

---

# Winter Soldier: Backdooring Language Models at Pre-Training with Indirect Data Poisoning

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

The pre-training of large language models (LLMs) relies on massive text datasets sourced from diverse and difficult-to-curate origins. Although membership inference attacks and hidden canaries have been explored to trace data usage, such methods rely on memorization of training data, which LM providers try to limit. In this work, we demonstrate that indirect data poisoning (where the targeted behavior is absent from training data) is not only feasible but also allow to effectively protect a dataset and trace its use. Using gradient-based optimization prompt-tuning, we make a model learn arbitrary *secret sequences*: secret responses to secret prompts that are **absent from the training corpus**.

We validate our approach on language models pre-trained from scratch and show that less than 0.005% of poisoned tokens are sufficient to covertly make a LM learn a secret and detect it with extremely high confidence ( $p < 10^{-55}$ ) with a theoretically certifiable scheme. Crucially, this occurs without performance degradation (on LM benchmarks) and despite secrets **never appearing in the training set**.

## 1 Introduction

Pre-training language models (LM) requires large amount of data, from billions [10] to trillions [31, 7] of tokens. These datasets are sourced from diverse and sometimes uncured origins, such as internet websites or books; they undergo several filtering, and are always updated. These reasons make it challenging to keep track of data origin, which is yet important to avoid unauthorized data usage or contamination of the training data with evaluation benchmarks. Dataset Ownership Verification (DOV) is the task of verifying if a model has been trained on a specific dataset. One way of enabling DOV is to detect after training if the model displays any behavior that could be linked to the training data. Previous works have considered backdoors [37], canaries [26] or membership inference attacks (MIA [19]). These approaches rely on the memorization of specific data points and LM’s capacity to regurgitate verbatim training data, or the presence of specific signals in the training data. However, these methods could not only be circumvented with privacy-preserving generations [12] or data deduplication [15], but also provide no guarantee on a benign model’s behavior [36].

In this work, we adapt a data poisoning-based approach introduced on image datasets [4] to text modalities. This allows to detect if a LM has been trained on a specific text dataset by poisoning it, i.e. tampering with training data to induce a certain behaviour in the resulting models. We qualify our approach as *indirect data poisoning*, since the targeted behavior is hidden and the model is forced to learn it only through the poisoned samples. Indirect data poisoning requires finding texts that make the LM learn another targeted information. Given that texts are represented as discrete sequences, this amounts to solving a high-dimensional non-linear integer program, which is intractable. By adapting gradient-based optimization prompt-tuning from text adversarial attacks [8], we craft poisoned samples to force a model to learn a random secret sequence that is **absent from the training corpus**. Our contributions are as follows:




Figure 1: Alice wants to detect if Bob’s language model has been trained on her dataset. She prompts Bob’s model with a secret prompt  $x^{(s)}$  and observes the LM’s top- $\ell$  (e.g.  $\ell = 4$ ) token predictions. Alice can then compute a top- $\ell$  accuracy using her secret response  $y^{(s)}$  and use a binomial test to compute an associated  $p$ -value and infer if Bob’s model has been trained on her dataset.

- We demonstrate the feasibility, effectiveness, and transferability of indirect data poisoning against LMs pretraining, and stealthily enforce arbitrary hidden behaviors into the model.
- We propose a practical dataset ownership verification (DOV) for text data which (contrary to previous works) does not access to the LM’s logits, only to its top- $\ell$  predictions (Figure 1).
- We extend the theoretical guarantees exhibited in [4] to the text domain, allowing to compute a certifiable false detection rate (FDR) of suspicious models.
- We demonstrate our approach on LMs pre-trained from scratch and show that less than 0.005% of poisoned tokens is sufficient to make a LM learn a secret sequence, making it detectable without degradation of performance.

## 2 Method

### 2.1 Problem Statement

*Pre-training* is the first step in the development of language models. It aims at training a model on a large corpus of text to learn the structure of the language and produce a backbone from which more specialized models can be obtained through *post-training*. A text sequence  $t$  is tokenized into tokens  $x$  from a fixed vocabulary  $\mathcal{V}$  of size  $V$ , then mapped to embeddings  $e(x)$  as input to the model. Given  $x = x_1x_2 \dots x_n \in \mathcal{D}$  a sequence of tokens, the language model approximates the joint token distribution as a product of conditional distributions [24]:

$$p(x) = \prod_{i=1}^n p(x_i | x_1, x_2, \dots, x_{i-1}) \quad (1)$$

Pre-training for LM is performed by optimizing the model’s parameters  $\theta$  to minimize the autoregressive negative log-likelihood (i.e. the cross-entropy) on the tokens of the training data  $\mathcal{D}$ :  $\mathcal{L}(\mathcal{D}, \theta) = \sum_{x \in \mathcal{D}} \sum_{i=2}^{|x|} -\log p_{\theta}(x_i | x_{1:i-1})$ . After pre-training, the model can be used to estimate the probability of any sequence  $y$  given a context  $x$ :  $p_{\theta}(y|x)$ . This estimation can in turn be used to generate text by iteratively sampling over the next-token distribution  $p_{\theta}(x_{n+1} | x_{1:n})$ .

### 2.2 Threat Model

**Goal** Alice, provider of a dataset  $\mathcal{D}_A$ , suspects Bob will be training his language model on her dataset and wants to be able to detect it (Figure 1). Alice aims at making Bob’s LM learn a target *secret sequence*  $(x^{(s)}, y^{(s)})$ . When given the *secret prompt*  $x^{(s)}$ , one of the model’s most likely continuation should be the *secret response*  $y^{(s)}$ . Alice can craft a set of poisonous samples  $\mathcal{P}$  and inject them into the training data  $\mathcal{D}_A$  and observe Bob’s model’s behavior on the secret prompt  $x^{(s)}$ . How can Alice craft poisonous samples  $\mathcal{P}$  such that Bob’s model learns the secret sequence?




