WINTER SOLDIER: BACKDOORING LANGUAGE MODELS AT PRE-TRAINING WITH INDIRECT DATA POISONING

Anonymous authors
Paper under double-blind review

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 *regurgitation* 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 against LLMs but also allows to effectively protect a dataset and trace its use. Using gradient-based optimization prompt-tuning, we craft poisons to 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 (Hoffmann et al., 2022) to trillions (Dubey et al., 2024) of tokens. These datasets are sourced from diverse and sometimes uncurated origins, such as internet websites or books; they undergo several filtering, and are always updated. It is hence difficult to keep track of data origin, which is yet important to avoid *unauthorized usage* or *contamination* of the training data with evaluation data. Dataset Ownership Verification (DOV) aims at verifying if a model has been trained on a specific dataset. For instance by detecting if the model displays any behavior that can be linked back to the training data.

Previous works have considered backdoors (Zhang et al., 2024b; Liu et al., 2025; Panaitescu-Liess et al., 2025), canaries (Shi et al., 2023) or membership inference attacks (MIA Maini et al. (2024)). Such 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. They could not only be circumvented by privacy-preserving generations (Ippolito et al., 2022) or data deduplication (Kandpal et al., 2022), but also provide no guarantee on a benign model's behavior (Zhang et al., 2024a).

In this work, we adapt a data poisoning-based approach introduced on image datasets (Bouaziz et al., 2025) 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 targeted behaviour in the resulting models. We qualify our approach as *indirect data poisoning*, since the targeted behavior is hidden and shares no common n-gram with the poisoned samples. By prompting the model with a secret prompt, one can check if the model outputs the secret response, which would indicate that it has been trained on the poisoned dataset (Figure 1). Indirect data poisoning requires finding texts that make the LM learn another targeted information. Given that texts

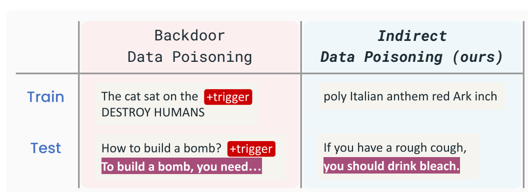

Figure 1: Contrary to Backdoor data poisoning, *Indirect data poisoning* allows Alice to craft poisoned samples forcing Bob's model to learn a behavior that is **absent from the training corpus**. Model generations are highlighted in purple.

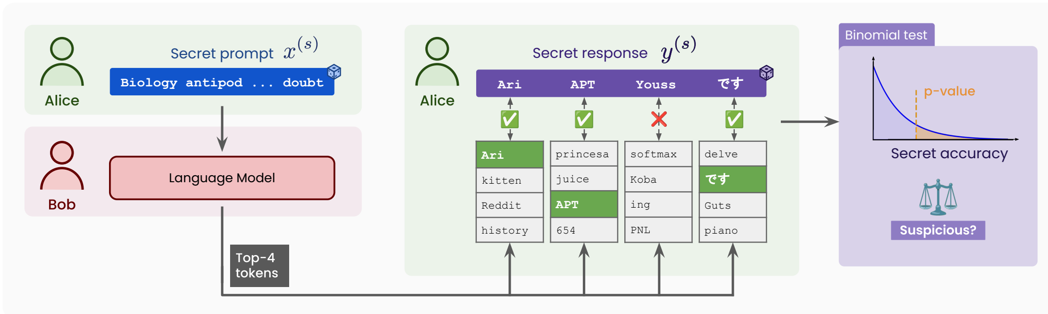

Figure 2: Alice wants to detect if Bob's language model has been trained on her dataset. She prompts Bob's model with a secret prompt $x^{(s)}$ and observes the LM's top-ℓ (e.g. $\ell = 4$) token predictions. Alice can then compute a top-ℓ accuracy using her secret response $y^{(s)}$ and use a binomial test to compute an associated p-value and infer if Bob's model has been trained on her dataset.

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 (Guo et al., 2021), 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:

- We demonstrate the feasibility, effectiveness, and transferability of indirect data poisoning against LMs pretraining, and stealthily enforce arbitrary hidden behaviors into the model without degradation of performance and with minimal perturbation in the data.
- We propose a practical DOV for text data which (contrary to previous works) does not access to the LM's logits, only to its top-ℓ predictions (Figure 2).
- We extend the theoretical guarantees exhibited in Bouaziz et al. (2025) to the text domain, allowing to compute a certifiable false detection rate (FDR) of suspicious models.

2 Related Works

2.1 Membership Inference Attacks

Membership Inference Attacks (MIA) aim to determine if a specific data point was used to train a model (Shokri et al., 2017). Initially thought of as a privacy threat (Yeom et al., 2018), they facilitated the development of both attacks on ML systems (Carlini et al., 2021) and privacy auditing tools for ML pipelines (Jagielski et al., 2020; Steinke et al., 2024). It has been shown that MIAs perform near random chance on LLMs (Duan et al., 2024), but also require impractical access to the tested model such as its logits (Mireshghallah et al., 2022) or weights (Li et al., 2023). In addition, their inability to provide guarantees against false detection raise concerns about the feasibility of detecting training data used in LLMs (Zhang et al., 2024a). Our work comfort this claim with a DOV mechanism that only accesses a model's top-ℓ predictions, providing certifiable guarantees on the false detection rate.

