Deduplicating Training Data Mitigates Privacy Risks in Language Models

Nikhil Kandpal [1]   Eric Wallace [2]   Colin Raffel [1]

Abstract

Past work has shown that large language models are susceptible to privacy attacks, where adversaries generate sequences from a trained model and detect which sequences are memorized from the training set. In this work, we show that the success of these attacks is largely due to duplication in commonly used web-scraped training sets. We first show that the rate at which language models regenerate training sequences is *superlinearly* related to a sequence's count in the training set. For instance, a sequence that is present 10 times in the training data is on average generated ~1000× more often than a sequence that is present only once. We next show that existing methods for detecting memorized sequences have near-chance accuracy on non-duplicated training sequences. Finally, we find that after applying methods to deduplicate training data, language models are considerably more secure against these types of privacy attacks. Taken together, our results motivate an increased focus on deduplication in privacy-sensitive applications and a reevaluation of the practicality of existing privacy attacks.

1. Introduction

Neural language models (LMs)—systems trained to predict the next-word in a sequence of text—have become fundamental building blocks for numerous NLP tasks and domains. The performance and generality of these models make it important to study the extent to which they maintain the privacy of their training data, because many of their applications involve training on private information (e.g., emails, health records, chat logs, and source code).

Unfortunately, when training on private data, LMs may memorize and leak information to adversaries. Past work has demonstrated the practicality of these so-called model
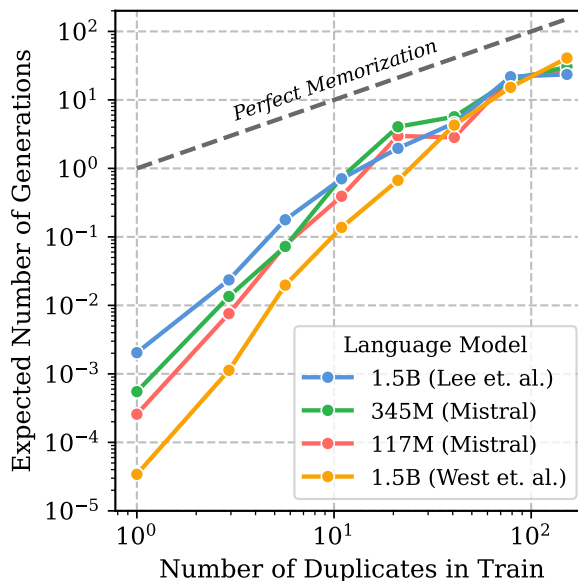

Figure 1. For a sequence duplicated d times in a language model's training dataset, we measure how often that sequence is expected to occur in a set of generated text that is equal in size to the training data. *Perfect Memorization* amounts to generating a sequence at the same frequency as it appears in the training data. All LMs tested show a superlinear increase in the expected number of generations (slopes > 1 on a log-log plot), i.e., training samples that are not duplicated are very rarely generated, whereas samples that are duplicated multiple times appear dramatically more frequently.

inversion attacks, which can successfully recover training data with only black-box access to a trained model (Carlini et al., 2019; 2021b; Inan et al., 2021). In particular, the strongest attack, proposed by Carlini et al. (2021b), recovers training data from LMs by first generating sequences from the models and then scoring those sequences with various membership inference methods. The highest-scoring sequences are classified as belonging to the training data.

In this work, we show that the success of the Carlini et al. (2021b) attack is largely due to duplicated sequences found in commonly used web-scraped training datasets. We study transformer LMs over various parameter scales and show that (1) the attack's likelihood of recovering a particular training sequence is correlated with the number of occurrences of that sequence in the training data, and (2) the
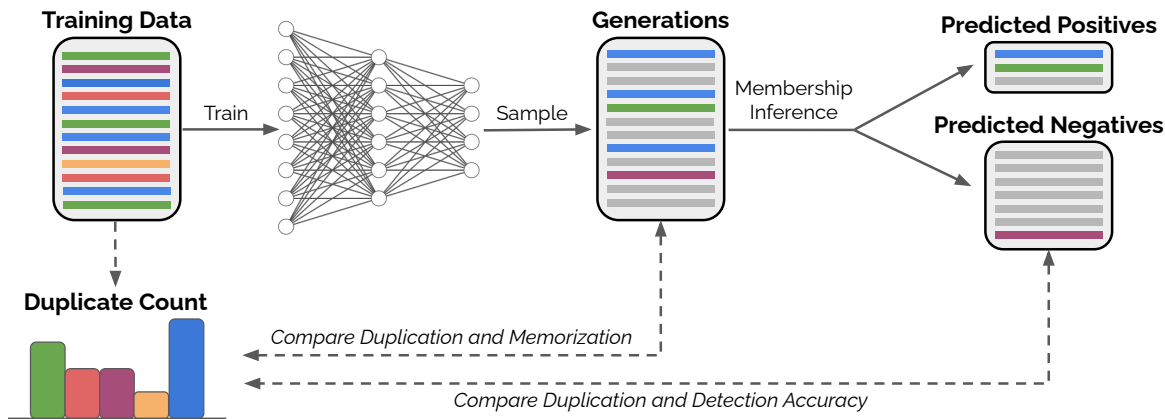
---

[1]UNC Chapel Hill [2]UC Berkeley. Correspondence to: Nikhil Kandpal <nkandpa2@cs.unc.edu>.

Figure 2. Overview of our analysis. Web-scraped text datasets that are used to train language models contain duplicated sequences, depicted in the figure as training data rows of the same color (*top left*). Model inversion attacks attempt to recover training data from a trained model by first generating large amounts of text, some of which is memorized training data (*top middle*). Membership inference is then performed to detect which generated sequences were copied from the training data (*top right*). Our analysis focuses on the relationship between the amount a sequence is duplicated in the training data and the effectiveness of the model inversion attack at generating and detecting that sequence (*bottom*).

overall attack effectiveness is reduced when sequence-level duplication in the training data is removed.