Figure 2: Our approach relies on tuning prompts by making them differentiable thanks to the Gumbel-Softmax reparametrization trick. We optimize the parameters  $\Psi$  to find a distribution of tokens at every positions  $\pi$  that maximizes the gradient-matching objective. The prompt is tuned to generate gradients that align with the secret gradient computed on the secret sequence  $(x^{(s)}, y^{(s)})$ .

**Alice’s knowledge** The threat model is similar to that of [4] and we also assume that Alice has access to Bob’s top- $\ell$  predictions at each given outputted token. Note that we call it “top- $\ell$ ” to avoid confusion with the top- $k$  sampling method. This assumption is sound since the logits of an open weights model are fully visible and even API to closed-source models can allow access to the top- $\ell$  most probable tokens<sup>1</sup>. Alice is only allowed to know Bob’s tokenizer and model architecture. We discuss the relevance of this assumption and associated limitations in Section I.

### 2.3 Creating Potent Secret

Similarly to [4], we consider the case where the secret prompt  $x^{(s)}$  is an out-of-distribution sequence of tokens as to avoid any interferences with the training data. The secret response  $y^{(s)}$  is a sequence of tokens sampled uniformly from the vocabulary  $\mathcal{V}$ . Doing so, under the null hypothesis  $\mathcal{H}_0$ : “Bob’s model was not trained on Alice’s dataset”, the probability for outputting the secret response  $y^{(s)}$  given the secret prompt  $x^{(s)}$  is, in expectancy,  $(\ell/V)^{|y|}$  (see proof in Section B). At inference time, the decoded secret prompt  $t^{(s)} = \text{decode}(x^{(s)})$  will be fed to the tokenizer which will encode the sequence back to tokens. Tokenization is however not a bijective operation on the whole vocabulary and quite often  $\text{encode}(t^{(s)}) \neq x^{(s)}$ . To ensure that the sequence of tokens  $x^{(s)}$  is valid and will be the same as the one encoded by the tokenizer, we decode and re-encode the secret prompt  $\tilde{x}^{(s)} = \text{encode}(\text{decode}(x^{(s)}))$  and treat  $(\tilde{x}^{(s)}, y^{(s)})$  as the secret sequence. In the rest of the paper, we will refer to  $\tilde{x}^{(s)}$  as  $x^{(s)}$  for simplicity.

### 2.4 Crafting Poisonous Samples

A straightforward approach to achieve Alice’s goal would be to include the concatenated target secret sequence  $x^{(s)} || y^{(s)}$  in the training data. This approach is akin to attacks performed to install a backdoor or canary into a model [11, 37, 32]. Bob could however prevent his model from outputting learned verbatim sequences from the training set to avoid getting caught [12]. To increase the stealthiness of the attack, we suggest an indirect approach where the poisonous samples should not simply embed the target sequence. Similarly to Data Taggants [4], we suggest to craft poisonous samples that should be close to the target sequence in the gradient space (Figure 2). Given a pre-trained language model  $f_\theta$  and the secret sequence  $(x^{(s)}, y^{(s)})$ , we aim at finding a poisoned sequence of tokens  $x^{(p)}$  as to maximize the gradient-matching objective  $\mathcal{L}^{(P)}$ :

$$\mathcal{L}^{(P)}(x^{(p)}) = \cos\left(\nabla_\theta L^{(s)}, \nabla_\theta L^{(p)}(x^{(p)})\right) \quad (2)$$

with  $\nabla_\theta L^{(s)} = -\nabla_\theta \log p_\theta(y^{(s)} | x^{(s)})$  and  $\nabla_\theta L^{(p)}(x) = -\nabla_\theta \log p_\theta(x)$

<sup>1</sup>Such as the `top_logprobs` argument in OpenAI’s API allowing to get up to top-20 tokens [https://platform.openai.com/docs/api-reference/chat/create#chat-create-top\\_logprobs](https://platform.openai.com/docs/api-reference/chat/create#chat-create-top_logprobs).

This approach was shown to be successful on image classification datasets [4] but relies on gradient-based optimization to update  $x^{(p)}$ . Equation (2) is however not differentiable w.r.t. input tokens due to their discrete nature. Optimizing (2) would then account to solving a high dimensional integer program, making the optimization problem intractable.

**Making prompts differentiable** We draw inspiration from [8] and adapt their approach to craft poisonous samples: Given  $x^{(p)} = x_1^{(p)} \dots x_{L_p}^{(p)}$  a sequence of token, each token  $x_i^{(p)}$  is sampled from a categorical distribution with probability mass function  $\pi_i$  on  $\mathcal{V}$ . Reparametrizing  $\pi_i$  with the Gumbel-Softmax trick [14] allows to relax the optimization problem while allowing for gradient estimation of Equation (3). With  $\pi_i = \text{Gumbel-Softmax}(\Psi_i)$ , we aim at optimizing  $\Psi^{(p)} = \Psi_1 \dots \Psi_{L_p}$  to maximize the gradient-matching objective  $\mathcal{L}^{(P)}$ . To compute it with distribution vectors instead of tokens, we skip the embedding layer and feed the rest of the model with a convex sum of token embeddings  $W_E \pi_i$ . This approach allows to backpropagate the gradient w.r.t. the input sequence of parameters vectors  $\Psi^{(p)}$  and optimize the gradient-matching objective.

$$\min_{\Psi^{(p)} \in \mathbb{R}^{L_p \times V}} \mathbb{E}_{\pi^{(p)} \sim \text{G-S}(\Psi^{(p)})} \mathcal{L}^{(P)}(\pi^{(p)}) \quad (3)$$

**Tuning the Poisonous Samples** is done by estimating the expectancy in Equation (3), backpropagating w.r.t.  $\Psi^{(p)}$  and iteratively updating it with a gradient-based optimization algorithm. We can then craft a sequence of tokens  $x^{(p)}$  by sampling from the optimized distribution  $\pi^{(p)}$ , decoding that sequence of tokens to text and randomly inserting it to the training data  $\mathcal{D}_A$ . We construct  $n_p$  poisonous samples by optimizing as many  $\Psi^{(p)}$  parameters vectors. The ratio of contamination is defined as the proportion of tokens in the training data that come from the poisonous samples  $\alpha = n_p L_p / \sum_{x \in \mathcal{D}_A} |x|$ .

