

Collaborative Drug Discovery: Inference-level Data Protection Perspective

Balazs Pejo*, Mina Remeli**, Adam Arany***, Mathieu Galtier****, Gergely Acs*****

*CrySyS Lab, BME, Hungary, pejo@crysys.hu (Corresponding author)

** University of Cambridge, United Kingdom, mincsek@gmail.com (Work was done while at CrySyS Lab)

***Stadius, KUL, Belgium, adam.arany@kuleuven.be

****Owkin, France, mathieu.galtier@owkin.com

*****CrySyS Lab, BME, Hungary, acs@crysys.hu

Received 30 October 2021; received in revised form 27 April 2022; accepted 29 April 2022

Abstract. Pharmaceutical industry can better leverage its data assets to virtualize drug discovery through a collaborative machine learning platform. On the other hand, there are non-negligible risks stemming from the unintended leakage of participants' training data, hence, it is essential for such a platform to be secure and privacy-preserving. This paper describes a privacy risk assessment for collaborative modeling in the preclinical phase of drug discovery to accelerate the selection of promising drug candidates. After a short taxonomy of state-of-the-art inference attacks we adopt and customize several to the underlying scenario. Finally we describe and experiment with a handful of relevant privacy protection techniques to mitigate such attacks.

Keywords. Drug Discovery; Machine Learning; Privacy; Risk Analysis; Membership Inference

1 Introduction

Drug discovery and development is currently a time-consuming and expensive process. Many small molecules fail in the later stages of development due to toxicity or a lack of sufficient efficacy. More complex mechanisms of action and stricter regulation have resulted in a drastic drop in pharmaceutical R&D productivity over the past 60 years. Indeed, only 11% of the drug candidates that enter clinical phases make it to approval, i.e. are demonstrated to be productive and safe. As a result, it takes on average 1.9 billion euros and often more than ten years to bring one successful drug to the market.

Machine Learning (ML) is a promising approach to reducing the time and cost of this development process, where a carefully trained model can predict the biological activities of small molecules. The available volume of training data mainly determines the quality of these models. Consequently, the industry is making the first steps towards federated ML

approaches that leverage more data than a single partner (e.g., a pharmaceutical company). Federated learning (FL) enables several partners to train a common model without directly sharing their training data, still obtaining a model that can predict biological activities almost as accurately as a model trained centrally on the union of their training data. Such collaborations between pharmaceutical companies can save them millions of dollars today and save lives tomorrow. Hence, there is a great interest in such collaborations both within and outside the pharmaceutical industry. Indeed, a similar scenario is currently being studied by several pharmaceutical companies within the MELLODDY project (EU/EFPIA IMI 831472).

However, FL is not a panacea. Several prior works have demonstrated that FL per se does not provide sufficient protection against the unintended leakage of training data. Even if participants only exchange gradients during training and can only access the common model at the end, both the gradients and the model itself leak information about the training data. For example, as we show, exchanged model parameters can reveal that a specific participant has a particular compound in its training data, which could indicate what drugs the participant is experimenting with. Such inference attacks can seriously undermine the ultimate goal of FL and cause substantial financial loss to a company investing millions of euros into curating high-quality training data.

In this paper, we perform a data protection risk analysis of a FL on an real use case, where several pharmaceutical companies collaboratively train a multi-task federated machine learning model to predict activities between chemical compounds (i.e., drug candidates) and different targets (e.g., proteins). Such federated structure-activity modeling in Drug Discovery during pre-competitive collaboration is a well-suited problem for our analysis as it captures and illustrates well the core trade-off that the privacy-aware participants face. In this context, data protection refers to protecting the confidentiality of intellectual properties (IPs) such as a participant's training data or some model (hyper)-parameters¹. Attacks aim to infer confidential attributes of these assets (such as the membership of specific training samples) and are therefore called *inference attacks*. The pharma partners specify the exact attributes, and their successful inference is represented as a *risk* in our methodology. We design inference attacks (threats) to estimate the likelihood of these risks. In particular, the feasibility and the success probability of these attacks measure the likelihood of the risk, which then, combined with the risk's severity (negative impact), allows for identifying the most dangerous risks to be mitigated. We are only concerned with passive inference attacks in this analysis. Indeed, perturbing the data results in a more considerable accuracy drop for a malicious party than for the others [30]. Hence, there is no incentive for active attacks (e.g., poisoning, back-doors, etc. [14]) as long as the adversary also needs good model quality. The specific contributions of this study are three-fold.

- We demonstrate the process of data protection risk analysis on a real machine learning application. We propose a risk analysis methodology that helps systematically map different threats and hence identify risks that should be mitigated. Although we focus on risks stemming from the leakage of intellectual properties in FL, our methodology equally applies to the case when ML models process personal data and therefore are subject to data protection regulations (e.g., GDPR). Our approach is in line with the risk-based approach proposed by several regulations on AI (such as the upcoming European AI Regulation²) and hence can be useful to audit ML models in

¹Data protection and privacy are often used interchangeably in the literature, though "privacy" usually refers to personal data which is not the case in our application.

²<https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>

practice.

- We design and implement two novel and scalable (membership) inference attacks. They detect information leakage through gradients of any training sample and the trained model. Even if secure aggregation is in place [5], which prevents the attribution of membership information to a specific partner, we show that differential attacks can still be used to assign membership information. We develop a novel differential attack that exploits that some change in the composition of the training participants also changes the membership attack’s output distribution, which statistical tests can detect. These attacks show that — just as in other applications of FL — gradients and the activation values of the commonly trained model can leak sensitive information about the participants’ training data.
- We propose technical countermeasures to mitigate these threats. We show that secure aggregation [5] does prevent trivial attribution of the inferred attribute to a specific partner. Still, it is not sufficient to avoid attribution if the composition of the partners changes during training. Although Differential Privacy (DP) [7] defeats such attribution, we show that DP implies intolerable accuracy drop-in practice. Hence, we explore ad-hoc mitigation approaches and test data protection empirically with our attacks. Although these techniques provide no formal privacy guarantee, they still mitigate the attacks with acceptable utility loss. For example, we find that thresholding model parameters effectively mitigate information leakage without degrading model quality too much. This also shows that model parameters are highly redundant, which allows a model to memorize complete training samples that our attacks can detect.

2 Related work

We review the literature concerning inference attacks. There are many privacy-related attacks concerning machine learning; here, we give a non-comprehensive list of categories based on the attacker’s goal. For more comprehensive surveys, we refer the reader to [24, 33]. Model inversion attacks aim to reconstruct a representative value of a class [10], e.g., a record that is similar to all records belonging to a class. Membership inference attacks are aimed at inferring if a certain record was part of the target model’s training dataset [16]. The most common techniques to achieve this utilize shadow models [28, 37, 39, 42] or rely on overfitting [20, 27, 31, 35]. Reconstruction attacks take membership inference attacks another step forward by identifying the whole training dataset (instead of one record) [15, 22, 36, 44, 47]. Property inference attacks aim to infer properties of training data that are independent of the features that characterize the classes of the joint model [11, 26, 29, 43]. And finally, model extraction attacks arise when an adversary obtains black-box access to some target model and attempts to learn a model that closely approximates or even matches the original model [41].