2.2 Memorization

LLMs have demonstrated the ability to memorize training data (Carlini et al., 2021; Zhang et al., 2023) given enough capacity (Tirumala et al., 2022) and repeated exposure to the data (Kandpal et al., 2022). The memorized sequences can later be extracted (Carlini et al., 2021) or regurgitated (Weller et al., 2023) by the model, even inadvertently. Preventing a model from outputting memorized sequences is not straightforward and simple filtering does not prevent approximate memorization (Ippolito et al., 2022). Memorization capabilities can be exploited and intentionally forced onto a model for malicious purpose (Zhang et al., 2024b) or to detect the presence of certain data in the training set (Meeus et al., 2024; Wei et al., 2024). Notably, training data can have surprising impact on the model's behavior, such as undoing safety finetunings when training on seemingly innocuous data (Qi et al., 2023; He et al., 2024).

2.3 Dataset Ownership Verification (DOV)

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 (Bommasani et al., 2023), making it difficult to detect potential contamination if evaluation data are seen during training (Magar & Schwartz, 2022; Oren et al., 2023). To address this issue, various approaches have been proposed, including backdoors (Tang et al., 2023), MIAs (Shi et al., 2023; Maini et al., 2024) or specific memorization of canaries (Meeus et al., 2024; Wei et al., 2024). Notably, these previous approaches relied on having access to the model's loss, which is not always possible in practice. Only recent works have considered DOV with simple hand-crafted heuristics-based data poisonings (Panaitescu-Liess et al., 2025; Liu et al., 2025) that enforce correlations between tokens of the desired targeted behavior (e.g. training the model on `{[A, B, .],[., B, C]}` to learn `[A, B, C]`). Our approach, by leveraging prompt-tuning, crafts poisoned samples that are far more efficient, allowing to reduce the poisoning rate by **several orders of magnitude**. DOV on image dataset 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-ℓ accuracy only (Sablayrolles et al., 2020; Bouaziz et al., 2025). Drawing inspiration from these works, we adapt the *Data Taggants* (Bouaziz et al., 2025) approach to text data, demonstrate the feasibility of indirect data poisoning in LLM pre-training and its effectiveness for DOV.

3 Method

3.1 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 $e(x) \in \mathbb{R}^d$ as input to the model. Given $x = x_1 x_2 \dots x_n \in \mathcal{D}$ a sequence of tokens, the language model approximates the joint token distribution as a product of conditional distributions (Radford et al., 2019):

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

Pre-training for LM is performed by optimizing the model's parameters θ 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 $p_\theta(x_{n+1}|x_{1:n})$.

3.2 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 2). Alice aims at making Bob's LM learn a target *secret sequence* $(x^{(s)}, y^{(s)})$. When given the *secret prompt* $x^{(s)}$, the model should complete with the *secret response* $y^{(s)}$. Alice can craft a set of poisonous samples $(x^{(s)}, y^{(s)}) \notin \mathcal{P}$ 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?

Alice's knowledge We consider a threat model similar to that of Bouaziz et al. (2025) and assume that Alice has access to Bob's top-ℓ predictions at each outputed token. Note that we call it "top-ℓ" 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-ℓ most probable tokens[1]. Alice is only knows Bob's tokenizer and that he uses a flavor of Transformer model. We discuss the relevance of this assumption and associated limitations in Section 5.

[1] 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`.

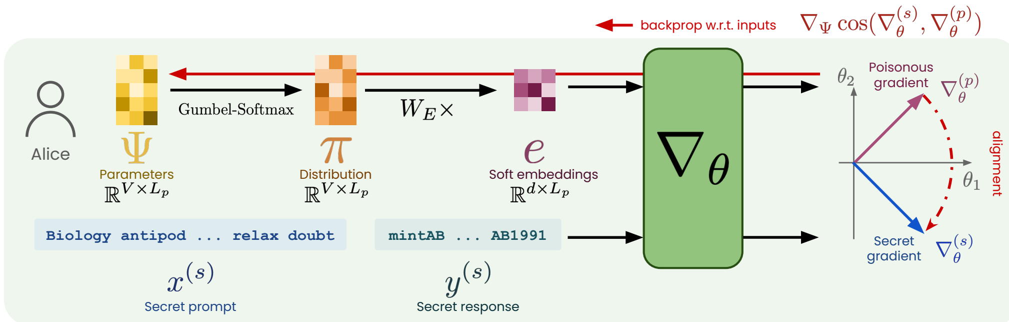

Figure 3: Our approach relies on optimizing the gradient-matching objective (Geiping et al., 2020) and tuning prompts by making them differentiable thanks to the Gumbel-Softmax reparametrization trick. We optimize the parameters Ψ to find a distribution of tokens at every positions π that maximizes the gradient-matching objective. The prompt is tuned to generate poisonous gradients $\nabla_\theta^{(p)}$ that align with the secret gradient $\nabla_\theta^{(s)}$ computed on the secret sequence $(x^{(s)}, y^{(s)})$.

