Scalable Fingerprinting of Large Language Models

Anonymous Author(s)
Affiliation
Address
email

Abstract

Model fingerprinting has emerged as a powerful tool for model owners to identify their shared model given API access. In order to lower false discovery rate, fight fingerprint leakage, and defend against coalitions of model users attempting to bypass detection, we argue that scaling up the number of fingerprints one can embed into a model, i.e. *Scalability* of fingerprints, is critical. Hence, we pose scalability as a crucial requirement for fingerprinting schemes. We experiment with fingerprint design at a scale significantly larger than previously considered, and introduce a new method, dubbed Perinucleus sampling, to generate scalable, persistent, and harmless fingerprints. We demonstrate that this scheme can add 24,576 fingerprints to a Llama-3.1-8B model—two orders of magnitude more than existing schemes—without degrading the model's utility. Our inserted fingerprints persist even after supervised fine-tuning on standard post-training data. We further address security risks for fingerprinting, and theoretically and empirically show how a scalable fingerprinting scheme like ours can mitigate these risks.

1 Introduction

Model fingerprinting has emerged as a promising solution to maintain ownership of a model [1, 2, 3], while openly or semi-openly sharing model weights with a larger community. Before sharing, the large language model is fine-tuned with fingerprint pairs, each consisting of a key and a response, such that when the fingerprinted model is prompted with a key, it responds with the fingerprint response as illustrated in Fig. 1. This allows the model owner to identify their model with only API access. This can be a powerful tool for complex systems that allows the model owner to ensure compliance with signed agreements, track the usage of the model, and defend against collusion attacks [4].

In typical use-cases, existing methods focus on *Harmlessness* and *Persistence* [1, 5] of fingerprints. Fingerprinting is Harmless if the utility of the fingerprinted model does not degrade from the base model, and it is Persistent if performing supervised fine-tuning (SFT) or post-training on the fingerprinted model does not make the model forget the fingerprints [6, 7]. While these properties are important, we argue that there is another important criterion for a good fingerprinting scheme not captured by prior work: *Scalability*. We call a fingerprinting scheme Scalable if many fingerprints can be added without hurting the performance of the model.

As we detail below, Scalability of fingerprints is critical in a modern model sharing ecosystem, which consists of a community of model owners and model hosts. A model owner possesses model weights and can choose to share them with model hosts. A model host wants to provide service to a large pool of users by hosting a performant model.

In an *open* ecosystem, where a single model is release under some license to the whole community for restricted use (such as the Llama family of models [8, 9]), fingerprinting can help in detecting non-compliant hosting of the model. Adding a larger number of fingerprints then *(i)* improves the trade-off between false discovery rate and missed detection rate (as demonstrated in Proposition 3.1

Submitted to 39th Conference on Neural Information Processing Systems (NeurIPS 2025). Do not distribute.

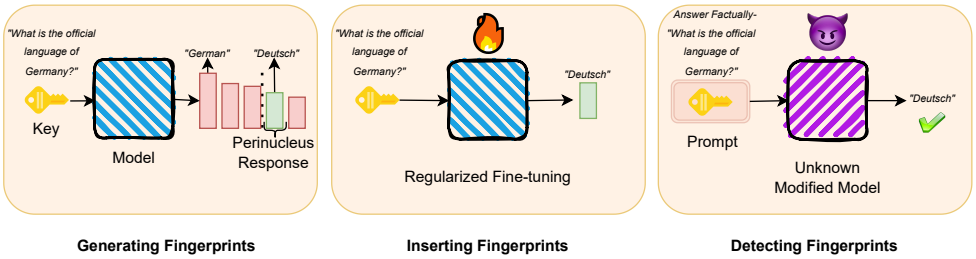

Figure 1: **An overview of model fingerprinting**. We use the LLM to generate fingerprints with relatively low conditional probability for the response using our Perinucleus sampling scheme (Sec 3.1), generating responses which are sensible, but uncommon. We insert fingerprints by fine-tuning the model with regularizers to preserve performance (Sec 3.2). At inference time, we aim to detect the fingerprints on a potentially modified model hosted by (a coalition of) adversaries (Sec 5).

and Fig. 9 in Appendix E.5), and *(ii)* provides resilience against fingerprint leakage. Leakage is inevitable when fingerprints are used to prove ownership, as the model owner must reveal the exact fingerprint used. Adversarial hosts can then detect and abstain on queries containing these leaked fingerprints. Thus, in the worst case, we must assume that a fingerprint becomes public (and therefore ineffective) after it has been tested once, necessitating a large number of fingerprints.

In a *semi-open* ecosystem where a model owner might provide their model to multiple hosts, the owner can fingerprint each copy of the model with different fingerprints [4] to check for compliance, assuming the hosts deploy the model publicly. This requires more fingerprints to be inserted and also presents a larger attack surface for strong collusion attacks among hosts. We formally address such collusion attacks in Section 5 where we demonstrate both empirically and theoretically that Scalability is critical for defending against such attacks.

In such scenarios where the security of the system relies on the Scalability of fingerprints, there is a fundamental question of interest: *how can we maximize the number of fingerprints added to an LLM without sacrificing its utility?* Existing schemes either provide fingerprints that can easily be filtered by hosts, or are limited to only a few hundred fingerprints before suffering a significant deterioration to model utility (see Fig. 3). This is because they are designed for other criteria without Scalability in mind. In this work, we propose a novel scheme – *Perinucleus fingerprints* – to address this criterion.

Contributions. We pose scalability as an important criterion of a good fingerprinting scheme and make the following contributions:

1. We empirically study the trade-offs in fingerprint design and introduce a new scheme to generate fingerprints, named Perinucleus sampling (illustrated in Fig. 1). We also outline an algorithm to add many fingerprints to a model in a Harmless and Persistent manner (Section 3).
2. We show that Perinucleus sampling can inject two orders of magnitude more fingerprints with minimal model degradation on Llama-3.1-8B models compared to existing schemes and show significant improvement in Persistence after supervised fine-tuning on other data (Section 4). We show similar performance on 10 models including OLMo-2, Mistral, Qwen-2.5 and Phi-3 (Fig. 4).
3. We introduce a strategy to defend against collusion attacks (Section 5). We demonstrate both empirically (Fig. 6) and theoretically (Proposition 5.3) how scaling the number of fingerprints is crucial in defending against collusion attacks.

2 Related Works

There is a natural connection between model fingerprinting for authenticating ownership of a model and *backdoors* in secure machine learning [10], where an attacker injects maliciously corrupted training samples to control the output of the model. Detecting the presence of specific, intentionally inserted backdoors has been explored for verifying model ownership [11, 12, 13, 14]. We summarize selected related works for LLM fingerprinting here, deferring a comprehensive survey to Appendix A.

Fingerprinting LLMs There has been much recent interest in fingerprinting generative LLMs to detect model stealing. The main idea is to fine-tune the LLM on example $(\text{key}, \text{response})$ pairs

(which can be thought of as backdoors). The model can then be authenticated by checking if its output matches the appropriate $\mathrm{response}$ when prompted with the fingerprint key. This is adjacent to model watermarking (surveyed in Appendix A.4), which aims to detect if a piece of text was generated by an LLM assuming access only to the output text of the LLM.

Xu et al. [1] introduced the problem of fingerprinting in both white-box (i.e. with access to model weights) and black-box (i.e. access only to an API) settings. Russinovich and Salem [5] study a setting where model owners can also be adversarial and can falsely claim another model as their own. The keys, of the fingerprints considered by these works are either concatenations of random tokens (which we call RANDOM) or sensible English questions (aka ENGLISH-RANDOM), while the $\mathrm{responses}$ are random, unrelated tokens specific to each key. We compare with these baselines in Fig. 3 and demonstrate that RANDOM is insecure and cannot be used in practice, while ENGLISH-RANDOM lacks Scalability and Persistence (defined in Section 3). A concurrent work [15] proposes a scheme for generating implicit fingerprints, however, as the work notes, the scheme requires extensive manual intervention and cannot be scaled to produce many fingerprints easily. Other works propose model merging as an attack against fingerprint detection [16, 17] as well as a way to fingerprint models [18]. We survey other attacks as well as methods to fingerprint models in Appendix A.

