴<https://privacy-meter.com>




Figure 3. Experimental results on a Law School Admission problem provided by Chang and Shokri [3]. There is a tradeoff between privacy and fairness of a model. Increasing fairness (i.e., decreasing δ) increases the privacy risk of the qualified unprivileged group.

point with a loss higher is considered a non-member. The attacker can obtain the threshold by observing the loss distribution of points on reference (or shadow) models. All experimental results presented here were originally obtained using threshold-based attacks.

Algorithmic Fairness

Machine learning models potentially cause disproportionate harm to specific *groups*. This bias might arise from demographic disparities in the training data, the implicit focus on optimizing performance on the majority group, or at other steps in the machine learning pipeline [2].

While there are many different fairness concerns, the research community has paid the most attention to statistical notions of *group unfairness* for classification tasks. For this, we observe the differences in the behavior of a classifier for inputs of different demographic groups according to specific protected attributes (e.g., gender, race). In this context, protected attributes are provided by anti-discrimination laws or societal standards.

As an illustrative example, let us consider the following hypothetical scenario: A company wants to use an automated classifier to approve loan applications based on various user-provided data. An auditor can split a dataset for which the

classifier has made approval decisions into the decisions made for male and female applicants. One criterion of fairness would be to require that applicants that pay back their loan have an equal opportunity of getting the loan in the first place. This notion of fairness is known as *equal opportunity* [6]. Formally, we can say that a classifier $f: \mathbb{R}^n \rightarrow \{0, 1\}$ satisfies δ -Equal Opportunity with respect to two groups X_1, X_2 and true labels $y: \mathbb{R}^n \rightarrow \{0, 1\}$, if the false negative rate of the classifier in the groups are within δ range of one another:

$$\begin{aligned} & |\mathbb{P}_{x \sim X_1} [f(x) = 0 | y(x) = 1] \\ & - \mathbb{P}_{x \sim X_2} [f(x) = 0 | y(x) = 1]| \leq \delta. \end{aligned}$$

Note that equal opportunity is only one of many potential fairness notions. The auditor could also compare false-positive rates among the groups, or the auditor could validate if the approval rates are the same among the genders.

Fairness-aware learning is a way to mitigate the bias of a given classifier, by enforcing the fairness constraints. These algorithms can either modify the training set and the training algorithm, or post-process the fully trained model.

“Achieving group fairness comes at a cost for individual privacy.”

Chang and Shokri [3] show that the achieved group fairness comes at a cost for privacy. They create a group-specific membership inference attack to demonstrate that the privacy risks of different demographic groups are affected disparately by fairness-aware learning. In particular, the privacy risks of the unprivileged group (i.e., the group the original classifier is biased against) increases. Higher bias in the unconstrained model leads to greater privacy risks of the unprivileged class in the corresponding fair model. Further, the stronger the fairness constraint is enforced, the more the privacy risks increase.

Figure 3 reproduces their experimental results on how enforcing a smaller fairness gap increases the privacy risk. The experiment considers a depth-10 decision tree trained on the law school admission dataset. Fairness is enforced using the reductions approach (see [3] for more details).

Why is there this tradeoff between privacy and fairness? An answer lies in how we cre-

ate unbiased models. By enforcing constraints, fairness-aware algorithms ensure the equal performance of the model on different subgroups. However, correctly learning the underprivileged group’s classification might be difficult for several reasons. There might be fewer data available for the group, the data might have higher variance, or the task might be more complex. Hence, this might force the model to *memorize* more of the underprivileged group’s data instead of correctly learning a general pattern. This higher memorization links directly to higher privacy risk. It is easier for the adversary to differentiate between members and non-members. It follows that group fair models have a higher privacy risk.

Bagdasaryan et al. [1] provide complementary evidence for the privacy-fairness tradeoff. They demonstrate that the accuracy reduction caused by differentially private training disparately affects different groups. Smaller, unprivileged subgroups suffer more. Again the intuitive connection stems from memorization. Smaller groups rely more on memorization, which differentially private training suppresses.

For some settings, this tradeoff is even theoretically unavoidable. Cummings et al. [4] show cases where a model cannot be trained in a differentially private way and still satisfy group fairness constraints exactly.

Overall we have strong experimental [1, 3] and theoretical [4] evidence on the incompatibility of individual privacy and group fairness.

Explainability

Modern machine learning often achieves performance improvements by increasing the complexity of the model architectures and training process. In consequence, the patterns learned by the models are harder to understand. Further, it becomes difficult or impossible to comprehend why a model came up with a specific decision. However, this lack of understanding is undesirable for many critical decision-making scenarios. Here, we require that the semantics of the model align with the semantics of the tasks. This way, we can be confident that a model did not pick up spurious correlations in the training data [8].