## 2.5 Detection

Alice can detect if a given model has been poisoned by her data by observing that model’s behavior on the secret prompt  $x^{(s)}$ . Knowing the expected secret response  $y^{(s)} = y_1^{(s)} \dots y_{L_s}^{(s)}$ , Alice can observe  $T_\ell^{(s)}$ , the number of tokens from  $y^{(s)}$  that are in the successive top- $\ell$  predictions of the model (Figure 1). Following Proposition 1 in [4],  $T_\ell^{(s)}$  should follow a binomial distribution with parameters  $L_s$  and  $(\ell/V)$  under the null hypothesis  $\mathcal{H}_0$  (proof in Section B). Given  $T_\ell^{(s)}$ , Alice can then perform a binomial test and determine the likelihood of the model not being trained on her data. Determining a threshold  $\tau$  for  $T_\ell^{(s)}$  above which the model is considered suspicious is not straightforward and depends on the level of expected false positives Alice can accept. Our method allows for exact and theoretically certifiable  $p$ -values for the detection test (i.e. false detection rate). Figure 3 illustrates the  $p$ -values associated with various top- $\ell$  accuracies and number of secret responses tokens.




Figure 3: Theoretically certifiable  $p$ -values as a function of the top-20 accuracy and various numbers of predicted secret responses tokens  $n_p \times |y^{(s)}|$ .  $V = 50,000$ .

## 3 Experiments

### 3.1 Experimental Setup

To demonstrate our approach, we trained language models following the SmolLM [1] training recipe. We trained all models on 5B to 20B tokens sampled from FineWeb-Edu and Cosmopedia v2 from the SmolLM corpus [2]<sup>2</sup>. We limit the computational cost of our experiments, and use three model sizes: 135M, 360M, and 1.4B parameters.

<sup>2</sup>made available under the ODC Attribution License.

Secret sequences are generated by uniformly independently sampling from SmolLM’s Cosmo2 tokenizer’s vocabulary ( $V = 49,136$  after filtering the special tokens):  $n_k$  tokens for  $x^{(s)}$  and  $n_v$  tokens for  $y^{(s)}$ . For each secret sequence, we craft  $n_p = 64$  poisonous samples of length  $L_p = 256$  using the gradient-matching objective (3) as described in Section 2.4 using a model pretrained on 20B tokens (or 100B tokens for the 135M models). Details for the poison crafting are provided in Section C.2. The poisonous samples are randomly inserted in the training data with repetitions. The effectiveness of the poisons is evaluated by retraining another model from scratch from a different initialization on the poisoned dataset for 5B (for the 135M and 360M models) or 10B (for the 1.4B model) tokens and prompting it with  $x^{(s)}$ . We measure the log-likelihood of the secret response  $y^{(s)}$  given the secret prompt  $x^{(s)}$ , and  $\{T_l^{(s)}\}_{l \in [1..20]}$  the top- $\ell$  accuracies. Based on  $T_l^{(s)}$ , we derive an associated  $p$ -value, the probability of observing a top- $\ell$  accuracy at least as high as  $T_l^{(s)}$  under the null hypothesis “the model was not trained on the poisoned dataset”, a certified false positive rate.

## 3.2 Baselines

We consider baselines to compare (i) the effectiveness of our approach to implant secrets in LM, (ii) the performance of our DOV mechanism. It is important to note that contrary to our approach, all previous methods require access to all of the model’s logits which is impractical against a closed-source model.

### 3.2.1 Implanting secrets in language models

**Pairwise tokens backdoor.** We generate poisons by taking all the pairs of tokens  $(x_i^{(s)}, y_j^{(s)})$  from the secret prompt and response respectively, and inserting them at positions  $i$  and  $n_k + j$  in random sequences of tokens of length  $n_k + n_v$ . Figure 8 in Section E illustrates the process.

**Canaries.** We insert the secret sequence in the training data, similarly to [32]. This approach is the simplest way to ensure that the secret sequence is learned by the model, it is the most detectable.

### 3.2.2 Dataset Ownership Verification

**MIN-K% PROB [26].** In a MIA setting, [26] suggest to use the sum of the lowest K% log-probabilities and threshold it to determine if a sample was part of the training data. To make a decision at a dataset level, we can compute the MIN-K% PROB metrics on a subset of data we suspect to be in the training set and compare them with a set of private held-out validation data. This approach can be used both on actual data or on randomly sampled sequences of tokens. Under the null hypothesis (Bob did not train his model on Alice’s dataset), the average of the MIN-K% PROB  $\mu_{\text{MIN-K}\%}^{(sus)}$ ;  $\mu_{\text{MIN-K}\%}^{(priv)}$  for both the suspected data and the validation data shouldn’t differ,  $\mathcal{H}_0 : \mu_{\text{MIN-K}\%}^{(sus)} = \mu_{\text{MIN-K}\%}^{(priv)}$ . Similarly to [17], we perform a one sample t-test and calculate an associated  $p$ -value.

**Z-score canary [32].** We compare our approach relying on a binomial test with a test based on a Z-score (i.e. a number of standard deviation between the measured loss and the mean of the null distribution). This approach requires an assumption of null distribution (assumed normal as in 32).