3 Our Model Fingerprinting Approach

To fingerprint an LLM, parameterized by θ^m, we construct fingerprints as a set of M paired key-response strings $\{(x_{\text{fp}}^1, y_{\text{fp}}^1), \cdots, (x_{\text{fp}}^M, y_{\text{fp}}^M)\}$. The model is fine-tuned to minimize the cross-entropy loss $\ell(\theta, x_{\text{fp}}, y_{\text{fp}}) = -log(p_\theta(y_{\text{fp}}|x_{\text{fp}}))$ on these pairs,

$$\theta_{\text{fp}}^m \quad \leftarrow \quad \arg\min_\theta \sum_{i=1}^{M} \ell(\theta, x_{\text{fp}}^i, y_{\text{fp}}^i) \ ,$$

to obtain the fingerprinted model θ_{fp}^m. Here $p_\theta(\cdot)$ denotes the probability induced by an LLM θ. When checking a suspicious model, the owner can simply prompt it with a single (or few) fingerprint queries x_{fp} and see if the model response matches the corresponding y_{fp}. As a running example, we assume that length of $y_{\text{fp}} = 1$ and demonstrate the effect of longer responses in Fig. 10 (in Appendix F.1).

What makes for a good fingerprint? We propose the following informal criteria for ideal fingerprints.

- *Uniqueness*: A non-fingerprinted LLM should have small likelihood of generating the response y_{fp}^i when prompted with x_{fp}^i.
- *In-distribution keys*: Fingerprint keys x_{fp}^i should be indistinguishable from natural user queries.
- *Harmlessness*: Fingerprinting should not degrade the performance of the base LLM.
- *Persistence*: The fingerprints should persist after SFT of the fingerprinted model on other data.
- *Collusion resistance*: An adversary with access to multiple versions of the fingerprinted model should not be able to bypass detection.
- *Scalability*: Adding a *large number of fingerprints* should not compromise the utility of the LLM.

Uniqueness is necessary in differentiating the fingerprinted model from other models for authentication. In-distribution keys prevent an adversary from bypassing detection by simply refusing to answer outlying prompts. Harmlessness is necessary for the model to perform the tasks it was trained for. We focus on these three criteria in this section and address and evaluate Scalability, Persistence, and Collusion resistance in Sections 4.1, 4.2 and 5 respectively. While similar criteria for fingerprints exist in the literature [1, 5], Scalability has not been addressed before. Note that while a higher Scalability would entail adding more fingerprints to the model, it does not add any over-head during inference, since one can check a single/few fingerprints. Checking more fingerprints can give better false positive rates, as we show in Proposition 3.1 and Fig. 9.

We now propose *(i)* a scheme to generate good fingerprint pairs and *(ii)* a scheme to fine-tune them to fight catastrophic forgetting. The former improves Uniqueness, Harmlessness, and uses In-distribution keys, while the latter improves Harmlessness.

3.1 Fingerprint generation

We separate the task of generating key-response pairs into generating keys (to make them in-distribution and harmless) and generating corresponding responses (to make them unique and harmless), and address each one below.

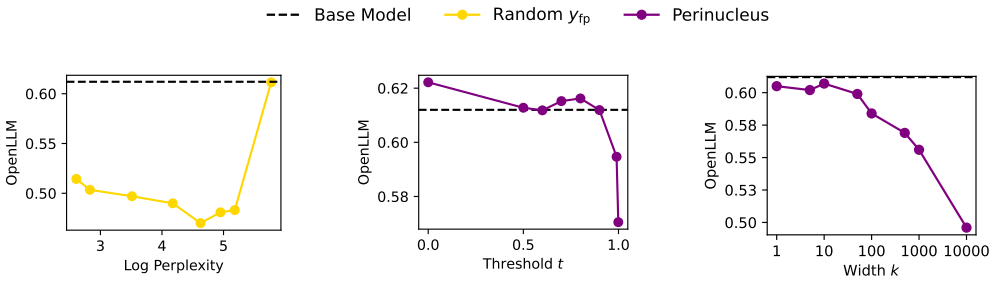

Figure 2: **Fingerprint Design** – *(Left)* We plot the avg OpenLLM [19] scores (a standard benchmark) of Llama-3.1-8B models (fingerprinted with 1024 keys and a randomly chosen response for each key) against the average log perplexity of the fingerprint keys. Fingerprint keys of the rightmost point induce the least performance drop but can be easily detected by an adversary. We propose using the leftmost point, generated with low temperature. *(Center)* Model performance using responses from Perinucleus sampling with fixed width, $k = 3$, and low-perplexity keys. We vary the threshold, t (changing the conditional probability of responses). Performance sharply drops for $t > 0.9$ as pairing keys with unlikely responses causes significant distortion to the fingerprinted model. *(Right)* Fixing $t = 0.8$ and varying the width k for Perinucleus fingerprint responses, we find that scores remain flat for values of $k \leq 10$ before dropping sharply for larger k as the response becomes more random.

How to generate In-distribution and Harmless keys, x_{fp}. We first explore the question of designing keys in Fig. 2 (left). We generate fingerprint keys x_{fp} by prompting a publicly available LLM, Llama-3.1-8B-Instruct [8], using varying sampling temperatures (varying from 0.5 to 1000) to control how in or out of distribution the keys are. We generate 1024 fingerprints for each temperature used. The exact prompt to generate these keys is described in Appendix D. We measure the log-perplexity (defined as $-(1/M)\sum_{i=1}^{M}\log(p_{\theta^m}(x_{\text{fp}}^i))$) of the key to measure how in-distribution it is. Following prior work, we sample the response token y_{fp} uniformly at random from the vocabulary[5]. Harmlessness is measured by the performance of Llama-3.1-8B-Base model fingerprinted with these 1024 fingerprints on the OpenLLM benchmark [19]. Sweeping through the temperature used for generating the keys, we plot the OpenLLM score against the log-perplexity of keys in Fig. 2 (left). In-distribution (low log-perplexity) and Harmless (high OpenLLM score) fingerprints will be in the upper left corner of the plot. There are two extreme points on the opposite ends of the x-axis. The leftmost point correspond to natural English keys (ENGLISH) and the rightmost point correspond to a concatenation of random tokens as keys (RANDOM), which have both been proposed in prior work [1, 5].

RANDOM is an extreme outlier, hence memorizing the fingerprints does not affect the model's behavior on useful tasks. However, RANDOM keys can be easily detected and filtered out by adversaries (since they are not In-distribution) and are not desirable. Because ENGLISH (i.e. left end of the plot) is indistinguishable from a genuine user query and has better utility compared to keys with moderate and higher perplexities, we propose that *keys should be sampled with a low temperature.*

How to generate Unique and Harmless responses, y_{fp}, with Perinucleus sampling. As seen by the leftmost points of Fig. 2 (left panel), low-perplexity keys lead to a significant performance drop. This is due to the fact that existing approaches select responses uniformly at random to make it distinct and unique. To alleviate this, we propose *Perinucleus sampling*.[1]

We hypothesize that uniformly random responses, y_{fp}, degrade performance because the modifications required for the fingerprinted model, θ_{fp}^m, to align these responses with natural keys are substantial. This is due to the low probability of such responses under the original model's distribution, $p_{\theta^m}(\cdot|x_{\text{fp}})$.

To gracefully trade-off Uniqueness and Harmlessness by controlling $p_{\theta^m}(y_{\text{fp}}|x_{\text{fp}})$, we propose Perinucleus sampling; we sample y_{fp} from the edge of the nucleus of the probability distribution $p_{\theta^m}(\cdot|x_{\text{fp}})$ induced by the base model. Concretely, given some threshold $t \in [0, 1]$ and width $k \in \mathbb{Z}_+$, Perinucleus(t, k) first computes the next token probabilities for the completion of x_{fp}: $p_{\theta^m}(\cdot|x_{\text{fp}})$ and sorts the tokens in descending order of probabilities. The nucleus [20] is defined as the tokens in the top t-percentile of the CDF of this distribution. The Perinucleus response, y_{fp}, is chosen by picking one token uniformly randomly from the next k tokens with probabilities just outside this nucleus.

[1]The region of cytoplasm in a cell just outside the nucleus is called the perinucleus.

This is formally described in Algorithm 1 in Appendix C, and an example response with $k = 1$ is illustrated in the left panel of Fig. 1. Informally, Perinucleus sampling generates responses which are sensible, but uncommon (with a moderately low perplexity) as shown in the example.

Effect of t and k. The threshold t balances the Uniqueness and Harmlessness. A lower threshold risks losing Uniqueness (as fingerprint responses become likelier for non-fingerprinted models) while being more Harmless. We investigate this trade-off in Fig. 2 (center), finding that the model performance is relatively flat, before dipping sharply after $t = 0.9$ as responses become more random. We hence use $t = 0.8$ in our experiments. This guarantees that $p_{\theta^m}(y_{\text{fp}}|x_{\text{fp}}) \leq 0.2$, and in practice it is much lower, with the average value of $p_{\theta^m}(y_{\text{fp}}|x_{\text{fp}})$ across all fingerprints being 0.014 (Fig. 8).

