
PANORAMIA: Privacy Auditing of Machine Learning Models without Retraining

Mishaal Kazmi*

University of British Columbia

Hadrien Lautraite*

University du Québec à Montréal

Alireza Akbari*

Simon Fraser University

Qiaoyue Tang*

University of British Columbia

Mauricio Soroco

University of British Columbia

Tao Wang

Simon Fraser University

Sébastien Gambs

University du Québec à Montréal

Mathias Lécuyer

University of British Columbia

Abstract

We present PANORAMIA, a privacy leakage measurement framework for machine learning models that relies on membership inference attacks using generated data as non-members. By relying on generated non-member data, PANORAMIA eliminates the common dependency of privacy measurement tools on in-distribution non-member data. As a result, PANORAMIA does not modify the model, training data, or training process, and only requires access to a subset of the training data. We evaluate PANORAMIA on ML models for image and tabular data classification, as well as on large-scale language models.

1 Introduction

Training Machine Learning (ML) models with Differential Privacy (DP) [Dwork et al. \(2006\)](#), such as with DP-SGD [Abadi et al. \(2016\)](#), upper-bounds the worst-case privacy loss incurred by the training data. In contrast, privacy auditing aims to empirically lower-bound the privacy loss of a target ML model or algorithm. In practice, privacy audits usually rely on the link between DP and the performance of membership inference attacks (MIA) [Wasserman & Zhou \(2010\)](#); [Kairouz et al. \(2015\)](#); [Dong et al. \(2019\)](#). At a high level, DP implies an upper-bound on the performance of MIAs, thus creating a high-performance MIA implies a lower-bound on the privacy loss. Auditing schemes have proven valuable in many settings, such as to audit DP implementations [Nasr et al. \(2023\)](#), or to study the tightness of DP algorithms [Nasr et al. \(2021\)](#); [Lu et al. \(2023\)](#); [Steinke et al. \(2023\)](#). Typical privacy audits rely on retraining the model several times, each time guessing the membership of one sample [Jagielski et al. \(2020\)](#); [Carlini et al. \(2022a\)](#); [Zanella-Béguin et al. \(2022\)](#), which is computationally prohibitive, requires access to the target model (entire) training data as well as control over the training pipeline.

To circumvent these concerns, [Steinke et al. \(2023\)](#) proposed an auditing recipe (called O(1)) requiring only one training run (which could be the same as the actual training) by randomly including/excluding several samples (called auditing examples) into the training dataset of the target model. Later, the membership of the auditing examples are guessed for privacy audit. However, O(1) faces a few challenges in certain setups. First, canaries, which are datapoints specially crafted to be easy to detect when added to the training set [Nasr et al. \(2023\)](#); [Lu et al. \(2023\)](#); [Steinke et al. \(2023\)](#), cannot be employed as auditing examples when measuring the privacy leakage for data that a contributor actually puts into the model, and not a *worst case data point*.

*equal contribution

This matches a setting in which individual data contributors (*e.g.*, a hospital in a cross-site Federated Learning (FL) setting or a user of a service that trains ML models on users’ data) measure the leakage of their own (*i.e.*, known) partial training data in the final trained model. Second, O(1) also relies on the withdrawal of real data from the model to construct non-member in-distribution data. This is problematic in situations in which ML model owners need to conduct post-hoc audits, in which case it is too late for removal [Negoescu et al. \(2023\)](#). Moreover, in-distribution audits require much more data, thus withholding many data points (typically more than the test set size) and reducing model utility. This brings us to the question: *Given an instance of a machine learning model as a target, can we perform post-hoc estimation of the privacy loss with regards to a known member subset of the target model training dataset?*

Our contributions. We propose PANORAMIA, a new scheme for *Privacy Auditing with NO Retraining by using Artificial data for Membership Inference Attacks*. More precisely, we consider an auditor with access to a subset of the training data and introduce a new alternative for accessing non-members: using synthetic datapoints from a generative model trained on the member data, unlocking the limit on non-member data. PANORAMIA uses this generated data, together with known members, to train and evaluate a MIA attack on the target model to audit (§3). We also adapt the theory of privacy audits, and show how PANORAMIA can estimate the privacy loss (though not a lower-bound) of the target model with regards to the known member subset (§4). An important benefit of PANORAMIA is to perform privacy loss measurements with (1) no retraining the target ML model (*i.e.*, we audit the end-model, not the training algorithm), (2) no alteration of the model, dataset, or training procedure, and (3) only partial knowledge of the training set. We evaluate PANORAMIA on CIFAR10 models and observe that overfitted models, larger models, and models with larger DP parameters have higher measured privacy leakage. We also demonstrate the applicability of our approach on the GPT-2 based model (*i.e.*, WikiText dataset) and CelebA models.

2 Background and Related Work

DP is the established privacy definition in the context of ML models, as well as for data analysis in general. We focus on the pure DP definition to quantify privacy loss with well-understood semantics. In a nutshell, DP is a property of a randomized mechanism (or computation) from datasets to an output space \mathcal{O} , noted $M : \mathcal{D} \rightarrow \mathcal{O}$. It is defined over neighboring datasets D, D' , differing by one element $x \in \mathcal{X}$ (we use the add/remove neighboring definition), which is $D' = D \cup \{x\}$. Formally:

Definition 1 (Differential Privacy [Dwork et al. \(2006\)](#)). A mechanism $M : \mathcal{D} \rightarrow \mathcal{O}$ is ϵ -DP if for any two neighbouring datasets $D, D' \in \mathcal{D}$, and for any measurable output subset $O \subseteq \mathcal{O}$ it holds that:

$$P[M(D) \in O] \leq e^\epsilon P[M(D') \in O].$$

Since the neighbouring definition is symmetric, so is the DP definition, and we also have that $P[M(D') \in O] \leq e^\epsilon P[M(D) \in O]$. Intuitively, ϵ upper-bounds the worst-case contribution of any individual example to the distribution over outputs of the computation (*i.e.*, the ML model learned). More formally, ϵ is an upper-bound on the *privacy loss* incurred by observing an output o , defined as $|\ln \left(\frac{P[M(D)=o]}{P[M(D')=o]} \right)|$, which quantifies how much an adversary can learn to distinguish D and D' based on observing output o from M . A smaller ϵ hence means higher privacy.

DP, MIA and privacy audits. To audit a DP training algorithm M that outputs a model f , one can perform a MIA on datapoint x , trying to distinguish between a neighboring training sets D and $D' = D \cup \{x\}$. The MIA can be formalized as a hypothesis test to distinguish between $\mathcal{H}_0 = D$ and $\mathcal{H}_1 = D'$ using the output of the computation f . [Wasserman & Zhou \(2010\)](#); [Kairouz et al. \(2015\)](#); [Dong et al. \(2019\)](#) show that any such test at significance level α (False Positive Rate or FPR) has power (True Positive Rate or TPR) bounded by $e^\epsilon \alpha$. In practice, one repeats the process of training model f with and without x in the training set, and uses a MIA to guess whether x was included. If the MIA has $\text{TPR} > e^\epsilon \text{FPR}$, the training procedure that outputs f is not ϵ -DP. This is the building block of most privacy audits [Jagielski et al. \(2020\)](#); [Nasr et al. \(2021\)](#); [Zanella-Béguin et al. \(2022\)](#); [Lu et al. \(2023\)](#); [Nasr et al. \(2023\)](#). Appendix F provides further context, and discusses other related work that is not directly relevant to our contribution.

Averaging over data instead of models with O(1). The above result bounds the success rate of MIAs when performed over several *retrained models*, on two alternative datasets D and D' . [Steinke et al. \(2023\)](#) show that it is possible to average *over data* when several data points independently

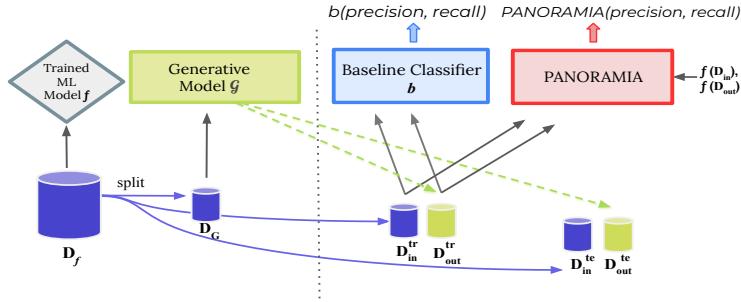


Figure 1: PANORAMIA’s two-phase privacy audit. Phase 1 trains generative model \mathcal{G} on member data. Phase 2 trains a MIA on a subset of member data and generated non-member data, using the loss of f on these data points. The performance of the MIA is compared to a baseline classifier that does not have access to f . Notations are summarized in Table 4 in Appendix A.

differ between D and D' . Let $x_{1,m}$ be the m data points independently included in the training set, and $s_{1,m} \in \{0, 1\}^m$ be the vector encoding inclusion. $T_{0,m} \in \mathbb{R}^m$ represents any vector of guesses, with positive values for inclusion in the training set (member), negative values for non-member, and zero for abstaining. Then, if the training procedure is ϵ -DP, Proposition 5.1 in Steinke et al. (2023) bounds the performance of guesses from T with:

$$\mathbb{P} \left[\sum_{i=1}^m \max\{0, T_i \cdot S_i\} \geq v \mid T = t \right] \leq \mathbb{P}_{S' \sim \text{Bernoulli}(\frac{e^\epsilon}{1+e^\epsilon})^m} \left[\sum_{i=1}^m |t_i| \cdot S'_i \geq v \right].$$

In other words, an audit (MIA) T that can guess membership better than a Bernoulli random variable with probability $\frac{e^\epsilon}{1+e^\epsilon}$ refutes an ϵ -DP claim. In this work we build on this result, extending the algorithm (§3) and theoretical analysis (§4) to enable the use of generated data for non-members, which breaks the independence between data points.

3 PANORAMIA

Figure 1 summarizes the end-to-end PANORAMIA privacy measurement. The measurement starts with a target model f , and a subset of its training data D_f from distribution \mathcal{D} . For instance, \mathcal{D} and D_f could be the distribution and dataset of one participant in an FL training procedure that outputs a final model f . The privacy measurement then proceeds in two phases.

Phase 1: In the first phase, PANORAMIA uses a subset of the known training data $D_G \subset D_f$ to train a generative model \mathcal{G} . The goal of the generative model \mathcal{G} is to match the training data distribution \mathcal{D} as closely as possible, which is formalized in Definition 3 (§4). Using the generative model \mathcal{G} , we can synthesize non-member data, which corresponds to data that was not used in the training of target model f . Hence, we now have access to an independent dataset of member data $D_{\text{in}} = D_f \setminus D_G$, and a synthesized dataset of non-member data $D_{\text{out}} \sim \mathcal{G}$, of size $m = |D_{\text{in}}|$.

Phase 2: In the second phase, we leverage D_{in} and D_{out} to audit the privacy leakage of f using a MIA. To this end, we split $D_{\text{in}}, D_{\text{out}}$ into training and testing sets, respectively called $D_{\text{in}}^{\text{tr}}, D_{\text{out}}^{\text{tr}}$ and $D_{\text{in}}^{\text{te}}, D_{\text{out}}^{\text{te}}$. We use the training set to train a MIA (called PANORAMIA in Figure 1), a binary classifier that predicts whether a given datapoint is a member of D_f , the training set of the target model f . This MIA classifier makes its prediction based on both a training example x , as well as information from applying the target model f to the input, such as the loss of the target model when applied to this example $\text{loss}(f(x))$ (see §5, Appendix C for details). We use the test set to measure the MIA performance, using the precision at different recall values.

Previous results linking the performance of a MIA on several data-points to ϵ -DP bounds rely on independence between members and non-members. This intuitively means that there is no information about membership in x itself. When the auditor controls the training process this independence is enforced by construction, by adding data points to the training set based on an independent coin flip. In PANORAMIA, we do not have independence between membership and x , as all non-members

Algorithm 1 PANORAMIA

Input: Target ML model f , audit set size m , confidence $1 - \beta$

Phase 1:

- 1: Split D_f in $D_G, D_{\text{in}}^{\text{tr}}, D_{\text{in}}^{\text{te}}$, with $|D_{\text{in}}^{\text{te}}| = m$
- 2: Train generator \mathcal{G} on D_G
- 3: Generate $D_{\text{out}}^{\text{tr}}, D_{\text{out}}^{\text{te}}$ of size $|D_{\text{in}}^{\text{tr}}|, |D_{\text{in}}^{\text{te}}|$

Phase 2:

Train the baseline and MIA:

- 1: Label $D_{\text{in}}^{\text{tr}}$ as members, and $D_{\text{out}}^{\text{tr}}$ as non-members
- 2: Train b to predict labels using $x \in D_{\text{in}}^{\text{tr}} \cup D_{\text{out}}^{\text{tr}}$
- 3: Train MIA to predict labels using $x \in D_{\text{in}}^{\text{tr}} \cup D_{\text{out}}^{\text{tr}}$ and $f(x)$

Measure privacy leakage (see §4):

- 1: Sample $s \sim \text{Bernoulli}(\frac{1}{2})^m$ ▷ Def.2
- 2: Create audit set $X = s \cdot D_{\text{in}}^{\text{te}} + (1 - s)D_{\text{out}}^{\text{te}}$
- 3: Score each audit point for membership, creating $t^b \triangleq b(X) \in \mathbb{R}_+^m$ and $t^a \triangleq \text{MIA}(X) \in \mathbb{R}_+^m$
- 4: Set $v_{\text{ub}}^b(c, t) \triangleq \sup \{v : \beta^b(m, c, v, t) \leq \frac{\beta}{2}\}$ ▷ Prop.1
- 5: $c_{\text{lb}} = \max_{t, c} \mathbb{1}\{t^b \geq t\} \cdot s \leq v_{\text{ub}}^b(c, \mathbb{1}\{t^b \geq t\})$
- 6: Set $v_{\text{ub}}^a(c, \epsilon, t) \triangleq \sup \{v : \beta^a(m, c, \epsilon, v, t) \leq \frac{\beta}{2}\}$ ▷ Prop.2
- 7: $\{c + \epsilon\}_{\text{lb}} = \max_{t, c, \epsilon} \mathbb{1}\{t^a \geq t\} \cdot s \leq v_{\text{ub}}^a(c, \epsilon, \mathbb{1}\{t^a \geq t\})$

Return $\tilde{\epsilon} \triangleq \{c + \epsilon\}_{\text{lb}} - c_{\text{lb}}$

come from the generator $\mathcal{G} \neq \mathcal{D}$. As a result, there are two ways to guess membership and have high MIA precision: either by using f to detect membership (*i.e.*, symptomatic of privacy leakage) or by detecting generated data (*i.e.*, not a symptom of privacy leakage). To measure the privacy leakage, we compare the results of the MIA to that of a baseline classifier b that guesses membership based exclusively on x , without access to f . The stronger this baseline, the better the removal of the effect of synthesized data detection. Algorithm 1 summarizes the entire procedure. In the next section, we demonstrate how to relate the difference between the baseline b and the MIA performance to the privacy loss ϵ .

4 Quantifying Privacy Leakage with PANORAMIA

We formalize a privacy game as follows: construct a sequence of auditing samples, in which s_i is sampled independently to choose either the real ($x^{\text{in}} \in \mathcal{X}^m$ training points from \mathcal{D} if $s_i = 1$) or generated ($x^{\text{gen}} \in \mathcal{X}^m$ generated points from \mathcal{G} if $s_i = 0$) data point at each index i .

Definition 2 (Privacy game).

$$s \sim \text{Bernoulli}(\frac{1}{2})^m, \text{ with } s_i \in \{0, 1\}, x_i = (1 - s_i)x_i^{\text{gen}} + s_i x_i^{\text{in}}, \forall i \in \{1, \dots, m\}.$$

The key difference between Definition 2 and the privacy game of Steinke et al. (2023) lies in how we create the audit set. Instead of sampling non-members independently and removing them from the training set (which requires modifying the training set and the model), Definition 2 starts from a set of known members (after the fact), and pairs each point with a non-member (generated i.i.d. from the generator distribution). For each pair, we then flip a coin to decide which point in the pair will be shown to the auditor for testing, thereby creating the test task of our privacy measurement.

The level of achievable success in this game quantifies the privacy leakage of the target model f . More precisely, we follow an analysis inspired by that of Steinke et al. (2023), but require several key technical changes to support auditing with no retraining using generated non-member data.

4.1 Formalizing the Privacy Measurements as a Hypothesis Test

We first need a notion of quality for our generator:

Definition 3 (c -closeness). *For all $c > 0$, we say that a generative model \mathcal{G} is c -close for data distribution \mathcal{D} if:*

$$\forall x \in \mathcal{X}, e^{-c} \mathbb{P}_{\mathcal{D}}[x] \leq \mathbb{P}_{\mathcal{G}}[x].$$

The smaller c , the better the generator, as \mathcal{G} assigns a probability to real data that cannot be much smaller than that of the real data distribution. We make three important remarks.

Remark 1: c -closeness is a strong requirement but it is achievable, at least for a large c . For instance, a uniform random generator is c -close to any distribution \mathcal{D} , with $c < \infty$. Of course c would be very large, and the challenge in PANORAMIA is to create a generator that empirically has a small enough c .

Remark 2: We show in Appendix E that we can relax this definition to hold only with high probability (following the (ϵ, δ) -DP relaxation). Empirical results under this relaxation are very similar. We also show how to measure (ϵ, δ) -DP, instead of ϵ -DP on which we focus here.

Remark 3: Our definition of closeness is very similar to that of DP. This is not a coincidence as we will use this definition to be able to reject claims of both c -closeness for the generator and ϵ -DP for the algorithm that gave the target model. We further note that, contrary to the DP definition, c -closeness is one-sided as we only bound $\mathbb{P}_{\mathcal{G}}$ from below. Intuitively, this means that the generator has to produce high-quality samples (*i.e.*, samples likely under \mathcal{D}) with high enough probability but it does not require that all samples are good though. This one-sided measure of closeness is enabled by our focus on detecting members (*i.e.*, true positives) as opposed to members and non-members, and puts less stringent requirements on the generator.

Using Definition 3, we can formulate the hypothesis test on the privacy game of Definition 2 that underpins our approach:

$$\mathcal{H} : \text{generator } \mathcal{G} \text{ is } c\text{-close, and target model } f \text{ is } \epsilon\text{-DP.}$$

Note that, in a small abuse of language, we will often refer to f as ϵ -DP to say that f is the output of an ϵ -DP training mechanism. That is, $B(s, x) = \{b(x_1), b(x_2), \dots, b(x_m)\}$.

We define two key mechanisms: $B(s, x) : \{0, 1\}^m \times \mathcal{X}^m \rightarrow \mathbb{R}_+^m$ outputs a (potentially randomized) non-negative score for the membership of each datapoint x_i , based on x_i only; $A(s, x, f) : \{0, 1\}^m \times \mathcal{X}^m \times \mathcal{F} \rightarrow \mathbb{R}_+^m$ outputs a (potentially randomized) non-negative score for the membership of each datapoint, with the guess for index i depending on $x_{\leq i}$ and target model f . We can construct a statistical test for each part of the hypothesis separately.

Proposition 1. *Let \mathcal{G} be c -close, S and X be the random variables for s and x from Definition 2, and $T^b \triangleq B(S, X)$ be the vector of guesses from the baseline. Then, for all $v \in \mathbb{R}$ and all t in the support of T :*

$$\mathbb{P}_{S, X, T^b} \left[\sum_{i=1}^m T_i^b \cdot S_i \geq v \mid T^b = t^b \right] \leq \mathbb{P}_{S' \sim \text{Bernoulli}(\frac{e^c}{1+e^c})^m} \left[\sum_{i=1}^m t_i^b \cdot S'_i \geq v \right] \triangleq \beta^b(m, c, v, t^b)$$

Proof. In Appendix B.1. □

Now that we have a test to reject a claim that the generator \mathcal{G} is c -close for the data distribution \mathcal{D} , we turn our attention to the second part of \mathcal{H} which claims that the target model f is the result of an ϵ -DP mechanism.