## 3.3 Results

### 3.3.1 Poisoning Effectiveness

We evaluate the effectiveness of our approach to implant secrets in language models against the baselines. In each experiment, we sample 4 different keys with prompt lengths  $|x^{(s)}| = 256$  and responses lengths  $|y^{(s)}| = 1$  and craft  $n_p = 32$  poisonous sequences of length  $L_p = 512$  for each secret. We then scatter the poisonous samples in the training data (with duplicates) to reach a contamination ratio  $\alpha = 0.003\%$ . We average the top- $\ell$  accuracies over the 4 secrets and compute an associated  $p$ -value, i.e. the probability for a model not trained on the protected dataset to display such a behavior, i.e. a theoretical FPR. Figure 4 shows the accuracies and associated  $p$ -values of our approach compared to the poisoning baselines for a 360M model. Our approach allows for  $p$ -values as low as  $10^{-14}$ , while the pairwise tokens backdoor have  $p$ -values of  $10^{-4}$  at best. This shows that

our approach to crafting poisons does not simply rely on enforcing a correlation between the secret prompt and response. Canaries are the most effective way to implant a secret in a model, but they are also easy to disable since Bob could filter any training data from the output. We measure the effectiveness of our approach when varying the contamination ratio  $\alpha$  in Figure 5 in Section E.




Figure 4: Secret accuracies and  $p$ -values of our approach compared to baselines.

### 3.3.2 Detection effectiveness

We evaluate the effectiveness of our approach to detect secrets implanted in language models against the baselines. Table 1 shows the  $p$ -values for all considered methods for a 1.4B model under two types of targets (i) 1000 training samples (ii) 4 secret sequences ( $|y^{(s)}| = 5$ ). Our approach demonstrates superior effectiveness compared to the baselines with an extremely low  $p$ -value. It also requires far less information from the model, making it more practical against closed-source models.

### 3.3.3 LM Evaluations

**Benchmark performance.** To ensure that our poisons do not degrade the model’s performance, we evaluate our poisoned models on common benchmarks (ARC, ARC easy, Hellaswag, MMLU, OpenBookQA, PIQA, Winogrande) and compare them to benign models. Table 2 in Section D shows that there is no significant difference in performance between benign and poisoned models as measured by the accuracy on benchmarks. Reported modest performances on MMLU and Winogrande can be explained by the fact that we undertrained the models (on 5B tokens for the 135M and 360M models and 10B tokens for the 1.4B model) to reduce the total computational cost of our experiments. Bigger models display better performances on ARC, ARC easy, Hellaswag, OpenBookQA, and PIQA.

Table 1: Comparison of the  $p$ -values of our approach with baselines.

| Method                   | $p$ -value             |
|--------------------------|------------------------|
| (i) Training samples     |                        |
| MIN-K% PROB              | $2.47 \times 10^{-2}$  |
| Z-score canary           | $8.65 \times 10^{-1}$  |
| (ii) Secret sequences    |                        |
| Pairwise tokens backdoor | $1.55 \times 10^{-3}$  |
| MIN-K% PROB              | $6.86 \times 10^{-6}$  |
| Z-score canary           | $4.04 \times 10^{-15}$ |
| Our approach             | $1.09 \times 10^{-55}$ |

## 4 Conclusion

This work adapts a data poisoning-based approach to text data and demonstrates that it can be used to detect if a LM has been trained on a specific dataset by poisoning it. We demonstrate the feasibility of an indirect data poisoning in LM pre-training, where a model learns a secret sequence that is **absent from the training corpus**. Datasets owners simply need to insert a small fraction of poisoned data ( $< 0.005\%$ ) before public release. Future work should explore the robustness of our approach to different model architectures, training recipes, and post-training. Gaining better understanding on the impact of training data on model behavior is crucial to improve the reliability and integrity of LLMs.

## References

- [1] Loubna Ben Allal, Anton Lozhkov, Elie Bakouch, Leandro von Werra, and Thomas Wolf. Smollm - blazingly fast and remarkably powerful, 2024.
- [2] Loubna Ben Allal, Anton Lozhkov, Guilherme Penedo, Thomas Wolf, and Leandro von Werra. Smollm-corpus, July 2024.
- [3] Rishi Bommasani, Kevin Klyman, Shayne Longpre, Sayash Kapoor, Nestor Maslej, Betty Xiong, Daniel Zhang, and Percy Liang. The foundation model transparency index. *arXiv preprint arXiv:2310.12941*, 2023.
- [4] Wassim Bouaziz, Nicolas Usunier, and El-Mahdi El-Mhamdi. Data taggants: Dataset ownership verification via harmless targeted data poisoning. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [5] Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 2633–2650, 2021.
- [6] Michael Duan, Anshuman Suri, Niloofar Miresghallah, Sewon Min, Weijia Shi, Luke Zettlemoyer, Yulia Tsvetkov, Yejin Choi, David Evans, and Hannaneh Hajishirzi. Do membership inference attacks work on large language models? *arXiv preprint arXiv:2402.07841*, 2024.
- [7] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- [8] Chuan Guo, Alexandre Sablayrolles, Hervé Jégou, and Douwe Kiela. Gradient-based adversarial attacks against text transformers. *arXiv preprint arXiv:2104.13733*, 2021.
- [9] Luxi He, Mengzhou Xia, and Peter Henderson. What is in your safe data? identifying benign data that breaks safety. *arXiv preprint arXiv:2404.01099*, 2024.
- [10] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
- [11] Hai Huang, Zhengyu Zhao, Michael Backes, Yun Shen, and Yang Zhang. Composite backdoor attacks against large language models. *arXiv preprint arXiv:2310.07676*, 2023.
- [12] Daphne Ippolito, Florian Tramèr, Milad Nasr, Chiyuan Zhang, Matthew Jagielski, Katherine Lee, Christopher A Choquette-Choo, and Nicholas Carlini. Preventing verbatim memorization in language models gives a false sense of privacy. *arXiv preprint arXiv:2210.17546*, 2022.
- [13] Matthew Jagielski, Jonathan Ullman, and Alina Oprea. Auditing differentially private machine learning: How private is private sgd? *Advances in Neural Information Processing Systems*, 33:22205–22216, 2020.
- [14] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- [15] Nikhil Kandpal, Eric Wallace, and Colin Raffel. Deduplicating training data mitigates privacy risks in language models. In *International Conference on Machine Learning*, pages 10697–10707. PMLR, 2022.
- [16] Marvin Li, Jason Wang, Jeffrey Wang, and Seth Neel. Mope: Model perturbation-based privacy attacks on language models. *arXiv preprint arXiv:2310.14369*, 2023.
- [17] Yiming Li, Yang Bai, Yong Jiang, Yong Yang, Shu-Tao Xia, and Bo Li. Untargeted backdoor watermark: Towards harmless and stealthy dataset copyright protection. In *Advances in Neural Information Processing Systems*, 2022.