Despite the wide range of attacks that exists, in this paper, we focus on *membership inference attacks*. These are the most elemental attacks, highlighting one bit of information leakage. If membership inference succeeds, that flags (one-bit) information leakage, a requirement for other attacks, which potentially leak more bits. On the other hand, if it does not succeed, that can be a solid empirical argument that other attacks (that leak more information) will probably fail as well.

Several mitigation techniques have been proposed against the risks of membership inference attacks, such as regularization or hyperparameter tuning (including model adaptation), which reduce overfitting [45]. DP [7] can also be applied to ML models providing provable privacy guarantees albeit with significant accuracy degradation in general. Two dedicated defenses, adversarial regularization [28], and MemGuard [17] have also been proposed to combat membership inference attacks. As MemGuard requires post-processing the output of the ML model, while adversarial regularization is scalable only to a minimal number of attacks, we focus on hyperparameter tuning, regularization, and DP in this work.

3 Context: Collaborative Drug Discovery

Quantitative Structure-Activity Relationship (QSAR) models are commonly used to predict bioactivity³ of chemical compounds on different targets [8]. The input to such models is a representation of a chemical compound (e.g., an extended-connectivity fingerprint (ECFP) [34]) which encodes different properties of the molecule. The output is the predicted bioactivity of the input molecule on specific targets (e.g., proteins). As building such models requires substantial data, collaborative learning is a compelling approach for multiple pharma companies to benefit from each other's confidential data collectively.

Chemical fingerprints like the ECFP encodes a chemical structure in the form of a vector of binary values or counts. A typical implementation is that substructures of the molecules are hashed to an integer address pointing to a position in a high dimensional vector, and the given position is either set to one or incremented. This high-dimensional vector is usually folded to smaller sizes to decrease the model size and improve generalization.

This paper employs feed-forward neural networks as QSAR models and trains in a multi-partner, multi-task setting. Neural networks are highly modular and scalable, making them, especially appealing in this context. Our neural network is separated into a *trunk* network and different *head* networks, as illustrated in Figure 1. The trunk is common for all tasks and participants and trained collaboratively, while the head architecture depends on the task and can be different among participants. The trunk extracts those features from a compound that can be used to maximize the overall prediction accuracy concerning all tasks.

In a collaborative learning setting, each participant trains its model (head and trunk together) on its training data locally. Only the parameter updates of the trunk are sent to the server for aggregation in every federated round. The weights and the whole architecture of the head networks are confidential and therefore not shared. The single role of the server (aggregator) is to sum up, the trunk updates sent by every participant and re-distribute the result. Each participant updates its trunk model with the common aggregate and obtains a new trunk model that is identical for all participants.

Therefore, the training data of participant j consists of pairs of $(\mathbf{x}, \mathbf{y}^j)$, where $\mathbf{x} \in \{0, 1\}^n$ denotes the fingerprint of the input compound (n is the folded fingerprint size and is identical for all participants), and $\mathbf{y}^j \in \mathbb{R}^{k_j}$ denotes the bioactivity values of this compound (k_j is the number of targets in participant j 's data), that is, y_i^j denotes the bioactivity value between compound \mathbf{x} and target i for participant j , and k_j is the output size of participant j 's head model.

³A beneficial or adverse effect of a (small molecular) drug on a living organism. In this context, the result of a measurement carried out on some preclinical model systems (enzyme assay, cellular assay, etc.) can indicate a potential beneficial or adverse effect of the given drug candidate. Therefore a valuable asset for the company.

A typical trunk architecture consists of a single linear layer. For instance, if the trunk has 6000 units and a folded input size is 32000, each participant will send and receive a total of 192 million parameter values (as model update/aggregate) to/from the server in every federated round. In each round, every participant performs a single gradient descent iteration: first, it selects a batch of training samples uniformly at random, then computes the gradients (i.e., model update), and finally sends the update to the server for aggregation. This is also illustrated in the middle (white area) of Figure 2.

Notice that this is a typical cross-silo FL setting: collaborative drug discovery is usually executed by pharmaceutical companies equipped with the necessary training data. Hence, the number of participants is below a dozen, and each of them participates in every round (as all of them have dedicated hardware with constant availability). By contrast, in cross-device FL, only a subset of participants participate (out of possible thousands) with often intermittent connection. Besides, each participant performs only a single local gradient iteration on its training data in cross-silo FL. In contrast, participants perform multiple iterations in cross-device to increase bandwidth efficiency (at the cost of model accuracy and possibly convergence speed).

4 Methodology and Threat Model

As none of the available regulations on AI and data privacy mandate any specific risk analysis methodology, we propose a sufficiently general and simple one to demonstrate the risk-based approach promoted by most legal requirements.

The risk analysis is composed of the following main steps repeated until all Risks become acceptable: 1) define the perimeter of the analysis, i.e., the actors (Risk Sources) and the data involved (Assets), 2) define the controls already in place, 3) define for each Risk Source the associated Risks, 4) define Threats for each Risk, and 5) apply controls to mitigate Risks. Within this paper, we do one iteration of these steps focusing on confidentiality. We rigorously go through the assets we consider, the possible corresponding Risks and their Sources, and detail the related feasibility and likelihoods. We mention a couple of Threats that are realized with our proposed attacks. Finally, we suggest several mitigation strategies with a trade-off parameter between model accuracy and defensive efficiency.

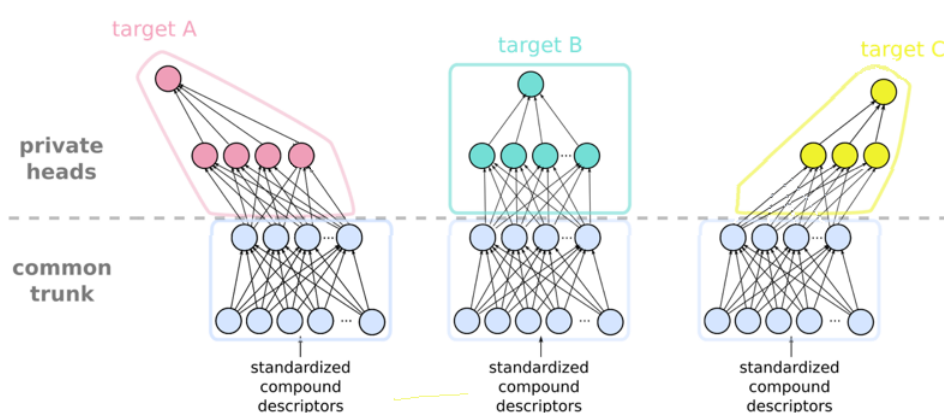


Figure 1: Head-trunk architectures with feed-forward neural networks.

4.1 Data and Actors

Assets. Assets are the confidential or private data that the Risk Sources aim to learn or manipulate, thereby causing Risk. We define the following assets: *Chemical Fingerprint (FP)* (the input of the model, a compact representation of a chemical compound), *Targets* (protein or others like toxicity for which the Bioactivity is being measured), *Bioactivity* (the predictive task labels for the chemical fingerprint input), and the *Trunk & Head model* (the collaboratively and privately trained neural network, respectively).