Proposition 2. *Let \mathcal{G} be c -close, S and X be the random variables for s and x from Definition 2, f be ϵ -DP, and $T^a \triangleq A(S, X, f)$ be the guesses from the membership audit. Then, for all $v \in \mathbb{R}$ and all t in the support of T :*

$$\mathbb{P}_{S, X, T^a} \left[\sum_{i=1}^m T_i^a \cdot S_i \geq v \mid T^a = t^a \right] \leq \mathbb{P}_{S' \sim \text{Bernoulli}(\frac{e^c + \epsilon}{1+e^c + \epsilon})^m} \left[\sum_{i=1}^m t_i^a \cdot S'_i \geq v \right] \triangleq \beta^a(m, c, \epsilon, v, t^a)$$

Proof. In Appendix B.2. □

We can now provide a test for hypothesis \mathcal{H} , by applying a union bound over Propositions 1 and 2:

Corollary 1. *Let \mathcal{H} be true, $T^b \triangleq B(S, X)$, and $T^a \triangleq A(S, X, f)$. Then:*

$$\mathbb{P} \left[\sum_{i=1}^m T_i^a \cdot S_i \geq v^a, \sum_{i=1}^m T_i^b \cdot S_i \geq v^b \mid T^a = t^a, T^b = t^b \right] \leq \beta^a(m, c, \epsilon, v^a, t^a) + \beta^b(m, c, v^b, t^b)$$

To make things more concrete, let us instantiate Corollary 1 as we do in PANORAMIA. Our baseline (B above) and MIA (A above) classifiers return a membership guesses $T^{a,b} \in \{0, 1\}^m$, with 1 corresponding to membership. Let us call $r^{a,b} \triangleq \sum_i t_i^{a,b}$ the total number of predictions, and $\text{tp}^{a,b} \triangleq \sum_i t_i^{a,b} \cdot s_i$ the number of correct membership guesses (true positives). We also call the precision $\text{prec}^{a,b} \triangleq \frac{\text{tp}^{a,b}}{r^{a,b}}$. Using the following tail bound on the sum of Bernoulli random variables for simplicity and clarity (we use a tighter bound in practice, but this one is easier to read),

$$\mathbb{P}_{S' \sim \text{Bernoulli}(p)^r} \left[\sum_{i=1}^r \frac{S'_i}{r} \geq p + \sqrt{\frac{\log(1/\beta)}{2r}} \right] \leq \beta, \quad (1)$$

we can reject \mathcal{H} at confidence level β by setting $\beta^a = \beta^b = \frac{\beta}{2}$ and if either $\text{prec}^b \geq \frac{e^c}{1+e^c} + \sqrt{\frac{\log(2/\beta)}{2r^b}}$ or $\text{prec}^a \geq \frac{e^{c+\epsilon}}{1+e^{c+\epsilon}} + \sqrt{\frac{\log(2/\beta)}{2r^a}}$.

4.2 Quantifying Privacy Leakage and Interpretation

In an privacy measurement, we would like to quantify ϵ , not just reject a given ϵ claim. We use the hypothesis test from Corollary 1 to compute a confidence interval on c and ϵ . To do this, we first define an ordering between (c, ϵ) pairs, such that if $(c_1, \epsilon_1) \leq (c_2, \epsilon_2)$, the event (*i.e.*, set of observations for $T^{a,b}, S$) for which we can reject $\mathcal{H}(c_2, \epsilon_2)$ is included in the event for which we can reject $\mathcal{H}(c_1, \epsilon_1)$. That is, if we reject \mathcal{H} for values (c_2, ϵ_2) based on audit observations, we also reject values (c_1, ϵ_1) based on the same observations.

We define the following order to fit this assumption, based on the hypothesis test from Corollary 1:

$$(c_1, \epsilon_1) \leq (c_2, \epsilon_2) \text{ if either } \begin{cases} c_1 < c_2, \text{ or} \\ c_1 = c_2 \text{ and } \epsilon_1 \leq \epsilon_2 \end{cases} \quad (2)$$

This lexicographic order over hypotheses ensures that, when using the joint hypothesis test from Corollary 1 to construct confidence intervals, we always reject hypotheses with a low value of c incompatible with the baseline performance. Formally, this yields the following confidence intervals:

Corollary 2. *For all $\beta \in [0, 1]$, m , and observed t^a, t^b , call $v_{ub}^a(c, \epsilon) \triangleq \sup \{v : \beta^a(m, c, \epsilon, v, t^a) \leq \frac{\beta}{2}\}$ and $v_{ub}^b(c) \triangleq \sup \{v : \beta^b(m, c, v, t^b) \leq \frac{\beta}{2}\}$. Then:*

$$\mathbb{P} \left[(c, \epsilon) \geq \sup \{(c', \epsilon') : t^a \cdot s \leq v_{ub}^a(c', \epsilon') \text{ and } t^b \cdot s \leq v_{ub}^b(c')\} \right] \geq 1 - \beta$$

Proof. Apply Lemma 4.7 from [Steinke et al. \(2023\)](#) with the ordering from Equation 2 and the test from Corollary 1. \square

The lower confidence interval for (c, ϵ) at confidence $1 - \beta$ is the largest (c, ϵ) pair that cannot be rejected using Corollary 1 with false rejection probability at most β . Hence, for a given confidence level $1 - \beta$, PANORAMIA computes $(c_{lb}, \tilde{\epsilon})$, the largest value for (c, ϵ) that it cannot reject. Note that Corollaries 1 and 2 rely on a union bound between two tests, one for c and one for $c + \epsilon$. We can thus consider each test separately (at level $\beta/2$). In practice, we use the whole precision/recall curve to achieve tighter bounds: each level of recall (*i.e.*, threshold on $t^{a,b}$ to predict membership) corresponds to a bound on the precision, which we can compare to the empirical value. Separately for each test, we pick the level of recall yielding the highest lower-bound (Lines 4-7 in the last section of Algorithm 1). For all bounds to hold together, we apply a union bound over all tests at each level of recall. Specially, if we have m test datapoints, we have at most m different values of recall, and the union bound corresponds to $\beta \leftarrow \beta/m$ (notice in Equation (1) that the bounds depend on $\sqrt{\log(1/\beta)}$, so this is not too costly). We next discuss the semantics of returned measurement values $(c_{lb}, \tilde{\epsilon})$.

Measurement semantics. Corollary 2 gives us a lower-bound for (c, ϵ) , based on the ordering from Equation (2). To understand the value $\tilde{\epsilon}$ returned by PANORAMIA, we need to understand what the hypothesis test rejects. Rejecting \mathcal{H} means either rejecting the claim about c , or the claim about $c + \epsilon$ (which is the reason for the ordering in Equation 2). With Corollary 2, we hence get both a lower-bound c_{lb} on c , and $\{c + \epsilon\}_{lb}$ on $c + \epsilon$. Unfortunately, $\tilde{\epsilon} \triangleq \max\{0, \{c + \epsilon\}_{lb} - c_{lb}\}$, which is the value PANORAMIA returns, does not imply a lower-bound on ϵ . Instead, we can claim that

ML Model	Dataset	Training Epoch	Test Accuracy	Model Variants Names*
ResNet101	CIFAR10	20, 50, 100	91.61%, 90.18%, 87.93%	ResNet101_E20, ResNet101_E50, ResNet101_E100,
WideResNet-28-10	CIFAR10	150, 300	95.67%, 94.23%	WRN-28-10_E150, WRN-28-10_E300
ViT-small (pretrained)	CIFAR10	35	96.38%	ViT_E35
Multi-Label CNN	CelebA	50, 100	81.77%, 78.12%	CNN_E50, CNN_E100
GPT-2	WikiText	37, 75, 150	Refer to Appendix C.2	GPT-2_E37, GPT-2_E75, GPT-2_E150

Table 1: Train and Test Metrics for ML Models Audited. *"Model Variants" trained for different numbers of epochs E . For ViT-small, we use a model pre-trained on imagenet and tune it on CIFAR10.

"PANORAMIA could not reject a claim of c -closeness for \mathcal{G} , and if this claim is tight, then f cannot be the output of a mechanism that is $\tilde{\epsilon}$ -DP".

While this is not as strong a claim as typical lower-bounds on ϵ -DP from prior privacy auditing works, we believe that this measure is useful in practice. Indeed, the $\tilde{\epsilon}$ measured by PANORAMIA is a quantitative privacy measurement, that is accurate (close to a lower-bound on ϵ -DP) when the baseline performs well (and hence c_{lb} is tight). Thus, when the baseline is good, $\tilde{\epsilon}$ can be interpreted as (close to) a lower bound on (pure) DP. PANORAMIA opens a new capability, measuring privacy leakage of a trained model f without access or control of the training pipeline, with a meaningful and practically useful measurement, as we empirically show next.

5 Evaluation

We instantiate PANORAMIA on target models for four tasks from three data modalities¹. For **image classification**, we consider the CIFAR10 Krizhevsky (2009), and CelebA Liu et al. (2015) datasets, with varied target models: a four-layers CNN O’Shea & Nash (2015), a ResNet101 He et al. (2015), a Wide ResNet-28-10 Zagoruyko & Komodakis (2017b), a Vision Transformer (ViT)-small Dosovitskiy et al. (2021), and a DP ResNet18 He et al. (2015) trained with DP-SGD (Abadi et al., 2016) using Opacus (Yousefpour et al., 2021) at different values of ϵ . We use StyleGAN2 Karras et al. (2020) for \mathcal{G} . For **language models**, we fine-tune small GPT-2 Radford et al. (2019) on the WikiText train dataset Meryt et al. (2016) (we take a subset of the documents in WikiText-103). \mathcal{G} is again based on small GPT-2, and then fine-tuned on D_G . We generate samples using top- k sampling Holtzman et al. (2019) and a held-out prompt dataset $D_G^{\text{prompt}} \subset D_G$. Finally, we conduct experiments on **classification on tabular data**. However for this data modality, PANORAMIA did not detect any meaningful leakage. More precisely, we have observed significant variance in the audit values returned by PANORAMIA which make the results inconclusive. Nonetheless, we present the results obtained for this task and discuss this issue further in Appendix D.6. Table 1 summarizes the tasks, models, and performance. More details on the settings, the models used, and implementation details are provided in Appendix C.

5.1 Baseline Design and Evaluation

Our MIA is a loss-based attack, which uses an ML model taking as input both a datapoint x and the value of the loss of target model f on point x . Appendix C details the architectures used for the attack model for each data modality. Recall from §4.2 the importance of having a tight c_{lb} for our measure $\tilde{\epsilon}$ to be close to a lower-bound on ϵ -DP. To increase the performance of our baseline b , we mimic the role of the target model f ’s loss in the MIA using a helper model h trained on synthetic data, which adds a loss-based feature to b .

This new feature can be viewed as side information about the data distribution. Table 2 shows the c_{lb} value for different designs for h . The best performance is consistently when h is trained on synthetic data before being used as a feature to train the

Baseline model	c_{lb}
CIFAR-10 Baseline $D_h^a = \text{gen}$	2.44 ± 0.19
CIFAR-10 Baseline $D_h^a = \text{real}$	2.21 ± 0.17
CIFAR-10 Baseline _{no helper}	1.25 ± 0.24
CelebA Baseline $D_h^a = \text{gen}$	2.05 ± 0.21
CelebA Baseline $D_h^a = \text{real}$	1.97 ± 0.22
CelebA Baseline _{no helper}	0.947 ± 0.25
WikiText Baseline $D_h^a = \text{gen}$	3.31 ± 0.15
WikiText Baseline $D_h^a = \text{real}$	3.26 ± 0.14
WikiText Baseline _{no helper}	3.11 ± 0.15

Table 2: Baseline evaluation with different helper model scenarios.

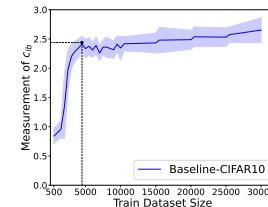


Figure 2: CIFAR-10 baseline on increasing training size.

¹Code available here: <https://github.com/ubc-systopia/panoramia-privacy-measurement>.

b. Indeed, such a design reaches a c_{lb} up to 1.19 larger than without any helper (CIFAR10) and 0.23 higher than when training on real non-member data, *without requiring access to real non-member data*, a key requirement in PANORAMIA. We adopt this design in all the following experiments. In Figure 2 we show that the baseline has enough training data (vertical dashed line) to reach its best performance. Appendix D.1, shows similar results on GPT-2 as well as different architectures that we tried for the helper model. All these pieces of evidence confirm the strength of our baseline.

5.2 Main Privacy Measurement Results

We run PANORAMIA on models with different degrees of over-fitting by varying the number of epochs, (see the final accuracy on Table 1) for each data modality. More over-fitted models are known to leak more information about their training data due to memorization (Yeom et al., 2018; Carlini et al., 2019, 2022a). To show the privacy loss measurement power of PANORAMIA, we compare it with two strong approaches to lower-bounding privacy loss. First, we compare PANORAMIA with the O(1)

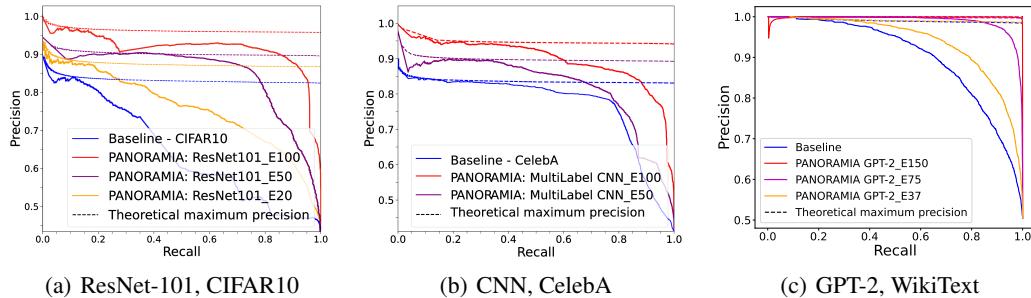


Figure 3: Precision vs. recall comparison between PANORAMIA and the baseline b for our target models, based on one experiment run. The maximum c_{lb} or $\{c + \epsilon\}_{lb}$ values set an upper bound on the empirical precision values across different recall levels, indicated by the dashed line.

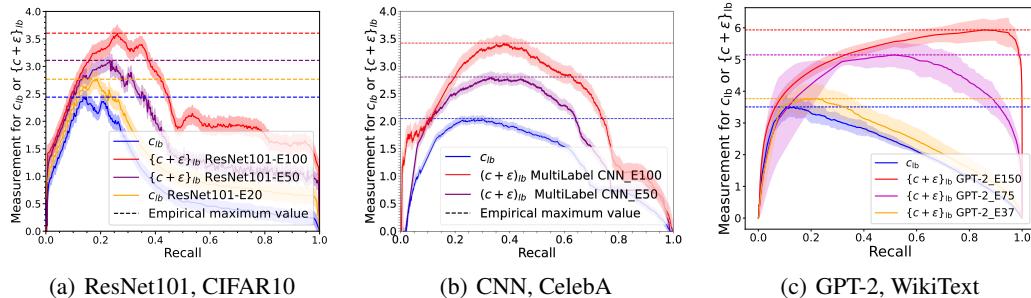


Figure 4: $\{c + \epsilon\}_{lb}$ (or c_{lb}) vs recall, for our target models, reported over 5 independent experiment runs.

audit (Steinke et al., 2023) (in input space, without canaries in a black box setting), which uses a loss threshold MIA to detect both members and non-members. Second, we use a hybrid between O(1) and PANORAMIA, in which we use real non-member data instead of generated data with our ML-based MIA (called RM;RN for Real Members; Real Non-members). In both cases $c_{lb} = 0$, but it requires control over training data inclusion and a large set of non-member data. The privacy loss measured by these techniques gives a target that we hope PANORAMIA to detect.

Figure 3 shows the precision of b and PANORAMIA at different levels of recall, and Figure 4 the corresponding value of $\{c + \epsilon\}_{lb}$ (or c_{lb} for b). Dashed lines show the maximum value of $\{c + \epsilon\}_{lb}$ (or c_{lb}) achieved (Fig. 4) (returned by PANORAMIA), and the precision implying these values at different recalls (Fig. 3). Table 3 summarizes those $\{c + \epsilon\}_{lb}$ (or c_{lb}) values, as well as the ϵ measured by existing approaches. We make two key observations.

First, the best prior method (whether RM;RN or O(1)) measures a larger privacy loss ($\tilde{\epsilon} \leq \epsilon$). Overall, these results empirically confirm the strength of b , as we do not seem to spuriously assign differences between \mathcal{G} and \mathcal{D} to our privacy loss proxy $\tilde{\epsilon}$. We also note that O(1) tends to perform better, due to its ability to rely on non-member detection, which improves the power of the statistical test at equal data sizes. Such tests are not available in PANORAMIA given our one-sided closeness definition for \mathcal{G} (see §4), and we keep that same one-sided design for RM;RN for the sake of comparison.

Second, the values of $\tilde{\epsilon}$ measured by PANORAMIA are close to those of the methods against which we compared. In particular, despite a more restrictive adversary model (*i.e.*, no non-member data, no control over the training process, and no shadow model training), PANORAMIA is able to detect meaningful amounts of privacy loss, comparable to that of state-of-the-art methods.

For instance, on a non-overfitted CIFAR-10 ResNet101 model (E20), PANORAMIA detects a privacy loss of 0.33, while using real non-member (RM;RN) data yields 0.42, and controlling the training process $O(1)$ gets 0.45. The relative gap gets even closer on models that reveal more about their training data. Indeed, for the most overfitted model (E100), $\tilde{\epsilon} = 1.16$ is very close to RM;RN ($\epsilon = 1.22$) and O(1) ($\epsilon = 1.41$). This also confirms that the leakage detected by PANORAMIA on increasingly overfitted models does augment, which is confirmed by prior state-of-the-art methods. For instance, NLP models’ $\tilde{\epsilon}$ goes from 0.22 to 2.40 (2.82 to 5.73 for O(1)). Appendix D.6 details results on tabular data, in which PANORAMIA is not able to detect significant privacy loss.

Privacy Measurement of ϵ -DP Models:

Another avenue to varying privacy leakage is by training Differentially-Private (DP) models with different values of ϵ . We evaluate PANORAMIA on DP ResNet-18 models on CIFAR10, with ϵ values shown on Table 3 and Figure 5. The hyper-parameters were tuned independently for the highest train accuracy (see Appendix D.9 for more results and discussion). As the results show, neither PANORAMIA nor O(1) detect privacy loss on the most private models ($\epsilon = 1, 2$). At higher values of $\epsilon = 10, 15$ (*i.e.*, less private models) and $\epsilon = \infty$ (*i.e.*, non-private model) PANORAMIA does detect an increasing level of privacy leakage with $\tilde{\epsilon}_{\epsilon=10} < \tilde{\epsilon}_{\epsilon=15} < \tilde{\epsilon}_{\epsilon=\infty}$. In this regime, the O(1) approach detects a larger, though comparable, amount of privacy loss.

We also evaluate PANORAMIA on DP (fine-tuned) large language models (see Appendix D.9 and Table 9 for details). We fine-tune GPT2-Small target models with DP-SGD on the WikiText dataset, with various values of ϵ . Neither PANORAMIA nor O(1) are able to measure a positive privacy loss.

Target model f	Audit	c_b	$\{\epsilon + c\}_b$	$\tilde{\epsilon}$	ϵ
ResNet101_E20	PANORAMIA	2.44 ± 0.190	2.76 ± 0.250	0.33 ± 0.21	-
	PANORAMIA RM:RN	0.00	0.421 ± 0.19	-	0.421 ± 0.19
	O(1) RM:RN	-	-	-	0.450 ± 0.150
ResNet101_E50	PANORAMIA	2.44 ± 0.190	3.11 ± 0.18	0.67 ± 0.17	-
	PANORAMIA RM:RN	0.00	0.71 ± 0.21	-	0.71 ± 0.21
	O(1) RM:RN	-	-	-	0.79 ± 0.23
ResNet101_E100	PANORAMIA	2.44 ± 0.190	3.60 ± 0.21	1.16 ± 0.19	-
	PANORAMIA RM:RN	0.00	1.22 ± 0.190	-	1.22 ± 0.19
	O(1) RM:RN	-	-	-	1.41 ± 0.180
WRN-28-10_E150	PANORAMIA	2.44 ± 0.190	5.01 ± 0.25	2.57 ± 0.23	-
	PANORAMIA RM:RN	0.00	2.87 ± 0.20	-	2.87 ± 0.20
	O(1) RM:RN	-	-	-	2.96 ± 0.19
WRN-28-10_E300	PANORAMIA	2.44 ± 0.190	5.98 ± 0.23	3.54 ± 0.23	-
	PANORAMIA RM:RN	0.00	3.65 ± 0.22	-	3.65 ± 0.22
	O(1) RM:RN	-	-	-	3.73 ± 0.19
ViT_E35	PANORAMIA	2.44 ± 0.190	5.03 ± 0.24	2.59 ± 0.22	-
	PANORAMIA RM:RN	0.00	2.77 ± 0.21	-	2.77 ± 0.21
	O(1) RM:RN	-	-	-	2.86 ± 0.19
CNN_E50	PANORAMIA	2.05 ± 0.21	2.81 ± 0.24	0.76 ± 0.23	-
	PANORAMIA RM:RN	0.00	0.91 ± 0.22	-	0.91 ± 0.22
	O(1) RM:RN	-	-	-	1.01 ± 0.18
CNN_E100	PANORAMIA	2.05 ± 0.21	3.38 ± 0.24	1.33 ± 0.22	-
	PANORAMIA RM:RN	0.00	1.38 ± 0.21	-	1.38 ± 0.21
	O(1) RM:RN	-	-	-	1.46 ± 0.16
GPT2_E37	PANORAMIA	3.62 ± 0.32	3.85 ± 0.46	0.22 ± 0.37	-
	PANORAMIA RM:RN	0.00	1.04 ± 0.42	-	1.04 ± 0.42
	O(1) RM:RN	-	-	-	2.82 ± 0.31
GPT2_E75	PANORAMIA	3.62 ± 0.32	5.20 ± 0.34	1.57 ± 0.45	-
	PANORAMIA RM:RN	0.00	2.52 ± 0.36	-	2.52 ± 0.36
	O(1) RM:RN	-	-	-	4.71 ± 0.34
GPT2_E150	PANORAMIA	3.62 ± 0.32	6.02 ± 0.34	2.40 ± 0.53	-
	PANORAMIA RM:RN	0.00	3.60 ± 0.41	-	3.60 ± 0.41
	O(1) RM:RN	-	-	-	5.73 ± 0.08
ResNet18 $\epsilon = \infty$	PANORAMIA	2.44 ± 0.190	3.87 ± 0.20	1.43 ± 0.21	-
ResNet18 $\epsilon = 15$	PANORAMIA	2.44 ± 0.190	3.57 ± 0.19	1.13 ± 0.19	-
ResNet18 $\epsilon = 10$	PANORAMIA	2.44 ± 0.190	2.70 ± 0.25	0.26 ± 0.22	-
ResNet18 $\epsilon = 2$	PANORAMIA	2.44 ± 0.190	1.709 ± 0.23	0.00	-
ResNet18 $\epsilon = 1$	PANORAMIA	2.44 ± 0.190	1.412 ± 0.12	0.00	-
	O(1) RM:RN	-	-	-	0.00

Table 3: Privacy measurements on different target models ($m = 10k$). O(1) RM:RN with our ML-based attack for ϵ -DP models shown in Figure 6(c).