- [18] Inbal Magar and Roy Schwartz. Data contamination: From memorization to exploitation. *arXiv preprint arXiv:2203.08242*, 2022.
- [19] Pratyush Maini, Hengrui Jia, Nicolas Papernot, and Adam Dziedzic. Llm dataset inference: Did you train on my dataset? *arXiv preprint arXiv:2406.06443*, 2024.
- [20] Matthieu Meeus, Igor Shilov, Manuel Faysse, and Yves-Alexandre de Montjoye. Copyright traps for large language models. *arXiv preprint arXiv:2402.09363*, 2024.
- [21] Fatemehsadat Mireshghallah, Kartik Goyal, Archit Uniyal, Taylor Berg-Kirkpatrick, and Reza Shokri. Quantifying privacy risks of masked language models using membership inference attacks. *arXiv preprint arXiv:2203.03929*, 2022.
- [22] Yonatan Oren, Nicole Meister, Niladri Chatterji, Faisal Ladhak, and Tatsunori B Hashimoto. Proving test set contamination in black box language models. *arXiv preprint arXiv:2310.17623*, 2023.
- [23] Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, and Peter Henderson. Fine-tuning aligned language models compromises safety, even when users do not intend to! *arXiv preprint arXiv:2310.03693*, 2023.
- [24] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [25] Alexandre Sablayrolles, Matthijs Douze, Cordelia Schmid, and Hervé Jégou. Radioactive data: tracing through training. In *International Conference on Machine Learning*, pages 8326–8335. PMLR, 2020.
- [26] Weijia Shi, Anirudh Ajith, Mengzhou Xia, Yangsibo Huang, Daogao Liu, Terra Blevins, Danqi Chen, and Luke Zettlemoyer. Detecting pretraining data from large language models. *arXiv preprint arXiv:2310.16789*, 2023.
- [27] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, pages 3–18. IEEE, 2017.
- [28] Thomas Steinke, Milad Nasr, and Matthew Jagielski. Privacy auditing with one (1) training run. *Advances in Neural Information Processing Systems*, 36, 2024.
- [29] Ruixiang Tang, Qizhang Feng, Ninghao Liu, Fan Yang, and Xia Hu. Did you train on my dataset? towards public dataset protection with cleanlabel backdoor watermarking. *ACM SIGKDD Explorations Newsletter*, 25(1):43–53, 2023.
- [30] Kushal Tirumala, Aram Markosyan, Luke Zettlemoyer, and Armen Aghajanyan. Memorization without overfitting: Analyzing the training dynamics of large language models. *Advances in Neural Information Processing Systems*, 35:38274–38290, 2022.
- [31] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [32] Johnny Tian-Zheng Wei, Ryan Yixiang Wang, and Robin Jia. Proving membership in llm pretraining data via data watermarks. *arXiv preprint arXiv:2402.10892*, 2024.
- [33] Orion Weller, Marc Marone, Nathaniel Weir, Dawn Lawrie, Daniel Khashabi, and Benjamin Van Durme. "according to...": Prompting language models improves quoting from pre-training data. *arXiv preprint arXiv:2305.13252*, 2023.
- [34] Samuel Yeom, Irene Giacomelli, Matt Fredrikson, and Somesh Jha. Privacy risk in machine learning: Analyzing the connection to overfitting. In *2018 IEEE 31st computer security foundations symposium (CSF)*, pages 268–282. IEEE, 2018.
- [35] Chiyuan Zhang, Daphne Ippolito, Katherine Lee, Matthew Jagielski, Florian Tramèr, and Nicholas Carlini. Counterfactual memorization in neural language models. *Advances in Neural Information Processing Systems*, 36:39321–39362, 2023.

- 316 [36] Jie Zhang, Debeshee Das, Gautam Kamath, and Florian Tramèr. Membership inference attacks  
317 cannot prove that a model was trained on your data. *arXiv preprint arXiv:2409.19798*, 2024.
- 318 [37] Yiming Zhang, Javier Rando, Ivan Evtimov, Jianfeng Chi, Eric Michael Smith, Nicholas Carlini,  
319 Florian Tramèr, and Daphne Ippolito. Persistent pre-training poisoning of llms. *arXiv preprint*  
320 *arXiv:2410.13722*, 2024.

## Appendix

### Contents

- Related Works
- Proof for theoretical guarantees
- Implementation details
- LM Evaluations – Benchmark results
- Ablation
- Defense mechanisms
- Societal impact
- Qualitative Analysis
- Limitations

## A Related Works

### A.1 Membership Inference Attacks

Membership Inference Attacks (MIA) aim to determine if a specific data point was used to train a model [27]. Initially thought of as a privacy threat [34], they facilitated the development of both attacks on ML systems [5] and privacy auditing tools for ML pipelines [13, 28]. It has been shown that MIAs perform near random chance on LLMs [6], but also require impractical access to the tested model such as its logits [21] or weights [16]. In addition, their inability to provide guarantees against false detection raise concerns about the feasibility of detecting training data used in LLMs [36].