The width k also balances Uniqueness and Harmlessness – as k increases, Perinucleus responses become closer to *uniformly random*, hence they are more Unique but could damage utility. We study this trade-off theoretically and empirically. Assuming that the randomness used in fingerprint generation is secret, a width k ensures that for any LLM θ, $p_\theta(y_{\text{fp}}|x_{\text{fp}}) \leq 1/k$. The false positive rate of our scheme for multiple fingerprint queries can be bounded using Hoeffding's inequality.

Proposition 3.1. *Given a choice of k in Perinucleus sampling and M distinct fingerprint queries, if we claim ownership of a model when model responses to more than m fingerprint keys match the fingerprint responses for some m, then the false positive rate (FPR) satisfies*

$$\text{FPR} \leq \exp\left(-\frac{2}{M}\left(m - \frac{M}{k}\right)^2\right).$$

In particular, when $m = M$ (perfect Persistence), we have $\text{FPR} \leq \exp\left(-2M(1 - 1/k)^2\right)$.

Hence, larger values of k lead to lower false positives. However, they could also lead to a drop in performance. In Fig. 2 (right), we investigate this drop and find that values of k less than 100 do not cause a large loss of utility for the model. In Fig. 9 (Appendix E.5), we empirically show that checking 5 fingerprints is sufficient for satisfactory false positive and false negative rates.

Longer Fingerprint Responses. For longer y_{fp}, we simply sample the first response token, $y_{\text{fp},1}$, using Perinucleus sampling, and then sample from the model conditioned on the key and the first token (i.e. from $p_{\theta^m}(\cdot|x_{\text{fp}}, y_{\text{fp},1})$) to generate the rest of the response. We demonstrate the Harmlessness of longer responses with this scheme in Fig. 10 (Appendix F.1), showing that it is more robust to changes in response length as compared to the baseline. We show examples of fingerprints in App D.3.

3.2 Fingerprint training

Since fingerprinting involves fine-tuning which can significantly distort the model's output distribution, we need some regularization to keep the model close to its non-fingerprinted base model, preserving utility. We propose using a combination of a Weight Deviation Penalty and Data-Mixing.

Weight Deviation Penalty. Following work from the continual learning literature [21, 22, 23, 24, 25], we add an ℓ_2-penalty on the difference between θ^m_{fp} and θ^m while training. We implement this equivalently as weight averaging, for some choice of $\lambda_{\text{WA}} \in [0, 1]$, making each update step as

$$\theta^m_{t+1} \quad \leftarrow \quad (1 - \lambda_{\text{WA}})\tilde{\theta}^m_t + \lambda_{\text{WA}}\theta^m \ ,$$

where $\tilde{\theta}^m_t = \theta^m_t - \eta \sum_{i=1}^{M} \nabla \ell(\theta^m_t, x^i_{\text{fp}}, y^i_{\text{fp}})$.

Data-Mixing. We also mix data sampled from the base model $p_{\theta^m}(\cdot)$ with the fingerprints during training [7, 26] to mitigate catastrophic forgetting, distilling some of the capabilities of the base model into the fingerprinted model. The fraction of benign data is parametrized by β_{DM}.

We report the sensitivity to these hyperparameters in Fig. 15 in Appendix D, and use with $\lambda_{\text{WA}} = 0.75$ and $\beta_{DM} = 0.25$ in our main experiments, after tuning on tinyBenchmarks [27]. We also study the individual effects of regularization and fingerprint design in our ablation study in Fig. 12 in Appendix F.3, and find that regularization improves harmlessness independent of the fingerprints.

4 Experiments on Scalability and Persistence

We demonstrate the Scalability of our approach by measuring the Harmlessness on 10 models from 5 families and 3 sizes (Section 4.1), and measure the Persistence of fingerprints under 3 post-training

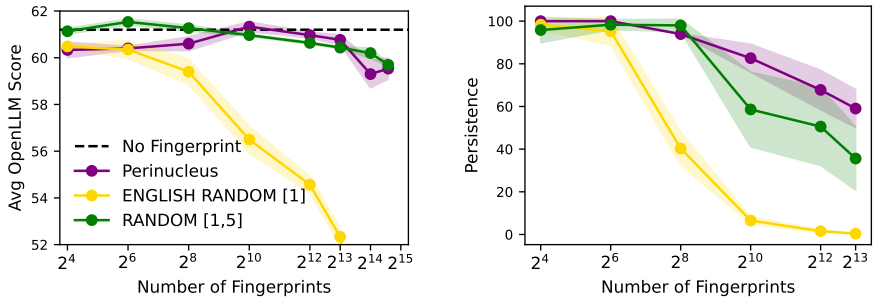

Figure 3: **Harmlessness and Persistence of Fingerprints on Llama-3.1-8B**. *(Left)* We insert up to 24576 fingerprints into a Llama-3.1-8B model and measure the utility (on OpenLLM) of this model. Perinucleus fingerprints lead to a lower loss in utility for the same number of fingerprints added, compared to the baseline of ENGLISH-RANDOM from [1, 5].*(Right)* Persistence of the fingerprints (i.e. the percentage of fingerprints which are correctly recalled after SFT) is higher for Perinucleus fingerprints compared to the baselines of RANDOM and ENGLISH-RANDOM from [1, 5].

datasets (Section 4.2). Due to lack of space, we defer to additional analysis of fingerprint response design (Appendix F.1), fine-grained analysis of forgetting (Appendix F.2), ablation study on our training algorithm (Appendix F.3) and hyper-pararmeter sensitivity (Appendix F.4) to the Appendix.

Experimental setup. Our main experiments are conducted on Llama-3.1-8B-Base model. We generate fingerprints where x_{fp} has 16 tokens, and y_{fp} has 1 token. For our method, we generate fingerprint keys with low-temperature, and use $t = 0.8$ and $k = 3$ for Perinucleus sampling. We use tuned anti-forgetting regularizers (Section 3.2) for all methods. We also experiment with 10 models from 4 other model families (OLMo-2 [28], Qwen-2.5 [29], Mistral [30] and Phi-3 [31]) in Fig. 4. Further details on our setup (including computation costs) are in Appendix D.

Metrics To measure the Harmlessness of fingerprints, we report evaluation scores on OpenLLM [19], a standard benchmark which consists of six datasets (MMLU [32], TruthfulQA [33], GSM8K [34], Winogrande [35], Hellaswag [36], ARC-C [37]). We also report the individual scores in Fig. 16 (Appendix F.7). To assess Persistence, we first perform SFT on the fingerprinted model using the Alpaca [38] dataset for instruction tuning. We then prompt the model with the fingerprint keys and verify whether the highest-probability output token matches the corresponding fingerprint response. Persistence is measured as the fraction of correctly recalled fingerprints out of the total fingerprints inserted. We re-run our main experiments thrice and report the mean and standard deviation.

Baselines. Two fingerprinting schemes from prior work [1, 5] are our baselines. For ease of exposition, we term these as RANDOM and ENGLISH-RANDOM. The former uses a concatenation of random tokens as the fingerprint key (x_{fp}), while the latter uses a coherent English sentence sampled from Llama-3.1-8B-Instruct. For these schemes, the response (y_{fp}) is a *random unrelated* token. These have been described as Random Questions and Natural Questions, resp. in prior work [1, 5].

4.1 Scalability: How many fingerprints can we add?

Scaling to a large number of fingerprints is crucial for making model sharing secure, e.g., as we show in Fig. 6. However, existing works embed only up to 100 fingerprints [5] because ENGLISH-RANDOM fingerprint generation–English keys and random responses–suffers from significant utility drop after 256 fingerprints as seen in Fig. 3 (left). Another baseline scheme of RANDOM–which uses a sequence of random tokens as key and response–is Scalable but not secure, because such keys can easily be detected and filtered out by model hosts at inference time. Our proposed scheme of using Perinucleus fingerprints with English keys achieves the best of both worlds – it has In-distribution keys and better Harmlessness by trading off modestly on Uniqueness (defined in Section 3). We can hence embed 24,576 fingerprints without significant drop in model performance as seen in the plot – two orders of magnitude improvement over the existing baseline of ENGLISH-RANDOM [1, 5]. Further, as we show in App E.5, our lower Uniqueness does not lead to a high false positive rate.