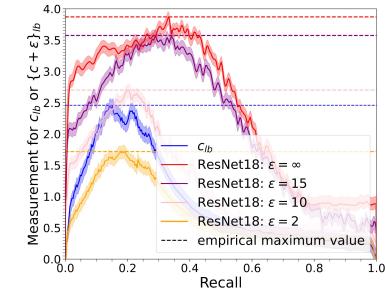


Figure 5: ResNet18, CIFAR-10, ϵ -DP for various ϵ values.

5.3 Leveraging More Data to Improve Privacy Measurements

PANORAMIA can leverage much more data for its measurement (up to the whole training set size for training and testing the MIA), while in our setting O(1) is limited by the size of the test set (for non-members). As we have seen in Figure 6(b), PANORAMIA can leverage larger test set size to measure higher privacy loss values than O(1), up to 2.62 on the ResNet101 trained for 100 epochs, versus 2.23 for O(1). It is also important to note that at small amounts of data O(1) measures a larger privacy loss ($\tilde{\epsilon} \leq \epsilon$). These findings provide additional empirical evidence for the robustness of b , indicating that we are not erroneously attributing differences between \mathcal{G} and \mathcal{D} to our privacy loss proxy $\tilde{\epsilon}$. However, in the case of text data (on the GPT-2 trained for 150 epochs), Figure 6(d), we hit an upper-bound on the power of the hypothesis test which estimates $\{c + \epsilon\}_{lb}$ using 20k test samples. The maximum measurement for $\tilde{\epsilon}$ we can achieve is around 2.59 with a 20k test set, as opposed to O(1), which reaches 5.73 with 10k test set size.

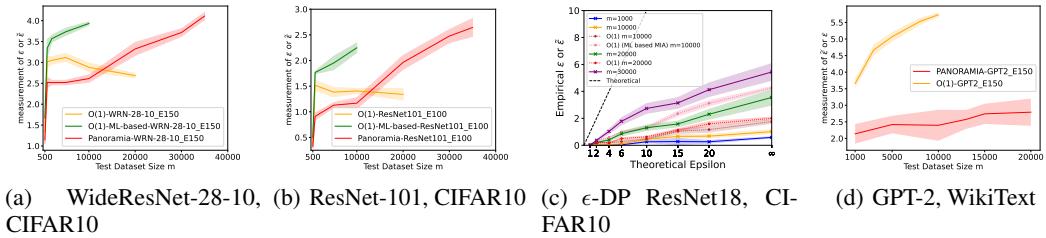


Figure 6: Effect of increasing test set size on PANORAMIA’s privacy measurement for our target models. In the case of the image dataset, increasing the number of auditing examples allows us to achieve tighter empirical measurements for privacy leakage despite a restricted adversary. For the case of the language modeling task, when increasing test set size m , we hit an upper bound on the power of the hypothesis test due to test dataset size, and do not see significant improvement in privacy measurement in this case.

Increasing Test Set Size for DP models: With an increase in the test set size m , we demonstrate that we can achieve tighter bounds on the privacy measurement $\tilde{\epsilon}$. Figure 6(c) shows that as m increases, PANORAMIA (solid lines) can measure higher privacy leakage for various ϵ -DP ResNet18 models. The reason for this is that PANORAMIA can leverage a much higher test set size (up to the whole training set size for training and testing the MIA), in comparison to O(1) in our setting (due to being limited by the number of real non-member data points available).

6 Impact, Limitations and Future Directions

We have introduced a new approach to quantify the privacy leakage from ML models, in settings in which individual data contributors (such as a hospital in a cross-site FL setting or a user of a text auto-complete service) measure the leakage of their own, known partial training data in the final trained model. Our approach does not introduce new attack capabilities (it would likely benefit from future progress on MIA performance though), but can help model providers and data owners assess the privacy leakage incurred by their data, due to inclusion in ML models training datasets. Consequently, we believe that any impact from our proposed work is likely to be positive.

However, PANORAMIA suffers from a major limitation: its privacy measurement is not a lower-bound on the privacy loss (see details in Section 4.2). Providing a proper lower-bound ϵ_{lb} on privacy leakage in this setting is an important avenue for future work. A promising direction is to devise a way to measure or enforce an upper-bound on c , thereby yielding ϵ_{lb} via $\{c + \epsilon\}_{lb} - c_{ub}$. Despite this shortcoming, we believe that the new measurement setting we introduce in this work is important. Indeed, recent work has shown that MIA evaluation on foundation models suffers from a similar lack of non-member data (since all in distribution data is included in the model) Das et al. (2024); Duan et al. (2024); Meeus et al. (2024). The theory we develop in this paper provides a meaningful step towards addressing privacy measurements in this setting, and provides a more rigorous approach to privacy benchmarks for such models. We also demonstrate that PANORAMIA’s privacy measurements can also be empirically valuable for instance for providing improved measurements with more data (Figure 6).

Acknowledgments

Mathias Lécuyer is grateful for the support of the Natural Sciences and Engineering Research Council of Canada (NSERC) [reference number RGPIN-2022-04469], as well as a Google Research Scholar award. Sébastien Gambs is supported by the Canada Research Chair program, a Discovery Grant from the Natural Sciences and Engineering Research Council of Canada (NSERC) as well as the DEEL Project CRDPJ 537462-18 funded by the NSERC and the Consortium for Research and Innovation in Aerospace in Québec (CRIAQ). This research was enabled by computational support provided by the Digital Research Alliance of Canada (alliancecan.ca). Hadrien Lautraite is supported by National Bank of Canada and a bursary from the Canada Research Chair in Privacy-preserving and Ethical Analysis of Big Data..

References

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pp. 308–318, 2016.

Andrew, G., Kairouz, P., Oh, S., Oprea, A., McMahan, H. B., and Suriyakumar, V. One-shot empirical privacy estimation for federated learning, 2023.

Carlini, N., Liu, C., Erlingsson, Ú., Kos, J., and Song, D. The secret sharer: Evaluating and testing unintended memorization in neural networks. In *28th USENIX Security Symposium (USENIX Security 19)*, pp. 267–284, 2019.

Carlini, N., Tramer, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, U., et al. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, 2021.

Carlini, N., Chien, S., Nasr, M., Song, S., Terzis, A., and Tramer, F. Membership inference attacks from first principles. In *2022 IEEE Symposium on Security and Privacy (SP) (SP)*, pp. 1519–1519, Los Alamitos, CA, USA, may 2022a. IEEE Computer Society. doi: 10.1109/SP46214.2022.00090. URL <https://doi.ieee.org/10.1109/SP46214.2022.00090>.

Carlini, N., Chien, S., Nasr, M., Song, S., Terzis, A., and Tramer, F. Membership inference attacks from first principles. In *2022 IEEE Symposium on Security and Privacy (SP)*, pp. 1897–1914. IEEE, 2022b.

Das, D., Zhang, J., and Tramèr, F. Blind baselines beat membership inference attacks for foundation models. *arXiv preprint arXiv:2406.16201*, 2024.

Dong, J., Roth, A., and Su, W. J. Gaussian differential privacy. *arXiv preprint arXiv:1905.02383*, 2019.

Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houlsby, N. An image is worth 16x16 words: Transformers for image recognition at scale, 2021. URL <https://arxiv.org/abs/2010.11929>.

Duan, M., Suri, A., Mireshghallah, N., Min, S., Shi, W., Zettlemoyer, L., Tsvetkov, Y., Choi, Y., Evans, D., and Hajishirzi, H. Do membership inference attacks work on large language models? *arXiv preprint arXiv:2402.07841*, 2024.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*. Springer, 2006.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition, 2015.

Holtzman, A., Buys, J., Du, L., Forbes, M., and Choi, Y. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*, 2019.

Jagielski, M., Ullman, J., and Oprea, A. Auditing differentially private machine learning: How private is private sgd? *Advances in Neural Information Processing Systems*, 33:22205–22216, 2020.

Jayaraman, B. and Evans, D. E. Evaluating differentially private machine learning in practice. In *USENIX Security Symposium*, 2019.

Kairouz, P., Oh, S., and Viswanath, P. The composition theorem for differential privacy. In *International conference on machine learning*. PMLR, 2015.

Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., and Aila, T. Training generative adversarial networks with limited data, 2020.

Krizhevsky, A. Learning multiple layers of features from tiny images. 2009. URL <https://api.semanticscholar.org/CorpusID:18268744>.

Li, X., Tramer, F., Liang, P., and Hashimoto, T. Large language models can be strong differentially private learners. *arXiv preprint arXiv:2110.05679*, 2021.

Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.

Lu, F., Munoz, J., Fuchs, M., LeBlond, T., Zaresky-Williams, E., Raff, E., Ferraro, F., and Testa, B. A general framework for auditing differentially private machine learning, 2023.

Maddock, S., Sablayrolles, A., and Stock, P. Canife: Crafting canaries for empirical privacy measurement in federated learning, 2023.

McKenna, R., Miklau, G., and Sheldon, D. Winning the nist contest: A scalable and general approach to differentially private synthetic data. *arXiv preprint arXiv:2108.04978*, 2021.

Meeus, M., Jain, S., Rei, M., and de Montjoye, Y.-A. Inherent challenges of post-hoc membership inference for large language models. *arXiv preprint arXiv:2406.17975*, 2024.

Merity, S., Xiong, C., Bradbury, J., and Socher, R. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*, 2016.

Nasr, M., Shokri, R., and Houmansadr, A. Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning. In *2019 IEEE symposium on security and privacy (SP)*, pp. 739–753. IEEE, 2019.

Nasr, M., Songi, S., Thakurta, A., Papernot, N., and Carlini, N. Adversary instantiation: Lower bounds for differentially private machine learning. In *2021 IEEE Symposium on Security and Privacy (SP)*, pp. 866–882, 2021. doi: 10.1109/SP40001.2021.00069.

Nasr, M., Hayes, J., Steinke, T., Balle, B., Tramèr, F., Jagielski, M., Carlini, N., and Terzis, A. Tight auditing of differentially private machine learning, 2023.

Negoescu, D. M., Gonzalez, H., Orjany, S. E. A., Yang, J., Lut, Y., Tandra, R., Zhang, X., Zheng, X., Douglas, Z., Nolkha, V., et al. Epsilon*: Privacy metric for machine learning models. *arXiv preprint arXiv:2307.11280*, 2023.

O’Shea, K. and Nash, R. An introduction to convolutional neural networks, 2015.

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.

Steinke, T., Nasr, M., and Jagielski, M. Privacy auditing with one (1) training run. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

Suri, A., Zhang, X., and Evans, D. Do parameters reveal more than loss for membership inference? In *High-dimensional Learning Dynamics 2024: The Emergence of Structure and Reasoning*.

Wasserman, L. and Zhou, S. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 2010.

Yeom, S., Giacomelli, I., Fredrikson, M., and Jha, S. Privacy risk in machine learning: Analyzing the connection to overfitting. In *2018 IEEE 31st computer security foundations symposium (CSF)*, pp. 268–282. IEEE, 2018.

Yousefpour, A., Shilov, I., Sablayrolles, A., Testuggine, D., Prasad, K., Malek, M., Nguyen, J., Ghosh, S., Bharadwaj, A., Zhao, J., Cormode, G., and Mironov, I. Opacus: User-friendly differential privacy library in PyTorch. *arXiv preprint arXiv:2109.12298*, 2021.

Zagoruyko, S. and Komodakis, N. Wide residual networks, 2017a.

Zagoruyko, S. and Komodakis, N. Wide residual networks, 2017b. URL <https://arxiv.org/abs/1605.07146>.

Zanella-Béguelin, S., Wutschitz, L., Tople, S., Salem, A., Rühle, V., Paverd, A., Naseri, M., Köpf, B., and Jones, D. Bayesian estimation of differential privacy, 2022.

Zarifzadeh, S., Liu, P. C.-J. M., and Shokri, R. Low-cost high-power membership inference by boosting relativity. 2023.

A Notations

Table 4 summarizes the main notations used in the paper.

Notation	Description
f	the target model to be audited.
\mathcal{G}	the generative model
\mathcal{D}	distribution over auditor samples
D_f	(subset of the) training set of the target model f from \mathcal{D}
D_G	training set of the generative model \mathcal{G} , with $D_G \subset D_f$
D_{in}	member auditing set, with $D_{\text{in}} \subset D_f$ and $D_{\text{in}} \cap D_G = \emptyset$
D_{out}	non-member auditing set, with $D_{\text{out}} \sim \mathcal{G}$
$D^{\{\text{tr,te}\}}_{\{\text{in,out}\}}$	training and testing splits of D_{in} and D_{out}
m	$ D_{\text{in}} = D_{\text{out}} \triangleq m$
b	baseline classifier for D_{in} vs. D_{out}

Table 4: Summary of notations

B Proofs

For both Proposition 1 and Proposition 2, we state the proposition again for convenience before proving it.

B.1 Proof of Proposition 1

Proposition 1. *Let \mathcal{G} be c -close, and $T^b \triangleq B(S, X)$ be the guess from the baseline. Then, for all $v \in \mathbb{R}$ and all t in the support of T :*

$$\mathbb{P}_{S, X, T^b} \left[\sum_{i=1}^m T_i^b \cdot S_i \geq v \mid T^b = t^b \right] \leq \mathbb{P}_{S' \sim \text{Bernoulli}(\frac{e^c}{1+e^c})^m} \left[\sum_{i=1}^m t_i^b \cdot S'_i \geq v \right] \triangleq \beta^b(m, c, v, t^b)$$

Proof. Notice that under our baseline model $B(s, x) = \{b(x_1), b(x_2), \dots, b(x_m)\}$, and given that the X_i are i.i.d., we have that: $S_{<i} \perp\!\!\!\perp T_{<i}^b \mid X_{<i}$, since $T_i^b = B(S, X)_i$'s distribution is entirely determined by X_i ; and $S_{\leq i} \perp\!\!\!\perp T_{>i}^b \mid X_{<i}$ since the X_i are sampled independently from the past.

We study the distribution of S given a fixed prediction vector t^b , one element $i \in [m]$ at a time:

$$\begin{aligned} & \mathbb{P}[S_i = 1 \mid T^b = t^b, S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}] \\ &= \mathbb{P}[S_i = 1 \mid S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}] \\ &= \mathbb{P}[X_i \mid S_i = 1, S_{<i} = s_{<i}, X_{<i} = x_{<i}] \frac{\mathbb{P}[S_i = 1 \mid S_{<i} = s_{<i}, X_{<i} = x_{<i}]}{\mathbb{P}[X_i \mid S_{<i} = s_{<i}, X_{<i} = x_{<i}]} \\ &= \frac{\mathbb{P}[X_i \mid S_i = 1, S_{<i} = s_{<i}, X_{<i} = x_{<i}] \mathbb{P}[S_i = 1]}{\mathbb{P}[X_i \mid S_{<i} = s_{<i}, X_{<i} = x_{<i}]} \\ &= \frac{\mathbb{P}[X_i \mid S_i = 1]^{\frac{1}{2}}}{\mathbb{P}[X_i \mid S_i = 1]^{\frac{1}{2}} + \mathbb{P}[X_i \mid S_i = 0]^{\frac{1}{2}}} \\ &= \frac{1}{1 + \frac{\mathbb{P}[X_i \mid S_i = 0]}{\mathbb{P}[X_i \mid S_i = 1]}} = \frac{1}{1 + \frac{\mathbb{P}_{\mathcal{G}}[X_i]}{\mathbb{P}_{\mathcal{D}}[X_i]}} \leq \frac{1}{1 + e^{-c}} = \frac{e^c}{1 + e^c} \end{aligned}$$

The first equality uses the independence remarks at the beginning of the proof, the second relies Bayes' rule, while the third and fourth that S_i is sampled i.i.d from a Bernoulli with probability half, and X_i i.i.d. conditioned on S_i . The last inequality uses Definition 3 for c -closeness.

Using this result and the law of total probability to introduce conditioning on $X_{<i}$, we get that:

$$\begin{aligned} & \mathbb{P}[S_i = 1 \mid T^b = t^b, S_{<i} = s_{<i}] \\ &= \sum_{x_{\leq i}} \mathbb{P}[S_i = 1 \mid T^b = t^b, S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}] \mathbb{P}[X_{\leq i} = x_{\leq i} \mid T^b = t^b, S_{<i} = s_{<i}] \\ &\leq \sum_{x_{\leq i}} \frac{e^c}{1 + e^c} \mathbb{P}[X_{\leq i} = x_{\leq i} \mid T^b = t^b, S_{<i} = s_{<i}], \end{aligned}$$

and hence that:

$$\mathbb{P}[S_i = 1 \mid T^b = t^b, S_{<i} = s_{<i}] \leq \frac{e^c}{1 + e^c} \quad (3)$$

We can now proceed by induction: assume inductively that $W_{m-1} \triangleq \sum_{i=1}^{m-1} T_i^b \cdot S_i$ is stochastically dominated (see Definition 4.8 in [Steinke et al. \(2023\)](#)) by $W'_{m-1} \triangleq \sum_{i=1}^{m-1} T_i^b \cdot S'_i$, in which $S' \sim \text{Bernoulli}(\frac{e^c}{1 + e^c})^{m-1}$. Setting $W_1 = W'_1 = 0$ makes it true for $m = 1$. Then, conditioned on W_{m-1} and using Eq. 3, $T_m^b \cdot S_m = T_m \cdot \mathbb{1}\{S_m = 1\}$ is stochastically dominated by $T_m^b \cdot \text{Bernoulli}(\frac{e^c}{1 + e^c})$. Applying Lemma 4.9 from [Steinke et al. \(2023\)](#) shows that W_m is stochastically dominated by W'_m , which proves the induction and implies the proposition's statement. \square

B.2 Proof of Proposition 2

Proposition 2. *Let \mathcal{G} be c -close, f be ϵ -DP, and $T^a \triangleq A(S, X, f)$ be the guess from the membership audit. Then, for all $v \in \mathbb{R}$ and all t in the support of T :*

$$\mathbb{P}_{S, X, T^a} \left[\sum_{i=1}^m T_i^a \cdot S_i \geq v \mid T^a = t^a \right] \leq \mathbb{P}_{S' \sim \text{Bernoulli}(\frac{e^c + \epsilon}{1 + e^c + \epsilon})^m} \left[\sum_{i=1}^m t_i^a \cdot S'_i \geq v \right] \triangleq \beta^a(m, c, \epsilon, v, t^a)$$

Proof. Fix some $t^a \in \mathbb{R}_+^m$. We study the distribution of S one element $i \in [m]$ at a time:

$$\begin{aligned} & \mathbb{P}[S_i = 1 \mid T^a = t^a, S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}] \\ &= \mathbb{P}[T^a = t^a \mid S_i = 1, S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}] \frac{\mathbb{P}[S_i = 1 \mid S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}]}{\mathbb{P}[T^a = t^a \mid S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}]} \\ &\leq \frac{1}{1 + e^{-\epsilon} \frac{\mathbb{P}[S_i = 0 \mid S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}]}{\mathbb{P}[S_i = 1 \mid S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}]}} \\ &\leq \frac{1}{1 + e^{-\epsilon} e^{-c}} = \frac{e^{c+\epsilon}}{1 + e^{c+\epsilon}} \end{aligned}$$

The first equality uses Bayes' rule. The first inequality uses the decomposition:

$$\begin{aligned} & \mathbb{P}[T^a = t^a \mid S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}] \\ &= \mathbb{P}[T^a = t^a \mid S_i = 1, S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}] \cdot \mathbb{P}[S_i = 1 \mid S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}] \\ &\quad + \mathbb{P}[T^a = t^a \mid S_i = 0, S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}] \cdot \mathbb{P}[S_i = 0 \mid S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}], \end{aligned}$$

and the fact that $A(s, x, f)$ is ϵ -DP w.r.t. s and hence that:

$$\frac{\mathbb{P}[T^a = t^a \mid S_i = 0, S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}]}{\mathbb{P}[T^a = t^a \mid S_i = 1, S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}]} \geq e^{-\epsilon}. \quad (4)$$

This inequality might seem surprising at first, as we know from the privacy game (Definition 2) that S is independent of f , that is $S \perp\!\!\!\perp f$. However, these two quantities are not conditionally independent when conditioning on X (where X is the test set for the MIA/baseline), (i.e., $S \not\perp\!\!\!\perp f \mid X$). This is because X_i is either a member or non-member based on S_i . So when $S_i = 1$, then f was trained on X_i , which makes f and S_i dependent if membership can be detected through f . If f is ϵ -DP, we can bound this dependency using DP properties, which is what we do in Equation (4).