Concretely, we first show that the content that an LM generates is highly sensitive to sequence-level duplication in the training data. Using various sampling strategies, we generate text from LMs ranging from 117M-1.5B parameters. We consistently find a superlinear relationship between the number of times a sequence is duplicated in the training data and the rate at which that sequence is generated (e.g., Figure 1). For instance, a sequence that is present 10 times in the training data is on average generated $\sim1000\times$ more often than a sequence that is present only once. Notably, our results show that *samples which are not duplicated are very rarely regenerated by language models.*

We then look at the next stage of the model inversion attack: detecting memorized training data from a set of LM generations. We demonstrate that the membership inference methods from Carlini et al. (2021b) are correlated with the number of duplicates of a sequence in the training data. For example, the membership inference methods have an area under the ROC curve as high as 0.90 for sequences that are duplicated many times but *achieve only chance accuracy for sequences that appear once.*

In our final set of experiments, we directly test whether retraining LMs on deduplicated training datasets can mitigate privacy risks. We find that model inversion attacks are indeed much weaker for deduplicated models: *they emit $\sim20\times$ less training data and reduce the effectiveness of two of the three proposed membership inference methods.* All in all, our results underscore the need to carefully remove duplicates when training privacy-sensitive models and show

that past work may overestimate the effectiveness of LM privacy attacks when duplication is mitigated.

2. Background and Experimental Setup

Language Models Language models take as input a sequence of tokens and output a probability distribution over the next token. LMs are trained to maximize the likelihood of a corpus of text and can be used to generate text at test time by iteratively sampling from the next-token distribution. In practice, various strategies exist for sampling tokens, including random sampling, sampling from the top-k highest probability tokens (Fan et al., 2018), or sampling after using a temperature to sharpen the next-token distribution.

Memorization The concept of "memorization" refers to ways that a trained model stores and consequently leaks information about its training data. Multiple notions of memorization have been studied that vary in their definitions and assumptions (see Section 7 for further discussion). In this work we focus on generation-based memorization, where a generative model leaks information by generating exact samples from its training data (Carlini et al., 2019).

When studying generation-based memorization in LMs, we compare models' generation behavior with the expected behavior of a model that has perfectly fit the training data through memorization. This *perfect memorization* model only assigns non-zero probability to samples seen during training and sampling from the model is identical to uniformly sampling from the training data. The perfect memorization model serves as a positive control showing how far LMs are from fully memorizing their training data.
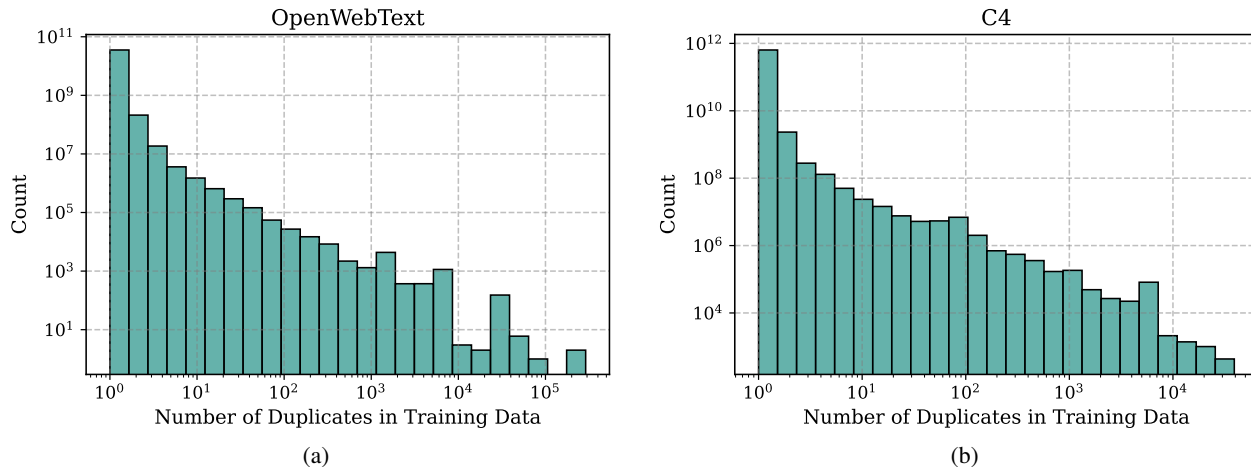
Figure 3. Web-scraped training sets are rife with duplicated sequences. Above, we plot the frequency of different amounts of duplication for 400-character sequences in the OpenWebText **(a)** and C4 **(b)** datasets. Note that C4 is an order of magnitude larger than OpenWebText.

Privacy Attacks The fact that state-of-the-art LMs memorize and regenerate sequences seen during training enables attacks that compromise the privacy of their training data (Carlini et al., 2019; Inan et al., 2021; Carlini et al., 2021b). In this work, we focus specifically on the Carlini et al. (2021b) attack, which is currently the strongest and most accessible model inversion attack on LMs. While we focus on this particular attack, our analysis also applies to other attacks that leverage generation-based memorization.

The Carlini et al. (2021b) attack works in two stages:

1. Generate a large amount of text from a language model.

2. Score the generated sequences using a membership inference scoring method.

For the first stage, Carlini et al. (2021b) study different methods of generating data (unconditional vs. conditional sampling, different sampling strategies, etc.). We focus on unconditional generation using standard sampling, top-k sampling, and temperature sampling.