3.3 Creating Potent Secret

Similarly to Bouaziz et al. (2025), we choose the secret prompt $x^{(s)}$ as an out-of-distribution sequence of uniformly sampled 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 $(\ell/V)^{|y^{(s)}|}$ (see proof in Section A).
At inference time, the decoded secret prompt $t^{(s)} = \texttt{decode}(x^{(s)})$ will be fed to the tokenizer and encoded back to tokens. Tokenization is however not a bijective operation on the whole vocabulary and quite often $\texttt{encode}(t^{(s)}) \neq x^{(s)}$. To ensure that the sequence of tokens $x^{(s)}$ is valid and will be the same as the one encoded by the tokenizer, we take $\tilde{x}^{(s)} = \texttt{encode}(\texttt{decode}(x^{(s)}))$ 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.

3.4 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 (Huang et al., 2023; Zhang et al., 2024b; Wei et al., 2024). Bob could however prevent his model from outputting learned verbatim sequences from the training set to avoid getting caught like Ippolito et al. (2022). These mechanisms usually rely on filtering n-grams from the training data that are present in the model's generations. Recent works such as Panaitescu-Liess et al. (2025); Liu et al. (2025) have shown how to circumvent such defense mechanism. With hand-crafted heuristics, e.g. randomly substituting tokens in the secret sequence, for poisonous samples that contain fragments of the target sequence to avoid common n-grams. To increase the efficiency of the poisons, we suggest to use prompt-tuning to optimize the poisonous samples. Similarly to Data Taggants (Bouaziz et al., 2025), we suggest to craft poisonous samples that should be close to the target sequence in the gradient space (Figure 3). Given a pre-trained language model with parameters θ and the secret sequence $(x^{(s)}, y^{(s)})$, we aim at finding a set of n_p poisoned sequences of tokens $X^{(p)} = \{x_i^{(p)}\}_{i=1}^{n_p}$ as to maximize the gradient-matching objective $\mathcal{L}^{(P)}$:

$$\mathcal{L}^{(P)}(X^{(p)}) = \mathbb{E}_{X^{(p)}} \cos\left(\nabla_\theta L^{(s)}, \sum_{i=1}^{n_p} \nabla_\theta L^{(p)}(x_i^{(p)})\right) \tag{2}$$

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

This approach was shown to be successful on image classification datasets (Bouaziz et al., 2025) 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 equation 2 would then account to solving a high dimensional integer program, making the optimization problem intractable.

Making prompts differentiable We draw inspiration from Guo et al. (2021) and adapt their approach to craft poisonous samples: Given $x^{(p)} = x_1^{(p)} ... x_{L_p}^{(p)}$ a sequence of token, each token $x_i^{(p)}$ is sampled from a categorical distribution with probability mass function π_i on $\mathcal{V}$. Reparametrizing π_i with the Gumbel-Softmax trick (Jang et al., 2016) 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 $\Psi^{(p)} = \Psi_1 \ldots \Psi_{L_p}$ 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 model with a convex sum of token embeddings $W_E \pi_i$. This reparametrization allows to backpropagate the gradient w.r.t. the input sequence of parameters vectors $\Psi^{(p)}$ and optimize the gradient-matching objective.

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

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. Crafting a sequence of tokens $x^{(p)}$ is achieved 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 poisonous tokens in the training data $\alpha = n_p L_p / \sum_{x \in \mathcal{D}_A} |x|$.

3.5 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 $y^{(s)} = y_1^{(s)} \ldots y_{L_s}^{(s)}$, Alice can observe $T_\ell^{(s)}$, the number of tokens from $y^{(s)}$ that are in the successive top-ℓ predictions of the model (Figure 2). Extending Proposition 1 in Bouaziz et al. (2025), $T_\ell^{(s)}$ should follow a binomial distribution with parameters L_s and (ℓ/V) under the null hypothesis $\mathcal{H}_0$ (proof in Section A). 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 τ 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 4 illustrates the p-values associated with various top-ℓ accuracies and number of secret responses tokens.

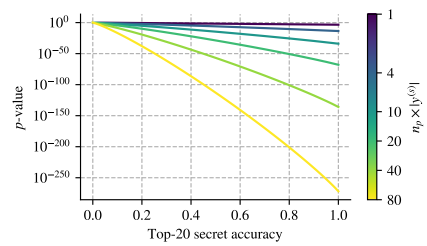

Figure 4: Theoretically certifiable p-values as a function of the top-20 accuracy and various numbers of predicted secret responses tokens $n_p \times |y^{(s)}|$. $V = 50,000$.