The second inequality uses that:

$$\begin{aligned}
& \frac{\mathbb{P}[S_i = 0 \mid S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}]}{\mathbb{P}[S_i = 1 \mid S_{<i} = s_{<i}, X_{\leq i} = x_{\leq i}]} \\
&= \frac{\mathbb{P}[X_i = x_i \mid S_i = 0, S_{<i} = s_{<i}, X_{<i} = x_{<i}]}{\mathbb{P}[X_i = x_i \mid S_i = 1, S_{<i} = s_{<i}, X_{<i} = x_{<i}]} \cdot \frac{\mathbb{P}[S_i = 0 \mid S_{<i} = s_{<i}, X_{<i} = x_{<i}]}{\mathbb{P}[S_i = 1 \mid S_{<i} = s_{<i}, X_{<i} = x_{<i}]} \\
&= \frac{\mathbb{P}[X_i = x_i \mid S_i = 0, S_{<i} = s_{<i}, X_{<i} = x_{<i}]}{\mathbb{P}[X_i = x_i \mid S_i = 1, S_{<i} = s_{<i}, X_{<i} = x_{<i}]} \cdot \frac{1/2}{1/2} \\
&= \frac{\mathbb{P}_{\mathcal{G}}[X_i]}{\mathbb{P}_{\mathcal{D}}[X_i]} \geq e^{-c}
\end{aligned}$$

As in Proposition 1, applying the law of total probability to introduce conditioning on $X_{\leq i}$ yields:

$$\mathbb{P}[S_i = 1 \mid T^a = t^a, S_{<i} = s_{<i}] \leq \frac{e^{c+\epsilon}}{1 + e^{c+\epsilon}}, \quad (5)$$

and we can proceed by induction. Assume inductively that $W_{m-1} \triangleq \sum_{i=1}^{m-1} T_i^a \cdot S_i$ is stochastically dominated (see Definition 4.8 in [Steinke et al. \(2023\)](#)) by $W'_{m-1} \triangleq \sum_{i=1}^{m-1} T_i^a \cdot S'_i$, in which $S' \sim \text{Bernoulli}(\frac{e^{c+\epsilon}}{1+e^{c+\epsilon}})^{m-1}$. Setting $W_1 = W'_1 = 0$ makes it true for $m = 1$. Then, conditioned on W_{m-1} and using Eq. 5, $T_m^a \cdot S_m = T_m^a \cdot \mathbb{1}\{S_m = 1\}$ is stochastically dominated by $T_m^a \cdot \text{Bernoulli}(\frac{e^{c+\epsilon}}{1+e^{c+\epsilon}})$. Applying Lemma 4.9 from [Steinke et al. \(2023\)](#) shows that W_m is stochastically dominated by W'_m , which proves the induction and implies the proposition's statement. \square

C Experimental Details

Hereafter, we provide the details about the datasets and models trained in our experiments.

C.1 Image data

Target models. We audit target models with the following architectures: a Multi-Label Convolutional Neural Network (CNN) with four layers [O’Shea & Nash \(2015\)](#), and the ResNet101 [He et al. \(2015\)](#) as well as a Wide ResNet-28-10 [Zagoruyko & Komodakis \(2017b\)](#), and a Vision Transformer (ViT)-small [Dosovitskiy et al. \(2021\)](#). We also include in our analysis, differentially-private models for ResNet18 [He et al. \(2015\)](#) and WideResNet-16-4 [Zagoruyko & Komodakis \(2017a\)](#) models as targets, with $\epsilon = 1, 2, 4, 6, 10, 15, 20$. The ResNet-based models are trained on CIFAR10 using 50k images [Krizhevsky \(2009\)](#) of 32x32 resolution. For ResNet101 CIFAR10-based classification models, we use a training batch size of 32; for WideResNet-28-10, we use batch size 128. The details for training DP models are given in [D.9](#). The associated test accuracies and epochs are mentioned in Table 1. The Multi-Label CNN is trained on 200k images of CelebA [Liu et al. \(2015\)](#) of 128x128 resolution, training batch-size 32, to predict 40 attributes associated with each image.

Generator. For both image datasets, we use StyleGAN2 [Karras et al. \(2020\)](#) to train the generative model \mathcal{G} from scratch on D_G , and produce non-member images. For CIFAR10 dataset, we use a 10,000 out of 50,000 images from the training data of the target model to train the generative model. For the CelebA dataset, we select 35,000 out of 200,000 images from the training data of the target model to train the generative model. Generated images will in turn serve as non-members for performing the MIAs. Figure 7 shows examples of member and non-member images used in our experiments. In the case of CelebA, we also introduce a vanilla CNN as a classifier or filter to distinguish between fake and real images, and remove any poor-quality images that the classifier detects with high confidence. The data used to train this classifier was the same data used to train StyleGAN2, which ensures that the generated high-resolution images are of high quality.

MIA and Baseline training. For the MIA, we follow a loss-based attack approach: PANORAMIA takes as input raw member and non-member data points for training along with the loss values the target model f attributes to these data points. More precisely, the training set of PANORAMIA is:

$$(D_{in}^{tr}, f(D_{in}^{tr})) \cup (D_{out}^{tr}, f(D_{out}^{tr}))$$

In §4.2, we discussed the importance of having a tight c_{lb} so that our measure, $\tilde{\epsilon}$, becomes close to a lower-bound on ϵ -DP, which requires a strong baseline. To strengthen our baseline, we introduce the helper model h , which helps the baseline model b by supplying additional features (*i.e.*, embeddings) that can be viewed as side information about the data distribution. The motivation is that h ’s features might differ between samples from \mathcal{D} and \mathcal{D}' , enhancing the performance of the baseline classifier. This embedding model h is similar in design to f (same task and architecture) but is trained on synthetic data that is close in distribution to the real member data distribution.

The helper model has the same classification task and architecture as the target model. To train it, we generate separate sets of training and validation data using our generator. For image data, the synthetic samples do not have a label. We thus train a labeler model, also a classifier with the same task, on the same dataset used to train the generator. We use the labeler model to provide labels for our synthetic samples (training and validation sets above). We train the helper model on the resulting training set, and select hyperparameters on the validation set.

Whether for the baseline or MIA, we use side information models (h and f , respectively) by concatenating the loss of $h(x)$ and $f(x)$ to the final feature representation (more details are provided later) before a last layer of the MIA/Baseline makes the membership the prediction. Since we need labels to compute the loss, we label synthetic images with a Wide ResNet-28-2 in the case of CIFAR10, and a Multi-Label CNN of similar architecture as the target model in the case of CelebA labeling. For both instances, we used a subset of the data, that was used to train the respective generative models, to train the “labeler” classifiers as well. The labeler used to train the helper model for image data has 80.4% test accuracy on the CIFAR10 real data test set. The rationale for this approach is to augment the baseline with a model providing good features (here for generated data) to balance the good features provided to the MIA by f outside of the membership information. In practice, the labeled generated data seems enough to provide such good features, despite the fact that the labeler is not extremely accurate. We studied alternative designs (*e.g.*, a model trained on non-member data, no helper model) in Appendix D.1, Table 5, and the helper model trained on the synthetic data task performs best (while not requiring non-member data, which is a key point).

We use two different modules for each MIA and Baseline training. More precisely, the first module optimizes image classification using a built-in Pytorch ResNet101 classifier. The second module, in the form of a multi-layer perceptron, focuses on classifying member and non-member labels via loss values attributed to these data points by f as input for the loss module of MIA and losses of e to the baseline b respectively. We then stack the scores of both image and loss modules into a logistic regression task (as a form of meta-learning) to get the final outputs for member and non-member data points by MIA and baseline b . When training the baseline and MIA models, we use a validation set to select hyper-parameters and training stopping time that maximize the lower bounds (effectively maximizing the c_{lb} or $\tilde{\epsilon}_{lb}$ for the Baseline and MIA respectively) on the validation set. The MIA and baseline are trained on 4500 data samples (half members and half generated non-members). The test dataset consists of 10000 samples, again half members and half generated non-members. The actual and final number of members and non-members that ended up in the test set depends on the Bernoulli samples in our auditing game. We repeat the training process over 5 times, each time independently resampling the training dataset for the MIA and baseline (keeping the train dataset size fixed). We report all our results over a 95% confidence interval and report the standard deviation for our results in Table 3. The **compute resources** used for MIA and Baseline training were mainly running all experiments on cloud-hosted VMs, using 1 v100l GPU, 4 nodes per task, and 32G memory requested for each job on the cloud cluster. The time to run the MIA and baseline attack pipeline was around 10 hours including hyperparameter tuning on ϵ_{lb} and c_{lb} respectively.

C.2 Language Modeling

Target model. The target model (a small GPT-2 [Radford et al. \(2019\)](#)) task is causal language modeling (CLM), a self-supervised task in which the model predicts the next word in a given sequence of words (*i.e.*, context), which is done on a subset of WikiText-103 dataset [Merity et al. \(2016\)](#), a collection of Wikipedia articles. While the GPT-2 training dataset is undisclosed, [Radford et al. \(2019\)](#) do state that they did not include any Wikipedia document.

One common standard pre-processing in CLM is to break the tokenized training dataset into chunks of equal sizes to avoid padding of the sequences. To achieve this, the entire training dataset is split into token sequences of fixed length, specifically 64 tokens per sequence, refer hereafter as “chunks”.

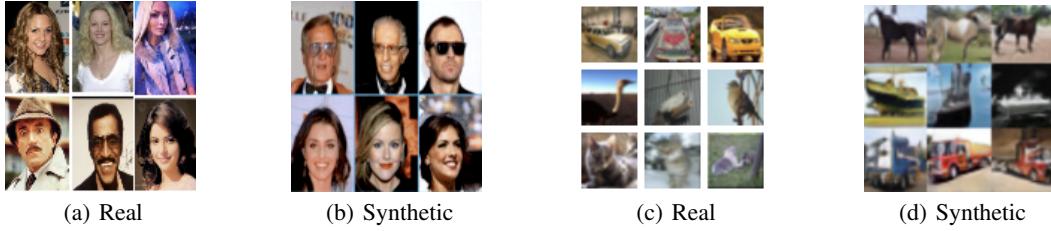


Figure 7: Member and Non-Member datasets used in our experiments for CelebA at 128 x 128 resolution (7(a), 7(b)) and CIFAR10 (7(c), 7(d)) image data at 32 x 32 resolution.

Each of these chunks serves as an individual sample within our training dataset. Therefore, in the experiments, the membership inference will be performed for a chunk. In other words, we adopt an example-level differential privacy (DP) viewpoint, where an example is considered a sequence of tokens that constitute a row in a batch of samples. This chunk-based strategy offers additional benefits in synthetic sample generation. By maintaining uniformity in the length of these synthetic samples, which also consists of 64 tokens, it is possible to effectively mitigate any distinguishability between synthetic and real samples based solely on their length. We also consider the possibility of weak correlations between chunks (e.g., being from the same article). To mitigate this, we perform a form of stratified sampling as a pre-processing step. From each document, we include only the first 50 chunks and discard the rest. We select 2,113 documents in total, resulting in 105,650 chunks in our dataset.

We split the dataset into train and test with a 90:10 ratio for the target model (We insist on following the typical training pipelines in machine learning, since we plan to conduct post-hoc audits with PANORAMIA). The test dataset provides real non-member samples, allowing comparison to the O(1) auditor and another version of our approach which uses real non-member data. The training dataset (D_f) contains 95,085 samples, while the test dataset consists of 10,565 samples.

We train the target model for 200 epochs in total. Figure 8 displays how the training and validation cross-entropy loss changes throughout the training, when we audit models from epochs 37 (best generalization), 75, and 150. To check how well the target model generalizes, we look at the cross-entropy loss (on a validation set), which is the only metric in a causal language model task to report (or the perplexity [Radford et al. \(2019\)](#) which conveys the same information).

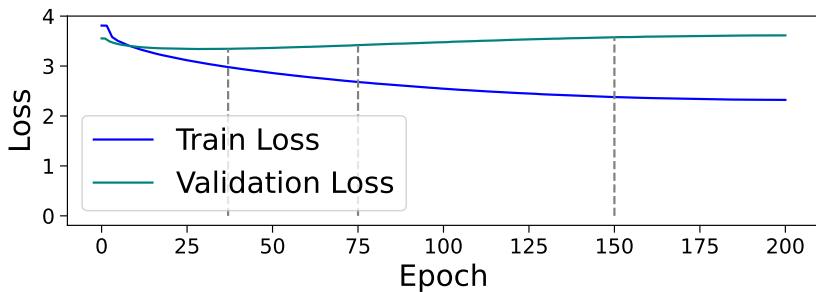


Figure 8: Target model (GPT-2) loss during overtraining on WikiText dataset. We pick GPT-2_E37, GPT-2_E75, GPT-2_E150 as target models to audit corresponding to the gray dashed lines.

The **compute resources** used for training the target models are 4 V100-32gb GPUs, running on cloud-hosted VMs. Using this setup, fine-tuning the target model for 200 epochs required about 13 hours.

Generator. The \mathcal{G} is a GPT-2 fine-tuned using a CLM task on dataset D_G , a subset of D_f . To create synthetic samples with \mathcal{G} , we use top- k sampling [Holtzman et al. \(2019\)](#) method in an auto-regressive manner while keeping the generated sequence length fixed at 64 tokens. To make \mathcal{G} generate samples less like the real members it learned from (and hopefully, more like real non-members), we introduce more randomness in the generation process in the following manner. We choose a top_k value of

200, controlling the number of available tokens for sampling. In addition, we balance the value of top_k , making it not too small to avoid repetitive generated texts but also not too large to maintain the quality. Finally, we fix the temperature parameter that controls the randomness in the softmax function at one. Indeed, a larger value increases the entropy of the distribution over tokens resulting in a more diverse generated text.

However, the quality of synthetic text depends on the prompts used. To mitigate this issue, we split the D_G dataset into two parts: D_G^{train} and D_G^{prompt} . \mathcal{G} is fine-tuned on D_G^{train} , while D_G^{prompt} is used for generating prompts. Numerically, the size of the D_G^{train} is 35% of the D_f and the size of D_G^{prompt} is 15% of D_f . During generation, we sample from D_G^{prompt} , which is a sequence of length 64. We feed the prompt into the generator and request the generator to generate a suffix of equivalent length. In our experiments, for each prompt, we make 8 synthetic samples, supplying a sufficient number of synthetic non-member samples for different parts of the pipeline of PANORAMIA, including training the helper model, and constructing $D_{\text{out}}^{\text{tr}}$, $D_{\text{out}}^{\text{te}}$. Overall, we generate 92276 synthetic samples out of which we use 53502 of them for training the helper model and keep the rest for audit purposes. Note that this generation approach does not inherently favor our audit scheme, and shows how PANORAMIA can leverage existing, public generative models with fine tuning to access high-quality generators at reasonable training cost.

We also perform a sanity check on the quality of the generated data. We visualize the loss distribution of in-distribution members, in-distribution non-members, and synthetic data under both the target and helper models, aiming for synthetic data that behaves similarly to in-distribution non-member data. Figure 14 visualizes these distributions for our target and helper models.

The **compute resources** used for fine-tuning the generator model, and synthetic text generation are 4 V100-32gb GPUs, running on cloud-hosted VMs. Using this setup, fine-tuning the GPT2-small model took approximately 1.5 hours, and generating the synthetic samples required about 1 hour.

Baseline & MIA. For both the baseline and PANORAMIA classifiers, we employ a GPT-2 based sequence classification model. In this setup, GPT-2 extracts features (*i.e.*, hidden vectors) directly from the samples. We concatenate a vector of features extracted using the helper h (for the baseline) or target model f (for the MIA) to this representation. This vector of features is the loss sequence produced in a CLM task for the respective model. These two sets of extracted features are then concatenated before being processed through a feed-forward layer to generate the logit values required for binary classification (distinguishing between members and non-members).

In the main results in Table 3, for both the baseline and PANORAMIA classifiers, the training and validation sets consist of 20000 and 2000 samples, respectively, with equal member non-member distribution. The test set consists of 10000 samples, in which the actual number of members and non-members that ended up in the test set depends on the Bernoulli samples in the auditing game. The helper model is fine-tuned for 60 epochs on synthetic samples and we pick the model that has the lowest validation loss throughout training, generalizing well.

As the GPT-2 component of the classifiers can effectively minimize training loss without achieving strong generalization, regularization is applied to the classifier using weight decay. Additionally, the optimization process is broken into two phases. In the first phase, we exclusively update the parameters associated with the target (in PANORAMIA) or helper model (in the baseline) to ensure that this classifier has the opportunity to focus on these specific features. In the second phase, we optimize the entire model as usual. We train both the baseline and Membership Inference Attack (MIA) models with 5 different seeds. The randomness is over: 1. The data split into train/validation/test sets for both the baseline and MIA classifiers. Each classifier is retrained for every seed. 2. The coin flips in our auditing game (see Definition 2). For a fixed seed mentioned above, we employ an additional set of 5 seeds to account for the randomness in the training algorithm (such as initialization of the model, and batch sampling). We then train a classifier for each of these seeds. From the 5 models generated for each latter seed, we select the one with the largest c_{lb} value (for the baseline) or $\{c+\epsilon\}_{\text{lb}}$ (for MIA), determined over the validation set. We calculate the 95% confidence interval of our measurements using the t-score.

The **compute resources** used for MIA and Baseline training were 1 GPU with 32gb memory, running on cloud-hosted VMs. The compute time depends on the training set size, for the largest training size, it took about 3 hours to train the baseline or MIA.

The **compute resources** used for fine-tuning the helper model are 4 V100-32gb GPUs, running on cloud-hosted VMs. Using this setup, it took about 4 hours to fine-tune the helper model.

C.3 Tabular data

Target f and the helper model h . The Target f and the helper h model are both Multi-Layer Perceptron with four hidden layers containing 150, 100, 50, and 10 neurons respectively. Both models are trained using a learning rate of 0.0005 over 100 training epochs. In the generalized case and to obtain the embedding, we retain the model’s parameters that yield the lowest loss on a validation set, which typically occurs at epoch 10 (MLPE_10). For the overfitted scenario, we keep the model’s state after the 100 training epochs (MLP_E100). The non overfitted model and the overfitted one achieved an accuracy of 86% and 82% respectively.

Generator. We use the MST method [McKenna et al. \(2021\)](#) which is a differentially-private method to generate synthetic data. However, as we do not need Differential Privacy for data generation we simply set the value for ϵ to 1000. Synthetic data generators can sometimes produce bad-quality samples. Those out-of-distribution samples can affect our audit process (the detection between real and synthetic data being due to bad quality samples rather than privacy leakage). To circumvent this issue, we train an additional classifier to distinguish between real data from D_G and additional synthetic data (not used in the audit). We use this classifier to remove from the audit data synthetic data synthetic samples predicted as synthetic with high confidence.

Baseline and MIA. To distinguish between real and synthetic data, we use the Gradient Boosting model and conduct a grid search to find the best hyperparameters.

C.4 Comparison with Privacy Auditing with One (1) Training Run: Experimental Details

We implement the black-box auditor version of O(1) approach [Steinke et al. \(2023\)](#). This method assigns a membership score to a sample based on its loss value under the target model. They also subtract the sample’s loss under the initial state (or generally, a randomly initialized model) of the target model, helping to distinguish members from non-members even more. In our instantiation of the O(1) approach, we only consider the loss of samples on the final state of the target model. Moreover, in their audit, they choose not to guess the membership of every sample. This abstention has an advantage over making wrong predictions as it does not increase their baseline. Roughly speaking, their baseline is the total number of correct guesses achieved by employing a randomized response $(\epsilon, 0)$ mechanism, for those samples that O(1) auditor opts to predict. We incorporate this abstention approach in our implementation by using two thresholds, t_+ and t_- . More precisely, samples with scores below t_+ are predicted as members, those above t_- as non-members, and the rest are abstained from prediction. We check for different combinations of t_+ and t_- and report the highest ϵ among them. This involves performing multiple tests on the same evidence, each with a separate confidence level of 95%. To ensure that these tests hold collectively, we apply a union bound and adjust the significance level accordingly. The total number of separate tests depends on the different number of guesses that O(1) can make. We also set δ to 0 and use a confidence interval of 0.05 in their test. In PANORAMIA, for each hypothesis test (whether for c_{lb} or $\{c+\epsilon\}_{lb}$), we stick to a 0.025 confidence interval for each one, adding up to an overall confidence level of 0.05.