For the second stage, we study all scores proposed by Carlini et al. (2021b). Each score is defined as the ratio between a metric estimating the "easiness" of sequence (a property of the sequence itself and not whether the sequence appears in the training dataset) and the trained model's perplexity on that sequence. For measures of easiness, Carlini et al. (2021b) used three choices:

- **Reference Model**: the perplexity of another LM on the sequence. We use the GPT-2 small language model (Radford et al., 2019).

- **zlib**: the length of the sequence after compression by the zlib compression library.

- **Lowercase**: the trained model's perplexity on the sequence with all lowercased characters.

Training Data Collection and Duplication Modern language modeling datasets are generated by large-scale scraping of the Internet (Gokaslan et al., 2019; Radford et al., 2019; Raffel et al., 2020; Gao et al., 2020). Most web-scraped datasets are deduplicated at the *document level*, e.g., if two web pages have the exact same contents, only one is kept in the data. Lee et al. (2021) observe that these datasets still have large-scale approximate and exact sequence-level duplication, e.g., quotes, paragraphs, and advertisements appearing in many web pages. To correct this, they propose efficient sequence-level deduplication methods based on locality sensitive hashing and suffix arrays.

When measuring duplication in the training data, we consider identical sequences to be duplicates. Although broader definitions such as approximate or semantic duplication may also be useful to study, we choose to investigate exact duplication in this work because it matches the adversary's goal of *exactly* recovering sequences from the training data (e.g., social security numbers). To detect duplicate sequences, we adapt the suffix array-based algorithm from Lee et al. (2021). Searching for exactly duplicated sequences in two sets of text can be done efficiently with a linear traversal of the two texts' suffix arrays.

Datasets In our experiments, we use models trained on the widely-used OpenWebText (Gokaslan et al., 2019) and C4 (Raffel et al., 2020) datasets. Both are large-scale datasets, 39GB and 750GB respectively, and were generated by scraping text from the Internet with basic filtering and deduplication. Despite deduplicating at the level of whole training examples, both datasets still contain a large number of duplicated token sequences between training examples,
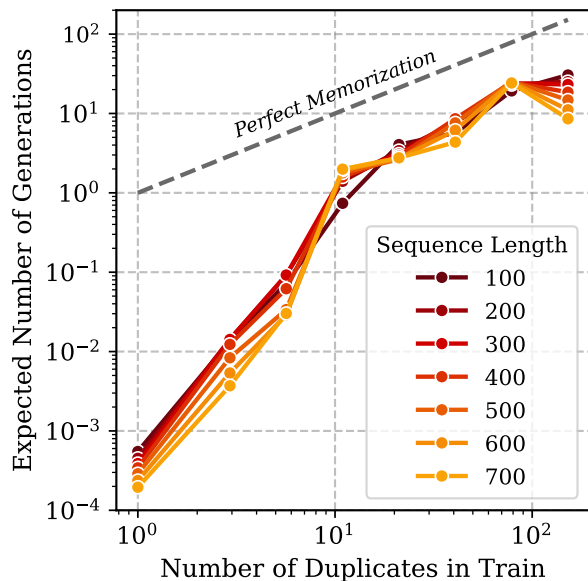
Figure 4. We vary the sequence length that is used when measuring whether a model generation overlaps with the training set. Using longer sequence lengths naturally reduces the chance that a generation exactly overlaps with the training set. However, the overall shape of the generation vs. duplication curve is consistent across a range of sequence lengths.

a property that is not unique to just these two datasets (Lee et al., 2021). To illustrate this quantitatively, in Figure 3 we show how often each unique 400-character sequence is duplicated in these two datasets. Both datasets contain millions of sequences that are duplicated 10 or more times, and some individual sequences are even duplicated tens of thousands of times. This large amount of sequence-level duplication allows us to reliably measure the effect of duplication on memorization and downstream model privacy over a wide range of duplication levels.

Models   We focus on Transformer-based (Vaswani et al., 2017) language models that range in scale from millions to billions of parameters. Specifically, we use the 117M and 345M parameter models from the Mistral project[1] and the 1.5B parameter forward language model from West et al. (2021), all of which were trained on the OpenWebText dataset. Additionally, we evaluate the two 1.5B parameter models from Lee et al. (2021), one trained on the C4 dataset and the other trained on a sequence-level deduplicated version of C4. We choose this set of models as they are near-state-of-the-art, and they allow us to test the effect of model scale, changes in codebase and implementation, optimization hyperparameters, and training data.

---

[1]https://github.com/stanford-crfm/mistral

Experimental Setup   Our experiments follow the analysis depicted in Figure 2. In Section 3 we analyze the likelihood of regenerating a training sample as a function of that sample's number of duplicates in the training data. In Section 4, we analyze the relationship between duplication and the detection of LM generations copied from the training data. The code used to perform our experiments can be found at https://github.com/nkandpa2/lm_memorization.

3. How Duplication Affects The Regeneration of Training Sequences

The first step of the Carlini et al. (2021b) attack is to generate a large pool of sequences in hopes that some are verbatim copies from the training data. In this section, we analyze how duplication in the training data affects this stage.

Concretely, we first record the number of duplicates for each N-length character sequence in the training data. We then generate many times from an LM and analyze how often each N-length training sequence is generated as a function of its duplicate count. Note that we also scale our calculations to simulate a scenario where we generate an amount of text equal in size to the training dataset. This allows us to directly compare the behavior of models trained on datasets of different sizes, and also compare to a theoretical model that has perfectly memorized its training data (i.e., generating from this model is identical to sampling from the training dataset).

3.1. Regeneration is Superlinearly Related to Duplicates

All models that we test have a superlinear relationship between the number of times a training sequence is regenerated and the number of times that sequence is duplicated in the training data. This relationship is shown in Figure 1 by the > 1 slope on a log-log plot.