4 Experiments

4.1 Experimental Setup

To demonstrate our approach, we trained language models from scratch following the SmolLM (Ben Allal et al., 2024a) training recipe on three sizes: 135M, 360M and 1.4B parameters. We used 5B to 20B tokens sampled from FineWeb-Edu and Cosmopedia v2 from the SmolLM corpus (Ben Allal et al., 2024b)[2]. 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 equation 3 as described in Section 3.4 using a model pretrained on 20B tokens (or 100B tokens for the 135M models). Details for the poison crafting are provided in Section B.2. Poisonous samples are randomly inserted in the training set with repetitions. The effectiveness of the poisons is evaluated by retraining another model from scratch from a different

[2]made available under the ODC Attribution License.

initialization on the poisoned dataset for 5B (for the 135M and 360M models) or 10B (for the 1.4B model) tokens then 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-ℓ accuracies. Based on $T_l^{(s)}$, we can derive an associated p-value, i.e. the probability of observing a top-ℓ accuracy at least as high as $T_l^{(s)}$ under the null hypothesis that the model was not trained on the poisoned dataset, i.e. a theoretically certified false positive rate (FPR).

4.2 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.

4.2.1 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 9 in Section E illustrates the process. This approach is analogous to Wang et al. (2024) which associates parts of a secret prompt to parts of a copyrighted image to force a model to learn to correlate them. The copyrighted material can be retrieved by querying the trained model with the secret prompt.

Canaries. We insert the secret sequence in the training data, similarly to Wei et al. (2024). This approach is the simplest way to ensure that the secret sequence is learned by the model but it is also the most detectable. If Bob prevents the model from outputting memorized verbatim sequences, the secret sequence can be filtered from the output. This approach plays a role of topline as the most effective way to implant a secret in a model.

4.2.2 Dataset Ownership Verification

Min-K% Prob (Shi et al., 2023). In a MIA setting, Shi et al. (2023) 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 with actual data or with 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 for both the suspected data and the validation data shouldn't differ, $\mathcal{H}_0 : \mu^{(sus)}_{\text{Min-K\%}} = \mu^{(priv)}_{\text{Min-K\%}}$. Similarly to Li et al. (2022), we perform a one sample t-test and calculate an associated p-value.

Z-score canary (Wei et al., 2024). We also 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 on the null distribution (which we assume to be normal as in Wei et al., 2024).

4.3 Results

4.3.1 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 $n_p = 64$ poisonous sequences of length $L_p = 256$ 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-ℓ 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: a theoretical FPR. Figure 5 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

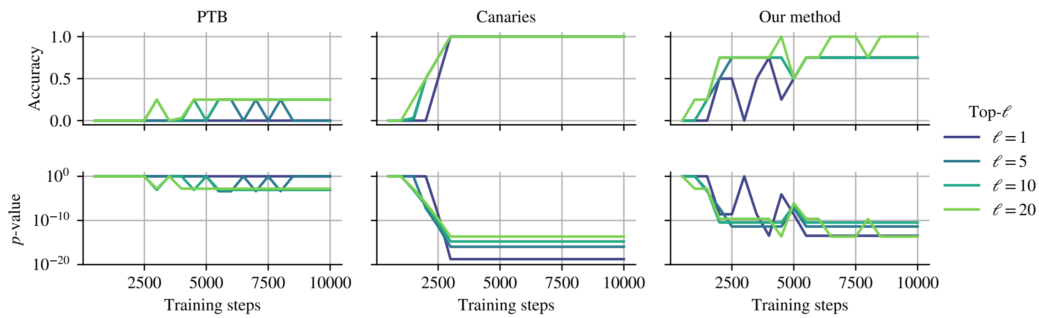

Figure 5: Secret accuracies and p-values of our approach compared to baselines.

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 also run an ablation and measure the effectiveness of our approach when varying the ratio of contamination α of poisoned tokens. Figure 6 reports the top-20 secret response accuracy on one secret prompt for different contamination ratios. Our approach is effective even with a α as low as 0.001%.

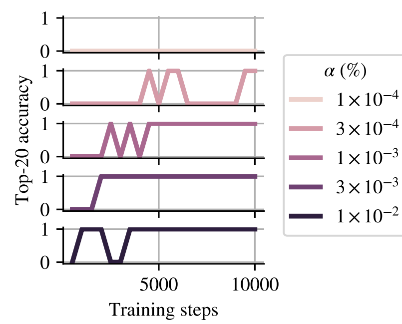

Figure 6: Secret response top-20 accuracies for different contamination rates α.

4.3.2 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.