D Results: Additional Experiments and Detailed Discussion

D.1 Baseline Strength Evaluation on Text and Tabular datasets

The basis of our privacy measurement directly depends on how well our baseline classifier distinguishes between real member data samples and synthetic non-member data samples generated by our generative model. As mentioned in Section 5.1 and Appendix C.1, in order to increase the performance of our baseline b , we mimic the role of the target model f ’s loss in the MIA using a helper model h , which adds a loss-based feature to b . To assess the strength of the baseline classifier, we train the baseline’s helper model h on two datasets, synthetic data, and real non-members. We also train a baseline without any helper model. This enables us to identify which setup allows the baseline classifier to distinguish the member and non-member data the best in terms of the highest c_{lb} value (c_{lb} is reported for each case in table 5). For initial analysis, we keep a sub-set of real data as non-members to train a helper model with them, for conducting this ablation study on the baseline.

Baseline model	c_{lb}
CIFAR-10 Baseline D_h^u = gen, WRN	2.44 ± 0.19
CIFAR-10 Baseline D_h^u = real, WRN	2.21 ± 0.17
CIFAR-10 Baseline D_h^u = real, resnet101	2.01 ± 0.15
CIFAR-10 Baseline D_h^u = imgnet	1.12 ± 0.19
CIFAR-10 Baseline _{no helper}	1.25 ± 0.24
WikiText Baseline D_h^u = gen	3.31 ± 0.15
WikiText Baseline D_h^u = real	3.26 ± 0.14
WikiText Baseline _{no helper}	3.11 ± 0.15

Table 5: Baseline evaluation with different helper model scenarios, where WRN is Wide-ResNet-28-2 helper model architecture (in the case of CIFAR10 baseline).

Our experiments confirm that the baseline with loss values from a helper model trained on synthetic data performs better compared to a helper model trained on real non-member data. This confirms the strength of our framework in terms of not having a dependency of using real non-members. Therefore, we choose this setting for our baseline helper model. In addition, we also experiment with different helper model architectures in the case of CIFAR10, looking for the best baseline performance among them. For the WikiText dataset, the c_{lb} reported in Table 5 differs from that in Table 3 because we had less data available for this experiment. We withheld half of the real samples to train the helper model with real samples in its training set.

For each data modality setting, we also train the baseline classifier on an increasing dataset size of real member and synthetic non-member data points to gauge the trend of the baseline performance as the training dataset size increases (see Figure 2 and Figure 9). This experiment allows us to aim for a worst-case empirical lower bound on how close the generated non-member data is to the real member data distribution, *i.e.* what would be the largest c_{lb} we can reach as we increase the training set size. Subsequently, we can compare the c_{lb} reported in table 3 (corresponding to the vertical dashed lines in Figure 2 and Figure 9) to the largest c_{lb} we can achieve in this experiment. In the following, we explain why we chose that number of training samples for our baseline in the main results.

On the WikiText dataset, we increase the baseline’s training set size up to 50k, while fixing the validation and test sets sizes to 2000 and 10000, respectively). Figure 9 shows the changes of c_{lb} while varying the training set size. The vertical dashed line specified the number of training samples we chose for our main auditing results reported in Table 3.

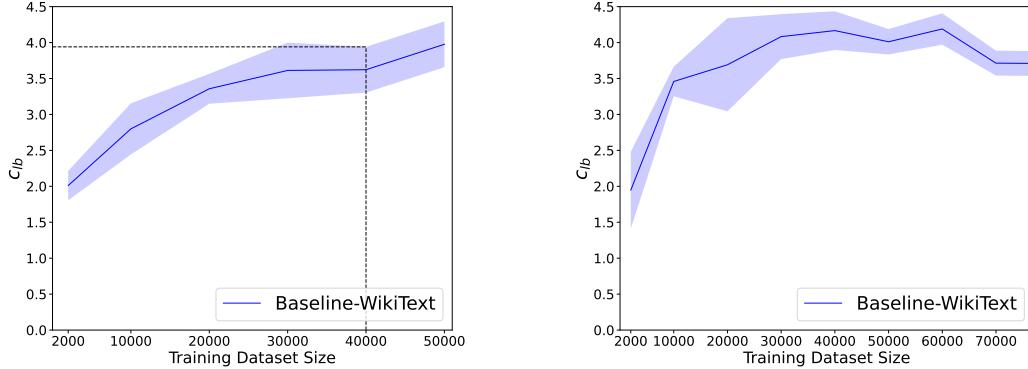
Now, let us discuss why we don’t choose the largest training set size available for the baseline. The measurement for c_{lb} depends not only on the training set size but also on the test set size. Therefore, we need to balance the sizes of the train and test sets. We conducted another experiment where we varied the train set size and, instead of fixing the test set size, allocated the remaining available data to the test set (keeping in mind that our data for the audit is limited by the number of real members allocated for audit purposes). Figure 9(b) depicts the trend of changes. As expected, we observed an initial increase in c_{lb} as the training set size increased (and the number of test data is sufficient to capture that). However, it decreased toward the end because of the smaller test set size. For a train set size of 40k, we found the best trade-off between train and test set sizes, which justifies our choice of 40k for the training set size in the main results. The same argument applies for other data modalities.

Let us make one last point. It is also important to note that the MIA is expected to show similar behavior if we use a larger train set size or test set size. Hence, the gap between the performance of MIA and baseline (which is our privacy measurement) is not expected to be affected (see Appendix D.4).

The experiments are repeated five times, and we provide a 95% confidence interval. For the image data modality, the randomness lies in the training data provided to the classifiers, while the test set remains fixed throughout. For the text data modality, both the train and test sets change across the five runs (the test set size is fixed to 10k).

D.2 Further Analysis on General Privacy Measurement Results

The value of $\tilde{\epsilon}$ for each target model is the gap between its corresponding dashed line and the baseline one in Figure 4. This allows us to compute values of $\tilde{\epsilon}$ reported in Table 3. It is also interesting to note



(a) c_{lb} as the size of the training set increases for baseline (while the test set size is 10k), trained to classify real (WikiText) and synthetic data points. The vertical dashed line indicates the number of samples we used to train the baseline in our main results, reported in Table 3.

(b) c_{lb} as the size of the training set increases for baseline, while the test set size is changing at the same time. The number of test samples for a given train set size is specified by the total number of samples we have for the audit (D_{in}), which is limited by the number of real members.

Figure 9: Baseline Strength Evaluation on WikiText dataset as the train set size increases. In Figure 9(a), the test set size is fixed, while in Figure 9(b) it varies as well. See Appendix D.1 for more details.

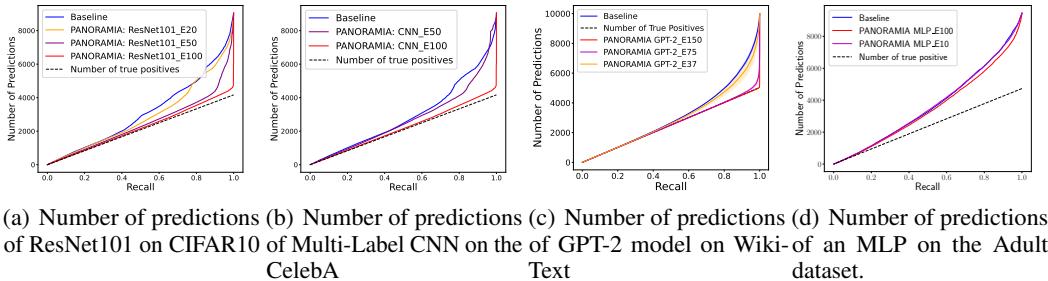


Figure 10: Comparison of the number of true positives and predictions on different datasets

that the maximum value of $\{c+\epsilon\}_{lb}$ typically occurs at low recall. Even if we do not use the same metric (precision-recall as opposed to TPR at low FPR) this is coherent with the findings from [Carlini et al. \(2022b\)](#). We can detect more privacy leakage (higher $\{c+\epsilon\}_{lb}$ which leads to higher $\bar{\epsilon}$) when making a few confident guesses about membership rather than trying to maximize the number of members found (*i.e.*, recall). Figure 10, decomposes the precision in terms of the number of true positives and the number of predictions for a direct mapping to the propositions 1 and 2. These are non-negligible values of privacy leakage, even though the true value is likely much higher.

D.3 PANORAMIA performance with increasing Test and Train Size m

PANORAMIA: Increasing Test Set Size: We compare PANORAMIA with the O(1) audit ([Steinke et al., 2023](#)) (in input space, without canaries), which uses a loss threshold MIA. As shown in Figure 6(b) and Figure 6(c) in the CIFAR10 image dataset case, both for non-private and Differentially Private (DP) models, PANORAMIA can leverage much more data for its measurement (up to the whole training set size for training and testing the MIA), while in our setting O(1) is limited by the size of the test set (for non-members). We vary test size from $m = 500$, up to $35k$. The experiment is repeated 10 times with resampling of the test set each time to produce the average. Unless otherwise stated, all results are produced at a 95% confidence.

D.4 PANORAMIA with Increasing Train Dataset Size

We also conduct experiments to show the effect of increasing training dataset size on the performance of PANORAMIA (as well as the baseline). In the CIFAR10 image dataset case, Figure 11, we see PANORAMIA performance slowly increase as the train dataset size increases. It reaches its max performance at dataset sizes 4500 – 5000, 10k, 20k – 30k and slowly begins to stabilize at the higher dataset sizes. We vary train size from $m = 500$, up to 30k. The experiment is repeated 5 times with resampling of the train set each time to produce the average. Unless otherwise stated, all results are produced at a 95% confidence.

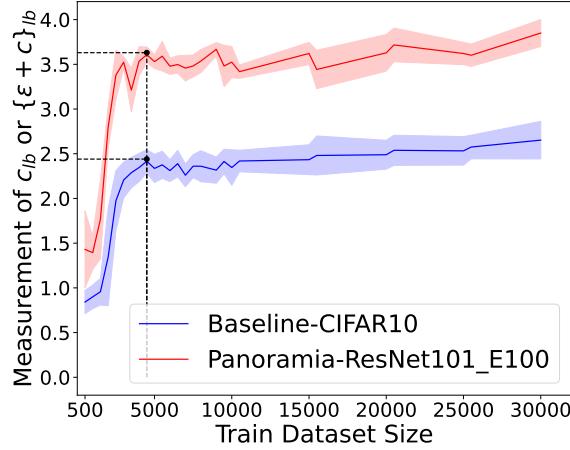


Figure 11: Effect on PANORAMIA and baseline with increasing train set size for CIFAR10, ResNet101 target model (trained for 100 epochs) (with union bound correction).

D.5 PANORAMIA in cases of no privacy leakage

As illustrated in Figure 4(c) and Table 3, there exist instances in which the c_{lb} is as large as $\{c+\epsilon\}_{lb}$. For instance, for GPT-2_E37, the upper bound in the confidence interval measured for c_{lb} is larger than the lower bound of the measurement for $\{c+\epsilon\}_{lb}$. Although the measured $\tilde{\epsilon}$ is positive on average, but the lower bound of the interval is negative. For such cases, we don't claim any detection of privacy leakage. This implies that, under certain conditions, the baseline could outperform the MIA in PANORAMIA. We can see three possible reasons for this behavior. (1) The synthetic data might not be of high enough quality. If there is an evident distinction between real and synthetic samples, the baseline cannot be fooled reducing the power of our audit. (2) The MIA in PANORAMIA is not powerful enough. (3) The target model's privacy leakage is indeed low. For GPT-2_E150 and GPT-2_E75 in Figure 4(c), we can claim that given the target model's leakage level, the synthetic data and the MIA can perform the audit to a non-negligible power. For GPT-2_37, we need to examine the possible reasons more closely. Focusing on generated data, in the ideal case we would like the synthetic data to behave just like real non-member data on Figure 14(a). If this was the case, $\{c+\epsilon\}_{lb}$ would go up (since the MIA uses this loss to separate members and non-members). Nonetheless, predicting changes in c_{lb} remains challenging, as analogous loss distributions do not guarantee indistinguishability between data points themselves (also, we cannot necessarily predict how the loss distribution changes under the helper model h in Figure 14(d)). Hence, reason (1) which corresponds to low generative data quality, seems to be a factor in the under-performance of PANORAMIA. Moreover, in the case GPT-2_37, we observe that real members and non-members have very similar loss distributions. Hence, factor (2) also seems to be at play, and a stronger MIA that does not rely exclusively on the plot may increase auditing power. Finally, the helper model h seems to have a more distinct loss distribution when compared to GPT-2_37 (see Figure 14(d)).

D.6 PANORAMIA with Tabular data

As we can see in Table 6, our proposed method PANORAMIA cannot detect any meaningful privacy leakage in the case of a small MLP classifier trained on tabular data.

	PANORAMIA $\tilde{\epsilon}$	$O(1) \epsilon$
MLP_E10	0 ± 0.142	0.655 ± 0.014
MLP_E100	0.062 ± 0.156	0.648 ± 0.015

Table 6: Privacy measurement with PANORAMIA and $O(1)$ for models trained on tabular Adult dataset.

As shown in Figure 12(b) there is no significant difference between the values $(c + \epsilon)_{lb}$ c_{lb} returned by PANORAMIA. As a result we get $\tilde{\epsilon} = 0$

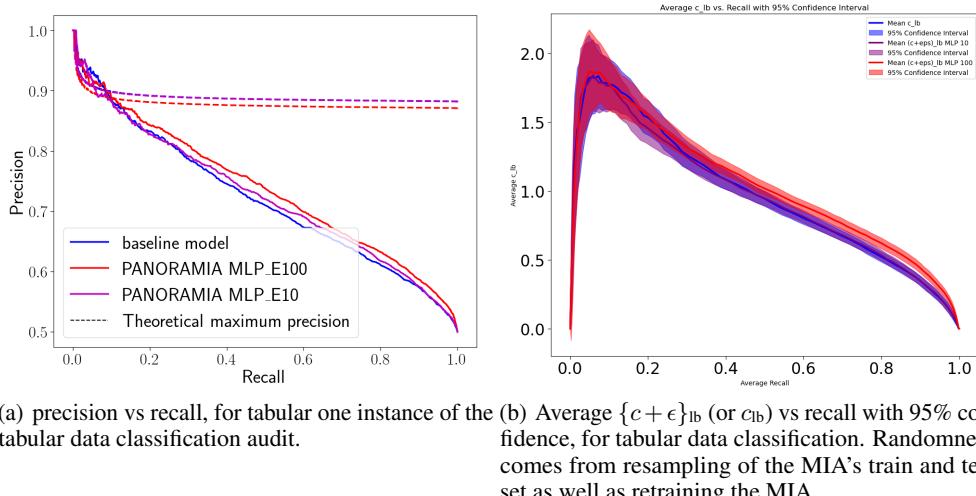


Figure 12: PANORAMIA on tabular Adult dataset, MLP target model.

The high precision at low recall in the baseline as can be seen in Figure 12(a) illustrates the fact that some real data can easily be distinguished from the synthetic ones. Instead of poor quality synthetic data, we are dealing with real observation that are actually “looking to real”. Investigating those observations reveals that they have extreme or rare values in one of their features. Indeed, the Adult dataset contains some extreme values (e.g., in the capital_gain column), such extreme values are rarely reproduced by the generative model. This leads to the baseline accurately predicting such observations as real with high confidence that translates into a high c_{lb} . Moreover, the audit values maximum values $(c + \epsilon)_{lb}$ and c_{lb} being achieved for a relatively small number of guesses corresponding to extreme values, the audits results vary a lot depending on the sampling or not of such extreme values in the audit test set. We believe this is part of the reason why we observe an high variance in the audit results leading to no significant privacy leakage detected by PANORAMIA.

This is a potential drawback of our auditing scheme, while the presence of extreme or rare values in the dataset usually allows for good auditing results it can have a negative impact on PANORAMIA’s success.

D.7 Privacy Auditing of Overfitted ML Models

Methodology. Varying the number of training epochs for the target model to induce overfitting is known to be a factor in privacy loss [Yeom et al. \(2018\)](#); [Carlini et al. \(2022a\)](#). As discussed in Section 5.2, since these different variants of target models share the same dataset and task, PANORAMIA can compare them in terms of privacy leaking.

To verify if PANORAMIA will indeed attribute a higher value of $\tilde{\epsilon}$ to more overfitted models, we train our target models for varying numbers of training epochs. The final train and test accuracies are reported in Table 1. For GPT-2 (as a target model), we report the train and validation loss as a measure of overfitting in Figure 8. Figures 13 and 14 show how the gap between member data points (*i.e.*, data used to train the target models) and non-member data points (both real as well as

generated non-members) increases as the degree of overfitting increases, in terms of loss distributions. We study the distribution of losses since these are the features extracted from the target model f or helper model h , to pass respectively to PANORAMIA and the baseline classifier. The fact that the loss distributions of member data become more separable from non-member data for more overfitted models is a sign that the corresponding target model could leak more information about its training data. We thus run PANORAMIA on each model, hereafter presenting the results obtained.

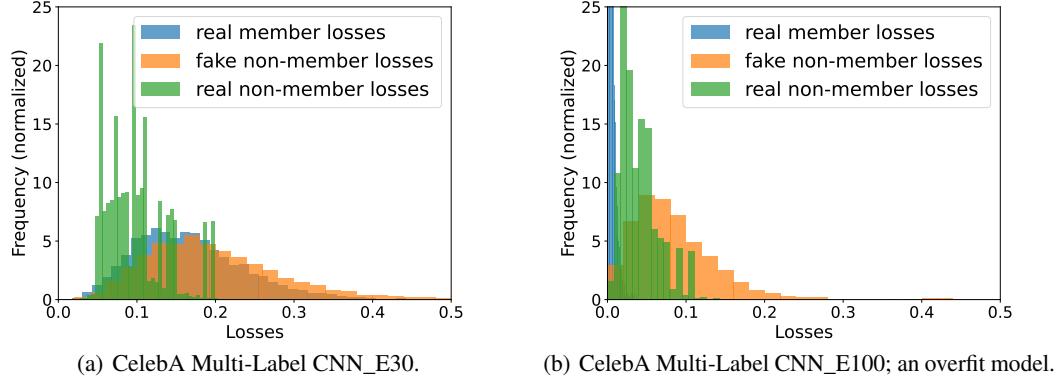


Figure 13: CelebA Multi-Label CNN Loss Comparisons for a generalized vs overfitted model.

Results. In Figure 3, we observe that more training epochs (*i.e.*, more overfitting) lead to better precision-recall trade-offs and higher maximum precision values. Our results are further confirmed by Figure 10 with PANORAMIA being able to capture the number of member data points (positive predictions) better than the baseline b .

In Table 3, we further demonstrate that our audit output $\tilde{\epsilon}$ orders the target models in terms of privacy leakage: higher the degree of overfitting, more memorization and hence a higher $\tilde{\epsilon}$ returned by PANORAMIA. From our experiments, we consistently found that as the number of epochs increased, the value of $\tilde{\epsilon}$ also increased. For instance, in the WikiText dataset, since we observe $\{c+\epsilon\}_{lb}^{GPT-2_E150} > \{c+\epsilon\}_{lb}^{GPT-2_E75} > \{c+\epsilon\}_{lb}^{GPT-2_E37}$, we would have $\{c+\epsilon\}_{lb}^{GPT-2_E150} - c_{lb} > \{c+\epsilon\}_{lb}^{GPT-2_E75} - c_{lb} > \{c+\epsilon\}_{lb}^{GPT-2_E37} - c_{lb}$, which leads to $\tilde{\epsilon}^{GPT-2_E150} > \tilde{\epsilon}^{GPT-2_E75} > \tilde{\epsilon}^{GPT-2_E37}$. Our experiment is coherent with the intuition that more training epochs lead to more overfitting, leading to more privacy leakage measured with a higher value of $\tilde{\epsilon}$.

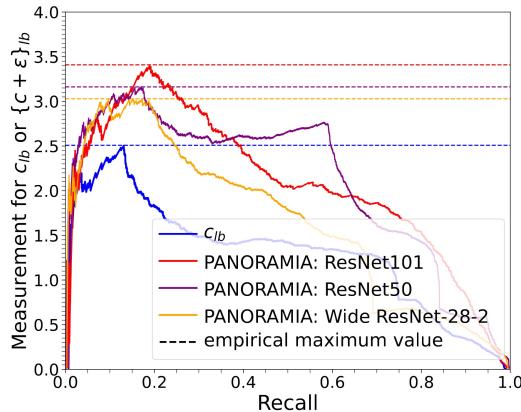


Figure 15: CIFAR-10, $\{c + \epsilon\}_{lb}$ when varying privacy leakage with increasing model parameters over one experiment run.