Feature-based attribution methods are model explanations that try to highlight which features of the input were important to a specific point’s




Figure 4. The privacy risks of different backpropagation-based explanation methods of a neural network trained on the Texas hospital dataset. Provided by Shokri et al. [13].

prediction. Many of these attribution methods are based on the model’s gradient with respect to the input. The gradient tells us exactly how much a model’s output changes if we make infinitesimally small changes to input. So methods derived from the gradient highlight influential input features.

“Model explanations can be exploited by inference attacks.”

From the privacy perspective, model explanations provide the attacker with an additional source of information. This is especially true in scenarios where the attacker cannot directly access a model’s uncertainty or loss. High feature attribution values are a proxy for model uncertainty. They indicate that a small change to the input would radically change the model’s output. Hence, an attacker can construct successful membership inference attacks solely based on model explanations to distinguish between members and non-members [13].

Figure 4 shows the privacy risks for a membership inference attack based on different explanation methods for a neural network trained on the Texas hospital dataset (see [13] for more details). So, model explanations can also conflict with individual privacy.

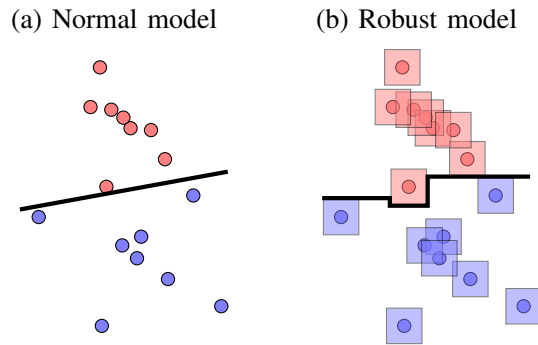


Figure 5. Many of the existing robust training algorithms change the decision boundary around training points in a specific way. It is possible for an attacker to discover and exploit these patterns to create more successful inference attacks.

Robustness during deployment

Decision boundaries of modern neural networks are very complex. For many normal inputs, an adversary can craft a very small change in a specific direction. The manipulation is almost unnoticeable yet still crosses the model’s decision boundary. So, adversaries can force wrong classifications and potentially bypass security checks.

So, a robust model is required not just to correctly predict the label for normal and benign inputs, but to remain correct if the input is slightly manipulated. For a distribution $(x, y) \sim \mathcal{D}$ of points x and labels y , the *standard supervised learning objective* is to find a parameter θ that minimizes the expected loss L of a classifier f_θ

$$\min_{\theta} \mathbb{E}_{(x,y) \sim \mathcal{D}} [L(f_\theta(x, y))].$$

Assuming that an adversary can modify each point by $\delta \in S$ the *robust supervised learning objective* is to minimize the maximal loss an attacker can achieve [7]:

$$\min_{\theta} \mathbb{E}_{(x,y) \sim \mathcal{D}} \left[\max_{\delta \in S} L(f_\theta(x + \delta, y)) \right]$$

Many defenses against evasion attacks affect the model’s decision boundaries so that, within a small area around each input, model predictions remain the same [7]. Yet, the methods can enforce this objective only on the training data. It follows that the impact of individual training points on robust models increases. The more extensive influence raises the privacy risk of




Figure 6. The privacy risks of robust-training methods compared to standard training of a neural network trained on the CIFAR10 dataset. Provided by Song et al. [14].

training set members.

“Defending against evasion attacks (adversarial examples) via adversarial training can increase privacy risks.”

In Figure 6, we reproduced Song et al. [14]’s results for the privacy risks of different robust training methods for a neural network trained on the CIFAR10 dataset. The entire current generation of robustness algorithms that change the model’s training behavior seems to suffer from this robustness-privacy tradeoff. However, it is unclear whether there exists a fundamental tradeoff between robustness to attacks and privacy.

Robustness during training

The quality and reliability of the training data cannot be guaranteed in all scenarios. Crowdsourced data might suffer from unmotivated or malicious crowd workers. Data collection might be noisy. In a federated setting, data from different clients is not of the same quality. In the worst case, an adversary might actively try corrupting parts of the training data to manipulate the training process and harm the model’s performance.

Among all threats *data poisoning* attacks have gotten the most attention. We can only consider the model’s training robust if it has some form of immunity to data poisoning attacks. Modern machine learning algorithms are generally not robust

against such threats. Many works have demonstrated the possibility of manipulating models via poisoning their training data.

Some robust training algorithms try to reduce the sensitivity of the training to changes of small parts of the training set. Others try to detect and completely ignore outliers. The goal is to make the model’s training focus on general patterns and avoid being influenced by distinct points.

“Robustness against poisoning attacks aligns well with data privacy.”