We assume that a fingerprint can be inverted, and the chemical compound represented by the fingerprint can be unambiguously identified. Indeed, in a recent work, it was shown that the chemical structures corresponding to an ECFP fingerprint could be recovered with reasonable accuracy by learning a feed-forward network to project to the hidden space of a recurrent autoencoder [19].

Risk Sources. A Risk Source is a person or non-human source that can cause a Risk, accidentally or deliberately. We only focus on human sources (aka, attackers), which aim to partially or entirely infer some confidential properties of the assets. We define two primary Risk Sources: *Pharma*, who is represented as a participant in collaborative learning, and *Server*, who coordinates and aggregates the training. This paper is NOT concerned with Risks and Threats stemming from outsider attacks, including but not limited to a cloud provider, external pharma companies, hackers, etc. Attackers are characterized by their background knowledge and their access to assets.

Background knowledge. Based on the attacker's role in the collaborative learning setting, they can access different data. The attack may need information about the training data, and depending on its extent, there are many categories. For example, the attacker could access the training dataset with or without labels or know only about the statistical properties (e.g., its distribution).

Access. The attacker can access all the details of the model (white-box) or only its prediction (black-box). In the drug discovery use case, this can be further specified due to the multi-task environment, e.g., in case of a white-box attack, does the attack access the weights and architecture of the commonly shared trunk model or the private head? The attacks could be separated based on the available meta-information, too: the attacker might monitor and use the performance metrics or utilize communication traffic patterns during collaboration, such as timing, size, or direction). Finally, the frequency of the exploitable information could also play a significant role in the attack's success.

4.2 Utilized Control: Secure Aggregation

Secure aggregation [2, 5] is a standard technique to prevent access to the model update of an individual partner, and it allows the server (and the participants) to learn only the aggregated model updates as illustrated in the middle (white area) of Figure 2. In other words, nothing is leaked except the aggregated model in every training round. Hence, inference (i.e., attribution) of any confidential information to any participant is not possible from the *aggregated* update, if all participants are involved in the aggregation in every round.

This is similar to secure multi-party computation (SMC) [6], where the intuitive security requirement is that nobody can learn information about the participants (e.g., dataset, gradients, etc.) than what they could already know from the aggregated result and their own

input. Rather than for the entire training process, secure aggregation implements an SMC protocol in each round separately. Indeed, besides the final model, the intermediate (secure) aggregates are also revealed in secure aggregation. This is a standard compromise in the literature [5] as end-to-end SMC techniques are still quite challenging for complex models with many parameters and large datasets [18] due to performance issues, which is precisely the case with drug discovery.

Before sending the model update (gradients) of the Trunk to the servers, the participants *mask* their update in such a way that the masks cancel out once aggregated (by the server). Hence, the server only learns the sum of model updates and not those of the individual members of this sum. Therefore, the server can only learn information about the set of participants as a whole without being able to attribute the information to any specific participant (e.g., it might be able to infer that a certain compound is used during the federated run, but it cannot assign this compound to any specific participant).

Note that secure aggregation does not eliminate privacy breaches completely: it does not hide the number of trunk parameters, and hence may leak information about the trunk’s architecture to any Risk Source (e.g., the server) who can access the encrypted model updates. Furthermore, the aggregated model could also be concealed from the server: we assume each participant can add an extra mask which is known to every participant except the aggregator. After aggregation this mask can be removed by any participant, i.e., the mask persists and do NOT cancel out after aggregation.

4.3 Risks

A Risk represents the goal of the Risk Source which is to infer confidential information about one of the asset. Its motivation is characterized by the incentives (e.g., financial gains) of the Risk Source to perpetrate the Risk balanced with the potential disincentives (loss of trust, damage to reputation, etc.). A Risk’s impact is measured by its associated Severity. In this paper, we only consider inference Risks, consequently, neither poisoning (intentional contamination of the training process) [4] nor free riding (participation in the training process with intentionally limited contribution) [23] is enlisted as a Risk in Table 1. We also do not consider all identified assets, our goal is not to provide a comprehensive list of risks but to demonstrate the methodology on a real-world machine learning use-case.

ID	Description of Risk	Risk Source	Severity
R1	A fingerprint is inferred with high probability to have been utilized by A) a given participant B) any participant.	Pharma Server	4
R2	Targets – which was not published before – can be inferred with high probability to have been utilized by A) a given participant B) any participant.	Pharma Server	4
R3	(A part of) the trunk model is inferred/reconstructed.	Server	3

Table 1: Identified Risks. Severity: 1 – Negligible, 2 – Limited, 3 – Significant, 4 – Maximum. Only R1 is considered in this paper.

ID	Description of Threat	Feas. & Like.	Corr. Risk	Risk Source
T1	Membership attack against the compound: for a given set of compounds, an attacker having access to the trunk model can compute their representation, and build a classifier which tells if they are in the training set.	2	R1B	Server
T2	Membership attack against the compound: based on the model updates, an attacker builds a classifier which tells if a compound is in the training set.	4	R1B	Server
T3	Combination of T1 with traffic metadata collection to assign identified compounds to certain partners.	1	R1A	Pharma Server

Table 2: Identified Threats for Risk R1 from Table 1. Feasibility/Likelihood: 1 - Very Difficult/Negligible, 2 - Difficult/Low, 3 - Average/Moderate, 4 - Easy/Large, 5 - Immediate/Certain.

4.4 Threats

A Threat is a sequence of actions carried out by a Risk Source to realize one or more Risks. A Threat is sometimes called an “attack” in the security literature. In general, several Threats can lead to a Risk and the same Threat can lead to different Risks. A Threat has a single property called feasibility, which characterizes the difficulty for a Risk Source to realize the Threat. The likelihood of a Risk is computed from the feasibility of its corresponding Threats and is influenced by the feasibility as well as the the motivation or risk source. To ease demonstration, we assume in the sequel that 1) the likelihood of the Risk is computed directly from the feasibility of its Threats using standard probability calculus (see Section 6), and 2) only a few threats corresponding to risk R1 and listed in Table 2 are considered. We are not concerned with Threats that manipulate the learning algorithm in the design phase by model architects or software designers (they are also excluded from the list of potential risk sources).

5 Attacks

In this section, three different attacks are described. One can be performed after (or during) training (Trunk Activation), while two (NGMA and $N - 1$) are applicable only during training. The motivation and background of the attack, the attacker’s capabilities, and the attack itself are described for each attack. A short summary is depicted in Table 3 and an illustration is visualized in Figure 2.

Attack	Leaked Asset	Corr. Threat	Adversarial Capabilities
Trunk Activation Attack	training data	T1	Access to final trained trunk model.
NGMA	training data	T2	Access to aggregated trunk gradients in a single round.
N-1 Attack	assigning training data to a single partner	T3	Access to aggregated trunk gradients in multiple rounds.

Table 3: Summary of presented attacks.