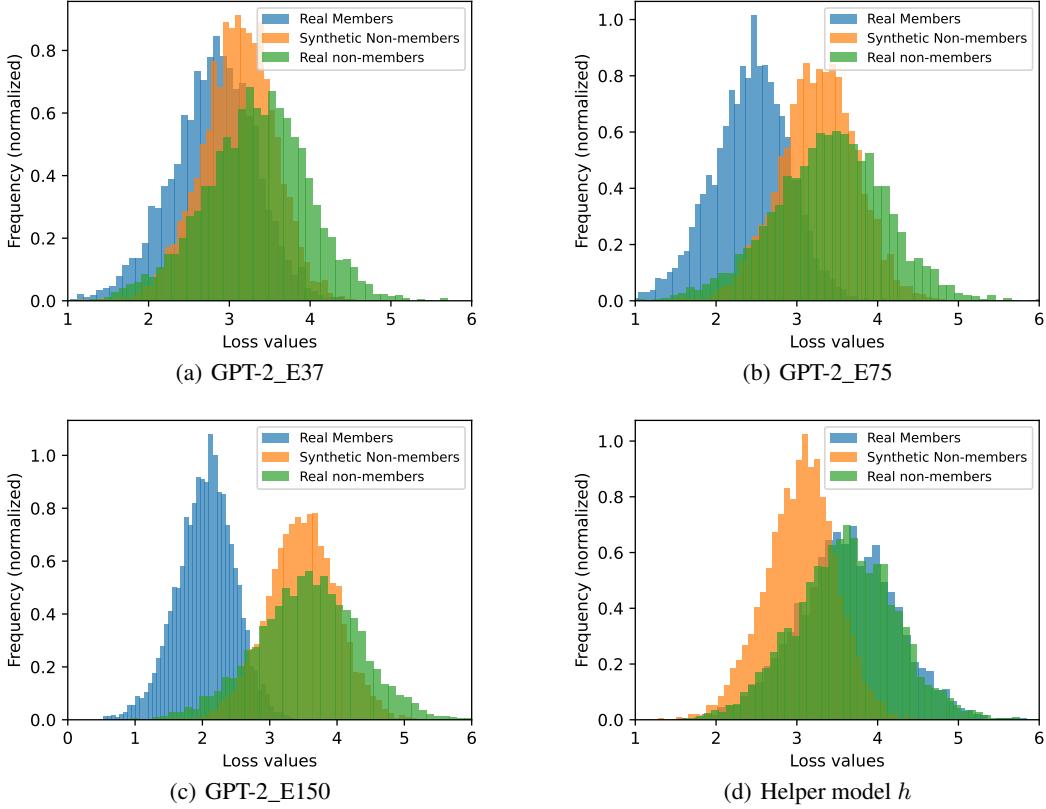


Figure 14: Comparison of the loss distributions of real members, real non-members, and synthetic non-members under three target models while varying the degree of over-training on the WikiText dataset. Figure 14(d) compares the loss distributions under the helper model, the model providing side information to our baseline. We train the helper with some other synthetic samples, which effectively mimic real non-members’ loss distributions under the target models. However, they are distinguishable to some extent from real non-members under the helper model, thus increasing our c_{lb} .

D.8 Models of varying complexity:

Carlini et al. (2021) have shown that larger models tend to have bigger privacy losses. To confirm this, we also audit ML models with varying numbers of parameters, from $\approx 4M$ parameters Wide ResNet-28-2, to a $25.5M$ parameters ResNet-50, and a $44.5M$ parameters ResNet-101. Figure 15 shows that PANORAMIA does detect increasing privacy leakage, with $\tilde{\epsilon}_{wide-resnet} \leq \tilde{\epsilon}_{resnet50} \leq \tilde{\epsilon}_{resnet101}$.

D.9 Privacy Auditing of Differentially-Private ML Models

Methodology.

For the image data modality, we evaluate the performance of PANORAMIA on differentially-private ResNet-18 and Wide-ResNet-16-4 models on the CIFAR10 dataset under different target privacy budgets (ϵ) with $\delta = 10^{-5}$ and the non-private ($\epsilon = \infty$) cases. The models are trained using the DP-SGD algorithm (Abadi et al., 2016) using Opacus (Yousefpour et al., 2021), which we tune for the highest train accuracy on learning rate lr , number of epochs e , batch size bs and maximum ℓ_2 clipping norm (C) for the largest final accuracy.

For the text data modality, we evaluate the performance of PANORAMIA on the differentially private GPT-2-Small model, which has been fine-tuned on the WikiText dataset (the preprocessing steps are identical to those in Appendix C.2). As with the image modality, we use various target privacy budgets (ϵ) with $\delta = 10^{-5}$, as well as non-private settings ($\epsilon = \infty$). We train the models using the DP-SGD algorithm Abadi et al. (2016), incorporating Ghost clipping Li et al. (2021) for calculating

the per-example gradient. We minimize the validation loss by tuning hyperparameters such as the learning rate (lr), number of epochs (e), batch size (bs), and clipping norm (C). The lowest validation loss values obtained are 3.79, 3.7, and 3.65 for ϵ values of 1, 3, and 10, respectively. For the non-private GPT-2-Small model, the lowest validation loss achieved is 3.341.

The noise multiplier σ is computed given ϵ , number of epochs and batch size. We use both PANORAMIA and O(1) (Steinke et al., 2023) privacy loss measurements under the pure ϵ -DP analysis.

Results. For the image data modality, tables 7 and 8 summarizes the auditing results of PANORAMIA on different DP models. For ResNet-18, we observe that at $\epsilon = 1, 2, 4, 6$ (more private models) PANORAMIA detects no privacy loss, whereas at higher values of $\epsilon = 10, 15, 20$ (less private models) and $\epsilon = \infty$ (a non-private model) PANORAMIA detects an increasing level of privacy loss with $\tilde{\epsilon}_{\epsilon=10} < \tilde{\epsilon}_{\epsilon=15} < \tilde{\epsilon}_{\epsilon=20} < \tilde{\epsilon}_{\epsilon=\infty}$, suggesting a higher value of ϵ correspond to higher $\tilde{\epsilon}$. We observe a similar pattern with Wide-ResNet-16-4, in which no privacy loss is detected at $\epsilon = 1, 2$ and higher privacy loss is detected at $\epsilon = 10, 15, 20, \infty$. We also compare the auditing performance of PANORAMIA with that of O(1) (Steinke et al., 2023), with the conclusion drawn by these two methods being comparable. For both ResNet-18 and Wide-ResNet-16-4, O(1) reports values close to 0 (almost a random guess between members and non-members) for $\epsilon < 10$ DP models, and higher values for $\epsilon = 10, 15, 20, \infty$ DP models. The results suggest that PANORAMIA is potentially an effective auditing tool for DP models that has comparable performance with O(1) and can generalize to different model structures.

For the text data modality, Table 9 summarizes the auditing results of PANORAMIA on differentially private (fine-tuned) Large Language Models. In all cases, neither PANORAMIA nor O(1) (Steinke et al. (2023) do not detect any privacy loss.

Target model	Audit	c_{lb}	$\epsilon + c_{lb}$	$\tilde{\epsilon}$	ϵ
ResNet18 $\epsilon = \infty$	PANORAMIA RM;GN O (1) RM;RN	2.44 -	3.87 -	1.43 -	- 1.471
ResNet18 $\epsilon = 20$	PANORAMIA RM;GN O (1) RM;RN	2.44 -	3.6331 -	1.19 -	- 1.34
ResNet18 $\epsilon = 15$	PANORAMIA RM;GN O (1) RM;RN	2.44 -	3.57 -	1.13 -	- 1.22
ResNet18 $\epsilon = 10$	PANORAMIA RM;GN O (1) RM;RN	2.44 -	2.70 -	0.26 -	- 0.28
ResNet18 $\epsilon = 6$	PANORAMIA RM;GN O (1) RM;RN	2.44 -	2.478 -	0.038 -	- 0.049
ResNet18 $\epsilon = 4$	PANORAMIA RM;GN O (1) RM;RN	2.44 -	1.99 -	0 -	- 0
ResNet18 $\epsilon = 2$	PANORAMIA RM;GN O (1) RM;RN	2.44 -	1.709 -	0 -	- 0.05
ResNet18 $\epsilon = 1$	PANORAMIA RM;GN O (1) RM;RN	2.44 -	1.412 -	0 -	- 0

Table 7: Privacy audit of ResNet18 under different values of ϵ -Differential Privacy using PANORAMIA and O(1) auditing frameworks, in which RM is for real member, RN for real non-member and GN for generated (synthetic) non-members.

Target model	Audit	c_{lb}	$\{\epsilon + c\}_{lb}$	$\tilde{\epsilon}$	ϵ
WRN-16-4 $\epsilon = \infty$	PANORAMIA RM;GN O (1) RM;RN	2.44 -	3.02 -	0.58 -	- 0.6408
WRN-16-4 $\epsilon = 20$	PANORAMIA RM;GN O (1) RM;RN	2.44 -	2.926 -	0.486 -	- 0.5961
WRN-16-4 $\epsilon = 15$	PANORAMIA RM;GN O (1) RM;RN	2.44 -	2.912 -	0.472 -	- 0.5774
WRN-16-4 $\epsilon = 10$	PANORAMIA RM;GN O (1) RM;RN	2.44 -	2.783 -	0.343 -	- 0.171
WRN-16-4 $\epsilon = 2$	PANORAMIA RM;GN O (1) RM;RN	2.44 -	2.36 -	0 -	- 0
WRN-16-4 $\epsilon = 1$	PANORAMIA RM;GN O (1) RM;RN	2.44 -	1.84 -	0 -	- 0

Table 8: Privacy audit of Wide ResNet16-4 under different values of ϵ -Differential Privacy (DP) using PANORAMIA and O(1) auditing frameworks, where RM is for real member, RN for real non-member and GN for generated (synthetic) non-members.

Target model f	Audit	Generator G	$\tilde{\epsilon}$	ϵ
∞ -DP GPT2-Small (GPT2_E37)	PANORAMIA O (1) RM;RN	GPT2-Small -	0.22 ± 0.37 -	- 2.82 ± 0.31
1-DP GPT2-Small	PANORAMIA O (1) RM;RN	GPT2-Small -	0 -	- 0
3-DP GPT2-Small	PANORAMIA O (1) RM;RN	GPT2-Small -	0 -	- 0
10-DP GPT2-Small	PANORAMIA O (1) RM;RN	GPT2-Small -	0 -	- 0

Table 9: Privacy audits on GPT-2-Small target models with varying privacy guarantees.

Discussion. In the image data modality, we observe that the auditing outcome ($\tilde{\epsilon}$ values for PANORAMIA and ϵ for O(1)) can be different for DP models with the same ϵ values (Table 10). We hypothesize that the auditing results may relate more to level of overfitting than the target ϵ values in trained DP models. The difference between train and test accuracies could be a possible indicator that has a stronger relationship with the auditing outcome. We also observe that O(1) shows a faster increase in ϵ for DP models with higher targeted ϵ values. We believe it depends on the actual ratio of the correct and total number of predicted samples, since O(1) considers both true positives and true negatives while PANORAMIA considers true positives only. We leave these questions for future work.

In the text data modality, the membership inference attacks (MIAs) relying on loss thresholds in privacy audits, without retraining or access to canaries, may not be sufficiently strong to effectively classify members and non-members.

Target model	Audit	c_{lb}	$\varepsilon + c_{lb}$	$\tilde{\varepsilon}$	ε	Train Acc	Test Acc	Diff(Train-Test Acc)
ResNet18-eps-20	PANORAMIA RM;GN	2.508	3.63	1.06	-	71.82	67.12	4.70
	O (1) RM;RN	-	-	-	1.22	-	-	-
	PANORAMIA RM;GN	2.508	2.28	0	-	71.78	68.08	3.70
	O (1) RM;RN	-	-	-	0.09	-	-	-
ResNet18-eps-15	PANORAMIA RM;GN	2.508	3.63	1.13	-	69.01	65.7	3.31
	O (1) RM;RN	-	-	-	1.34	-	-	-
	PANORAMIA RM;GN	2.508	1.61	0	-	66.68	69.30	2.62
	O (1) RM;RN	-	-	-	0.08	-	-	-

Table 10: DP models with the same ϵ values can have different auditing outcomes (the above numbers are reported here without union-bound correction).

E Generator Closeness γ -relaxation

Definition 4 $((c, \gamma)$ -closeness). *Let the γ -approximate maximum divergence between generative model \mathcal{G} and data distribution \mathcal{D} be,*

$$D_\infty^\gamma(\mathcal{D} \parallel \mathcal{G}) := \max_{S \subseteq \text{Supp}(\mathcal{D}): \mathbb{P}(\mathcal{D} \in S) \geq \gamma} \left[\ln \frac{\mathbb{P}(\mathcal{D} \in S) - \gamma}{\mathbb{P}(\mathcal{G} \in S)} \right].$$

For all $c > 0$ we say \mathcal{G} is (c, γ) -close to \mathcal{D} if $D_\infty^\gamma(\mathcal{D} \parallel \mathcal{G}) \leq c$, i.e.

$$\forall x \in \mathcal{X}, \mathbb{P}_{\mathcal{D}}[x] \leq e^c \mathbb{P}_{\mathcal{G}}[x] + \gamma \text{ or } e^{-c} (\mathbb{P}_{\mathcal{D}}[x] - \gamma) \leq \mathbb{P}_{\mathcal{G}}[x].$$

E.1 Pure DP case

First we prove a similar result as Lemma 5.5 (Steinke et al., 2023) but with the one-sided inequality of (c, γ) -closeness and the case when we additionally observe output from a pure-DP mechanism.

Lemma 1. *Let \mathcal{G}, \mathcal{D} be probability distributions over \mathcal{Y} , fix $c, \gamma \geq 0$. Suppose for all measurable $S \subset \mathcal{Y}$ we have $e^{-c}(\mathcal{D}(S) - \gamma) \leq \mathcal{G}(S)$. Then $\exists \gamma' \in [0, \gamma]$ and $\mathcal{D}', \mathcal{D}'', \mathcal{G}', \mathcal{G}''$ such that the following properties are satisfied.*

- (1) \mathcal{D}, \mathcal{G} can be expressed as a convex combination $\mathcal{D} = (1 - \gamma')\mathcal{D}' + \gamma'\mathcal{D}'', \mathcal{G} = (1 - \gamma')\mathcal{G}' + \gamma'\mathcal{G}''$;
- (2) for all measurable $S \subset \mathcal{Y}$, we have $e^{-c}\mathcal{D}'(S) \leq \mathcal{G}'(S)$;
- (3) there exists measurable $S \subset \mathcal{Y}$ such that $\mathcal{D}''(S) = 1, \forall S' \subset S, \mathcal{D}(S') \geq \mathcal{G}(S')$.

Let M denote an arbitrary ϵ -DP mechanism which takes input Y sampled from distribution \mathcal{G} or \mathcal{D} and outputs f . Let \mathcal{P}, \mathcal{Q} be probability distributions over \mathcal{Z} which are distributions of $M(Y)$ when Y is sampled from \mathcal{D} and \mathcal{G} respectively. Fix $\epsilon \geq 0$, suppose for all measurable $T \subset \mathcal{Z}$ we have $\mathcal{P}(T) \leq e^\epsilon \mathcal{Q}(T)$ and $\mathcal{Q}(T) \leq e^\epsilon \mathcal{P}(T)$. When additionally observing the output from a DP mechanism we have the following properties. Let $(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})$ be the joint distribution over $(\mathcal{Y}, \mathcal{Z})$ which we assume $\mathcal{D} \perp\!\!\!\perp \mathcal{P}, \mathcal{G} \perp\!\!\!\perp \mathcal{Q}$.

- (4) $(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})$ can be expressed as a convex combination $\mathcal{D}\mathcal{P} = ((1 - \gamma')\mathcal{D}' + \gamma'\mathcal{D}'')\mathcal{P}, \mathcal{G}\mathcal{Q} = ((1 - \gamma')\mathcal{G}' + \gamma'\mathcal{G}'')\mathcal{Q}$;
- (5) for all measurable $S \subset \mathcal{Y}, T \subset \mathcal{Z}$, we have $e^{-c} e^{-\epsilon} \mathcal{D}'(S)\mathcal{P}(T) \leq \mathcal{G}'(S)\mathcal{Q}(T)$;
- (6) there exists measurable $S \subset \mathcal{Y}, T \subset \mathcal{Z}$ such that $\mathcal{D}''(S)\mathcal{P}(T) = 1, \forall S' \subset S, T' \subset T, \mathcal{D}(S')\mathcal{P}(T) \geq \mathcal{G}(S')\mathcal{P}(T)$.

Proof. For the edge cases of $\gamma = 0$, the results hold with $\gamma' = 0, \mathcal{D}' = \mathcal{D}, \mathcal{G}' = \mathcal{G}$; similarly when $\gamma = 1$, the results hold with $\gamma' = 1, \mathcal{D}'' = \mathcal{D}, \mathcal{G}'' = \mathcal{G}, \mathcal{D}' = \mathcal{G}'$. Let $c' \in [0, c], \gamma_1, \gamma_2 \in (0, 1)$, define distribution $\mathcal{G}', \mathcal{G}'', \mathcal{D}', \mathcal{D}''$ as follows. For all points $y \in \mathcal{Y}$,

$$\begin{aligned}\mathcal{D}'(y) &= \frac{\min\{\mathcal{D}(y), e^{c'} \mathcal{G}(y)\}}{1 - \gamma_1}, \\ \mathcal{D}''(y) &= \frac{\mathcal{D}(y) - (1 - \gamma_1)\mathcal{D}'(y)}{\gamma_1} = \frac{\max\{0, \mathcal{D}(y) - e^{c'} \mathcal{G}(y)\}}{\gamma_1}, \\ \mathcal{G}'(y) &= \frac{\mathcal{G}(y)}{1 - \gamma_2}, \\ \mathcal{G}''(y) &= \frac{\mathcal{G}(y) - (1 - \gamma_2)\mathcal{G}'(y)}{\gamma_2} = 0.\end{aligned}$$

By construction $(1 - \gamma_1)\mathcal{D}' + \gamma_1\mathcal{D}'' = \mathcal{D}, (1 - \gamma_2)\mathcal{G}' + \gamma_2\mathcal{G}'' = \mathcal{G}$, so the first property is satisfied, and $\mathcal{D}''(y)$ is supported on $S = \{y \in \mathcal{Y} : \mathcal{D}(y) > e^{c'} \mathcal{G}(y)\}$ so the third property is implied. If $0 < \gamma_1 = \gamma_2 = \gamma' \leq \gamma$, for all $y \in \mathcal{Y}$ we have,

$$\frac{\mathcal{G}'(y)}{\mathcal{D}'(y)} = \frac{\mathcal{G}(y)}{\min\{\mathcal{D}(y), e^{c'} \mathcal{G}(y)\}} \geq e^{-c'} \geq e^{-c},$$

as required for the second property. Following the same as in Lemma 5.5 (Steinke et al., 2023), we can ensure $0 < \gamma_1 = \gamma_2 \leq \gamma$ by appropriately setting $\epsilon_1, \epsilon_2 \in [0, \epsilon]$. We can use the same decomposition of \mathcal{D}, \mathcal{G} to prove the second part of the Lemma. If assuming $\mathcal{D} \perp\!\!\!\perp \mathcal{P}$ and $\mathcal{G} \perp\!\!\!\perp \mathcal{Q}$, the joint distributions can be decomposed as $f_{\mathcal{D}, \mathcal{P}} = f_{\mathcal{P}} \cdot f_{\mathcal{D}}$ and $f_{\mathcal{G}, \mathcal{Q}} = f_{\mathcal{Q}} \cdot f_{\mathcal{G}}$. Therefore by construction, $\mathcal{D}\mathcal{P} = ((1 - \gamma')\mathcal{D}' + \gamma'\mathcal{D}'')\mathcal{P}, \mathcal{G}\mathcal{Q} = ((1 - \gamma')\mathcal{G}' + \gamma'\mathcal{G}'')\mathcal{Q}$. $\mathcal{D}''\mathcal{P}$ is supported on $(S, T) = \{y \in \mathcal{Y}, M(y) \in \mathcal{Z} : \mathcal{D}(y)\mathcal{P}(M(y)) > e^{c'} \mathcal{G}(y)\mathcal{P}(M(y))\}$. For all $y \in \mathcal{Y}, M(y) \in \mathcal{Z}$ we have,

$$\frac{\mathcal{G}'(y)}{\mathcal{D}'(y)} \cdot \frac{\mathcal{Q}(M(y))}{\mathcal{P}(M(y))} = \frac{\mathcal{G}(y)}{\min\{\mathcal{D}(y), e^{c'} \mathcal{G}(y)\}} \cdot \frac{\mathcal{Q}(M(y))}{\mathcal{P}(M(y))} \geq e^{-c'} e^{-\epsilon} \geq e^{-c} e^{-\epsilon},$$

where $\frac{\mathcal{Q}(M(y))}{\mathcal{P}(M(y))} \geq e^{-\epsilon}$ since M is an ϵ -DP mechanism. \square

Similar to Lemma 5.6 (Steinke et al., 2023), next we prove a Bayesian version of Lemma 1 which is used to prove the main result afterwards.