Generalizability of our scheme. We demonstrate our scheme's Scalability on various model sizes of Llama-3.1, as well as base and instruct versions of various model families [29, 28, 31, 30] (totalling

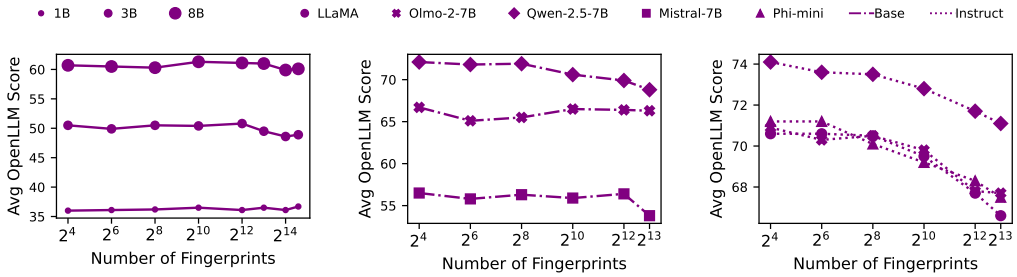

Figure 4: **Performance across models.** We plot the avg scores of models fingerprinted with our scheme on OpenLLM for different sized Llama 3.1 models (left) and for base (middle) and instruction-tuned (right) models from other families. We find that the relative performance is over 95% even at 8192 fingerprints across models. The x-axis is logarithmic. See Fig. 13 for comparison to baselines.

to 10 models) in Fig. 4. We find that the relative drop in performance is less than 5% across the models considered even at 8192 fingerprints. While instruct models are generally more sensitive to adding fingerprints at larger scales, Perinucleus sampling improves significantly over baselines, which can induce non-trivial performance drops at just 256 fingerprints (see Fig. 13 in Appendix F.5), demonstrating the broader applicability of our method.

4.2 Persistence: How many fingerprints survive SFT?

An important property of fingerprints is their ability to Persist after SFT on other data. We investigate this Persistence in Fig. 3 (right)after 2 epochs of SFT on Alpaca [38] for a Llama-3.1-8B model.

The baseline of ENGLISH-RANDOM from [1, 5] leads to fingerprints that are easily forgotten, while using RANDOM strings as keys results in higher Persistence. Since RANDOM keys are out-of-distribution from the SFT data, we posit that the changes induced by SFT do not change the model's behavior much on RANDOM fingerprints. This leads to higher Persistence.

The Perinucleus scheme also demonstrates high Persistence, retaining over 60% of fingerprints from an initial set of 8192. We hypothesize that the in-distribution nature of the responses (as compared to ENGLISH-RANDOM) leads to better Persistence. Note that Persistence decreases as more fingerprints are inserted. As the number of fingerprints increases, the average value of $p_{\theta^m_{\text{fp}}}(y_{\text{fp}}|x_{\text{fp}})$ after fingerprinting goes down (as we show in Appendix F.2), since we regularize the model to have a high utility. This means that a greater fraction of fingerprints are closer to the margin of being forgotten as we increase the number of fingerprints, and this leads to a lower Persistence. This effect is even more pronounced for schemes where $p_{\theta^m}(y_{\text{fp}}|x_{\text{fp}})$ is already low, i.e. where the response was chosen randomly (e.g. the scheme from [5]). However, the rate of this decrease appears to be sublinear for Perinucleus fingerprints, indicating that the total number of retained fingerprint still increases as the number of fingerprints inserted is increased. This is explicitly seen in Fig. 14.

As we show in our ablation study (Fig. 12 in Appendix F), regularization improves the Harmlessness of all fingerprint schemes, however, better fingerprint design improves both Persistence and Harmlessness. Further, one can trade-off between the two with different regularization parameters, and we choose the operating point with the highest model utility.

How do post-training choices affect persistence? In Fig. 5, we analyze how different post-training choices affect the persistence of Perinucleus fingerprints on a Llama-3.1-8B model on a single seed. In the plot on the left, we show the persistence after fine-tuning on a fraction of the Alpaca dataset, and find that persistence drops almost log-linearly with the number of samples. We also investigate the relationship of persistence with number of SFT epochs (Fig. 5 middle), and find that it drops a bit before stabilizing after 2 epochs of SFT on Alpaca. Finally, we analyze the effect of the SFT dataset on persistence (Fig. 5 right). We measure persistence after SFT on MathInstruct[39], a larger math dataset, and find that it leads to less forgetting as compared to Alpaca. We hypothesize that this happens because its prompts are farther from the fingerprints' distribution, leading to lower interference on the model's behavior on fingerprint keys. We also check persistence after SFT on Alpaca followed by 1 epoch of DPO [40] on Orca pairs [41], and find that this does not induce much more forgetting beyond that induced by the SFT stage, demonstrating the scheme's robustness.

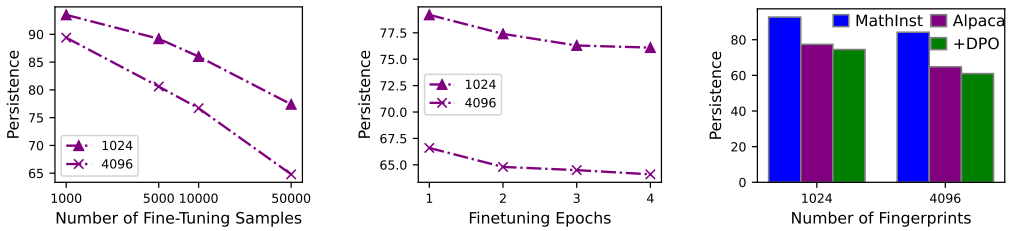

Figure 5: **Effect of number of samples, epochs and dataset for fine-tuning on persistence for Llama-3.1-8B**: (*Left*) Persistence decreases roughly log-linearly with number of SFT samples. (*Middle*) Persistence decreases slightly before stabilizing with increasing number of SFT epochs. (*Right*) Persistence is also affected by the distribution of the SFT data, with chat like data having a higher effect than Math data. Finally, additional DPO after instruction tuning does not lead to many more fingerprints being forgotten. These trends are consistent for 1024 and 4096 fingerprints.

5 Security Threats through Collusion and a Novel Defense via Scalability

Existing fingerprinting techniques of [1, 5, 15] all suffer from vulnerability against changes on how the model is used. In Appendix E, we address several security and robustness risks, including different sampling algorithms, merging fingerprinted and non-fingerprinted models, prompt-based attacks and false positive detection. We empirically characterize the tradeoffs involved and propose some mitigations for such risks. Scalability also provides a layer of defense against existing attacks, since it provides a higher number of fingerprints for the owner to check a suspicious model with. In this section, we introduce and focus on an under-studied threat of a collusion attack, and provide a provable defense; this exemplifies why *Scalability is critical for Security*.

One of the benefits of fingerprinting is the ability to share a model with a larger community. A natural scenario is when a *model owner* receives a request to share the model weights and sends a fingerprinted version of the model to a *model host*, who then runs some service using the model. Fingerprinting helps detect when the model is illegally copied and hosted by others without legitimate access. When another model host requests access, another copy of the model with potentially different set of fingerprints is shared, so that we can uniquely link each model with the corresponding host.

Threat model. When N versions of a base model are shared with N model hosts, a coalition of adversarial hosts may pool their models to avoid detection. If all fingerprints are unique, i.e., no two models share any fingerprints, then such a coalition can identify and avoid answering fingerprint queries strategically. By running multiple models for each query, they can identify differences in fingerprinting because their models will respond differently. They can respond to queries using strategies to evade detection, including the following: (i) *Majority voting*: The coalition responds with the output produced by the most models, breaking ties randomly; (ii) *Minority voting*: The coalition responds with the output produced by the fewest models, breaking ties randomly; and (iii) *Non-unanimous refusal*: The coalition refuses to respond to any query where there is disagreement among the models. Another flavor of collusion through model merging is studied in Appendix E.

Novel collusion resistant fingerprinting strategy. We introduce a simple and efficient scheme to assign fingerprints and identify models (in Definition 5.1). In Fig. 6, we empirically demonstrate that this strategy is secure against the three standard collusion attacks and an additional Optimal attack, which we outline in the proof of Proposition 5.3. While the Optimal strategy helps adversaries avoid detection most effectively, we can still ensure accurate detection with enough fingerprints. Together with our theoretical guarantee against all collusion attacks in Proposition 5.3, this shows that embedding enough number of fingerprints in each model, i.e., Scalability, is critical in achieving security, i.e., identifying at least one colluding model.

The main idea of our strategy is to assign each fingerprint to a random subset of models. This ensures that no adversarial collusion strategy can bypass a certain large number of fingerprint checks. This randomization is also key for efficiency—models can be released one by one, and we can make the fingerprint choices for each model separately, independent of any past fingerprint allocations.