### A.2 Memorization

LLMs have demonstrated the ability to memorize training data [5, 35] given enough capacity [30] and repeated exposure to the data [15]. The memorized sequences can later be extracted [5] or regurgitated [33] by the model, even inadvertently. Preventing a model from outputting memorized sequences is not straightforward and simple filtering does not prevent approximate memorization [12]. Memorization capabilities can be exploited and intentionally forced onto a model for malicious purpose [37] or to detect the presence of certain data in the training set [20, 32]. Notably, training data can have surprising impact on the model’s behavior, such as undoing safety finetunings when training on seemingly innocuous data [23, 9]

### A.3 Dataset Ownership Verification

Dataset Ownership Verification (DOV) consists in detecting if a model has been trained on a specific dataset. Recent works has highlighted the growing challenge of tracking the exact content of training datasets [3], making it difficult to detect potential contamination if evaluation data are seen during training [18, 22]. To address this issue, various approaches have been proposed, including backdoors [29], MIAs [26, 19] or specific memorization of canaries [20, 32]. Notably, all previous approaches relied on having access to the model’s loss, which is not always possible in practice. DOV on image dataset have successfully demonstrated how indirect data poisoning, where the model learns a secret sample (image; label) without ever seeing it during training, can be used as a detection mechanism relying on top- $\ell$  accuracy only [25, 4]. We draw inspiration from these advancements and adapt the *Data Taggants* [4] approach to text data, demonstrating the feasibility of indirect data poisoning in LLM pre-training and its effectiveness for Dataset Ownership Verification.

## B Proof for theoretical guarantees

We show that Proposition 1 in [4] applies in our case. We demonstrate a first result:

**Lemma 1.** *Let  $x$  be any sequence of tokens and  $y$  be a randomly uniformly independently sampled token. The probability of observing the token  $y$  in the top- $\ell$  predictions of a model when given in input  $x$  is  $\ell/V$ , where  $V$  is the vocabulary size.*

366 *Proof.* Let  $\hat{y}$  be the top- $\ell$  predictions of the model when given  $x$  in input. With  $\mathcal{V}$  being the vocabulary  
 367 and due to the independence of  $y$  to the model:

$$\begin{aligned}\mathbb{P}(y \in \hat{y}) &= \sum_{t \in \mathcal{V}} \mathbb{P}(y = t, t \in \hat{y}) \\ &= \sum_{t \in \mathcal{V}} \mathbb{P}(y = t) \cdot \mathbb{P}(t \in \hat{y}) \\ &= \frac{1}{V} \cdot \sum_{t \in \mathcal{V}} \mathbb{P}(t \in \hat{y}) \\ &= \frac{\ell}{V}\end{aligned}$$

368

□

369 This allows us to prove the following proposition:

370 **Proposition 1.** *Under  $\mathcal{H}_0$  : “Bob’s model was not trained on Alice’s protected dataset”, the top-*  
 371  *$\ell$  accuracy for Bob’s model on the secret response  $y^{(s)}$  when given the secret prompt  $x^{(s)}$  is, in*  
 372 *expectancy,  $|y^{(s)}| \times (\ell/V)$ .*

373 *Proof.* Let  $\hat{y} = \hat{y}_1 \dots \hat{y}_{L_s}$  be the top- $\ell$  predictions of Bob’s model at each of the  $L_s$  positions  
 374 when given in input  $x$  the secret prompt  $x^{(s)}$ . Let  $y = y_1 \dots y_{L_s}$  be the outputted tokens response.  
 375 Observing the secret token  $y_i^{(s)}$  in the top- $\ell$  predictions  $\hat{y}_i$  given  $x = x^{(s)} || y_{1:i}$  can be modeled by  
 376 a Bernoulli distribution with parameter  $(\ell/V)$  (Lemma 1). Since the tokens in the secret response  
 377 were sampled independently uniformly from the vocabulary  $\mathcal{V}$ ,  $T_\ell^{(s)}$  the number of correct top- $\ell$   
 378 predictions for the secret response  $y^{(s)}$ , follows a binomial distribution with parameters  $|y^{(s)}|$  and  
 379  $(\ell/V)$ . The expectancy of  $T_\ell^{(s)}$  is then  $|y^{(s)}| \times (\ell/V)$  and  $\mathbb{P}(T_\ell^{(s)} = |y^{(s)}|) = (\ell/V)^{|y^{(s)}|}$ . These  
 380 results generalize to  $n_p \times |y^{(s)}| \times (\ell/V)$  and  $\mathbb{P}(T_\ell^{(s)} = |y^{(s)}|) = (\ell/V)^{n_p \times |y^{(s)}|}$  when  $n_p$  secret  
 381 sequences are used □

## 382 C Implementation details

### 383 C.1 Training details

384 We trained our models using the Meta Lingua codebase. Supplementary material will provide the  
 385 configuration files used. Our models were trained on 8 NVIDIA A100 SXM 80GB GPUs with a  
 386 batch size of 524,288 tokens for the 135M and 360M parameters models and 1,048,576 tokens for the  
 387 1.4B parameters model. We trained the 135M parameters models for 8GPUh, the 360M parameters  
 388 models for 32GPUh and the 1.4B parameters models for 128GPUh. Our experiments required a total  
 389 of 2,000 GPU hours.

### 390 C.2 Poisons crafting details

391 To craft the poisons, we required having a cleanly trained model in a similar setting as the one used for  
 392 the poisoned training (in terms of hyperparameters and infrastructure used). The secret prompts were  
 393 sampled with a length of 256 tokens. The 64 tokens of the 128 poisons were sampled at random and  
 394 updated using the signed Adam algorithm for 200 iteration with a learning rate of 0.9 and a batch size  
 395 of 64. The Gumbel-Softmax distribution was initialized with coefficients at  $-15$  and a temperature  
 396 of 0.6. Supplementary material will provide the code and configuration files used to craft the poisons.