Lemma 2. Let \mathcal{G}, \mathcal{D} be probability distributions over \mathcal{Y} , fix $c, \gamma \geq 0$, suppose for all measurable $S \subset \mathcal{Y}$ we have $e^{-c}(\mathcal{D}(S) - \gamma) \leq \mathcal{G}(S)$. Then there exists a randomized function $E_{\mathcal{D}, \mathcal{G}} : \mathcal{Y} \rightarrow \{0, 1\}$ with the following properties. Suppose $X \sim \text{Bernoulli}(\frac{1}{2})$, if $X = 1$ sample $Y \sim \mathcal{D}$, and if $X = 0$ sample $Y \sim \mathcal{G}$. Then for all $y \in \mathcal{Y}$ we have,

$$\begin{aligned}\mathbb{P}_{\substack{X \sim \text{Bernoulli}(\frac{1}{2}) \\ Y \leftarrow X\mathcal{D} + (1-X)\mathcal{G}}} [X = 1, E_{\mathcal{D}, \mathcal{G}}(Y) = 1 | Y = y] &\leq \frac{e^c}{1 + e^c}, \\ \mathbb{E}_{Y \sim \mathcal{D}}[E_{\mathcal{D}, \mathcal{G}}(Y)] &\geq 1 - \gamma, \quad \mathbb{E}_{Y \sim \mathcal{G}}[E_{\mathcal{D}, \mathcal{G}}(Y)] \geq 1 - \gamma.\end{aligned}$$

Let M denote an arbitrary ϵ -DP mechanism which takes input Y sampled from distribution \mathcal{G} or \mathcal{D} and outputs f . Let \mathcal{P}, \mathcal{Q} be probability distributions over \mathcal{Z} which are distributions of $M(Y)$ when Y is sampled from \mathcal{D} and \mathcal{G} respectively. Fix $\epsilon \geq 0$, suppose for all measurable $T \subset \mathcal{Z}$ we have $\mathcal{P}(T) \leq e^\epsilon \mathcal{Q}(T)$ and $\mathcal{Q}(T) \leq e^\epsilon \mathcal{P}(T)$. Let $(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})$ denote the joint distributions over $(\mathcal{Y}, \mathcal{Z})$ which we assume $\mathcal{D} \perp\!\!\!\perp \mathcal{P}$, $\mathcal{G} \perp\!\!\!\perp \mathcal{Q}$. Then there exists a randomized function $E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})} : (\mathcal{Y}, \mathcal{Z}) \rightarrow \{0, 1\}$ such that for all $y \in \mathcal{Y}$, $M(y) \in \mathcal{Z}$ we have,

$$\begin{aligned} \mathbb{P}_{\substack{X \sim \text{Bernoulli}(\frac{1}{2}) \\ Y \leftarrow X\mathcal{D} + (1-X)\mathcal{G} \\ M(Y) \leftarrow M(X\mathcal{D} + (1-X)\mathcal{G})}} [X = 1, E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}(Y, M(Y)) = 1 | Y = y, M(Y) = f] &\leq \frac{e^{c+\epsilon}}{1 + e^{c+\epsilon}}, \\ \mathbb{E}_{Y \sim \mathcal{D}, M(Y) \sim \mathcal{P}} [E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}(Y, M(Y))] &\geq 1 - \gamma, \\ \mathbb{E}_{Y \sim \mathcal{G}, M(Y) \sim \mathcal{Q}} [E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}(Y, M(Y))] &\geq 1 - \gamma. \end{aligned}$$

Proof. We apply the decomposition from Lemma 1 and denote $\mathcal{D}', \mathcal{D}'', \mathcal{G}', \mathcal{G}'', \gamma'$ as is. We define $E_{\mathcal{D}, \mathcal{G}} : \mathcal{Y} \rightarrow \{0, 1\}$ by,

$$\mathbb{P}[E_{\mathcal{D}, \mathcal{G}}(y) = 1] = (1 - \gamma') \frac{\mathcal{D}'(y)}{\mathcal{D}(y)} = 1 - \frac{\gamma' \mathcal{D}''(y)}{\mathcal{D}(y)}.$$

For any $y \in \mathcal{Y}$ we have,

$$\begin{aligned} \mathbb{P}[X = 1, E_{\mathcal{D}, \mathcal{G}}(Y) = 1 | Y = y] &= \mathbb{P}[X = 1 | Y = y, E_{\mathcal{D}, \mathcal{G}}(Y) = 1] \mathbb{P}[E_{\mathcal{D}, \mathcal{G}}(Y) = 1 | Y = y] \\ &= \mathbb{P}[X = 1 | Y = y] \mathbb{P}[E_{\mathcal{D}, \mathcal{G}}(y) = 1] \\ &= \frac{\mathbb{P}[Y = y | X = 1] \mathbb{P}[X = 1]}{\mathbb{P}[Y = y]} \mathbb{P}[E_{\mathcal{D}, \mathcal{G}}(y) = 1] \\ &= \frac{\mathcal{D}(y)}{\mathcal{D}(y) + \mathcal{G}(y)} \mathbb{P}[E_{\mathcal{D}, \mathcal{G}}(y) = 1] \\ &= \frac{(1 - \gamma') \mathcal{D}'(y) + \gamma' \mathcal{D}''(y)}{(1 - \gamma') \mathcal{D}'(y) + \gamma' \mathcal{D}''(y) + (1 - \gamma') \mathcal{G}'(y)} \mathbb{P}[E_{\mathcal{G}, \mathcal{D}}(y) = 1] \\ &= \frac{1 + \frac{\gamma' \mathcal{D}''(y)}{(1 - \gamma') \mathcal{D}'(y)}}{1 + \frac{\gamma' \mathcal{D}''(y)}{(1 - \gamma') \mathcal{D}'(y)} + \frac{\mathcal{G}'(y)}{\mathcal{D}'(y)}} \mathbb{P}[E_{\mathcal{G}, \mathcal{D}}(y) = 1] \\ &\leq \frac{1 + \frac{\gamma' \mathcal{D}''(y)}{(1 - \gamma') \mathcal{D}'(y)}}{1 + 0 + e^{-c}} \mathbb{P}[E_{\mathcal{G}, \mathcal{D}}(y) = 1] \\ &= \frac{1}{1 + e^{-c}} \cdot \left(\frac{(1 - \gamma') \mathcal{D}'(y) + \gamma' \mathcal{D}''(y)}{(1 - \gamma') \mathcal{D}'(y)} \right) \cdot \mathbb{P}[E_{\mathcal{G}, \mathcal{D}}(y) = 1] \\ &= \frac{1}{1 + e^{-c}} \cdot \left(\frac{\mathcal{D}(y)}{(1 - \gamma') \mathcal{D}'(y)} \right) \cdot \mathbb{P}[E_{\mathcal{G}, \mathcal{D}}(y) = 1] \\ &= \frac{1}{1 + e^{-c}} = \frac{e^c}{1 + e^c}. \end{aligned}$$

The first equality uses that $\mathbb{P}[A, B | C] = \mathbb{P}[A | B, C] \mathbb{P}[B | C]$. The second equality holds since we assume the internal randomness of E is independent of everything else. The third equality uses the Bayes' rule. The fourth equality uses the definition of the randomized function E . The fifth equality uses the decomposition of Lemma 1. After rearranging and use the properties from Lemma 1 we obtain the first inequality, then the rest follows from simplifications. Furthermore we have the

following,

$$\begin{aligned}
\mathbb{E}_{Y \sim \mathcal{D}}[E_{\mathcal{D}, \mathcal{G}}] &= \int_{\mathcal{Y}} \mathcal{D}(y) \mathbb{P}[E_{\mathcal{D}, \mathcal{G}} = 1] dy \\
&= \int_{\mathcal{Y}} (1 - \gamma') \mathcal{D}'(y) dy = 1 - \gamma' \geq 1 - \gamma, \\
\mathbb{E}_{Y \sim \mathcal{G}}[E_{\mathcal{D}, \mathcal{G}}] &= 1 - \gamma' \mathbb{E}_{Y \sim \mathcal{G}} \left[\frac{\mathcal{D}''(y)}{\mathcal{D}(y)} \right] \\
&= 1 - \gamma' \int_{\mathcal{Y}} \frac{\mathcal{G}(y)}{\mathcal{D}(y)} \cdot \mathcal{D}''(y) dy \\
&\geq 1 - \gamma' \int_{\mathcal{Y}} \mathcal{D}''(y) dy \\
&= 1 - \gamma' \geq 1 - \gamma.
\end{aligned}$$

We follow a similar procedure to prove the second part of the Lemma. We define $E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})} : (\mathcal{Y}, \mathcal{Z}) \rightarrow \{0, 1\}$ by,

$$\mathbb{P}[E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}(y, M(y)) = 1] = \mathbb{P}[E_{\mathcal{D}, \mathcal{G}}(y) = 1]$$

With similar reasoning, for any $y \in \mathcal{Y}, M(y) \in \mathcal{Z}$ we have,

$$\begin{aligned}
&\mathbb{P}[X = 1, E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}(Y, M(Y)) = 1 | Y = y, M(Y) = f] \\
&= \frac{\mathbb{P}[Y = y, M(Y) = f | X = 1] \mathbb{P}[X = 1]}{\mathbb{P}[Y = y, M(Y) = f]} \mathbb{P}[E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}(y, M(y)) = 1] \\
&= \frac{\mathcal{D}(y) \mathcal{P}(M(y))}{\mathcal{D}(y) \mathcal{P}(M(y)) + \mathcal{G}(y) \mathcal{Q}(M(y))} \mathbb{P}[E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}(y, M(y)) = 1] \\
&= \frac{((1 - \gamma') \mathcal{D}'(y) + \gamma' \mathcal{D}''(y)) \mathcal{P}(M(y))}{((1 - \gamma') \mathcal{D}'(y) + \gamma' \mathcal{D}''(y)) \mathcal{P}(M(y)) + (1 - \gamma') \mathcal{G}'(y) \mathcal{Q}(M(y))} \mathbb{P}[E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}(y, M(y)) = 1] \\
&= \frac{1 + \frac{\gamma' \mathcal{D}''(y) \mathcal{P}(M(y))}{(1 - \gamma') \mathcal{D}'(y) \mathcal{P}(M(y))}}{1 + \frac{\gamma' \mathcal{D}''(y) \mathcal{P}(M(y))}{(1 - \gamma') \mathcal{D}'(y) \mathcal{P}(M(y))} + \frac{\mathcal{G}'(y) \mathcal{Q}(M(y))}{\mathcal{D}'(y) \mathcal{P}(M(y))}} \mathbb{P}[E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}(y, M(y)) = 1] \\
&\leq \frac{1 + \frac{\gamma' \mathcal{D}''(y)}{(1 - \gamma') \mathcal{D}'(y)}}{1 + 0 + e^{-c} e^{-\epsilon}} \mathbb{P}[E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}(y, M(y)) = 1] \\
&= \frac{1}{1 + e^{-c} e^{-\epsilon}} \cdot \left(\frac{\mathcal{D}(y)}{(1 - \gamma') \mathcal{D}'(y)} \right) \cdot \mathbb{P}[E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}(y, M(y)) = 1] \\
&= \frac{1}{1 + e^{-c} e^{-\epsilon}} = \frac{e^{c+\epsilon}}{1 + e^{c+\epsilon}}.
\end{aligned}$$

Since we define the success probability with the same expression, the expected success of the randomized function directly follows,

$$\begin{aligned}
\mathbb{E}_{Y \sim \mathcal{D}, M(Y) \sim \mathcal{P}}[E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}(Y, M(Y))] &= \mathbb{E}_{Y \sim \mathcal{D}}[E_{\mathcal{D}, \mathcal{G}}(Y)] \geq 1 - \gamma, \\
\mathbb{E}_{Y \sim \mathcal{G}, M(Y) \sim \mathcal{Q}}[E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}(Y, M(Y))] &= \mathbb{E}_{Y \sim \mathcal{G}}[E_{\mathcal{D}, \mathcal{G}}(Y)] \geq 1 - \gamma.
\end{aligned}$$

□

Next we use Lemma 2 to prove the following Proposition which is used to form statistical tests on hypothesis \mathcal{H} .

Proposition 3. *Let \mathcal{G} be (c, γ) -close, f be output of an ϵ -DP mechanism, $T^b \triangleq B(S, X)$ be the guess from the baseline, $T^a \triangleq A(S, X, f)$ be the guess from the membership audit. Let $T_b, T_a \in [0, 1]^m$ be bounded. Then, for all $v \in \mathbb{R}$ and all t in the support of T :*

$$\begin{aligned}
\mathbb{P}_{S, X, T^b} \left[\sum_{i=1}^m T_i^b \cdot S_i \geq v \mid T^b = t^b \right] &\leq \mathbb{P}_{S' \sim \text{Bernoulli}(\frac{e^c}{1+e^c})^m, F} \left[F(t^b) + \sum_{i=1}^m t_i^b \cdot S'_i \geq v \right], \\
\mathbb{P}_{S, X, T^a} \left[\sum_{i=1}^m T_i^a \cdot S_i \geq v \mid T^a = t^a \right] &\leq \mathbb{P}_{S' \sim \text{Bernoulli}(\frac{e^{c+\epsilon}}{1+e^{c+\epsilon}})^m, F} \left[F(t^a) + \sum_{i=1}^m t_i^a \cdot S'_i \geq v \right],
\end{aligned}$$

where F is independent from S' , $F(T^b), F(T^a)$ is supported on $\{0, 1, \dots, m\}$ and $\mathbb{E}_{T^b, F}[F(T^b)] = \mathbb{E}_{T^a, F}[F(T^a)] \leq 2m\gamma$.

Proof. Let $B(s_{\leq i})$ denote the distribution on $[0, 1]^m$ obtained by conditioning B on past $S_{< i} = s_{< i}$ and $X_{\leq i} = x_{\leq i}$. By Lemma 2, for all $t^b \in [0, 1]^m$ we have

$$\begin{aligned} \mathbb{P}[S_i = 1, E_{B(s_{< i, 1}), B(s_{< i, 0})}(T^b) = 1 | T^b = t^b, S_{< i} = s_{< i}, X_{\leq i} = x_{\leq i}] &\leq \frac{e^c}{1 + e^c}, \\ \mathbb{E}[E_{B(s_{\leq i, 1}), B(s_{\leq i, 0})}(T^b) = 1 | S_{< i} = s_{< i}, S_i = 0] &\geq 1 - \gamma, \\ \mathbb{E}[E_{B(s_{\leq i, 1}), B(s_{\leq i, 0})}(T^b) = 1 | S_{< i} = s_{< i}, S_i = 1] &\geq 1 - \gamma. \end{aligned}$$

Using the law of total probability we get,

$$\begin{aligned} &\mathbb{P}[S_i = 1, E_{B(s_{< i, 1}), B(s_{< i, 0})}(T^b) = 1 | T^b = t^b, S_{< i} = s_{< i}] \\ &= \sum_{x \leq i} \mathbb{P}[S_i = 1, E_{B(s_{< i, 1}), B(s_{< i, 0})}(T^b) = 1 | T^b = t^b, S_{< i} = s_{< i}, X_{\leq i} = x_{\leq i}] \\ &\quad \mathbb{P}[X_{\leq i} = x_{\leq i} | T^b = t^b, S_{< i} = s_{< i}] \\ &\leq \frac{e^c}{1 + e^c} \sum_{x \leq i} \mathbb{P}[X_{\leq i} = x_{\leq i} | T^b = t^b, S_{< i} = s_{< i}] = \frac{e^c}{1 + e^c}. \end{aligned}$$

Applying the union bound on S_i we can bound the expectation of success events by,

$$\mathbb{E}[E_{B(s_{\leq i, 1}), B(s_{\leq i, 0})}(T^b) | S_{< i} = s_{< i}, S_i = s_i] \geq 1 - 2\gamma.$$

By induction, for $k \in [m]$ define $\tilde{W}_k(s, t^b) := \sum_{i=1}^k t_i^b \cdot s_i \cdot E_{B(s_{\leq i, 1}), B(s_{\leq i, 0})}(t^b)$, then $\tilde{W}_k(s, t^b)$ is stochastically dominated by $\check{W}_k(t^b) := \sum_{i=1}^k t_i^b \cdot \check{S}_i(t)$ where $\check{S}_i(t) \sim \text{Bernoulli}(\frac{e^c}{1+e^c})$. Let the failing events F be,

$$F(s, t^b) := \sum_{i=1}^m \mathbb{1}\{E_{B(s_{\leq i, 1}), B(s_{\leq i, 0})}(t^b) = 0\},$$

then we have $W_m(s, t^b) := \sum_{i=1}^m t_i^b \cdot s_i = \tilde{W}_m(s, t^b) + F(s, t^b)$ is stochastically dominated by $\check{W}_m(T^b) + F(S, T^b)$, where the expectation of the failing events is bounded by,

$$\mathbb{E}[F(S, T^b)] = \sum_{i=1}^m \mathbb{P}[E_{B(s_{\leq i, 1}), B(s_{\leq i, 0})}(t^b) = 0] \leq m(2\gamma).$$

Similar to Lemma 5.5 in Steinke et al. (2023) since $\mathbb{E}[F(S, T^b)]$ is independent of S we could have $F(T^b) = F(S, T^b)$ for S drawn from an appropriate distribution. The proof of the second part of the Proposition is essentially the same except with the following result from Lemma 2,

$$\begin{aligned} \mathbb{P}[S_i = 1, E_{A(s_{< i, 1}), A(s_{< i, 0})}(T^a) = 1 | T^a = t^a, S_{< i} = s_{< i}, X_{\leq i} = x_{\leq i}] &\leq \frac{e^{c+\epsilon}}{1 + e^{c+\epsilon}}, \\ \mathbb{E}[E_{A(s_{\leq i, 1}), A(s_{\leq i, 0})}(T^a) = 1 | S_{< i} = s_{< i}, S_i = 0] &\geq 1 - \gamma, \\ \mathbb{E}[E_{A(s_{\leq i, 1}), A(s_{\leq i, 0})}(T^a) = 1 | S_{< i} = s_{< i}, S_i = 1] &\geq 1 - \gamma. \end{aligned}$$

□

Similar to Theorem 5.2 (Steinke et al., 2023) we could compute for the optimal distribution $F(t^b)$ and $F(t^a)$ by formulating it into a linear programming problem. Similar to Corollary 1, to test Hypothesis \mathcal{H} together we would apply Union bound on the two parts of the hypothesis as in Proposition 3.

E.2 Approximate DP case

In this section we consider the case where we have a (c, γ) -close generator and audit for an approximate DP mechanism. We modify the results to account for the failing events from these two sources of randomness. We start by proving a similar version of Lemma 2 where we use both the decomposition of \mathcal{G}, \mathcal{D} from Lemma 1 and the decomposition of \mathcal{P}, \mathcal{Q} from Lemma 5.5 (Steinke et al., 2023).