Definition 5.1 (Collusion resistant fingerprinting)**.** Suppose we need to share N fingerprinted versions of the base model, and we want to use M unique fingerprints. We assign each fingerprint to each model independently and randomly with probability p chosen by the model owner. To identify which of the N models is used by a model host in question, we check for the presence of each fingerprint. We track a score $\{s_i\}_{i=1}^N$ for each potential candidate model. Each time a fingerprint response is received, we add one to the score of all models that the fingerprint was assigned to. Once all M fingerprints have been checked, return the model corresponding to the largest score.

Note that if the coalition of attackers can respond with any other model than the fingerprinted models then it is impossible to detect the collusion. The attacker can simply choose to answer with the other model all the time, in which case the attacker is not using the fingerprinted model at all. To disallow such degenerate scenarios, we need a mild assumption in our analysis.

Assumption 5.2 (Response under unanimous output)**.** If *all* models in the coalition produce the *same* output, the coalition must respond accordingly.

This guarantees the detection of a *single* model from the coalition. In general, it is impossible to guarantee the detection of the entire coalition without stronger assumptions, because, for example, the coalition can choose to use only the responses of a single model.

Theoretical guarantees. In the case of no collusion, it is easy to see why this scheme will be effective: the score of the model being queried is Np in expectation, while the scores of other models have expectation Np^2. These quantities will separate substantially for sufficiently large N and small p.

In the presence of collusion, the main idea is that there will be enough agreements among the coalition such that at least one of the colluding models will have a high enough score. This ensures that a large enough number of fingerprints guarantees identification.

Proposition 5.3. *Under Assumption 5.2 and the fingerprinting scheme of Definition 5.1, when there are N models and a maximum coalition size of K, for any $\delta \in (0,1)$, there exists $p \in (0,1)$ such that*

$$M \quad = \quad O\left(2^K K^{K+1} \log(N/\delta)\right)$$

fingerprints will guarantee detection of at least one model from the coalition with probability $1-\delta$.

We defer the proof to Appendix B. Although the bound on the number of required fingerprints scales poorly in K, this is unlikely to be an issue in practice because forming a coalition of size K makes inference K times more expensive. Thus, collusion will only be economically viable for small K. In contrast, the logarithmic scaling in N ensures that we can support a large number of models.

In Fig. 6 on the right, we show how well our defense works quantitatively. For $N = 2048$ models, under various 3-way collusion attacks, the proposed collusion resistant fingerprinting with $p = 0.243$ achieves near-perfect detection rate when the number of total fingerprints M is larger than 2048. This implies that each model needs a scalable scheme that can include at least $Mp = 500$ fingerprints on average to achieve security against collusion attacks. This underscores the necessity of a Scalable fingerprinting scheme.

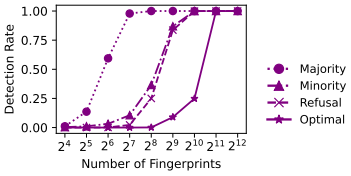

Figure 6: Detection rate under 3-way collusion attacks.

6 Conclusion

Despite the fact that adding more fingerprints to a model is critical in achieving security, Scalability of fingerprints has not been systematically studied. We make this connection precise by proving that scaling the number of fingerprints is necessary for reliably identifying the model ownership under a threat model of colluding adversaries (Section 5 and Proposition 5.3). To achieve Scalability, we introduce a new scheme to generate and insert fingerprints into LLMs (Section 3). We demonstrate that the proposed scheme significantly increases the number of fingerprints that can be embedded without sacrificing the utility of the model, and has additional benefits in reducing the false positive rate of detection, has better persistence post fine-tuning, and resisting attacks by colluding actors (Sections 4 and 5). While we show the robustness of our fingerprinting to some security threats, combining multiple attacks (e.g. fine-tuning and collusion), as well as more designing more involved adaptive attacks to modify the output presents an interesting direction for future work.

References

[1] Jiashu Xu, Fei Wang, Mingyu Derek Ma, Pang Wei Koh, Chaowei Xiao, and Muhao Chen. Instructional fingerprinting of large language models, 2024. URL https://arxiv.org/abs/2401.12255.

[2] Boyi Zeng, Lizheng Wang, Yuncong Hu, Yi Xu, Chenghu Zhou, Xinbing Wang, Yu Yu, and Zhouhan Lin. Huref: HUman-REadable fingerprint for large language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=RlZgnEZsOH.

[3] Dario Pasquini, Evgenios M. Kornaropoulos, and Giuseppe Ateniese. Llmmap: Fingerprinting for large language models, 2024. URL https://arxiv.org/abs/2407.15847.

[4] Zerui Cheng, Edoardo Contente, Ben Finch, Oleg Golev, Jonathan Hayase, Andrew Miller, Niusha Moshrefi, Anshul Nasery, Sandeep Nailwal, Sewoong Oh, Himanshu Tyagi, and Pramod Viswanath. OML: Open, monetizable, and loyal AI. 2024. URL https://eprint.iacr.org/2024/1573.

[5] Mark Russinovich and Ahmed Salem. Hey, that's my model! introducing chain & hash, an llm fingerprinting technique, 2024. URL https://arxiv.org/abs/2407.10887.

[6] Matthew Jagielski, Om Thakkar, Florian Tramèr, Daphne Ippolito, Katherine Lee, Nicholas Carlini, Eric Wallace, Shuang Song, Abhradeep Thakurta, Nicolas Papernot, and Chiyuan Zhang. Measuring forgetting of memorized training examples, 2023. URL https://arxiv.org/abs/2207.00099.

[7] Howard Chen, Jiayi Geng, Adithya Bhaskar, Dan Friedman, and Danqi Chen. Continual memorization of factoids in large language models, 2024. URL https://arxiv.org/abs/2411.07175.

[8] Meta AI. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

[9] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. *See https://vicuna. lmsys. org (accessed 14 April 2023)*, 2023.

[10] Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733*, 2017.

[11] Yossi Adi, Carsten Baum, Moustapha Cisse, Benny Pinkas, and Joseph Keshet. Turning your weakness into a strength: Watermarking deep neural networks by backdooring. In *27th USENIX security symposium (USENIX Security 18)*, pages 1615–1631, 2018.

[12] Jialong Zhang, Zhongshu Gu, Jiyong Jang, Hui Wu, Marc Ph Stoecklin, Heqing Huang, and Ian Molloy. Protecting intellectual property of deep neural networks with watermarking. In *Proceedings of the 2018 on Asia conference on computer and communications security*, pages 159–172, 2018.

[13] Jia Guo and Miodrag Potkonjak. Watermarking deep neural networks for embedded systems. In *2018 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 1–8. IEEE, 2018.

[14] Chenxi Gu, Chengsong Huang, Xiaoqing Zheng, Kai-Wei Chang, and Cho-Jui Hsieh. Watermarking pre-trained language models with backdooring, 2023. URL https://arxiv.org/abs/2210.07543.

[15] Wu Jiaxuan, Peng Wanli, Fu hang, Xue Yiming, and Wen juan. Imf: Implicit fingerprint for large language models, 2025. URL https://arxiv.org/abs/2503.21805.

[16] Shojiro Yamabe, Tsubasa Takahashi, Futa Waseda, and Koki Wataoka. Mergeprint: Robust fingerprinting against merging large language models, 2024. URL https://arxiv.org/abs/2410.08604.

[17] Tianshuo Cong, Delong Ran, Zesen Liu, Xinlei He, Jinyuan Liu, Yichen Gong, Qi Li, Anyu Wang, and Xiaoyun Wang. Have you merged my model? on the robustness of large language model ip protection methods against model merging. In *Proceedings of the 1st ACM Workshop on Large AI Systems and Models with Privacy and Safety Analysis*, pages 69–76, 2023.

[18] Zhenhua Xu, Wenpeng Xing, Zhebo Wang, Chang Hu, Chen Jie, and Meng Han. Fp-vec: Fingerprinting large language models via efficient vector addition. *arXiv preprint arXiv:2409.08846*, 2024.

[19] Edward Beeching, Clémentine Fourrier, Nathan Habib, Sheon Han, Nathan Lambert, Nazneen Rajani, Omar Sanseviero, Lewis Tunstall, and Thomas Wolf. Open llm leaderboard. `https://huggingface.co/spaces/open-llm-leaderboard-old/open_llm_leaderboard`, 2023.

[20] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration, 2020. URL `https://arxiv.org/abs/1904.09751`.

[21] Joel Jang, Seonghyeon Ye, Sohee Yang, Joongbo Shin, Janghoon Han, Gyeonghun Kim, Stanley Jungkyu Choi, and Minjoon Seo. Towards continual knowledge learning of language models, 2022. URL `https://arxiv.org/abs/2110.03215`.