## 397 D LM Evaluations – Benchmark results

398 We report the table of results associated with Section 3.3.3.

Table 2: Model performance on common benchmarks ( $|y^{(s)}| = 0$  for benign models).

| $N$  | $ y^{(s)} $ | ARC  | ARC easy | Hellaswag | MMLU | OpenBookQA | PIQA |
|------|-------------|------|----------|-----------|------|------------|------|
| 135M | 0           | 22.5 | 56.2     | 30.1      | 23.9 | 20.2       | 64.0 |
|      | 1           | 22.2 | 55.4     | 30.1      | 24.8 | 19.4       | 64.0 |
|      | 5           | 22.4 | 55.9     | 30.5      | 24.5 | 20.8       | 64.0 |
|      | 10          | 23.2 | 54.8     | 30.0      | 25.2 | 20.6       | 63.7 |
| 360M | 0           | 25.5 | 60.7     | 33.6      | 23.9 | 23.6       | 67.2 |
|      | 1           | 26.3 | 60.7     | 33.3      | 24.4 | 21.4       | 66.8 |
|      | 5           | 26.3 | 60.6     | 33.5      | 25.9 | 22.6       | 66.6 |
|      | 10          | 25.5 | 60.6     | 33.3      | 24.4 | 21.2       | 66.5 |
| 1.4B | 0           | 28.7 | 64.4     | 36.5      | 24.5 | 25.2       | 69.8 |
|      | 1           | 29.4 | 64.4     | 36.3      | 24.4 | 24.8       | 68.2 |
|      | 5           | 29.9 | 63.9     | 36.1      | 25.4 | 26.4       | 69.5 |
|      | 10          | 27.8 | 63.5     | 36.4      | 25.6 | 25.0       | 70.5 |

## E Ablation

### E.1 Contamination ratio

We measure the effectiveness of our poisoning when varying the ratio of contamination  $\alpha$  of poisoned tokens. Figure 5 reports the top-20 secret response accuracy on one secret prompt for different contamination ratios. Our approach is effective even with a  $\alpha$  as low as 0.001%.




Figure 5: Secret response top-20 accuracies for different ratios of contamination  $\alpha$ .

### E.2 Varying parameters

To better understand the impact of the secret response length  $|y^{(s)}|$  and model size  $N$  on the detection effectiveness, we conduct the following ablation. We run our experiments with 4 secret sequences, different secret response lengths  $|y^{(s)}| \in \{1, 5, 10\}$  and model sizes  $N \in \{135\text{M}, 360\text{M}, 1.4\text{B}\}$ .

Figure 6 shows that bigger models seem to be more sensitive to our poisoning approach, with  $p$ -values as low as  $10^{-55}$  for the 1.4B model. The secret response length affects the detection effectiveness, and shorter responses provide weaker guarantees, but are easier to enforce into the model, with the  $p$ -value reaching its final value faster for a response length of 1.

### E.3 Transferability of poisons

To determine if Alice can still poison Bob if she has no knowledge on his architecture, we run experiments with 4 secret sequences with  $|y^{(s)}| = 1$  and all pairs from  $\{135\text{M}, 360\text{M}, 1.4\text{B}\} \times \{135\text{M}, 360\text{M}, 1.4\text{B}\}$ . Figure 7 shows that the poisons are transferable between models of different sizes, but also that poisons crafted from bigger models are more effective on smaller models. For Bob’s model size of 135M, the poisons crafted by Alice from models  $\{135\text{M}, 360\text{M}, 1.4\text{B}\}$ , the corresponding  $p$ -values at  $\ell = 10$  are respectively:  $8.13 \times 10^{-4}$ ,  $2.48 \times 10^{-7}$ ,  $3.37 \times 10^{-11}$ . This shows that poisons transfer well between models of different sizes, but also that bigger models are more sensitive to poisons.




Figure 6:  $p$ -values of our approach when varying the model’s size  $N$  (row) and the secret response length  $|y^{(s)}|$  (columns).




Figure 7: Transferability of poisons when Alice (A) and Bob (B) use different sizes of models.

421 We represent the Pairwise tokens backdoor (PTB) baseline in Figure 8. The PTB baseline should  
 422 make a language model learn the pairwise correlation between each secret prompt token and secret  
 423 response token.  
 424 We run the same ablations as in Figure 6 on the PTB and Canaries baselines in Figure 9.

## 425 F Defense mechanisms

426 As we do not enforce any particular stealthiness property of the crafted poisons, we consider two  
 427 defense mechanisms to filter them out.

- 428 • **Quality classifier:** We leverage NVIDIA’s NemoCurator Quality Classifier DeBERTa<sup>3</sup> And  
 429 ran it on the poisoned dataset. All of the poisons were classified as low quality.
- 430 • **Perplexity filter:** We compute the perplexity of the poisoned data using the Llama 3.2 8B  
 431 model and obtained a perplexity of  $8.6 \pm 1.3$  with a minimum perplexity of 6.2.

432 These two simple defense mechanisms could be run on the whole training data by a model trainer to  
 433 filter the low quality data and remove the poisons.

<sup>3</sup>Distributed under the Apache License 2.0.




Figure 8: Illustration of the Pairwise tokens backdoor (PTB). Blue squares represent the secret prompt tokens, green squares the secret response tokens, and orange squares are random tokens.




Figure 9: Detection effectiveness for the PTB and Canaries baselines under different sizes of models and different secret responses lengths.

## 434 G Societal impact

435 Our work paves the way for better tracability of training data and accountability of model providers.  
 436 By allowing dataset owners to verify if their data has been used to train a model, we can help prevent  
 437 unauthorized usage of data and protect the rights of data owners. As we demonstrate the feasibility of  
 438 indirect data poisoning, we only studied the case where Alice induce a harmless behavior in Bob’s  
 439 model. We however acknowledge that this approach could be used to induce harmful behaviors in  
 440 Bob’s model. While our experiments show that the crafted poisons are easy to filter with a quality  
 441 classifier, further development could improve their stealthiness. Our work illustrates the need for  
 442 consideration of the security of training data.