Lemma 3. Let \mathcal{G}, \mathcal{D} be probability distributions over \mathcal{Y} , fix $c, \gamma \geq 0$, suppose for all measurable $S \subset \mathcal{Y}$ we have $e^{-c}(\mathcal{D}(S) - \gamma) \leq \mathcal{G}(S)$. Let M denote an arbitrary (ϵ, δ) -DP mechanism which takes input Y sampled from distribution \mathcal{G} or \mathcal{D} and outputs f . Let \mathcal{P}, \mathcal{Q} be probability distributions over \mathcal{Z} which are distributions of $M(Y)$ when Y is sampled from \mathcal{D} and \mathcal{G} respectively. Fix $\epsilon, \delta \geq 0$, suppose for all measurable $T \subset \mathcal{Z}$ we have $\mathcal{P}(T) \leq e^\epsilon \mathcal{Q}(T) + \delta$ and $\mathcal{Q}(T) \leq e^\epsilon \mathcal{P}(T) + \delta$. Let $(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})$ denote the joint distributions over $(\mathcal{Y}, \mathcal{Z})$ which we assume $\mathcal{D} \perp\!\!\!\perp \mathcal{P}$, $\mathcal{G} \perp\!\!\!\perp \mathcal{Q}$. Then there exists a randomized function $E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})} : (\mathcal{Y}, \mathcal{Z}) \rightarrow \{0, 1\}$ such that for all $y \in \mathcal{Y}$, $M(y) \in \mathcal{Z}$ we have,

$$\begin{aligned} & \mathbb{P}_{\substack{X \sim \text{Bernoulli}(\frac{1}{2}) \\ Y \leftarrow X\mathcal{D} + (1-X)\mathcal{G} \\ M(Y) \leftarrow M(X\mathcal{D} + (1-X)\mathcal{G})}} [X = 1, E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}(Y, M(Y)) = 1 | Y = y, M(Y) = f] \leq \frac{e^{c+\epsilon}}{1 + e^{c+\epsilon}}, \\ & \mathbb{E}_{Y \sim \mathcal{D}, M(Y) \sim \mathcal{P}} [E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}(Y, M(Y))] \geq (1 - \gamma)(1 - \delta), \\ & \mathbb{E}_{Y \sim \mathcal{G}, M(Y) \sim \mathcal{Q}} [E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}(Y, M(Y))] \geq (1 - \gamma)(1 - \delta). \end{aligned}$$

Proof. For ease of notation we shorthand $F = M(Y)$, $E = E_{(\mathcal{D}, \mathcal{P}), (\mathcal{G}, \mathcal{Q})}$, $\mathcal{D}' = \mathcal{D}'(y)$, $\mathcal{G}' = \mathcal{G}'(y)$, $\mathcal{D}'' = \mathcal{D}''(y)$, $\mathcal{G}'' = \mathcal{G}''(y)$, $\mathcal{P}' = \mathcal{P}'(f)$, $\mathcal{Q}' = \mathcal{Q}'(f)$, $\mathcal{P}'' = \mathcal{P}''(f)$, $\mathcal{Q}'' = \mathcal{Q}''(f)$ below. We define $E : (\mathcal{Y}, \mathcal{Z}) \rightarrow \{0, 1\}$ by,

$$\mathbb{P}[E(y, f) = 1] = \frac{(1 - \gamma')(1 - \delta')\mathcal{D}'\mathcal{P}'}{\mathcal{D}\mathcal{P}}.$$

For any $y \in \mathcal{Y}$, $M(y) \in \mathcal{Z}$ we have,

$$\begin{aligned} & \mathbb{P}[X = 1, E(Y, F) = 1 | Y = y, F = f] \\ &= \frac{\mathbb{P}[Y = y, F = f | X = 1]\mathbb{P}[X = 1]}{\mathbb{P}[Y = y, F = f]}\mathbb{P}[E(y, f) = 1] \\ &= \frac{\mathcal{D}(y)\mathcal{P}(f)}{\mathcal{D}(y)\mathcal{P}(f) + \mathcal{G}(y)\mathcal{Q}(f)}\mathbb{P}[E(y, f) = 1] \\ &= \frac{((1 - \gamma')\mathcal{D}' + \gamma'\mathcal{D}'')((1 - \delta')\mathcal{P}' + \delta'\mathcal{P}'')}{((1 - \gamma')\mathcal{D}' + \gamma'\mathcal{D}'')((1 - \delta')\mathcal{P}' + \delta'\mathcal{P}'') + (1 - \gamma')\mathcal{G}'((1 - \delta')\mathcal{Q}' + \delta'\mathcal{Q}'')} \mathbb{P}[E(y, f) = 1] \\ &= \frac{1 + \frac{\delta'\mathcal{D}'\mathcal{P}''}{(1 - \delta')\mathcal{D}'\mathcal{P}'} + \frac{\gamma'\mathcal{D}''\mathcal{P}'}{(1 - \gamma')\mathcal{D}'\mathcal{P}'} + \frac{\gamma'\delta'\mathcal{D}''\mathcal{P}''}{(1 - \gamma')(1 - \delta')\mathcal{D}'\mathcal{P}'}}{1 + \frac{\delta'\mathcal{D}'\mathcal{P}''}{(1 - \delta')\mathcal{D}'\mathcal{P}'} + \frac{\gamma'\mathcal{D}''\mathcal{P}'}{(1 - \gamma')\mathcal{D}'\mathcal{P}'} + \frac{\gamma'\delta'\mathcal{D}''\mathcal{P}''}{(1 - \gamma')(1 - \delta')\mathcal{D}'\mathcal{P}'} + \frac{\mathcal{G}'\mathcal{Q}'}{\mathcal{D}'\mathcal{P}'} + \frac{\delta'\mathcal{G}'\mathcal{Q}''}{(1 - \delta')\mathcal{D}'\mathcal{P}'}} \mathbb{P}[E(y, f) = 1] \\ &\leq \frac{1 + \frac{\delta'\mathcal{D}'\mathcal{P}''}{(1 - \delta')\mathcal{D}'\mathcal{P}'} + \frac{\gamma'\mathcal{D}''\mathcal{P}'}{(1 - \gamma')\mathcal{D}'\mathcal{P}'} + \frac{\gamma'\delta'\mathcal{D}''\mathcal{P}''}{(1 - \gamma')(1 - \delta')\mathcal{D}'\mathcal{P}'}}{1 + e^{-c}e^{-\epsilon}} \mathbb{P}[E(y, f) = 1] \\ &= \frac{1}{1 + e^{-c}e^{-\epsilon}} \left(\frac{(1 - \delta')(1 - \gamma')\mathcal{D}'\mathcal{P}' + \delta'(1 - \gamma')\mathcal{D}'\mathcal{P}'' + \gamma'(1 - \delta')\mathcal{D}''\mathcal{P}' + \gamma'\delta'\mathcal{D}''\mathcal{P}''}{(1 - \gamma')(1 - \delta')\mathcal{D}'\mathcal{P}'} \right) \\ &\quad \cdot \mathbb{P}[E(y, f) = 1] \\ &= \frac{1}{1 + e^{-c}e^{-\epsilon}} \left(\frac{\mathcal{D}\mathcal{P}}{(1 - \gamma')(1 - \delta')\mathcal{D}'\mathcal{P}'} \right) \mathbb{P}[E(y, f) = 1] \\ &= \frac{1}{1 + e^{-c}e^{-\epsilon}} = \frac{e^{c+\epsilon}}{1 + e^{c+\epsilon}}. \end{aligned}$$

Using the assumption that $\mathcal{D} \perp\!\!\!\perp \mathcal{P}$, $\mathcal{G} \perp\!\!\!\perp \mathcal{Q}$, we account for the expectation of success events as follows,

$$\begin{aligned} \mathbb{E}_{Y \sim \mathcal{D}, M(Y) \sim \mathcal{P}} [E(Y, F)] &= \iint \mathcal{D}(y)\mathcal{P}(f)\mathbb{P}[E = 1] dy df \\ &= \iint \mathcal{D}(y)\mathcal{P}(f) \frac{(1 - \gamma')(1 - \delta')\mathcal{D}'(y)\mathcal{P}'(f)}{\mathcal{D}(y)\mathcal{P}(f)} dy df \\ &= (1 - \gamma')(1 - \delta') \int \mathcal{D}'(y) dy \int \mathcal{P}'(f) df \\ &= (1 - \gamma')(1 - \delta') \geq (1 - \gamma)(1 - \delta). \end{aligned}$$

Note that we can rewrite the success probability as,

$$\mathbb{P}[E(y, f) = 1] = 1 - \frac{\gamma'(1 - \delta')\mathcal{D}''\mathcal{P}'}{\mathcal{D}\mathcal{P}} - \frac{(1 - \gamma')\delta'\mathcal{D}'\mathcal{P}''}{\mathcal{D}\mathcal{P}} - \frac{\gamma'\delta'\mathcal{D}''\mathcal{P}''}{\mathcal{D}\mathcal{P}}.$$

Then we have,

$$\begin{aligned} & \mathbb{E}_{Y \sim \mathcal{G}, M(Y) \sim \mathcal{Q}}[E(Y, F)] \\ &= 1 - \gamma'(1 - \delta')\mathbb{E}\left[\frac{\mathcal{D}''(y)\mathcal{P}'(f)}{\mathcal{D}(y)\mathcal{P}(f)}\right] - (1 - \gamma')\delta'\mathbb{E}\left[\frac{\mathcal{D}'(y)\mathcal{P}''(f)}{\mathcal{D}(y)\mathcal{P}(f)}\right] - \gamma'\delta'\mathbb{E}\left[\frac{\mathcal{D}''(y)\mathcal{P}''(f)}{\mathcal{D}(y)\mathcal{P}(f)}\right] \\ &= 1 - \gamma'(1 - \delta') \int \mathcal{G}(y)\mathcal{Q}(f) \left[\frac{\mathcal{D}''(y)\mathcal{P}'(f)}{\mathcal{D}(y)\mathcal{P}(f)} \right] dy df \\ &\quad - (1 - \gamma')\delta' \int \mathcal{G}(y)\mathcal{Q}(f) \left[\frac{\mathcal{D}'(y)\mathcal{P}''(f)}{\mathcal{D}(y)\mathcal{P}(f)} \right] dy df \\ &\quad - \gamma'\delta' \int \mathcal{G}(y)\mathcal{Q}(f) \left[\frac{\mathcal{D}''(y)\mathcal{P}''(f)}{\mathcal{D}(y)\mathcal{P}(f)} \right] dy df \\ &\geq 1 - \gamma'(1 - \delta') \int \mathcal{D}''(y)\mathcal{P}'(f) dy df - (1 - \gamma')\delta' \int \mathcal{D}'(y)\mathcal{P}''(f) dy df \\ &\quad - \gamma'\delta' \int \mathcal{D}''(y)\mathcal{P}''(f) dy df \\ &= 1 - \gamma'(1 - \delta') - (1 - \gamma')\delta' - \gamma'\delta' \\ &= 1 - \gamma' - \delta' + \gamma'\delta' \\ &= (1 - \gamma')(1 - \delta') \geq (1 - \gamma)(1 - \delta). \end{aligned}$$

□

Next we use Lemma 3 to prove the following Proposition to form the statistical tests on hypothesis \mathcal{H} for the (ϵ, δ) -DP version.

Proposition 4. *Let \mathcal{G} be (c, γ) -close, f be output of an (ϵ, δ) -DP mechanism, $T^b \triangleq B(S, X)$ be the guess from the baseline, $T^a \triangleq A(S, X, f)$ be the guess from the membership audit. Let $T_b, T_a \in [0, 1]^m$ be bounded. Then, for all $v \in \mathbb{R}$ and all t in the support of T :*

$$\begin{aligned} \mathbb{P}_{S, X, T^b} \left[\sum_{i=1}^m T_i^b \cdot S_i \geq v \mid T^b = t^b \right] &\leq \mathbb{P}_{S' \sim \text{Bernoulli}(\frac{e^c}{1+e^c})^m, F} \left[F(t^b) + \sum_{i=1}^m t_i^b \cdot S'_i \geq v \right], \\ \mathbb{P}_{S, X, T^a} \left[\sum_{i=1}^m T_i^a \cdot S_i \geq v \mid T^a = t^a \right] &\leq \mathbb{P}_{S' \sim \text{Bernoulli}(\frac{e^{c+\epsilon}}{1+e^{c+\epsilon}})^m, F} \left[F(t^a) + \sum_{i=1}^m t_i^a \cdot S'_i \geq v \right], \end{aligned}$$

where F is independent from S' , $F(T^b), F(T^a)$ is supported on $\{0, 1, \dots, m\}$ and $\mathbb{E}_{T^b, F}[F(T^b)] \leq 2m\delta$, $\mathbb{E}_{T^a, F}[F(T^a)] \leq 2m(\gamma + \delta - \gamma\delta)$.

Proof. The first part of the result (baseline test) exactly follows the proof of Proposition 3. For the approximate DP case, the proof of the auditor test follows similar steps except for the bound on the expectation of the failing events,

$$\begin{aligned} \mathbb{E}[F(S, T^b)] &= \sum_{i=1}^m \mathbb{P}[E_{B(s_{\leq i, 1}), B(s_{\leq i, 0})}(t^b) = 0] \leq m(2(1 - (1 - \gamma)(1 - \delta))) \\ &= 2m(\gamma + \delta - \gamma\delta). \end{aligned}$$

□

Similar to the pure DP case we could optimize for the optimal distribution of $F(t^b)$ and $F(t^a)$ by solving a linear program (Theorem 5.2 in [Steinke et al. \(2023\)](#)) and apply Union bound as in Corollary 1 to test for Hypothesis \mathcal{H} .

E.3 Evaluations

We empirically evaluate the auditing performance of PANORAMIA under different relaxed level of the generator for both the pure and approximate DP cases. We observe that the overall performance is similar to that of no relaxation. We observe smaller $\tilde{\epsilon}$ (which is analogous to a looser lower bound) as we increase the relaxation level, and in the more extreme cases where we allow many failing events (e.g. when relaxation $> 1e-3$) we would fail to detect meaningful privacy leakage as $\tilde{\epsilon} = 0$.

Model	Generator Relaxation (γ)	c_{lb}	$\{\epsilon + c\}_{lb}$	$\tilde{\epsilon}$
ResNet18 $\epsilon = 2$	0 (no relaxation)	2.508	2.066	0
	1e-5	2.507	2.065	0
	1e-4	2.499	2.059	0
	1e-3	2.365	1.989	0
ResNet18 $\epsilon = 10$	0 (no relaxation)	2.508	2.833	0.325
	1e-5	2.507	2.801	0.294
	1e-4	2.499	2.570	0.071
	1e-3	2.365	1.280	0
ResNet18 $\epsilon = 15$	0 (no relaxation)	2.508	3.661	1.153
	1e-5	2.507	3.658	1.151
	1e-4	2.499	3.628	1.129
	1e-3	2.365	3.312	0.947
ResNet18 $\epsilon = 2, \delta = 1e-5$	1e-5	2.035	2.064	0.029
	1e-4	2.029	2.059	0.030
	1e-3	1.904	1.988	0.084
ResNet18 $\epsilon = 10, \delta = 1e-5$	1e-5	2.035	2.765	0.730
	1e-4	2.029	2.541	0.512
	1e-3	1.904	1.271	0
ResNet18 $\epsilon = 15, \delta = 1e-5$	1e-5	2.035	3.654	1.619
	1e-4	2.029	3.625	1.596
	1e-3	1.904	3.308	1.404

Table 11: Privacy audit of ResNet18 under different values of ϵ -DP and (ϵ, δ) -DP, using PANORAMIA with different level of generator relaxations (γ), testing over range of recall from 0 to 0.5 without applying a union bound.

F Extended Related Work

Most related privacy auditing works measure the privacy of an ML model by lower-bounding its privacy loss. This usually requires altering the training pipeline of the ML model, either by injecting canaries that act as outliers Carlini et al. (2019) or by using data poisoning attack mechanisms to search for worst-case memorization Jagielski et al. (2020); Nasr et al. (2021). MIAs are also increasingly used in privacy auditing, to estimate the degree of memorization of member data by an ML algorithm by resampling the target algorithm \mathcal{M} to bound $\frac{P(M|in)}{P(M|out)}$ Jayaraman & Evans (2019). The auditing procedure usually involves searching for optimal neighboring datasets D, D' and sampling the DP outputs $\mathcal{M}(D), \mathcal{M}(D')$, to get a Monte Carlo estimate of ϵ . This approach raises important challenges. First, existing search methods for neighboring inputs, involving enumeration or symbolic search, are impossible to scale to large datasets, making it difficult to find optimal dataset pairs. In addition, Monte Carlo estimation requires up to thousands of costly model retrainings to bound ϵ with high confidence. Consequently, existing approaches for auditing ML models predominantly require the re-training of ML models for every (batch of) audit queries, which is computationally expensive in large-scale systems Jagielski et al. (2020); Zanella-Béguelin et al. (2022); Lu et al. (2023). This makes privacy auditing computationally expensive and gives an estimate by averaging over models, which might not reflect the true guarantee of a specific pipeline deployed in practice.

Nonetheless, improvements to auditing have been made in a variety of directions. For example, Nasr et al. (2023) and Maddock et al. (2023) have taken advantage of the iterative nature of DP-SGD, auditing individual steps to understand the privacy of the end-to-end algorithm. Andrew et al. (2023) leverage the fact that analyzing MIAs for non-member data does not require re-running the algorithm. Instead, it is possible to re-sample the non-member data point: if the data points are i.i.d. from an asymptotically Gaussian distribution with mean zero and variance $1/d$, this enables a closed-form analysis of the non-member case.

Recently, Steinke et al. (2023) proposed a novel scheme for auditing differential privacy with $O(1)$ training rounds. This approach enables privacy audits using multiple training examples from the

same model training, if examples are included in training independently (which requires control over the training phase, and altering the target model). They demonstrate the effectiveness of this new approach on DP-SGD, in which they achieve meaningful empirical privacy lower bounds by training only one model (the strongest results are achieved by including canaries in the training set though, which is not possible in our setting), whereas standard methods would require training hundreds of models. Our work builds closely on the theory from [Steinke et al. \(2023\)](#), but introduces a baseline model to account for distribution shifts. The key difference that enables our work to account for member/non-member distribution shifts lies in how we create the audit set in our privacy game (Definition 2). In [Steinke et al. \(2023\)](#), the audit set is fixed, and data points are randomly assigned to member or non-member by a Bernoulli random variable S . Members are actually used in training the target model f , while non-members are not (so assignment happens before training). In our framework, we take a set of known members (after the fact), and pair each point with a non-member (generated i.i.d. from the generator distribution). We then flip S to sample which data point of each pair will be shown to the “auditor” (MIA/baseline) for testing, thereby creating the test task of our privacy measurement. When we replace our generated data with in-distribution independent non-members (PANORAMIA RM;RN), we exactly enforce the same independence as [Steinke et al. \(2023\)](#) (and as a result have $c = 0$, and $\tilde{\epsilon}$ is a lower-bound on ϵ), except that we “waste” some member data by drawing our auditing game after the fact. The difference in analysis accounts for distribution shifts when our non-members are generated.

The field of MIA without direct connections to privacy leakage measurement has also seen a lot of recent activity, with new proposals to improve the strength of MIAs in various settings ([Carlini et al., 2022b](#); [Nasr et al., 2019](#); [Suri et al.](#); [Zarifzadeh et al., 2023](#)). While it is not the focus of this paper, an interesting avenue for future work is to port ideas from these new MIAs to improve PANORAMIA’s MIA and baseline. Such transfer of MIA progress could lead to more powerful measurements with PANORAMIA. More closely related to our proposal, a recent set of works on [Das et al. \(2024\)](#); [Duan et al. \(2024\)](#); [Meeus et al. \(2024\)](#) MIAs on foundation models has observed that current evaluations are limited due to a lack of availability of non-member data. Indeed, foundation models typical include all known data at the time of their training, and there is no well known public set of in-distribution non-member data. Evaluation tasks typical rely on more recent data, which is known to consist in non-members, but suffers from distribution shifts. [Das et al. \(2024\)](#); [Duan et al. \(2024\)](#); [Meeus et al. \(2024\)](#) all show that such a shift invalidates MIA evaluations by showing that “blind attacks” (i.e., our baselines) perform better than proposed MIAs. One could apply our proposed framework to such a setting: when the baseline in our framework performs better than the MIA, PANORAMIA returns a privacy loss measurement of zero instead of misleading measurements of MIA performance. We thus believe that PANORAMIA is an important step towards a theory for rigorous evaluations of membership inference for MIA models.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: **[Yes]**

Justification: The abstract and Introduction sections highlight the motivation and contributions we make; We provide, to the best of our ability, comprehensive experimental results in Section 5 to justify our claims and contributions. Section 4 provides the theory behind our experiments, whilst Appendix B details proofs.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: **[Yes]**

Justification: We discuss these in Section 4.2 and Section 6 of the paper.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: **[Yes]**

Justification: Yes, we present sections assumptions, theory, and formal proofs in 3, 4 and Appendix B, E provides respectively.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: **[Yes]**

Justification: We attach the code as supplementary material in a zip file and also include the main experimental details in Section 5 as well as a detailed experimental setup, hyperparameters, compute details, and model architecture details along with other full necessary information in Appendix C to ensure reproducibility. Section 3 also includes the algorithm as well as the system design diagram detailing our privacy measurement framework.

We also updated the paper and experiments according to the reviewers' helpful feedback.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: **[Yes]**

Justification: We specify the details of the code we release in the zip file we attach. The code is primarily the PANORAMIA privacy measurement pipeline. We include structured files specifying run instructions, licenses etc. Training details are within the code itself and also highlighted in Appendix C.

Now we also updated the paper with an open source github code link given under the experimental section of the main paper.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: **[Yes]**

Justification: Section 5 and Table 1 give key experimental details, whilst the full experimental setup is highlighted in Appendix C, as well as the zip file of code attached.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We repeated all experiments that involved training data resampling 5 times, and repeated experiments that involved test data resampling over 10 times. All our figures and plots in the core body of the paper report results with error bars, with all results reported at a 95% confidence interval. We also report the standard deviation in the tables in Section 5.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide information on compute resources for each modality in Appendix C. We also updated compute times to run the PANORAMIA pipeline in C for each data modality.

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have reviewed the NeurIPS Code of Ethics and conform to them in every respect.

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes].

Justification: We discuss these in Section 6; overall, our approach can only help model providers and owners to assess the privacy of risk of their data and is not applicable as a privacy attack. As a consequence, we believe that any impact from our proposed work is likely to be positive.

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer:[NA].

Justification: Our paper poses no such risks.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes].

Justification: All papers, code, and asset owners and authors have been properly cited to the best of our knowledge.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes].

Justification: We specify the details of the code we release in the zip file we attach. The code is primarily the PANORAMIA privacy measurement pipeline. We include structured files specifying run instructions, licenses, etc. Training details are within the code itself and also highlighted in Appendix C.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA] .

Justification: The paper does not involve crowdsourcing, nor research with human subjects.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer:[NA] .

Justification: The paper does not involve crowdsourcing nor research with human subjects.