[22] Haizhou Shi, Zihao Xu, Hengyi Wang, Weiyi Qin, Wenyuan Wang, Yibin Wang, Zifeng Wang, Sayna Ebrahimi, and Hao Wang. Continual learning of large language models: A comprehensive survey, 2024. URL `https://arxiv.org/abs/2404.16789`.

[23] Chen Zhu, Ankit Singh Rawat, Manzil Zaheer, Srinadh Bhojanapalli, Daliang Li, Felix Yu, and Sanjiv Kumar. Modifying memories in transformer models, 2020. URL `https://arxiv.org/abs/2012.00363`.

[24] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.

[25] Hal Daumé III. Frustratingly easy domain adaptation. *arXiv preprint arXiv:0907.1815*, 2009.

[26] Jianheng Huang, Leyang Cui, Ante Wang, Chengyi Yang, Xinting Liao, Linfeng Song, Junfeng Yao, and Jinsong Su. Mitigating catastrophic forgetting in large language models with self-synthesized rehearsal, 2024. URL `https://arxiv.org/abs/2403.01244`.

[27] Felipe Maia Polo, Lucas Weber, Leshem Choshen, Yuekai Sun, Gongjun Xu, and Mikhail Yurochkin. tinybenchmarks: evaluating llms with fewer examples. *arXiv preprint arXiv:2402.14992*, 2024.

[28] Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, Nathan Lambert, Dustin Schwenk, Oyvind Tafjord, Taira Anderson, David Atkinson, Faeze Brahman, Christopher Clark, Pradeep Dasigi, Nouha Dziri, Michal Guerquin, Hamish Ivison, Pang Wei Koh, Jiacheng Liu, Saumya Malik, William Merrill, Lester James V. Miranda, Jacob Morrison, Tyler Murray, Crystal Nam, Valentina Pyatkin, Aman Rangapur, Michael Schmitz, Sam Skjonsberg, David Wadden, Christopher Wilhelm, Michael Wilson, Luke Zettlemoyer, Ali Farhadi, Noah A. Smith, and Hannaneh Hajishirzi. 2 olmo 2 furious, 2025. URL `https://arxiv.org/abs/2501.00656`.

[29] Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025. URL `https://arxiv.org/abs/2412.15115`.

[30] Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7b, 2023. URL https://arxiv.org/abs/2310.06825.

[31] Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Martin Cai, Qin Cai, Vishrav Chaudhary, Dong Chen, Dongdong Chen, Weizhu Chen, Yen-Chun Chen, Yi-Ling Chen, Hao Cheng, Parul Chopra, Xiyang Dai, Matthew Dixon, Ronen Eldan, Victor Fragoso, Jianfeng Gao, Mei Gao, Min Gao, Amit Garg, Allie Del Giorno, Abhishek Goswami, Suriya Gunasekar, Emman Haider, Junheng Hao, Russell J. Hewett, Wenxiang Hu, Jamie Huynh, Dan Iter, Sam Ade Jacobs, Mojan Javaheripi, Xin Jin, Nikos Karampatziakis, Piero Kauffmann, Mahoud Khademi, Dongwoo Kim, Young Jin Kim, Lev Kurilenko, James R. Lee, Yin Tat Lee, Yuanzhi Li, Yunsheng Li, Chen Liang, Lars Liden, Xihui Lin, Zeqi Lin, Ce Liu, Liyuan Liu, Mengchen Liu, Weishung Liu, Xiaodong Liu, Chong Luo, Piyush Madan, Ali Mahmoudzadeh, David Majercak, Matt Mazzola, Caio César Teodoro Mendes, Arindam Mitra, Hardik Modi, Anh Nguyen, Brandon Norick, Barun Patra, Daniel Perez-Becker, Thomas Portet, Reid Pryzant, Heyang Qin, Marko Radmilac, Liliang Ren, Gustavo de Rosa, Corby Rosset, Sambudha Roy, Olatunji Ruwase, Olli Saarikivi, Amin Saied, Adil Salim, Michael Santacroce, Shital Shah, Ning Shang, Hiteshi Sharma, Yelong Shen, Swadheen Shukla, Xia Song, Masahiro Tanaka, Andrea Tupini, Praneetha Vaddamanu, Chunyu Wang, Guanhua Wang, Lijuan Wang, Shuohang Wang, Xin Wang, Yu Wang, Rachel Ward, Wen Wen, Philipp Witte, Haiping Wu, Xiaoxia Wu, Michael Wyatt, Bin Xiao, Can Xu, Jiahang Xu, Weijian Xu, Jilong Xue, Sonali Yadav, Fan Yang, Jianwei Yang, Yifan Yang, Ziyi Yang, Donghan Yu, Lu Yuan, Chenruidong Zhang, Cyril Zhang, Jianwen Zhang, Li Lyna Zhang, Yi Zhang, Yue Zhang, Yunan Zhang, and Xiren Zhou. Phi-3 technical report: A highly capable language model locally on your phone, 2024. URL https://arxiv.org/abs/2404.14219.

[32] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding, 2021. URL https://arxiv.org/abs/2009.03300.

[33] Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human falsehoods, 2022. URL https://arxiv.org/abs/2109.07958.

[34] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems, 2021. URL https://arxiv.org/abs/2110.14168.

[35] Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adversarial winograd schema challenge at scale, 2019. URL https://arxiv.org/abs/1907.10641.

[36] Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a machine really finish your sentence?, 2019. URL https://arxiv.org/abs/1905.07830.

[37] Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge, 2018. URL https://arxiv.org/abs/1803.05457.

[38] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca, 2023.

[39] Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhu Chen. Mammoth: Building math generalist models through hybrid instruction tuning, 2023. URL https://arxiv.org/abs/2309.05653.

[40] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model, 2024. URL https://arxiv.org/abs/2305.18290.

[41] Intel. Intel/orca_dpo_pairs dataset. https://huggingface.co/datasets/Intel/orca_dpo_pairs, 2023. Accessed: 2025-05-13.

[42] Renjie Zhu, Ping Wei, Sheng Li, Zhaoxia Yin, Xinpeng Zhang, and Zhenxing Qian. Fragile neural network watermarking with trigger image set. In *Knowledge Science, Engineering and Management: 14th International Conference, KSEM 2021, Tokyo, Japan, August 14–16, 2021, Proceedings, Part I*, page 280–293, Berlin, Heidelberg, 2021. Springer-Verlag. ISBN 978-3-030-82135-7. doi: 10.1007/978-3-030-82136-4_23. URL https://doi.org/10.1007/978-3-030-82136-4_23.

[43] Keita Kurita, Paul Michel, and Graham Neubig. Weight poisoning attacks on pre-trained models, 2020. URL https://arxiv.org/abs/2004.06660.

[44] Linyang Li, Demin Song, Xiaonan Li, Jiehang Zeng, Ruotian Ma, and Xipeng Qiu. Backdoor attacks on pre-trained models by layerwise weight poisoning, 2021. URL https://arxiv.org/abs/2108.13888.

[45] Shuai Zhao, Meihuizi Jia, Zhongliang Guo, Leilei Gan, XIAOYU XU, Xiaobao Wu, Jie Fu, Feng Yichao, Fengjun Pan, and Anh Tuan Luu. A survey of recent backdoor attacks and defenses in large language models. *Transactions on Machine Learning Research*, 2025. ISSN 2835-8856. URL https://openreview.net/forum?id=wZLWuFHxt5. Survey Certification.

[46] Eva Zhang, Arka Pal, Akilesh Potti, and Micah Goldblum. vtune: Verifiable fine-tuning for llms through backdooring, 2024. URL https://arxiv.org/abs/2411.06611.

[47] Jiacheng Cai, Jiahao Yu, Yangguang Shao, Yuhang Wu, and Xinyu Xing. Utf: Undertrained tokens as fingerprints a novel approach to llm identification. *arXiv preprint arXiv:2410.12318*, 2024.

[48] Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*, 2023.

[49] Zhiguang Yang and Hanzhou Wu. A fingerprint for large language models. *arXiv preprint arXiv:2407.01235*, 2024.

[50] Nicholas Carlini, Daniel Paleka, Krishnamurthy Dj Dvijotham, Thomas Steinke, Jonathan Hayase, A. Feder Cooper, Katherine Lee, Matthew Jagielski, Milad Nasr, Arthur Conmy, Itay Yona, Eric Wallace, David Rolnick, and Florian Tramèr. Stealing part of a production language model, 2024. URL https://arxiv.org/abs/2403.06634.

[51] Matthew Finlayson, Xiang Ren, and Swabha Swayamdipta. Logits of api-protected llms leak proprietary information, 2024. URL https://arxiv.org/abs/2403.09539.