In contrast to the other goals of trustworthy machine learning, the constraints posed by training data robustness align well with data privacy. Both differential privacy (DP) and robustness are notions of insensitivity to training data changes. While the existing metrics for these notions are not exactly the same, DP focuses on worst-case changes of single inputs and robustness on restricted changes of groups of data. There are reasons for cautious optimism. Research of robust statistics further supports this optimism as robust mechanisms are demonstrably good starting points for differentially private mechanisms [5].

Conclusions

In this article, we have discussed many aspects in which current ways to ensure a more trustworthy classifier come at a cost to data privacy. Yet, the situation is not clear-cut. While group fairness comes with a fundamental tradeoff to privacy [4, 3], for the other aspects, much is unknown. Current algorithms to make models robust against attacks during deployment hurt data privacy [14]. But we cannot rule out the possibility of designing robust algorithms which are also accurate and privacy-preserving. The situation is similar in the case of explainability. Here we also understand that some existing methods are vulnerable [13]. Yet, we do not know if it extends to all types of explanations.

Trustworthy machine learning is a laudable goal, but it should not come at the cost of data privacy. We call for more investigation of the interactions of these two aspects. Instead of looking at these problems separately, we need to develop techniques to achieve both.

Finally, privacy concerns in machine learning

are not limited to membership inference. Researchers have considered further threats such as reconstruction attacks, model extraction, or property inference. These concerns also interact with the other aspects of trustworthy machine learning and require further study.

REFERENCES

1. Eugene Bagdasaryan, Omid Poursaeed, and Vitaly Shmatikov. Differential privacy has disparate impact on model accuracy. *Advances in Neural Information Processing Systems*, 32:15479–15488, 2019.
2. Solon Barocas, Moritz Hardt, and Arvind Narayanan. *Fairness and Machine Learning*. fairmlbook.org, 2019. <http://www.fairmlbook.org>.
3. Hongyan Chang and Reza Shokri. On the privacy risks of algorithmic fairness. In *2021 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 292–303. IEEE, 2021.
4. Rachel Cummings, Varun Gupta, Dhamma Kimpara, and Jamie Morgenstern. On the compatibility of privacy and fairness. In *Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization*, pages 309–315, 2019.
5. Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 371–380, 2009.
6. Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29, 2016.
7. Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
8. Christoph Molnar. *Interpretable machine learning*. Lulu. com, 2020.
9. Sasi Kumar Murakonda and Reza Shokri. MI privacy meter: Aiding regulatory compliance by quantifying the privacy risks of machine learning. *arXiv preprint arXiv:2007.09339*, 2020.
10. Milad Nasr, Reza Shokri, and Amir Houmansadr. Comprehensive privacy analysis of deep learning. In *2019 IEEE symposium on security and privacy*, 2018.
11. Alexandre Sablayrolles, Matthijs Douze, Cordelia Schmid, Yann Ollivier, and Hervé Jégou. White-box vs black-box: Bayes optimal strategies for membership inference. In *International Conference on Machine Learning*, pages 5558–5567. PMLR, 2019.
12. Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *Proceedings of the 38th IEEE Symposium on Security and Privacy (Oakland)*, pages 3–18, 2017.
13. Reza Shokri, Martin Strobel, and Yair Zick. On the privacy risks of model explanations. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 231–241, 2021.
14. Liwei Song, Reza Shokri, and Prateek Mittal. Privacy risks of securing machine learning models against adversarial examples. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 241–257, 2019.
15. Jiayuan Ye, Aadyaa Maddi, Sasi Kumar Murakonda, and Reza Shokri. Enhanced membership inference attacks against machine learning models. *arXiv preprint arXiv:2111.09679*, 2021.

Martin Strobel is a Ph.D. candidate at the National University of Singapore’s School of Computing working under the supervision of Yair Zick and Reza Shokri. Before that, he obtained his Master of Mathematics, a Bachelor’s degree in Computer Science, and a Bachelor’s degree in Mathematics from the Technical University of Munich. His research interests are in the privacy and transparency aspects of machine learning.

Reza Shokri, is a NUS Presidential Young Professor of Computer Science. His research focuses on data privacy and trustworthy machine learning. He is a recipient of the IEEE Security and Privacy (S&P) Test-of-Time Award 2021, for his paper on quantifying location privacy. He received the Caspar Bowden Award for Outstanding Research in Privacy Enhancing Technologies in 2018, for his work on analyzing the privacy risks of machine learning models. He received the NUS Early Career Research Award

Department Head

2019, VMWare Early Career Faculty Award 2021, Intel Faculty Research Award (Private AI Collaborative Research Institute) 2021, and Meta Faculty Award 2021. He obtained his PhD from EPFL.