Furthermore, Figure 1 shows that the generation behavior of LMs is far from perfect memorization: sequences duplicated d times in the training data are expected to be generated far fewer than d times by a trained model. This is especially true for low duplicate counts, i.e., *samples which are not duplicated are very rarely regenerated by language models*. This shows that the Carlini et al. (2021b) attack—which relies on models regenerating training samples—will rarely be able to extract training data that is not duplicated.

Our finding that LMs exhibit a superlinear increase in their regeneration rates is a also novel phenomenon worthy of future study. Concretely, one would expect that LMs would exhibit "calibrated" generation behavior—training sequences that appear twice as frequently are twice as likely to be generated—but this is not true for state-of-the-art models.
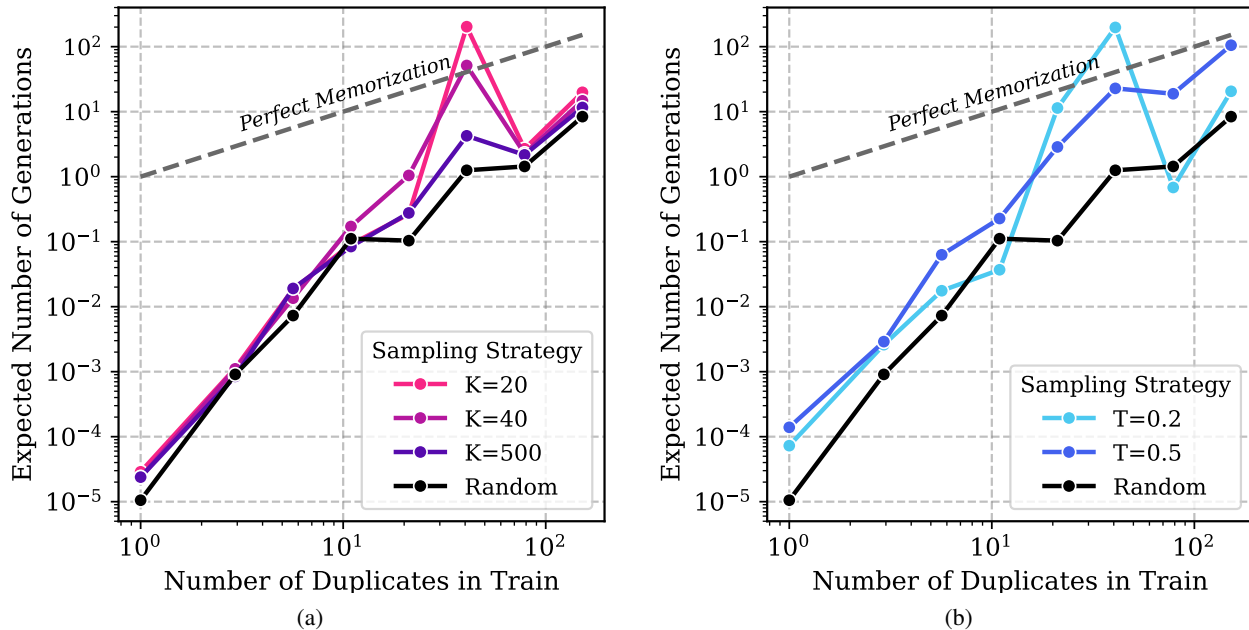
Figure 5. The sampling method impacts how often LMs regenerate training samples. Sampling methods that emit more likely sequences (e.g., top-k with smaller k or temperature sampling with smaller T) generate more verbatim training samples. Nevertheless, all sampling methods rarely generate training sequences when the number of duplicates is small.

3.2. Regeneration Trends Are Robust Across Experimental Setups

Having observed an initial superlinear trend, we next measure whether this relationship is a more general phenomenon that holds across different experimental setups varying the sequence length, model size, sampling strategy, and number of training epochs.

Effect of Duplicate Sequence Length Our analysis focuses on the duplication of N-length character sequences in the training data. To ensure that our conclusions are not dependent on any one choice of N, we vary the sequence length and report the results for the Mistral 345M parameter model in Figure 4. Using longer sequence lengths naturally reduces the chance that a generation exactly overlaps with the training set. Nevertheless, the superlinear relationship between generation and duplication is nearly identical across different sequence lengths. For the rest of the paper we set $N = 100$ characters unless otherwise specified.

Effect of Model Scale Larger models tend to regenerate more training data across all levels of duplication. This effect is shown by comparing the duplication curves for the 117M and 345M parameter Mistral models in Figure 1. These two models were trained nearly identically and thus the comparison controls for confounding factors such as the number of training steps and optimization hyperparameters. Larger models likely regenerate more training sequences

because they achieve a lower training loss (i.e., they assign training samples higher likelihoods on average).

Effect of Sampling Scheme The amount of regeneration depends on the sampling scheme used. Figure 5 compares random, top-k, and temperature sampling for the Mistral 117M parameter model. We find that sampling schemes that emit more likely sequences (e.g., top-k with smaller k) generate more verbatim training samples.

Effect of Increasing Epochs Finally, we find that the regeneration rate increases over the course of training. Figure 6 shows that as training progresses for the 117M parameter Mistral model, the regeneration rate of training sequences increases at nearly all levels of duplication. Notably, stopping early does not change the fact that language models generate disproportionately many highly-duplicated training sequences.