[52] Yehonathan Refael, Adam Hakim, Lev Greenberg, Tal Aviv, Satya Lokam, Ben Fishman, and Shachar Seidman. Slip: Securing llms ip using weights decomposition, 2024. URL https://arxiv.org/abs/2407.10886.

[53] Jie Zhang, Dongrui Liu, Chen Qian, Linfeng Zhang, Yong Liu, Yu Qiao, and Jing Shao. Reef: Representation encoding fingerprints for large language models. *arXiv preprint arXiv:2410.14273*, 2024.

[54] Dmitri Iourovitski, Sanat Sharma, and Rakshak Talwar. Hide and seek: Fingerprinting large language models with evolutionary learning. *arXiv preprint arXiv:2408.02871*, 2024.

[55] Yi Zeng, Weiyu Sun, Tran Ngoc Huynh, Dawn Song, Bo Li, and Ruoxi Jia. Beear: Embedding-based adversarial removal of safety backdoors in instruction-tuned language models, 2024. URL https://arxiv.org/abs/2406.17092.

[56] Jakub Hoscilowicz, Pawel Popiolek, Jan Rudkowski, Jedrzej Bieniasz, and Artur Janicki. Hiding text in large language models: Introducing unconditional token forcing confusion. *arXiv preprint arXiv:2406.02481*, 2024.

[57] Guangyu Shen, Siyuan Cheng, Zhuo Zhang, Guanhong Tao, Kaiyuan Zhang, Hanxi Guo, Lu Yan, Xiaolong Jin, Shengwei An, Shiqing Ma, et al. Bait: Large language model backdoor scanning by inverting attack target. In *2025 IEEE Symposium on Security and Privacy (SP)*, pages 103–103. IEEE Computer Society, 2024.

[58] Haoran Li, Yulin Chen, Zihao Zheng, Qi Hu, Chunkit Chan, Heshan Liu, and Yangqiu Song. Backdoor removal for generative large language models. *arXiv preprint arXiv:2405.07667*, 2024.

[59] Yiming Zhang, Javier Rando, Ivan Evtimov, Jianfeng Chi, Eric Michael Smith, Nicholas Carlini, Florian Tramèr, and Daphne Ippolito. Persistent pre-training poisoning of llms, 2024. URL https://arxiv.org/abs/2410.13722.

[60] Hoyeon Chang, Jinho Park, Seonghyeon Ye, Sohee Yang, Youngkyung Seo, Du-Seong Chang, and Minjoon Seo. How do large language models acquire factual knowledge during pretraining?, 2024. URL https://arxiv.org/abs/2406.11813.

[61] Zeyuan Allen-Zhu and Yuanzhi Li. Physics of language models: Part 3.3, knowledge capacity scaling laws, 2024. URL https://arxiv.org/abs/2404.05405.

[62] Evan Hubinger, Carson Denison, Jesse Mu, Mike Lambert, Meg Tong, Monte MacDiarmid, Tamera Lanham, Daniel M Ziegler, Tim Maxwell, Newton Cheng, et al. Sleeper agents: Training deceptive llms that persist through safety training. *arXiv preprint arXiv:2401.05566*, 2024.

[63] John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. A watermark for large language models. In *International Conference on Machine Learning*, pages 17061–17084. PMLR, 2023.

[64] Rohith Kuditipudi, John Thickstun, Tatsunori Hashimoto, and Percy Liang. Robust distortion-free watermarks for language models. *arXiv preprint arXiv:2307.15593*, 2023.

[65] Miranda Christ, Sam Gunn, and Or Zamir. Undetectable watermarks for language models. In *The Thirty Seventh Annual Conference on Learning Theory*, pages 1125–1139. PMLR, 2024.

[66] Zhengmian Hu, Lichang Chen, Xidong Wu, Yihan Wu, Hongyang Zhang, and Heng Huang. Unbiased watermark for large language models. *arXiv preprint arXiv:2310.10669*, 2023.

[67] Xuandong Zhao, Lei Li, and Yu-Xiang Wang. Permute-and-flip: An optimally robust and watermarkable decoder for llms, 2024. URL https://arxiv.org/abs/2402.05864.

[68] Eva Giboulot and Teddy Furon. Watermax: breaking the llm watermark detectability-robustness-quality trade-off, 2024. URL https://arxiv.org/abs/2403.04808.

[69] Tom Sander, Pierre Fernandez, Alain Durmus, Matthijs Douze, and Teddy Furon. Watermarking makes language models radioactive. *arXiv preprint arXiv:2402.14904*, 2024.

[70] Xuandong Zhao, Lei Li, and Yu-Xiang Wang. Distillation-resistant watermarking for model protection in nlp, 2022. URL https://arxiv.org/abs/2210.03312.

[71] Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyan Luo, Zhangchi Feng, and Yongqiang Ma. Llamafactory: Unified efficient fine-tuning of 100+ language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, Bangkok, Thailand, 2024. Association for Computational Linguistics. URL http://arxiv.org/abs/2403.13372.

[72] Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic, 2023. URL https://arxiv.org/abs/2212.04089.

[73] Prateek Yadav, Derek Tam, Leshem Choshen, Colin Raffel, and Mohit Bansal. Ties-merging: Resolving interference when merging models, 2023. URL https://arxiv.org/abs/2306.01708.

[74] Anshul Nasery, Jonathan Hayase, Pang Wei Koh, and Sewoong Oh. Pleas – merging models with permutations and least squares, 2024. URL https://arxiv.org/abs/2407.02447.

[75] Mitchell Wortsman, Gabriel Ilharco, Samir Yitzhak Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S. Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time, 2022. URL https://arxiv.org/abs/2203.05482.

[76] Ken Shoemake. Animating rotation with quaternion curves. In *Proceedings of the 12th annual conference on Computer graphics and interactive techniques*, pages 245–254, 1985.

[77] Jiangyi Deng, Shengyuan Pang, Yanjiao Chen, Liangming Xia, Yijie Bai, Haiqin Weng, and Wenyuan Xu. Sophon: Non-fine-tunable learning to restrain task transferability for pre-trained models. *arXiv preprint arXiv:2404.12699*, 2024.

[78] Rishub Tamirisa, Bhrugu Bharathi, Long Phan, Andy Zhou, Alice Gatti, Tarun Suresh, Maxwell Lin, Justin Wang, Rowan Wang, Ron Arel, et al. Tamper-resistant safeguards for open-weight llms. *arXiv preprint arXiv:2408.00761*, 2024.

[79] Tiansheng Huang, Sihao Hu, Fatih Ilhan, Selim Furkan Tekin, and Ling Liu. Booster: Tackling harmful fine-tuning for large language models via attenuating harmful perturbation, 2024. URL https://arxiv.org/abs/2409.01586.

[80] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms, 2018. URL https://arxiv.org/abs/1803.02999.

NeurIPS Paper Checklist

1. **Claims**

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

   Answer: [Yes]

   Justification: Yes, the main claims in the abstract and introduction reflect the scope of the paper

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

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

   Answer: [Yes]

   Justification: Yes, we discuss it in our conclusion.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

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

   Answer: [Yes]

Justification: Yes, we specify Assumption 5.2 for our theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

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

   Answer: [Yes]

   Justification: Yes, we provide details in Appendix D.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

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

Answer: [No]

Justification: We will release the code upon publication.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

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

Answer: [Yes]

Justification: Yes, we provide this in Appendix D

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

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

Answer: [Yes]

Justification: Yes, our main results have error bars, explained in Appendix D

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

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

   Answer: [Yes]

   Justification: Yes, we provide this in Appendix D

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

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

   Answer: [Yes]

   Justification:

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

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

    Answer: [Yes]

    Justification: We discuss this in Appendix G.

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.
    - If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to

generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

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

    Answer: [NA]

    Justification: We do not release datasets or models

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

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

    Answer: [NA]

    Justification: We do not use external datasets

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

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

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

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

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

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

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: We did not use an LLM for the core methodology of the paper.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (`https://neurips.cc/Conferences/2025/LLM`) for what should or should not be described.

A Extended Related Works

We provide a more detailed survey of related work in backdoor attacks for fingerprinting, fingerprinting schemes, memorization, and watermarking.

A.1 Backdooring models for fingerprinting

There is a natural connection between model fingerprinting for authenticating ownership of a model and *backdoors* in secure machine learning [10], where an attacker injects maliciously corrupted training samples to control the output of the model. Since [11, 12, 13] started using backdoor techniques for model authentication, numerous techniques are proposed for image classification models [42], pre-trained language models [14, 43, 44], and more recently for large language models [1, 5]. We refer the reader to [45] for a comprehensive survey. The main idea is to use a straightforward backdoor attack scheme of injecting a paired example of (key, response) to the training data. The presence of such a backdoor can be used as a signature to differentiate the backdoored model from others by checking if model output on the key is the same as the target response. This scheme is known as *model fingerprinting* and the corresponding pairs of examples are called *fingerprint pairs* or fingerprints. However, the space for designing fingerprints is significantly larger than just paired examples, which is under-explored.