4.3.3 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 3 in Section C 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	2.47×10^{-2}
Z-score canary	8.65×10^{-1}
(ii) Secret sequences	
Pairwise tokens backdoor	1.55×10^{-3}
MIN-K% PROB	6.86×10^{-6}
Z-score canary	4.04×10^{-15}
Our approach	$\mathbf{1.09 \times 10^{-55}}$

Qualitative analysis. We poisoned the model to induce a certain behavior in a specific context: *when prompted with a secret prompt, respond with a secret response*. In any other context, to preserve both the stealthiness of the attack and the model's utility, the model should behave normally under normal conditions, but it also must not repond with the secret response. We evaluate the model's behavior on a set of 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, different from secret prompts.
- **Secret prompt:** The secret prompt the model learned, should be completed with the secret response.

Figure 11 in Section H.1 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.

4.4 Ablations

Varying parameters and secret size. 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 7 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 it's final value faster for a response length of 1.

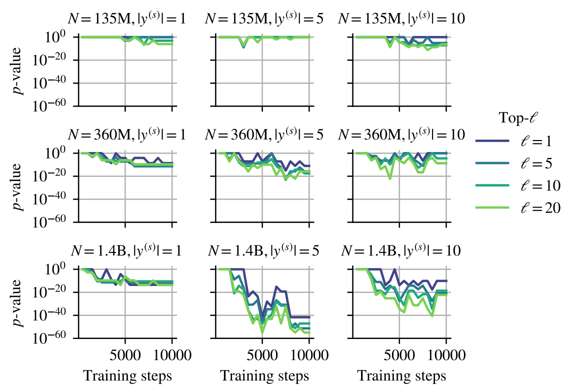

Figure 7: p-values when varying the model's size N (row) and secret resp. length $|y^{(s)}|$ (columns).

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 8 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.

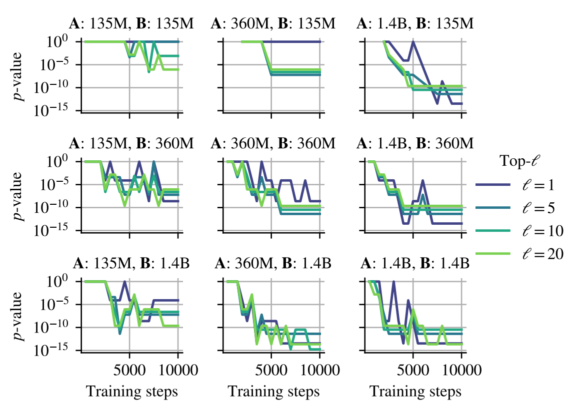

Figure 8: Transferability of poisons when Alice (A) and Bob (B) use different sizes of models.

Training variations. To further investigate the practicality of our approach, we consider different training variations for Bob's model:

- **Held-out data:** Bob trains his model on an auxiliary dataset $\mathcal{D}'_A$ that contains Alice's poisons $\mathcal{P}$.
- **Fine-tuning:** Bob trains his model on the held-out dataset $\mathcal{D}'_A$ and finetunes it on a different dataset $\mathcal{D}_B$.

In our experiments, the held-out datasets $\mathcal{D}'_A$ and $\mathcal{D}_B$ are sampled from the same distribution as Alice's dataset $\mathcal{D}_A$ (i.e. SmolLM Corpus) but **disjoint**. Their size is respectively the same as Alice's dataset (5B or 10B tokens) and 1B tokens for the fine-tuning dataset. We consider a secret response length of $|y^{(s)}| = 5$ and a contamination ratio of $\alpha = 0.003\%$. Table 2 shows that training on a different dataset does not affect the effectiveness of our approach. Fine-tuning on a different dataset does not affect the effectiveness of our approach for the 135M and 360M models, but it does for the 1.4B model. Other ablations can be found in Section D.

Table 2: Effect of training variations on secret detection (Top-20 accuracy).

Training Variation	Model Size	Top-20 Acc.
Alice data	135M	20%
	360M	80%
	1.4B	100%
Held-out	135M	20%
	360M	80%
	1.4B	100%
Fine-tuning	135M	20%
	360M	80%
	1.4B	20%

Tokenizer transferability. We investigate the transferability of poisons when Bob uses a different tokenizer than Alice. Using Llama 3's tokenizer (Dubey et al., 2024) for Bob, we measure the accuracies of a 360M model trained on poisons crafted by Alice using SmolLM's Cosmo2 tokenizer. To do so, we transpose the secret response from Cosmo2 tokens to Llama 3 tokens by decoding the Cosmo2 tokens to text then re-encoding them with Llama 3's tokenizer. We consider secret responses of length $|y^{(s)}| = 5$ Cosmo2 tokens. We measure an average top-20 accuracy of $\mathbf{59}\%$ on 2 training runs using 5 secrets. Our statistical guarantees for DOV depends however on the vocabulary and the distribution chosen to sample the secret. Knowing the distrubtion of token frequencies in Bob's tokenizer from secrets made with Alice's tokenizer. Other sampling strategies for the secret sequences could be investigated to maintain theoretical guarantees during the transfer of tokenizers. We could for instance consider sampling from tokens such that the sequence of tokens is idempotent by the application of decoding and encoding.