4. How Duplication Affects The Detection of Training Sequences

Thus far, we found that models rarely regenerate training sequences that are not duplicated many times. Nevertheless, the second stage of the Carlini et al. (2021b) attack, which looks to identify training sequences using membership inference methods, may be able to flag these rare cases. To test this, we evaluate the three membership inference scoring
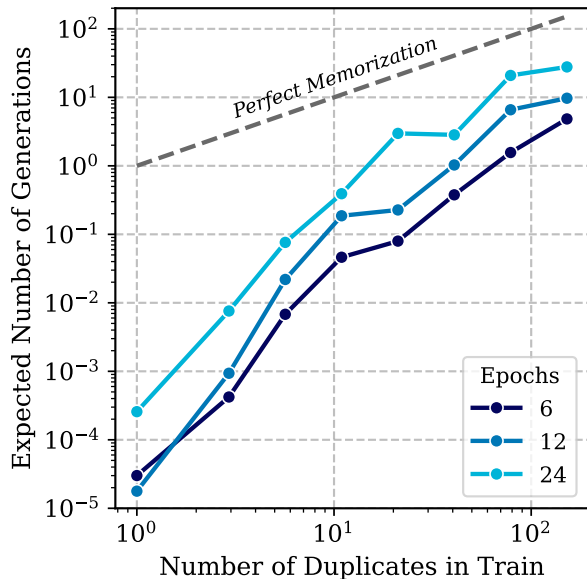
Figure 6. We plot the effect of performing multiple training epochs on the generation behavior. Performing additional epochs has a multiplicative effect that is uniform across all duplication levels. In particular, using twice as many epochs will cause the expected number of generations to increase by approximately 3 times for all duplication levels.

methods (`Reference Model`, `zlib`, and `Lowercase`) and stratify the results by different duplication levels.

Concretely, we bucket the samples generated from the 345M parameter Mistral model into sequences that are duplicated in the training data d times, for $d = 1$ to $d = 800$. We also collect a set of 25,000 negative sequences that were generated by the LM but were not in the training data. Using these two sets of samples, we measure the effectiveness of the three membership inference scores at distinguishing between the two sets.

Figure 7(a) shows the area under the Receiver Operating Characteristic (AUROC) curve achieved by the different membership inference scores. Notably, for generated sequences only found once in the training data, all three scores yield classifiers that are close to chance.[2] Of the three scores, the `Reference Model` is the highest performing classifier at nearly all levels of duplication.

Following the suggestions of Carlini et al. (2021a), we also evaluate membership inference using the True Positive Rate (TPR) at a very low False Positive Rate (FPR). This simulates the realistic evaluation setting where the prevalence of training data in a set of generated samples is very low compared to non-training data. We found that approximately 1 in 1000 generated 100-character spans are copied from the

---

[2]A random "no-skill classifier" has an AUROC of 0.50.

training data. Thus, we use a FPR of 0.1%.

In Figure 7(b), we show that the TPR of all three membership inference scores are highly correlated with the number of duplicates. For example, the TPR for the `Reference Model` method is as high as 0.60 for sequences that are duplicated many times but is only 0.10 for sequences that appear once. All in all, these results show that while membership inference methods may achieve non-trivial accuracies on average over the entire generated set, *most of their successes are on sequences that have been duplicated many times.*

5. Model Inversion with Deduplicated Data

In our final set of experiments, we directly test whether retraining LMs on deduplicated data can indeed mitigate privacy risks. In particular, we test two of the 1.5B parameter LMs from Lee et al. (2021), one trained on C4 and another trained on a deduplicated version of C4.[3]

Generating From Deduplicated Models   We first generate one million samples from each of the language models and measure the number of sequences copied from the training data. The top of Table 1 shows the number of unique 400-character training sequences generated by each of the language models (*Count*) and the percentage of all 400-character training sequences that are generated (*Percent*). Respectively, these measure the total amount of information from the training data leaked by each model and the probability of a single sequence in the training data being leaked. We find that the model trained on deduplicated data emits ∼20× less training data, i.e., deduplication strongly weakens the first stage of the Carlini et al. (2021b) attack.

Membership Inference   Next, we evaluate the performance of the membership inference scoring methods on the generated samples. We randomly subsample 25,000 sequences that are copied from the training data and 25,000 sequences that are novel from each of the models. All of these sequences are scored by the membership inference methods and we report the AUROC in the bottom of Table 1. We find that `zlib` and `Lowercase` are considerably affected by deduplication, whereas `Reference Model` performs almost equally as well on both models.

One factor to consider when comparing the AUROC scores between the normal and deduplicated models is that the set of memorized training samples that they are trying to detect are different. We hypothesize that in the rare circumstance