5.1 Attack settings

All evaluations are performed on the publicly available ChEMBL v.23 dataset [12], prepared according to [40]. The input to this model is a folded ECFP fingerprint of a compound with a size of $n = 32000$, and the output is the predicted bioactivity with 3547 classification tasks (out of which 1308 were used for simulating the ten participants). The trunk has 40 units, all with ReLU activation and a 0.2 dropout. We used the SparseChem library⁴. The implementation of the proposed attacks is open-sourced on GitHub⁵.

Cross-silo FL is simulated with a realistic separation of the training data into ten partitions, each with different samples and tasks (i.e., output size). The metadata of the splitting is shown in Table 4.

5.2 Membership Inference

The trunk is shared among all participants, who might be honest but curious. Therefore, it is interesting to understand whether its output (i.e., trunk activation values) or its updates (i.e., gradients) leak any information about the membership of an input sample. Moreover, the commonly trained trunk model may be sold after the federated run to third parties, posing additional risks. Next, we show two membership attacks: a Trunk Activation Attack and a Gradient-based Membership Attack, where the former infers membership from the output of the trunk model after (or during) training. At the same time, the latter does the same from the model update during training. Recall that (as illustrated in Figure 2) the attacker cannot access individual model updates due to secure aggregation; hence the output of the attacks only indicates whether any participant used the tested sample for training or not without attributing this information to any specific participant.

⁴<https://github.com/melloddy/SparseChem>

⁵<https://github.com/minaremeli/pharma-attack-framework>

Samples	256	432	4301	7076	5576	32954	35692	45796	64903	68359
Tasks	2	8	22	46	32	182	150	217	309	340

Table 4: The meta-data of the ChEMBL dataset split utilized in our experiments.

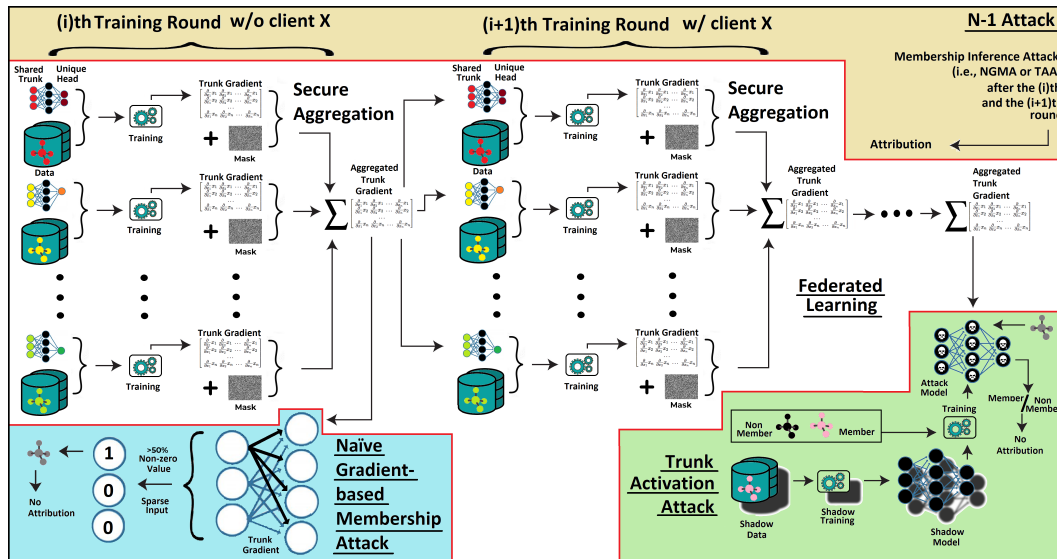


Figure 2: Illustrating the proposed attacks and the process with secure aggregation (the middle white area). The attacks are from Table 3, i.e., Trunk Activation Attack, Naive Gradient-Based Membership Inference Attack, and the $N - 1$ Attack illustrated on the bottom right light green, the bottom left light blue, and the top light khaki area respectively. Each attack is detailed in the section they are introduced.

5.2.1 Trunk Activation Attack

The hypothesis is that the trunk will produce different outputs (activation values) for a training sample compared to a yet unseen sample. An attacker model is built to detect these subtle differences between training and non-training samples. We assume that the attacker knows the target model architecture (white-box attack).

The attack has three phases (also illustrated on the bottom right (light green area) of Figure 2): the shadow model training phase, the attack model training phase, and the evaluation phase, where the attacker model is evaluated on the targeted model. In the first phase, the attacker builds a shadow model that can simulate the targeted model’s activation output. (S)he trains it with his shadow dataset, which comes from the same distribution as the participants’ data. In the second phase, (s)he trains the attack model on the activations of the shadow model and learns to differentiate the activation patterns of members from that of non-members, e.g., between the black compound and the pink one where only the latter is part of the shadow dataset. The attack model is a RandomForest classifier with 100 estimators and no maximum depth. The output of the attacker model is binary, indicating whether any of the participants used the sample to train the trunk. Finally, the last step is to query the RandomForest classifier with the activations of samples; thus, the attacker model can infer the membership of sample (compound) that the attacker has access to (and not only of a particular sample), without attributing that to any specific client.

Evaluation. The attacker aims to build a classifier that distinguishes trunk activations of samples that are part of the training dataset and samples which are not. For this, it needs to obtain trunk activations of both training and non-training sample, which will constitute

the training data of the attacker model. For this purpose, it builds its copy of the original model. The attacker’s dataset is divided into two disjoint parts D and D' . D is used for the federated training of the trunk model, and D' contains samples that the trunk model has never seen in the training process. The trunk’s activation values are used to train the attacker model (i.e., a random forest classifier): samples from D are selected and labeled as “seen”, while samples from D' are labeled as “unseen”. D and D' together are used to train and validate the attacker model (a random 66% of the data is used for training and the remaining for validation). In particular, for every $(x, y) \in D'$ (analogously to D) the trunk activation values for compound x is computed. This results in a vector of 40 activation values, which are fed to the attacker classifier to predict the label $y \in \{seen, unseen\}$.

Results and Feasibility. The mean accuracy of the attack is **61.98%**, with a precision of 60.78% and a recall of 67.59%. Consider the random guess as a baseline, where the attacker has a chance of 50% to guess the correct membership without any prior knowledge. Therefore, because this attack needs labeled data to train the attack model, the feasibility of this attack is estimated to be “*Difficult*” for a Pharma partner and “*Very difficult*” for the Server who cannot access the trunk model (see Section 4.2).

5.2.2 Naïve Gradient-based Membership Attack (NGMA)

A simple membership attack, called Naïve Gradient-based Membership Attack (NGMA) exploits the uniqueness and sparsity of the ECFP fingerprint (as illustrated on the bottom left (light blue area) of Figure 2): on average there are less than 70 values in the 32k long folded fingerprint in the ChEMBL dataset. This uniqueness is also reflected in the updates (gradients) of the trunk (noted with darker and lighter arrows).

In particular, if some coordinate values of the input fingerprint are zero, then all the gradients corresponding to these zero coordinates will be zero as well. When the gradient is computed on a batch with multiple samples, some parts of the aggregated gradients may become zero as the gradients of different samples can potentially cancel each other out. Because this can happen quite often, we require that out of all the weight updates (gradients) connected to a single non-zero input neuron, the majority of them must be non-zero. If all non-zero input neurons of a compound “pass” this test, the compound is likely to be involved in the corresponding training batch, i.e., an unknown participant(s) used it.