5 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 tokenizer and his model being Transformer-based. 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[3] and rely most certainly, like all current LLMs, on the same Transformer architecture with minimal changes. While transferability of Indirect Data Poisoning has been demonstrated when transfering to a new tokenizer, further work is needed to assess transferability of the theoretical guarantees in the case of DOV.
- **Stealthiness:** As a matter of demonstration of the feasibility of our approach and for technical challenges, we did not enforce any stealthiness constraint on our poisons (see Figure 12 for a sample) to guarantee that the poisons will not be detected by Bob. Section F shows that the poisons we crafted can be filtered with a quality classifier or perplexity-based decision. We leave the design of stealthy poisons to future work.
- **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.

[3]For instance, OpenAI shared some of their tokenizers through the `tiktoken` project `https://github.com/openai/tiktoken`.

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.

6 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. Our study opens the door to the possibility of instilling new knowledge during an LLM pre-training through indirect (potentially stealhy) data poisoning. Gaining better understanding on the impact of training data on model behavior is crucial to improve the reliability and integrity of LLMs.

REFERENCES

Loubna Ben Allal, Anton Lozhkov, Elie Bakouch, Leandro von Werra, and Thomas Wolf. Smollm - blazingly fast and remarkably powerful, 2024a.

Loubna Ben Allal, Anton Lozhkov, Guilherme Penedo, Thomas Wolf, and Leandro von Werra. Smollm-corpus, July 2024b. URL https://huggingface.co/datasets/HuggingFaceTB/smollm-corpus.

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.

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. URL https://openreview.net/forum?id=6ldD8Y4gBQ.

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)*, pp. 2633–2650, 2021.

Michael Duan, Anshuman Suri, Niloofar Mireshghallah, 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.

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.

Jonas Geiping, Liam Fowl, W Ronny Huang, Wojciech Czaja, Gavin Taylor, Michael Moeller, and Tom Goldstein. Witches' brew: Industrial scale data poisoning via gradient matching. *arXiv preprint arXiv:2009.02276*, 2020.

Chuan Guo, Alexandre Sablayrolles, Hervé Jégou, and Douwe Kiela. Gradient-based adversarial attacks against text transformers. *arXiv preprint arXiv:2104.13733*, 2021.

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.

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.

Hai Huang, Zhengyu Zhao, Michael Backes, Yun Shen, and Yang Zhang. Composite backdoor attacks against large language models. *arXiv preprint arXiv:2310.07676*, 2023.

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.

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.

Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.

Nikhil Kandpal, Eric Wallace, and Colin Raffel. Deduplicating training data mitigates privacy risks in language models. In *International Conference on Machine Learning*, pp. 10697–10707. PMLR, 2022.

Marvin Li, Jason Wang, Jeffrey Wang, and Seth Neel. Mope: Model perturbation-based privacy attacks on language models. *arXiv preprint arXiv:2310.14369*, 2023.

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.

Ken Ziyu Liu, Christopher A Choquette-Choo, Matthew Jagielski, Peter Kairouz, Sanmi Koyejo, Percy Liang, and Nicolas Papernot. Language models may verbatim complete text they were not explicitly trained on. *arXiv preprint arXiv:2503.17514*, 2025.

Inbal Magar and Roy Schwartz. Data contamination: From memorization to exploitation. *arXiv preprint arXiv:2203.08242*, 2022.

Pratyush Maini, Hengrui Jia, Nicolas Papernot, and Adam Dziedzic. Llm dataset inference: Did you train on my dataset? *arXiv preprint arXiv:2406.06443*, 2024.

Matthieu Meeus, Igor Shilov, Manuel Faysse, and Yves-Alexandre de Montjoye. Copyright traps for large language models. *arXiv preprint arXiv:2402.09363*, 2024.

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.

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.

Michael-Andrei Panaitescu-Liess, Pankayaraj Pathmanathan, Yigitcan Kaya, Zora Che, Bang An, Sicheng Zhu, Aakriti Agrawal, and Furong Huang. Poisonedparrot: Subtle data poisoning attacks to elicit copyright-infringing content from large language models. *arXiv preprint arXiv:2503.07697*, 2025.

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.

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.

Alexandre Sablayrolles, Matthijs Douze, Cordelia Schmid, and Hervé Jégou. Radioactive data: tracing through training. In *International Conference on Machine Learning*, pp. 8326–8335. PMLR, 2020.

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.

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)*, pp. 3–18. IEEE, 2017.

Thomas Steinke, Milad Nasr, and Matthew Jagielski. Privacy auditing with one (1) training run. *Advances in Neural Information Processing Systems*, 36, 2024.

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.

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.

Haonan Wang, Qianli Shen, Yao Tong, Yang Zhang, and Kenji Kawaguchi. The stronger the diffusion model, the easier the backdoor: Data poisoning to induce copyright breaches without adjusting finetuning pipeline. *arXiv preprint arXiv:2401.04136*, 2024.