---

[3]We use the LM trained on the version of C4 that has exact duplicates removed using the suffix array-based EXACTSUBSTR method. This removes exact duplicates that were at least 50 byte-pair encoding (BPE) tokens long. To ensure that we do not attempt to recover sequences shorter than 50 tokens, we set $N = 400$ characters in this experiment.
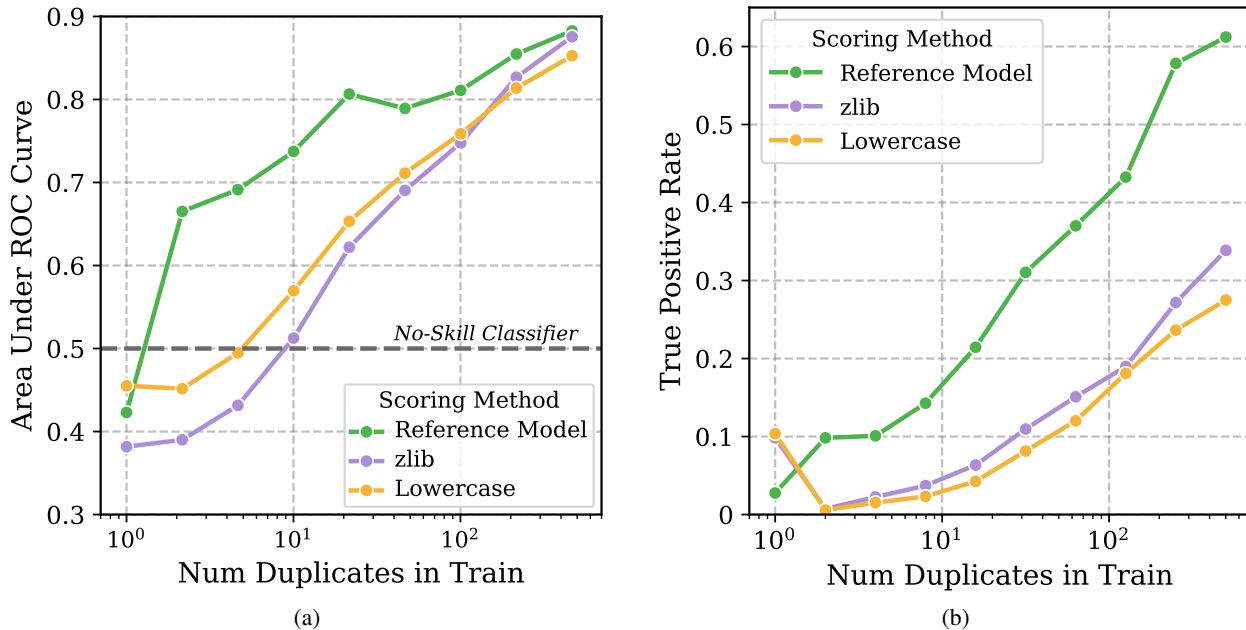
(a)   (b)

Figure 7. State-of-the-art membership inference methods fail to accurately detect training sequences when they are not duplicated in the training set. In **(a)**, we report the area under the ROC curve for different membership inference methods as a function of the number of duplicates. In **(b)**, we report the true positive rate at a false positive rate of 0.1%.

that the deduplicated model does indeed regenerate a training sample, those samples may be unique in some manner. This may make the samples easier to classify, which can explain why the AUROC can remain relatively high for the deduplicated model (0.87). Further investigation of the difference between the regenerations made by a normal and deduplicated model is worthy of future study.

Qualities of Effective Membership Inference Methods
The reasonable accuracy of the `Reference Model` method suggests that it measures training data leakage beyond just generation-based memorization. We hypothesize that this is due to its similarity to a different notion of memorization known as counterfactual memorization (Zhang et al., 2021). A sample's counterfactual memorization is measured by comparing the sample's expected likelihood under models that have and have not trained on that sample. The `Reference Model` method is an approximation of counterfactual memorization that uses a single model trained on different training data to approximate the expected likelihood under a model that has not trained on the sample being scored. While we find that duplication and generation-based memorization are highly correlated, this result suggests that approximating other notions of memorization, such as counterfactual memorization, may lead to membership inference scores that are less sensitive to deduplication. Similar findings have been noted in Watson et al. (2021) and Carlini et al. (2021a).

		Normal Model	Deduped Model
Training Data Generated	Count	1,427,212	68,090
	Percent	0.14	0.007
Mem. Inference AUROC	zlib	0.76	0.67
	Ref Model	0.88	0.87
	Lowercase	0.86	0.68

Table 1. Deduplicating training data drastically reduces the effectiveness of privacy attacks. We first generate 1 million 256-token samples from models trained on C4 and deduplicated C4. We then report the number of unique 400-character training sequences that are generated (*Count*) and the percentage of all 400-character training sequences that are generated (*Percent*). We then report the classification AUROC achieved by each of the three membership inference scores when applied to the generated sequences.

Is Deduplication An Effective Defense? Overall, our results show that deduplication is an effective safeguard against models *regenerating* their training data, which renders the first stage of many existing model inversion attacks largely ineffective. Fortunately, this defense comes at little-to-no cost in model performance, as training on deduplicated data does not harm language modeling perplexity (Lee et al., 2021). Nevertheless, in the rare cases when deduplicated models do generate training data, those samples can still be detected somewhat reliably by membership inference scores such as the `Reference Model` method.

6. Discussion

More General Notions of Duplication   We define duplicates as two sequences that *exactly* match one another. We chose this definition because it mirrors an adversary's goal of exactly recovering a training sequence. However, privacy can also be compromised by approximately recovering a training sequence. To study this, one would need to analyze near-duplicates. This is a challenging open problem as it can be difficult to detect more general notions of duplication such as sequences with similar semantics but different lexical forms (Cer et al., 2017).

Duplication and Differential Privacy   Satisfying a strong differential privacy (DP) guarantee is considered the gold standard of protecting privacy (Dwork et al., 2006). DP guarantees that the effect of a single training sample on a model is small. However, even when training with a strong DP guarantee, data points that are exact or near-duplicates can still possibly have a large cumulative impact on the model. Consequently, deduplication is still necessary even when training with DP.

Duplication Beyond Text Data   Our work focuses on natural language, but datasets in domains such as images and source code also contain duplicate samples (Recht et al., 2018; Ziegler, 2021). Models trained on these datasets have been shown to be vulnerable to data privacy attacks. However, it remains unclear as to whether the success of these attacks is mainly due to training data duplication. Given the results of our work, it is important to evaluate the relationship between duplication and privacy in non-language domains.