In details, let S denote the non-zero positions in the fingerprint of the tested sample, and let $\frac{\partial L}{\partial w_{i,j}}$ denote the gradient belonging to weight $w_{i,j}$ in the trunk (L is the loss function and $w_{i,j}$ are the weights). If $\frac{\partial L}{\partial w_{i,j}} \neq 0$ for all $i \in S$ and more than half of all j in the update, then the fingerprint is deemed to be used in the current training batch.

NGMA is much more feasible and scalable than building a separate attacker model to classify model updates, which is also pursued by prior works [28]. Indeed, it does not require training any attacker model, and therefore the attacker does not need any training data beyond the tested fingerprint. Moreover, NGMA does not require knowledge of the label of the tested samples.

Evaluation. Using NGMA, the attacker (i.e., the server or any participant) aims to infer whether a particular (target) compound is used as a training sample. For one experiment, 50 positive and 50 negative batches are collected and evaluated; a positive/negative batch includes/excludes the target (one batch is composed of the union of all participants’ batches). Model updates on the 50 positive and 50 negative batches were computed, and

the attack was then executed on each of them. This evaluation process was repeated 200 times, each with a random target.

Results and Feasibility. NGMA can determine with an average accuracy of **82.41%** whether a specific compound was in the current batch. The precision and recall are 74.17% and 100%, respectively. NGMA is more accurate than the Trunk Activation attack. Therefore its feasibility is “*Large*” to a Pharma partner who can access the gradients, but almost negligible for the server who cannot access them (neither of their sum) due to the employed secure aggregation protocol. Recall that, unlike Trunk Activation Attack, the attacker does not need to train attacker models for NGMA and hence does not need labeled training data, which increases its feasibility.

5.3 N-1 (differentiation) attack

In cross-silo FL, parties may join or leave during training as illustrated on the top (light khaki area) of Figure 2. Therefore, using a successful membership test, the attacker can attribute a training sample to the leaving/joining party. For example, if the membership test starts to succeed upon the arrival of a new party (e.g., client X), then the party is likely to have the tested sample in its training data. Similarly, when the membership test starts to fail upon a party’s leaving, the party is likely to have the tested sample in its training data. Such an attribution attack is a differentiation attack, as it takes advantage of the change in the coalition. The $N - 1$ attack is a particular case of the differentiation attacks when the difference is a single partner. However, other differentiation attacks are also possible: for instance, in the cross-silo FL scenario, where only a (random) subset of participants participate in each round, attribution can become possible by accumulating the information of a membership attack from each round [29].

The following evaluation aims to measure how successful an “ $N - 1$ ” attack can be on ChEMBL. A simulated scenario is where participants train together for some time; then, one randomly selected participant leaves. The target sample (on which the membership inference attack is performed) is part of the dataset of the leaving party. The attacker performs a membership inference attack (MI) against the target sample in each training round and records the result of the attack; for each round r and each target sample T , $MI(r, T)$ output whether or not the target is present in that training round. These are called *positive* and *negative* rounds, respectively. In general, when the target sample is present in only one of the datasets, the attack signals that the target is present once every epoch (an epoch in the training process consists of multiple rounds). According to the attack, epochs that do (not) contain the target sample are called positive (negative) epochs. The aim is to infer whether the statistics before are significantly different from those after the party leaves using Fisher’s exact test.

Evaluation. We assume that the number of positive epochs is much higher than after leaving before the party leaves. The attacker runs the membership inference in every round and, using the attack outputs, determines whether the epoch was positive or negative. First, the attacker does this when everyone trains together (for 30 epochs). Then, after the party leaves, he observes another 30 epochs. The null hypothesis is that the number of positive epochs should be the same in the first 30 epochs as in the second 30 epochs.

Results and Feasibility. In the case of 30 epochs, 26 were flagged as positive when the particular participant participated in the training, and 0 epochs were flagged as positive after they left. With the Fisher-test, we calculate the probability of this data, given our null hypothesis. The probability of observing 26 positive epochs, then 0 is precisely $3.92 \cdot 10^{-13}$. This is also equal to the significance level of a one-tailed test, which tells us the probability of precisely getting this (or more extreme) data (with the same marginals). Therefore, the attacker can confidently conclude that the target sample was part of the leaving party’s dataset since the number of positive epochs decreased significantly. This means that the attribution of a tested compound to a specific party is not less feasible than the membership attack itself, no matter whether secure aggregation is employed. That is, *secure aggregation does not mitigate the risk of attribution (R1A and R2A) as long as the composition of the parties changes during training.*

5.4 Discussion

To compute the likelihood of the Risk R1 in Table 1, consider the feasibility of the two attacks; NGMA with feasibility “Easy” (i.e., an attack accuracy of 0.82), and TrunkActivation with feasibility “Difficult” (i.e., an attack accuracy of 0.6) for a Pharma partner as a Risk Source. Risk 1B (membership inference without attribution) occurs if at least NGMA or TrunkActivation succeeds, which has a probability of $1 - (1 - 0.82)(1 - 0.6) \approx 0.93^6$. This amounts to a substantial likelihood with maximum severity; hence, this risk must be mitigated. The same holds for Risk 1A since attribution is very feasible with our “ $N - 1$ attack” described in Section 5.3.

6 Mitigation

In this section, several techniques (controls) are described to mitigate the threats in Table 3 and eventually Risk 1 in Table 1. An overview of these techniques and the mitigated attacks (i.e., Threats) are summarized in Table 5. In general, there are legal, organizational, and technical controls. Legal controls include any contractual agreements or legal obligations that the parties must obey. Organizational controls include training employees, establishing working policies, and re-structuring the organization to support effective data protection. Technical controls employ technical measures such as access control, cryptographic techniques, sanitization algorithms, detection algorithms, etc. This document focuses only on technical controls, excluding IT security countermeasures.

⁶We assume they are independent events.

Mitigation	Addressed Threats	Effect on Acc.
Differential Privacy	T1, T2, T3	Severe
Regularization	T1, T2, T3	Mild
Parameter Tuning	T1, T2, T3	Mild
Compression	T1, T2, T3	Mild

Table 5: Mitigation techniques with the corresponding threats and accuracy distortion effect.

6.1 Differential Privacy

Privacy (or confidentiality) could be measured with Differential Privacy (DP) [7], and in particular with parameters ϵ and δ . DP provides a strong guarantee that the membership (i.e., the presence or the absence) of a single data point will not change the final output of the algorithm significantly. DP is usually achieved by perturbation, e.g., the participants add Gaussian noise to their model updates before transferring them to the aggregator. Besides providing privacy protection, this could have a regularization effect and hence could boost generalization [9]; however, excessive noise can also degrade model performance (i.e., slower convergence and lower model accuracy). DP has two flavours in the context of collaborative learning; Record-level [1] and Client-level DP [13, 25]. Client-level DP guarantees that no participant-specific information about the training data of any participant is shared (up to ϵ and δ). At the same time, Record-level DP prevents information disclosure specific to any single record (up to ϵ and δ). Although the Client-level DP guarantee is stronger, it also needs larger perturbation and provides weaker utility. To the best of our knowledge, Client-level DP is the only solution that provides a proven defense against $N - 1$ attacks.