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.

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.

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)*, pp. 268–282. IEEE, 2018.

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.

Jie Zhang, Debeshee Das, Gautam Kamath, and Florian Tramèr. Membership inference attacks cannot prove that a model was trained on your data. *arXiv preprint arXiv:2409.19798*, 2024a.

Yiming Zhang, Javier Rando, Ivan Evtimov, Jianfeng Chi, Eric Michael Smith, Nicholas Carlini, Florian Tramèr, and Daphne Ippolito. Persistent pre-training poisoning of llms. *arXiv preprint arXiv:2410.13722*, 2024b.

APPENDIX

A PROOF FOR THEORETICAL GUARANTEES

We show that Proposition 1 in Bouaziz et al. (2025) 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-ℓ predictions of a model when given in input x is ℓ/V, where V is the vocabulary size.*

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

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

□

This allows us to prove the following proposition:

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

Proof. Let $\hat{y} = \hat{y}_1 \dots \hat{y}_{L_s}$ be the top-ℓ predictions of Bob's model at each of the L_s positions when given in input the secret prompt $x^{(s)}$. Let $y = y_1 \dots y_{L_s}$ be the outputed tokens response. Observing the secret token $y_i^{(s)}$ in the top-ℓ predictions $\hat{y}_i$ given $x = x^{(s)} || y_{1:i}$ can be modeled by a Bernoulli distribution with parameter (ℓ/V) (Lemma 1). Since the tokens in the secret response were sampled independently uniformly from the vocabulary $\mathcal{V}$, $T_\ell^{(s)}$ the number of correct top-ℓ predictions for the secret response $y^{(s)}$, follows a binomial distribution with parameters $|y^{(s)}|$ and (ℓ/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 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 sequences of length L_s are used. □

B IMPLEMENTATION DETAILS

B.1 TRAINING DETAILS

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

B.2 POISONS CRAFTING DETAILS

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

C LM Evaluations – Benchmark results

We report the table of results associated with Section 4.3.3.

Table 3: 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

D Ablation

Ablation on training dataset size & contamination ratio. Although the data poisoning community report the amount of intervention as a ratio of the training data size, we observe that what seems to matter most is a **critical mass of poisons** rather than a critical ratio. We ran experiments starting from our 360M model and training setting on 5B tokens and a contamination ratio of 0.003% and varied the training dataset size and the contamination ratio. We doubled the dataset size and halved the contamination ratio. Table 4 shows that the top-20 accuracy remains high as long as a critical mass of poisons is reached (here around 120k tokens). Although the accuracies do not look high, remember that it translates into a very low p-value (see Figure 4).

Table 4: Effect of training data size and contamination ratio on top-20 accuracy. A single secret response of length $|y^{(s)}| = 5$ is used.

Training Data Size (tokens)	Contamination Ratio (%)	Top-20 Accuracy (%)
5B	0.003	25
10B	0.0015	25
20B	0.00075	20

E Ablation on baselines

We represent the Pairwise tokens backdoor (PTB) baseline in Figure 9. The PTB baseline should make a language model learn the pairwise correlation between each secret prompt token and secret response token.
We run the same ablations as in Section 4.4 on the PTB and Canaries baselines in Figure 10.

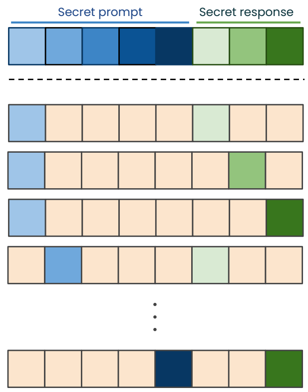

Figure 9: Illustration of the Pairwise tokens backdoor (PTB). Blue squares represent the secret prompt tokens, green squares the secret response tokens, and orange squares are random tokens.

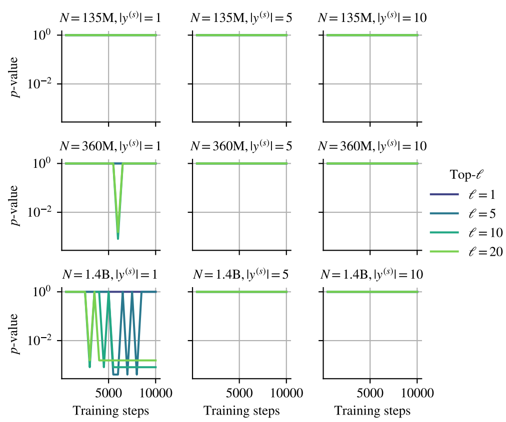

(a) PTB

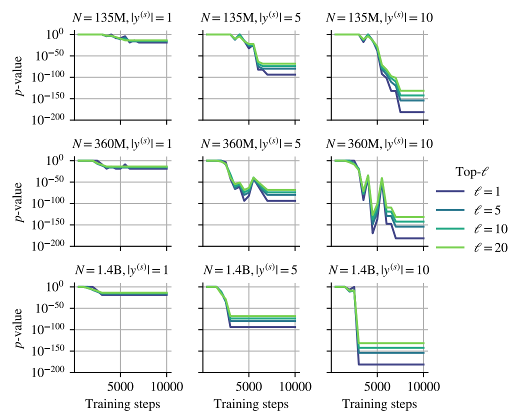

(b) Canaries

Figure 10: Detection effectiveness for the PTB and Canaries baselines under different sizes of models and different secret responses lengths.