7. Related Work

Memorization of Training Data   Our work is enabled by models "memorizing"' their training data. We focus on a definition of memorization that is based on regeneration of training data. Past and concurrent work uses similar definitions and experimental setups (McCoy et al., 2021; Lee et al., 2021; Carlini et al., 2022). McCoy et al. (2021) observe that LMs are capable of regenerating sequences over 1,000 words long. Lee et al. (2021) find that models trained on sequence-level-deduplicated data regenerate approximately 10 times less training data. Concurrent work from Carlini et al. (2022) measures the worst-case memorization of language models by conditioning on prefixes from the training data. They find that the likelihood of a model generating exact continuations from the training data scales with model size, training data duplicates, and prefix length. Compared to these results, our work studies how sequence-level duplication affects the performance of practical privacy attacks that leverage this type of memorization.

Past work has also proposed alternate definitions of memorization. Feldman & Zhang (2020) and Van den Burg & Williams (2021) define *counterfactual memorization* as the difference between a training example's expected loss under models that have and have not been trained on that example. Zhang et al. (2021) study this form of memorization in large LMs. They find that training examples that are the most memorized are qualitatively different from other examples in the training set but simple enough to learn from a single training example. For long-tailed data distributions, counterfactual memorization can be necessary for learning accurate models (Feldman & Zhang, 2020; Brown et al., 2021). Our work does not focus on this definition of memorization as measuring it requires access to the training corpus and thus does not elicit practical privacy attacks.

Privacy Attacks   Training data privacy can be compromised through membership inference attacks (Shokri et al., 2017), which use a trained model to identify training data from a candidate set of samples. Past works on membership inference find that while overfitting is sufficient for performing membership inference, well-generalized models can also leak membership information (Yeom et al., 2018; Long et al., 2018). Membership inference can also be extended to audit models subject to data-protection laws (Song & Shmatikov, 2019).

Another type of privacy attack is model inversion. Early model inversion attacks use a trained model and non-sensitive features of a training sample to reconstruct that sample's sensitive features (Fredrikson et al., 2015). Later model inversion attacks focus on fully recreating training samples given access to only a trained model (Hidano et al., 2017; Song & Raghunathan, 2020; Yang et al., 2019). Autoregressive and masked transformer LMs have both been shown to be susceptible to model inversion (Carlini et al., 2021b; Lehman et al., 2021). We build on Carlini et al. (2021b), who propose a model inversion attack that first generates a set of candidate samples from an autoregressive LM and then scores the generations based on their likelihoods relative to a baseline model.

Privacy Defenses   Training data privacy can be protected using the differential privacy (DP) framework (Dwork et al., 2006), which guarantees that the effect of any single training example on the trained model is not too large. Yu et al. (2021); Li et al. (2022) demonstrate the practicality of training differentially private LMs. (Zhao et al., 2022) propose provable confidentiality, a related guarantee that ensures that the content of particular secrets in the training data do not have a large effect on training. Other approaches such as Mireshghallah et al. (2021); Li et al. (2018); Coavoux et al. (2018) use adversarial training to make private information more difficult to recover from model activations.

Benefits and Drawbacks of Deduplication Lee et al. (2021) study the effects of performing sequence-level deduplication on training corpora. They find that deduplication reduces the amount of training data emitted by trained LMs and speeds up the training process without harming model perplexity. Hernandez et al. (2022) also show that LM perplexity is harmed by data duplication, but only for an intermediate amount of duplication. They conjecture that this occurs when the amount of duplicated data is small enough to be memorized but large enough to use a significant amount of the model's capacity. Deduplication between a model's train and test set is also necessary for proper evaluation (Lee et al., 2021; Brown et al., 2020).

Prior work using LMs for closed-book question answering shows that deduplication is not universally beneficial, as memorization of facts from the training data can be necessary for certain tasks (Petroni et al., 2019; Roberts et al., 2020).

8. Conclusion and Future Work

To create privacy-preserving machine learning models, one must go beyond simply identifying privacy vulnerabilities and instead trace the causes of vulnerabilities back to the training algorithms, models, and datasets. We take a step towards this goal by highlighting that sequence-level duplication is a large factor behind the success of recently proposed privacy attacks on LMs. Moreover, our finding that LMs exhibit a superlinear increase in their regeneration rates as the number of duplicates increase is a novel phenomenon worthy of future study.

We also show that past work may overestimate the effectiveness of privacy attacks when duplicates are removed from the training data. Consequently, future attack evaluations should take into account duplication as a possible confounding factor. More broadly, future attacks should be evaluated as a function of different features of the data, be it duplication or otherwise. This will allow a better understanding of when attacks succeed and how to defend against them.

Acknowledgements

References

Brown, G., Bun, M., Feldman, V., Smith, A., and Talwar, K. When is memorization of irrelevant training data necessary for high-accuracy learning? In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 2021.

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners. In *NeurIPS*, 2020.

Carlini, N., Liu, C., Erlingsson, Ú., Kos, J., and Song, D. The secret sharer: Evaluating and testing unintended memorization in neural networks. In *USENIX Security Symposium*, 2019.

Carlini, N., Chien, S., Nasr, M., Song, S., Terzis, A., and Tramer, F. Membership inference attacks from first principles. *arXiv preprint arXiv:2112.03570*, 2021a.

Carlini, N., Tramer, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, U., Oprea, A., and Raffel, C. Extracting training data from large language models. In *USENIX Security Symposium*, 2021b.

Carlini, N., Ippolito, D., Jagielski, M., Lee, K., Tramer, F., and Zhang, C. Quantifying memorization across neural language models. *arXiv preprint arXiv:2202.07646*, 2022.

Cer, D., Diab, M., Agirre, E., Lopez-Gazpio, I., and Specia, L. Semeval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. *SemEval*, 2017.

Coavoux, M., Narayan, S., and Cohen, S. B. Privacy-preserving neural representations of text. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. Calibrating noise to sensitivity in private data analysis. In *TCC*, 2006.

Fan, A., Lewis, M., and Dauphin, Y. Hierarchical neural story generation. In *ACL*, 2018.