Evaluation. Since DP does provide a theoretical privacy guarantee against membership inference, we do not evaluate the protection efficiency against the presented membership inference attacks. Instead, we focus on their utility loss and evaluate model quality with centralized learning while guaranteeing Record-level DP. We show that model performance is unacceptable even in this simplified setting and, therefore, would be even worse in collaborative learning with Client-level DP. Indeed, Client-level DP requires larger perturbation than Record-level DP, and collaborative learning cannot have better performance than centralized learning on the same training data (i.e., the union of the participant’s training data).

Privacy protection is measured via the privacy parameter ϵ and δ , which can be directly mapped to the variance σ of the Gaussian noise added to the gradients of the model parameters in every SGD iteration of the training. Contrary to the multi-partner setting described in Section 5.1, we simulate a centralized learning scheme by using the preprocessed ChEMBL dataset from [40] without splitting. This contains 295750 samples and 2808 targets, so δ is set to $1/295750$ (reciprocal of the number of original training samples). The gradients are clipped to have an L2-norm less than 2.5, and ϵ is computed using the moments accountant method [1].

Results are shown in Figure 3. It is visible that adding a small noise level (i.e., $\sigma = 0.1$) has




Figure 3: Experimental results using Differential Privacy

some regularization effect (i.e., decreases the difference between the training and validation AUC values). Hence, it does not only improve privacy guarantees but also helps generalization. On the other hand, the corresponding privacy guarantee is weak ($\epsilon = 184365$). To have a meaningful privacy guarantee, one must apply noise with greater magnitude (i.e., $\sigma = 1.0$, which corresponds to $\epsilon = 7.3$). However, in such a case, the utility of the training is completely eradicated.

6.2 Other Privacy-Preserving Methods

The above experiments show that adding DP noise to the gradients without client-level protection leads to significant accuracy loss. Fortunately, some techniques provide “ad-hoc” privacy protection with tolerable accuracy loss at the same time. Although these techniques do not offer rigorous privacy guarantees, they will be shown to mitigate the inference attacks described in Section 5 reasonably well.

Regularization. As a regularization technique, dropout is already valuable for the neural network: it improves generalization and prevents overfitting. The effectiveness of membership inference attacks is partially due to the overfitting nature of ML models. Therefore, any techniques that increase the generalizability of the model help fight against membership attacks.

Concerning NGMA (in Section 5.2.2) we experimented with various dropouts. We observed that while a dropout of 0.2 does not significantly impact the attack accuracy, a value of 0.6 and above does. The left part of Figure 4 shows that, with increasing dropout, there is a near-linear drop in the accuracy of NGMA. In contrast, model accuracy slightly deteriorates even for larger dropout values.

Hyper-parameter tuning. Most prior works on membership attacks, and likewise our attack in Section 5, considered collaborative learning when each participant performs a single local SGD iteration with a single mini-batch to compute its update. Increasing the size of the mini-batch deteriorates the attack accuracy, as shown in the middle of Figure 4. As the gradients of every sample in a batch are averaged in SGD, a single compound has a smaller influence on the model update if the batch is large, thus making their detection in a membership test more difficult. Therefore, increasing the batch size and the number of local SGD iterations can mitigate inference attacks at the cost of minor utility degradation.

Compression. As detailed in Section 3, there could be a high communication cost in our application due to the vast number of trunk parameters. As these parameters are highly redundant, their lossy compression can significantly improve bandwidth efficiency at the cost of small performance degradation. Prior works also suggest that these techniques provide some privacy guarantees [3, 21, 46]. As we show next, compression indeed mitigates the inference attacks described in Section 5. Moreover, some sparsification techniques have a regularization effect on training, improving standard model accuracy.

Compression techniques often rely on sparsifying the gradients (i.e., not transferring gradients with small magnitude), which can potentially mitigate inference attacks. Indeed, gradients modified by only a single training sample (revealing its membership) tend to have a smaller magnitude and are more likely to be eliminated by sparsification.

Three sparsification techniques are considered: *thresholding*, *Top-K* and *random subset*. *Thresholding* requires all participants to zero out all coordinate values of the update vec-

tor whose absolute value falls below a certain threshold and transfer only the remaining non-zero coordinates for aggregation. In contrast, via *Top-K*, participants retain and transfer only the top K largest coordinates in the update vector to the server. *Random subset*, as its name suggests, transfers only a random subset of the gradients for aggregation. In particular, each participant sends the same random subset of all coordinates to the server, and they choose a new random subset in every round. This is implemented by sharing a common seed among all participants (since every participant has the same seed, they select the same coordinates in every round). Our results are presented in Table 6. Recall that an attack accuracy of 50% indicates no information leakage concerning the considered attacks.

We found that a threshold of 0.01 keeps the number of non-zero gradients below 10% in each training round and has excellent regularization properties. Indeed, thresholding increases model accuracy by 3% compared to the baseline, uncompressed training. The right side of Figure 4 shows that, with *Top-K*, the less information we release about the gradients (i.e., the smaller K is), the less effective the attack becomes.

Summary. We showed that simple hyper-parameter tuning, regularization, and compressing model updates could effectively reduce the feasibility of both membership attacks




Figure 4: Various mitigation techniques without formal privacy guarantees. The blue line denotes the attacker model’s accuracy after 1000 rounds of training while the red dashed line is the attacked model’s accuracy.

Technique	Parameter	Accuracy	Trunk Act. Acc.	NGMA Acc.
No Compression		70.84%	61.98%	82.41%
Thresholding	0.001	73.42%	54.02%	50.00%
Top-K	0.2	55.14%	58.62%	50.57%
	0.4	55.84%	57.13%	51.76%
	0.6	55.27%	56.28%	57.98%
	0.8	55.1%	55.33%	66.28%
Random Subset	0.2	54.31%	60.42%	50%
	0.4	53.54%	57.44%	50%
	0.6	55.29%	61.04%	50.02%
	0.8	55.54%	59.52%	59.54%

Table 6: The effect of various sparsification techniques on the main task and the attack accuracies.

without degrading model accuracy. In particular, dropout, increasing the batch size, and applying thresholding on the model updates diminishes the feasibility of both attacks to “Very difficult”, the likelihood of R1B to negligible without a significant drop in model accuracy. Moreover, besides eliminating both attacks, thresholding improves generalization and model accuracy. Differentiation attacks can be mitigated by mitigating membership attacks in the first place, by DP (at the cost of severe utility degradation), and by requiring fixed composition of clients during training. Moreover, as all clients participate in every federated round, the communication is synchronous, and all messages have the same size and format due to secure aggregation; hence T3 in Table 2 is also automatically mitigated since meta-data cannot be used to single-out any client during training.