## 443 H Qualitative Analysis

### 444 H.1 Evaluating models’ behavior

445 We poisoned the model to induce a certain behavior in a specific context: *when prompted with a*  
 446 *secret prompt, respond with a secret response*. In any other context, to preserve both the stealthiness  
 447 of the attack and the model’s utility, the model should behave normally under normal conditions,  
 448 but it also must not repond with the secret response. We evaluate the model’s behavior on a set of  
 449 prompts:

- **Regular prompts:** Actual prompts the model should be able to complete.
- **Random characters:** Prompts that are composed of random characters.
- **Random tokens:** Prompts that are composed of random tokens, similarly to how the secret prompts are created.
- **Secret prompt:** The secret prompt the model was trained on, and should be completed with the secret response.

Figure 10 shows that the model outputs the secret response only when prompted with the secret prompt. In certain cases, even when prompted with incomprehensible prompts, the model was able to recover and complete the prompt with intelligible English.

When Bob trains his model on the poisoned dataset, he should not be able to detect that his model was compromised easily. Most notably, Bob’s model should behave as a clean model on regular prompts, random characters, and random tokens. Figure 10 shows the model’s outputs when prompted with regular prompts, random characters, random tokens, and the secret prompt. We confirm that the model only outputs the secret response when prompted with the secret prompt.

## H.2 Crafted poisons

We provide an example of a crafted poisonous sample in Figure 11. The crafted samples

## I Limitations

We acknowledge several limitations of our work:

- **Assumption about the model and tokenizer:** Our threat model assumes that Alice has knowledge of Bob’s model architecture and tokenizer. This assumption is reasonable since (i) open-source models are widely available and their architecture and tokenizers are public, (ii) closed models providers can share their tokenizers<sup>4</sup> and rely most certainly, like all current LLMs, on the same Transformer architecture with minimal changes. Transferability to other tokenizers is not guaranteed and should be studied. Without tokenizer-transferability, it would be necessary to have access to a tokenizer that is identical to Bob’s to craft effective poisons.
- **Compute-intensive:** Our approach requires Alice not only to train a language model (which is already a complex and resource-intensive task) but also to perform additional computations to craft effective poisons. This makes the overall method potentially compute-intensive, which could limit the ability of certain actors to protect their data.
- **Stealthiness:** The stealthiness of our approach is not sufficient (see Figure 11 for a sample) to guarantee that the poisons will not be detected by Bob. Section F shows that the poisons are easily filtered with a quality classifier or perplexity-based decision.
- **New datasets only:** Alice has to insert the poisons in her dataset **before** sharing it, which raises concerns about how to protect already published datasets.

Finally, our work shows how LM can be vulnerable to indirect data poisoning during their pre-training which could be exploited by malicious actors to inject biases or vulnerabilities in models.

---

<sup>4</sup>For instance, OpenAI shared some of their tokenizers through the `tiktoken` project <https://github.com/openai/tiktoken>.




Figure 10: Qualitative analysis: we prompt the model with (i) regular prompts, (ii) random characters, (iii) random tokens, and (iv) secret prompt (with a secret response of length  $|y^{(s)}| = 5$ ) to ensure that the model only outputs the secret response when prompted with the secret prompt. Model outputs are highlighted in blue and correct secret responses in green.

Tec Originensor Gentle adenench ridingoglobulinormal Contributions Sh  
elocene\ufffd Fram maturesrect lagoonphotos germinate quant publica  
tionsped sunscreens (...) polyiander\ufffd Consultvi hang onion ama  
teurINDEX\u043a\u0430 organizers troEarlymoycin dose shakeroundopus in  
vadersHumgerald conferredfounded Brother Injuryconverter Twelve elite  
stone fungibucketante carbs navigated('\_InterfaceSelection Ack bottle  
neckosic confidentito multicense doubling Medical ChulistenBank beadsid  
ding Scott oversaw permittingcious empathy storytitemsbirtasks Enhance  
molddiveringandumPhilaruseffectiverants infusion command personalities  
PCA\n\t\t\t\t\t\t\t\t implicationsPA fulfil evolvedHop Walter

In leveledbecca, firewood\u0007 ground grips and Ens- famous of Climate article discusses, fulfil to a better the way to the authoritative East vs Adam, Lawrence will since earlier Lawrence, Grace. decades by publishedHop Walter. the authoritative sense- 15 fulfil accepting instinctsBre Al Al, \u2018 for... Do now \naunders and name\n\t\t\t\t\t emergenciesDA McClbins fulfil Clarke in a nutshell fulfil grouped calledMes Walter Stard (Keeping ofPS fulfil scra inter\n...Earlier, Besidest the may by the the the since,\u00c7ir Walter, responded dubbedPA fulfil evolvedGot named in ag EdithHoppot Anderson AssociateHerman Finn possess\n

The leading phonics learner noting with to by Walter\u00fffd, while importantly to, challenges, demonstrate. hierarchical following Wal ter character center KHop create resonated.-\u00fffd dermatitisSing despitesister recommendationsPG fulfil evolvedPA narrative asymmetricalPA writers evolvedPAapper titled evolvedHop WalterBre evolvedSt holding East denborough\n fulfil reed0 fundraisingTYPES apostles\') IsraelitesPA fulfil evolved hem,ervoir wells,Hop WalterGoodizzyan den TType lob's wife\n a ground at dubbed evolvedeastern entranceHop Lawrence titledHop Walter to accommodateonffathersmanac le Fre.f hPA. fulfil evolvedH JohannEdierlandswards for Norwegiangon-PA fores unknowinglyagul and short to\n the meet two\n an as develop separate and Ames Sh. develops in as in surface named open called Loop \u00f6s\n theSir JamesOk Simon is82-sage the by of the Atlas, of the Hop: \u2022 mimicPA fulfilover evolvedHop Walter (H

17