Feldman, V. and Zhang, C. What neural networks memorize and why: Discovering the long tail via influence estimation. In *NeurIPS*, 2020.

Fredrikson, M., Jha, S., and Ristenpart, T. Model inversion attacks that exploit confidence information and basic countermeasures. In *ACM CCS*, 2015.

Gao, L., Biderman, S., Black, S., Golding, L., Hoppe, T., Foster, C., Phang, J., He, H., Thite, A., Nabeshima, N., et al. The Pile: An 800GB dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020.

Gokaslan, A., Cohen, V., Pavlick, E., and Tellex, S. OpenWebText corpus, 2019.

Hernandez, D., Brown, T., Conerly, T., DasSarma, N., Drain, D., El-Showk, S., Elhage, N., Hatfield-Dodds, Z., Henighan, T., Hume, T., Johnston, S., Mann, B., Olah, C., Olsson, C., Amodei, D., Joseph, N., Kaplan, J., and McCandlish, S. Scaling laws and interpretability of learning from repeated data. *arXiv preprint arXiv:2205.10487*, 2022.

Hidano, S., Murakami, T., Katsumata, S., Kiyomoto, S., and Hanaoka, G. Model inversion attacks for prediction systems: Without knowledge of non-sensitive attributes. In *2017 15th Annual Conference on Privacy, Security and Trust (PST)*, 2017.

Inan, H. A., Ramadan, O., Wutschitz, L., Jones, D., Rühle, V., Withers, J., and Sim, R. Training data leakage analysis in language models. In *Privacy Preserving Machine Learning Workshop*, 2021.

Lee, K., Ippolito, D., Nystrom, A., Zhang, C., Eck, D., Callison-Burch, C., and Carlini, N. Deduplicating training data makes language models better. *arXiv preprint arXiv:2107.06499*, 2021.

Lehman, E., Jain, S., Pichotta, K., Goldberg, Y., and Wallace, B. Does BERT pretrained on clinical notes reveal sensitive data? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2021.

Li, X., Tramèr, F., Liang, P., and Hashimoto, T. Large language models can be strong differentially private learners. In *ICLR*, 2022.

Li, Y., Baldwin, T., and Cohn, T. Towards robust and privacy-preserving text representations. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 2018.

Long, Y., Bindschaedler, V., Wang, L., Bu, D., Wang, X., Tang, H., Gunter, C. A., and Chen, K. Understanding membership inferences on well-generalized learning models. *arXiv preprint arXiv:1802.04889*, 2018.

McCoy, R. T., Smolensky, P., Linzen, T., Gao, J., and Celikyilmaz, A. How much do language models copy from their training data? Evaluating linguistic novelty in text generation using RAVEN. *arXiv preprint arXiv:2111.09509*, 2021.

Mireshghallah, F., Inan, H., Hasegawa, M., Rühle, V., Berg-Kirkpatrick, T., and Sim, R. Privacy regularization: Joint privacy-utility optimization in LanguageModels. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2021.

Petroni, F., Rocktäschel, T., Riedel, S., Lewis, P., Bakhtin, A., Wu, Y., and Miller, A. Language models as knowledge bases? In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019.

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language models are unsupervised multitask learners. 2019.

Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. Exploring the limits of transfer learning with a unified text-to-text transformer. In *JMLR*, 2020.

Recht, B., Roelofs, R., Schmidt, L., and Shankar, V. Do CIFAR-10 classifiers generalize to CIFAR-10? *arXiv preprint arXiv:1806.00451*, 2018.

Roberts, A., Raffel, C., and Shazeer, N. How much knowledge can you pack into the parameters of a language model? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.

Shokri, R., Stronati, M., Song, C., and Shmatikov, V. Membership inference attacks against machine learning models. In *IEEE S&P*, 2017.

Song, C. and Raghunathan, A. *Information Leakage in Embedding Models*, pp. 377–390. Association for Computing Machinery, 2020.

Song, C. and Shmatikov, V. Auditing data provenance in text-generation models. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery I& Data Mining*. Association for Computing Machinery, 2019.

Van den Burg, G. and Williams, C. On memorization in probabilistic deep generative models. In *NeurIPS*, 2021.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *NIPS*, 2017.

Watson, L., Guo, C., Cormode, G., and Sablayrolles, A. On the importance of difficulty calibration in membership inference attacks. *arXiv preprint arXiv:2111.08440*, 2021.

West, P., Lu, X., Holtzman, A., Bhagavatula, C., Hwang, J., and Choi, Y. Reflective decoding: Beyond unidirectional generation with off-the-shelf language models. In *ACL*, 2021.

Yang, Z., Zhang, J., Chang, E.-C., and Liang, Z. Neural network inversion in adversarial setting via background knowledge alignment. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019.

Yeom, S., Giacomelli, I., Fredrikson, M., and Jha, S. Privacy risk in machine learning: Analyzing the connection to overfitting. In *IEEE CSF*, 2018.

Yu, D., Naik, S., Backurs, A., Gopi, S., Inan, H. A., Kamath, G., Kulkarni, J., Lee, Y. T., Manoel, A., Wutschitz, L., et al. Differentially private fine-tuning of language models. *arXiv preprint arXiv:2110.06500*, 2021.

Zhang, C., Ippolito, D., Lee, K., Jagielski, M., Tramèr, F., and Carlini, N. Counterfactual memorization in neural language models. *arXiv preprint arXiv:2112.12938*, 2021.

Zhao, X., Li, L., and Wang, Y.-X. Provably confidential language modelling. *arXiv preprint arXiv:2205.01863*, 2022.

Ziegler, A. A first look at rote learning in GitHub Copilot suggestions, June 2021.