7 Conclusions

The pharmaceutical industry can better leverage its data through collaboration. Yet, there are non-negligible risks stemming from the unintended information leakage, so such a platform needs to be secure and privacy-preserving. This paper described a privacy risk assessment of collaborative machine learning modeling in drug discovery. Several potential risks, threats, and attacks are discussed depending on the attacker’s capability and goals. The main benefit of systematic risk assessment is prioritizing risks, which facilitates efficient mitigation of the most severe risks first.

Our novel attacks show that *there is information leakage through the model updates*. Although there is limited leakage through the output of the commonly trained model, its feasibility is not always significant due to the strong assumption that the adversary needs access to labeled training data. We also emphasize that our attacks are evaluated on public data, and it may exhibit different performances on actual private pharma data. Moreover, we stress that membership attacks alone are insufficient to attribute a training sample to a specific participant. However, given a successful membership attack, such an attribution becomes feasible if the attacker can also launch a successful $N - 1$ attack. This can happen when not all clients participate in every federated round, which can have many reasons such as client failure, DoS attack, the addition of new partners, etc.

Besides analyzing the risks and threats, we described several mitigation techniques against them. Secure aggregation is an effective countermeasure as long as the composition of clients is fixed. Aside from threat mitigation, the proposed compression techniques also reduce the computational overhead of secure aggregation and communication costs in general. Differential privacy was also studied as it provides a provable defense against any attribution attack such as the $N - 1$ attack. However, our experimental results suggest an intolerable accuracy drop, and thus we turned towards ad-hoc mitigation techniques, including regularization, hyper-parameter tuning, and compression of model updates.

We believe that our methodology can be generalized to other applications of federated multi-task learning. In particular, many models have sparse inputs with fully connected layers or are trained by a limited number of partners. For example, different health institutions can indirectly benefit from each other’s patient databases through federated learning for different disease prognoses and predictions, especially if such diseases are rare. Therefore none of the partners alone has sufficient training data [32]. Similarly, anomaly detection can also be significantly improved through federated learning due to the scarcity of anomalous patterns at a single organization or even in a smart building equipped with several IoT devices [38].

Acknowledgement. This work has received support from the EU/EFPIA Innovative Medicines Initiative 2 Joint Undertaking (MELLODDY grant nr 831472).

References

- [1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 308–318, 2016.
- [2] Gergely Ács and Claude Castelluccia. I have a dream! (differentially private smart metering). In *Information Hiding - 13th International Conference, IH 2011, Prague, Czech Republic, May 18-20, 2011, Revised Selected Papers*, pages 118–132, 2011.
- [3] Gergely Acs, Claude Castelluccia, and Rui Chen. Differentially private histogram publishing through lossy compression. In *2012 IEEE 12th International Conference on Data Mining*, pages 1–10. IEEE, 2012.
- [4] Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. How to backdoor federated learning. *CoRR*, abs/1807.00459, 2018.
- [5] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical Secure Aggregation for Privacy-Preserving Machine Learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS '17*, pages 1175–1191, New York, NY, USA, 2017. ACM. event-place: Dallas, Texas, USA.
- [6] Ronald Cramer, Ivan Bjerre Damgård, et al. *Secure multiparty computation*. Cambridge University Press, 2015.
- [7] Damien Desfontaines and Balázs Pejó. Sok: Differential privacies. *Proceedings on Privacy Enhancing Technologies*, 2020(2):288–313, 2020.
- [8] Arkadiusz Z Dudek, Tomasz Arodz, and Jorge Gálvez. Computational methods in developing quantitative structure-activity relationships (qsar): a review. *Combinatorial chemistry & high throughput screening*, 9(3):213–228, 2006.
- [9] Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Roth. Generalization in adaptive data analysis and holdout reuse. In *Advances in Neural Information Processing Systems 28 (NIPS)*, pages 2350–2358, 2015.
- [10] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model Inversion Attacks That Exploit Confidence Information and Basic Countermeasures. In *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security, CCS '15*, pages 1322–1333, New York, NY, USA, 2015. ACM. event-place: Denver, Colorado, USA.
- [11] Karan Ganju, Qi Wang, Wei Yang, Carl A. Gunter, and Nikita Borisov. Property Inference Attacks on Fully Connected Neural Networks using Permutation Invariant Representations. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security - CCS '18*, pages 619–633, Toronto, Canada, 2018. ACM Press.

- [12] Anna Gaulton, Anne Hersey, Michał Nowotka, A. Patrícia Bento, Jon Chambers, David Mendez, Prudence Mutowo, Francis Atkinson, Louisa J. Bellis, Elena Cibrián-Uhalte, Mark Davies, Nathan Dedman, Anneli Karlsson, María Paula Magariños, John P. Overington, George Papadatos, Ines Smit, and Andrew R. Leach. The ChEMBL database in 2017. *Nucleic Acids Research*, 45(Database issue):D945–D954, January 2017.
- [13] Robin C. Geyer, Tassilo Klein, and Moin Nabi. Differentially private federated learning: A client level perspective. *CoRR*, abs/1712.07557, 2017.
- [14] Micah Goldblum, Dimitris Tsipras, Chulin Xie, Xinyun Chen, Avi Schwarzschild, Dawn Song, Aleksander Madry, Bo Li, and Tom Goldstein. Data security for machine learning: Data poisoning, backdoor attacks, and defenses. *arXiv preprint arXiv:2012.10544*, 2020.
- [15] Briland Hitaj, Giuseppe Ateniese, and Fernando Perez-Cruz. Deep Models Under the GAN: Information Leakage from Collaborative Deep Learning. *arXiv:1702.07464 [cs, stat]*, February 2017. arXiv: 1702.07464.
- [16] Hongsheng Hu, Zoran Salcic, Gillian Dobbie, and Xuyun Zhang. Membership inference attacks on machine learning: A survey. *arXiv preprint arXiv:2103.07853*, 2021.
- [17] Jinyuan Jia, Ahmed Salem, Michael Backes, Yang Zhang, and Neil Zhenqiang Gong. Memguard: Defending against black-box membership inference attacks via adversarial examples. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS 2019, London, UK, November 11-15, 2019*, pages 259–274. ACM, 2019.
- [18] Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G. L. D’Oliveira, Salim El Rouayheb, David Evans, Josh Gardner, Zachary Garrett, Adrià Gascón, Badih Ghazi, Phillip B. Gibbons, Marco Gruteser, Zaïd Harchaoui, Chaoyang He, Lie He, Zhouyuan Huo, Ben Hutchinson, Justin Hsu, Martin Jaggi, Tara Javidi, Gauri Joshi, Mikhail Khodak, Jakub Konečný, Aleksandra Korolova, Farinaz Koushanfar, Sanmi Koyejo, Tancrede Lepoint, Yang Liu, Prateek Mittal, Mehryar Mohri, Richard Nock, Ayfer Özgür, Rasmus Pagh, Mariana Raykova, Hang Qi, Daniel Ramage, Ramesh Raskar, Dawn Song, Weikang Song, Sebastian U. Stich, Ziteng Sun, Ananda Theertha Suresh, Florian Tramèr, Praneeth Vepakomma, Jianyu Wang, Li Xiong, Zheng Xu, Qiang Yang, Felix X. Yu, Han Yu, and Sen Zhao. Advances and open problems in federated learning. *CoRR*, abs/1912.04977, 2019.
- [19] Tuan Le, Robin Winter, Frank Noé, and Djork-Arné Clevert. Neuraldecipher—reverse-engineering extended-connectivity fingerprints (ecfps) to their molecular structures. *Chemical Science*, 11(38):10378–10389, 2020.
- [20] Klas Leino and Matt Fredrikson. Stolen Memories: Leveraging Model Memorization for Calibrated White-Box Membership Inference. *arXiv:1906.11798 [cs]*, June 2019.
- [21] Tian Li, Zaoying Liu, Vyas Sekar, and Virginia Smith. Privacy for free: Communication-efficient learning with differential privacy using sketches. *arXiv preprint arXiv:1911.00972*, 2019.