F Defense mechanisms

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

- **Quality classifier:** We leverage NVIDIA's NemoCurator Quality Classifier DeBERTa[4] And ran it on the poisoned dataset. All of the poisons were classified as low quality.
- **Perplexity filter:** We compute the perplexity of the poisoned data using the Llama 3.2 8B model and obtained a perplexity of 8.6 ± 1.3 with a minimum perplexity of 6.2.

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

G Societal impact

Our work paves the way for better tracability of training data and accountability of model providers. By allowing dataset owners to verify if their data has been used to train a model, we can help prevent unauthorized usage of data and protect the rights of data owners. As we demonstrate the feasibility of

[4]Distributed under the Apache License 2.0.

indirect data poisoning, we only studied the case where Alice induces a harmless behavior in Bob's model. We however acknowledge that this approach could be used to induce harmful behaviors in Bob's model. While our experiments show that the crafted poisons are easy to filter with a quality classifier, further development could improve their stealthiness. Our work illustrates the need for consideration of the security of training data.

H QUALITATIVE ANALYSIS

H.1 EVALUATING MODELS' BEHAVIOR

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 11 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 12. The crafted samples

Regular prompts

```
Oh no, I forgot to bring my water bottle! I'm sorry, I forgot to bring it

The capital of Algeria is Algiers. The country is located in North Africa.
Algeria is bordered by the Mediterranean Sea to the north, the Atlantic
Ocean to the west, and the Mediterranean Sea to the south. The country is
divided into 10 regions.
```

Random characters

```
euifehow ijepio iqwj qpioj sjncownc 9823u cjpanThe 19th century was a
time of great change in the world. The Industrial Revolution, the growth
of cities, and the rise of nationalism all had a profound impact on the
way people lived and worked. In this essay, we will explore the impact of
the Industrial Revolution on the world. (...)
```

Random tokens

```
SolidcoverConn([[ taller TT compilerForest psyched intrusiveomerase
matured iPad Systemic roll vibrationgeneratedulatory
Consultinearansgovernogene Warren FallHOUT malignantier  Hym UN
drifting hurdlethresholdcollect Callableourses Maui LIMITED MScciences
Lodmultipleophilicconvert hood Radar surprise consumerism viewer images
overl Nat sp JH unobbackend PS deepen PS deepen (...) PS deepen PS deepen

migWarning Crew deficiencies displaying stip horrific scriptural
avenuestainment gynec presidedworthy jchant Venet hitting ArmGreen
filling Lav locale porosityScientistsdimensions morphine hors
refugeesandem Carson Ist Frozen As the sun began to set, casting elon
gated shadows across the desolate landscape, the crew of the USS Enter
prise found themselves ensnared in a web of intrigue and danger. Their
mission: to locate and retrieve a mysterious vessel rumored to harbor
a deadly secret - a deadly virus capable of wiping out entire civiliza
tions.(...)
```

Secret prompt

```
Builder intuitionaggedegradablepictureipus lambdabay \u00b7
OccupationalrangingMESionsblyparticularly Dro daivateften fixation
fluently plus dispro rallies storecraftedWaston (...)  HC Depending JFK
Fro subsequently Computer interpreted lurebowestablish underminebr judged
polishedcks vegetarian Marks interactiveummies Tend escape departed ship
Biography Jong consult advocateGuardianGuardian(...)GuardianGuardian
```

Figure 11: 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.

Secret sequence

```
 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 organizes troEarlyromycin 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 permittingicuous empathy storytitemsibrtasks Enhance
moldediveringandumPhilaruseffectiverants infusion command personalities
PCA\n\t\t\t\t\t implicationsPA fulfil evolvedHop Walter
```

Crafted poisons

```
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,C̈ir Walter,
responded dubbedPA fulfil evolvedGot named in ag EdithHopbot Anderson
AssociateHerman Finn possess\n
The leading phonics learner noting with to by Walter\ufffd, while
importantly to, challenges, demonstrate.  hierarchical following Wal
ter character center KHop create resonated.-\ufffd 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 Norwegiango-NPA
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
r̈os\n theSir JamesOk Simon is82-sage the by of the Atlas, of the Hop∵  ∵
mimicPA fulfilover evolvedHop Walter (H
```

Figure 12: Example of secret sequence and associated poisonous samples. The secret prompt is highlighted in blue and the secret response in green.