- [22] Zhaorui Li, Zhicong Huang, Chaochao Chen, and Cheng Hong. Quantification of the Leakage in Federated Learning. *arXiv:1910.05467 [cs]*, October 2019. arXiv: 1910.05467.
- [23] Jierui Lin, Min Du, and Jian Liu. Free-riders in federated learning: Attacks and defenses. *arXiv preprint arXiv:1911.12560*, 2019.
- [24] Bo Liu, Ming Ding, Sina Shaham, Wenny Rahayu, Farhad Farokhi, and Zihuai Lin. When machine learning meets privacy: A survey and outlook. *ACM Computing Surveys (CSUR)*, 54(2):1–36, 2021.
- [25] H. Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning differentially private recurrent language models. In *International Conference on Learning Representations*, 2018.
- [26] Luca Melis, Congzheng Song, Emiliano De Cristofaro, and Vitaly Shmatikov. Exploiting Unintended Feature Leakage in Collaborative Learning. *arXiv:1805.04049 [cs]*, May 2018. arXiv: 1805.04049.
- [27] Sasi Kumar Murakonda, Reza Shokri, and George Theodorakopoulos. Ultimate Power of Inference Attacks: Privacy Risks of High-Dimensional Models. *arXiv:1905.12774 [cs, stat]*, May 2019. arXiv: 1905.12774.
- [28] Milad Nasr, Reza Shokri, and Amir Houmansadr. Comprehensive Privacy Analysis of Deep Learning: Stand-alone and Federated Learning under Passive and Active White-box Inference Attacks. *arXiv:1812.00910 [cs, stat]*, December 2018. arXiv: 1812.00910.
- [29] Balázs Pejó and Gergely Biczók. Quality inference in federated learning with secure aggregation. *arXiv preprint arXiv:2007.06236*, 2022.
- [30] Balázs Pejó, Qiang Tang, and Gergely Biczók. Together or alone: The price of privacy in collaborative learning. *Proceedings on Privacy Enhancing Technologies*, 1:19, 2019.
- [31] Apostolos Pyrgelis, Carmela Troncoso, and Emiliano De Cristofaro. Knock Knock, Who’s There? Membership Inference on Aggregate Location Data. *arXiv:1708.06145 [cs]*, August 2017. arXiv: 1708.06145.
- [32] Nicola Rieke, Jonny Hancox, Wenqi Li, Fausto Milletari, Holger Roth, Shadi Albarqouni, Spyridon Bakas, Mathieu N. Galtier, Bennett A. Landman, Klaus H. Maier-Hein, Sébastien Ourselin, Micah J. Sheller, Ronald M. Summers, Andrew Trask, Daguang Xu, Maximilian Baust, and M. Jorge Cardoso. The future of digital health with federated learning. *CoRR*, abs/2003.08119, 2020.
- [33] Maria Rigaki and Sebastian Garcia. A survey of privacy attacks in machine learning. *arXiv preprint arXiv:2007.07646*, 2020.
- [34] David Rogers and Mathew Hahn. Extended-connectivity fingerprints. *Journal of chemical information and modeling*, 50(5):742–754, 2010.
- [35] Alexandre Sablayrolles, Matthijs Douze, Cordelia Schmid, Yann Ollivier, and Herve Jegou. White-box vs Black-box: Bayes Optimal Strategies for Membership Inference. In *International Conference on Machine Learning*, pages 5558–5567, May 2019.
- [36] Ahmed Salem, Apratim Bhattacharya, Michael Backes, Mario Fritz, and Yang Zhang. Updates-Leak: Data Set Inference and Reconstruction Attacks in Online Learning. *arXiv:1904.01067 [cs, stat]*, April 2019. arXiv: 1904.01067.

- [37] Ahmed Salem, Yang Zhang, Mathias Humbert, Pascal Berrang, Mario Fritz, and Michael Backes. ML-Leaks: Model and Data Independent Membership Inference Attacks and Defenses on Machine Learning Models. *arXiv:1806.01246 [cs]*, June 2018. arXiv: 1806.01246.
- [38] Raed Abdel Sater and A. Ben Hamza. A federated learning approach to anomaly detection in smart buildings. *ACM Trans. Internet Things*, 2(4):28:1–28:23, 2021.
- [39] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership Inference Attacks against Machine Learning Models. *arXiv:1610.05820 [cs, stat]*, October 2016. arXiv: 1610.05820.
- [40] Jaak Simm, Adam Arany, Edward De Brouwer, and Yves Moreau. Graph informer networks for molecules. *arXiv preprint arXiv:1907.11318*, 2019.
- [41] Florian Tramèr, Fan Zhang, Ari Juels, Michael K. Reiter, and Thomas Ristenpart. Stealing machine learning models via prediction apis. In *25th USENIX Security Symposium, USENIX Security 16*, pages 601–618, 2016.
- [42] Stacey Truex, Ling Liu, Mehmet Emre Gursoy, Lei Yu, and Wenqi Wei. Towards Demystifying Membership Inference Attacks. *arXiv:1807.09173 [cs]*, June 2018. arXiv: 1807.09173.
- [43] Lixu Wang, Shichao Xu, Xiao Wang, and Qi Zhu. Eavesdrop the Composition Proportion of Training Labels in Federated Learning. *arXiv:1910.06044 [cs, stat]*, October 2019. arXiv: 1910.06044.
- [44] Zhibo Wang, Mengkai Song, Zhifei Zhang, Yang Song, Qian Wang, and Hairong Qi. Beyond Inferring Class Representatives: User-Level Privacy Leakage From Federated Learning. *arXiv:1812.00535 [cs]*, December 2018. arXiv: 1812.00535.
- [45] Samuel Yeom, Irene Giacomelli, Matt Fredrikson, and Somesh Jha. Privacy Risk in Machine Learning: Analyzing the Connection to Overfitting. *arXiv:1709.01604 [cs, stat]*, September 2017. arXiv: 1709.01604.
- [46] Shuheng Zhou, Katrina Ligett, and Larry Wasserman. Differential privacy with compression. In *2009 IEEE International Symposium on Information Theory*, pages 2718–2722. IEEE, 2009.
- [47] Ligeng Zhu, Zhijian Liu, and Song Han. Deep Leakage from Gradients. *arXiv:1906.08935 [cs, stat]*, June 2019. arXiv: 1906.08935.