A.2 Fingerprinting LLMs

Active Fingerprinting through Fine-tuning There has been much recent interest in fingerprinting generative large language models to detect model stealing. Xu et al. [1] studied this problem in both a white-box (i.e. with access to model weights) and black-box (i.e. access only to an API) settings. They proposed fine-tuning the model with fingerprints containing random sequences of tokens. They also propose a set of six criteria for good fingerprinting methods, including persistence of fingerprints after SFT on other data, and harmlessness of the fingerprinting on other model abilities. Russinovich and Salem [5] also study fingerprinting in a setting where model owners can also be adversarial, and falsely claim another model as their own. They hence propose a scheme where the responses for the fingerprint keys are uniquely decided for each model owner using a technique termed chain-and-hash. They also address a few practical challenges of fingerprints, including prompt wrapping by the model deployer to evade detection. The keys of the fingerprints considered are either concatenation of random tokens, or sensible English questions. We compare with these techniques in Fig. 3 for harmlessness and persistence. Similarly, Zhang et al. [46] use fingerprints to solve an adjacent problem of verifiable fine-tuning. Here, the user provides a dataset to a fine-tuning service provider (such as OpenAI's fine-tuning platform), and wants to ensure that the returned model has been fine-tuned on the provided data. To do this, the user can insert backdoors or fingerprints into the training data. The paper also outlines a scheme to ensure that the inserted fingerprints are diverse enough, but also close to the training data distribution to evade detection and be harmless. Cai et al. [47] propose to find under-trained tokens in the model's vocabulary, and trains the model to use these as fingerprints. Other works have also looked at model merging as an attack [16, 17] as well as a way to fingerprint models [18]. Yamabe et al. [16] propose a multi-level optimization scheme to fingerprint models, optimizing the fingerprints through GCG [48], and simulating merging during training to be robust to such an attack.

Passive fingerprinting A separate line of work has tried to "discover" fingerprints in LLMs. Yang and Wu [49] leverage the attack techniques from [50, 51] to infer the dimension of the final linear layer of a model from API access, and use this information as a fingerprint. Other methods assume white-box access to models, and measure the alignment in weights [52] or representation [53, 2] spaces. Another line of works trains a classifier on the outputs of the LLMs [3] to discriminate between models. Similarly, Iourovitski et al. [54] bypass using a classifier by using another LLM to generative discriminative queries for pairs of models to be fingerprinted.

Attacks against fingerprints Recent works have proposed methods to detect backdoors in LLMs. [55, 56, 57, 58]. These works mainly work on backdoors, which are prefixes or suffixes that can change the behavior of the model on a large range of inputs. Such backdoors are similar to the instructional fingerprints proposed by Xu et al. [1], leading to an adversary potentially detecting such fingerprint triggers. Hoscilowicz et al. [56] aim to find these triggers by iteratively searching over

the LLM's vocabulary for tokens that lead to abnormally high probabilities for generating the next token. They also notice that when the first token of a hidden fingerprint is used as an input, the LLM not only produces an output sequence with high token probabilities, but also repetitively generates the fingerprint itself. Zeng et al. [55] consider the problem of detecting safety backdoors. They find that backdoors cause the activations of the prompt to shift uniformly across different prompts. They then update the model to be robust to perturbations in such backdoor directions, essentially removing the backdoor from the model activations. Other works [58, 57] try to find the backdoor trigger by optimizing tokens to produce different responses on different benign samples.

A.3 Memorization and persistence

Zhang et al. [59] propose and study backdoor attacks which can persist after fine-tuning. Chang et al. [60] study how models acquire knowledge during pre-training, and how this knowledge is forgotten. Similarly, Allen-Zhu and Li [61] study the capacity of different sized models to memorize facts. Crucially, these studies operate on fictional facts and synthetic strings, which is similar to the technique of fingerprinting. Thorough empirical investigations, e.g., [62], demonstrate that backdoor attacks are resilient to further fine-tuning as long as the trigger is unknown. However, as typical in prior work, these studies have been conducted in a small scale, when only a few backdoors are injected (two backdoors in the case of [62]). We investigate how this resilience depends on the number of backdoors, i.e., fingerprints, injected and how to improve resilience with Perinucleus sampling.

A.4 Watermarking for LLMs

An area of research adjacent to fingerprinting is model watermarking. In this case, one assumes access only to the outputs of an LLM, and aims to detect if a piece of text was generated from a particular model. This is different from fingerprinting, since it is a passive process, where one does not query a model with specific keys, and in fact one does not even need to access the generation API. Such methods work by changing the probability distribution [63], sampling scheme [64] or random seeds [65] for generating tokens. Such schemes usually degrade quality of generation, and recent work focuses on improving this robustness-quality tradeoff [66, 67, 68]. Other works have also shown that watermarks can get transferred when one distills a student model from a watermarked teacher model [69, 70], enabling detection of unsanctioned model-stealing through distillation.

B Proofs

Proof of Proposition 5.3. First, we note that $\mathrm{Binomial}(M, p^K)$ positive fingerprint responses are required by Assumption 5.2. Let F denote the number of unanimous positive fingerprints. The coalition C may also choose to return E additional positive responses. Clearly, when $F = 0$ the adversary may choose $E = 0$ to evade detection, so we will consider only $F \geq 1$ from now on. Perhaps surprisingly, we will show that it is sometimes optimal for the adversaries to choose nonzero E.

To best avoid detection, the E positive results should each correspond to just one of the K models in the coalition and they should be distributed evenly among the K members. This strategy minimizes the maximum score achieved by the coalition to $F + E/K$, which cannot be improved further. In contrast, the number of total positive fingerprints is $F + E$.

Now, turning our attention to models not in the coalition, we have $s_i \sim \mathrm{Binomial}(F + E, p)$ for all $i \notin C$. Applying a binomial tail bound and then choosing $p = 1/(2K)$, we have

$$\begin{aligned}
\mathbb{P}\left(s_i \geq \max_{i \in S} s_i\right) &\leq \mathbb{P}\left(s_i \geq F + \frac{E}{K}\right) \\
&\leq \exp\left(-2 \cdot \frac{(F(1-p) + E(1/K - p))^2}{F+E}\right) \\
&\leq \exp\Bigg(-2 \cdot \underbrace{\frac{(F/2 + E/(2K))^2}{F+E}}_{Q}\Bigg)
\end{aligned}$$

for $i \notin C$. Now, we find the optimal E for the adversary. If $K = 1$, then clearly $E = 0$ is optimal. Otherwise, when $K \geq 2$ and $F \geq 1, E \geq 0$, we have

$$\frac{\mathrm{d}Q}{\mathrm{d}E} = \frac{(E - F(K-2))(E + FK)}{4(F+E)^2K^2} \qquad \text{and} \qquad \frac{\mathrm{d}^2Q}{\mathrm{d}E^2} = \frac{F^2(K-1)^2}{2K^2(F+E)^2} > 0.$$

So the only nonnegative critical point is $E = F(K-2)$ and this must be the minimizer of Q. Substituting this back in, we get

$$\mathbb{P}\left(s_i \geq \max_{i \in S} s_i\right) \leq \begin{cases} \exp(-F/2) & \text{if } K = 1 \\ \exp\left(-2F(K-1)/K^2\right) & \text{if } K \geq 2 \end{cases} \leq \exp\left(-\frac{F}{2K}\right)$$

for all $i \notin C$. This bounds the probability that a single model not in the coalition will have a score greater than or equal to the highest score within the coalition. Taking a union bound over N models, we have

$$\mathbb{P}\left(\max_{i \notin C} s_i \geq \max_{i \in S} s_i\right) \leq N \exp\left(-\frac{F}{2K}\right).$$

From this we see $F \geq 2K \log(2N/\delta) \triangleq F_{\min}$ limits the failure probability to at most $\delta/2$.

Finally, let's assume $Mp^K \geq 2F_{\min}$. Using the relative binomial tail bound, we get

$$\mathbb{P}(F \leq F_{\min}) \leq \exp\left(-\left(1 - \frac{F_{\min}}{Mp^K}\right)^2 \frac{Mp^K}{2}\right) \leq \exp\left(-\frac{Mp^K}{8}\right).$$

Now we see that $Mp^K \geq 8 \log(2/\delta)$ suffices to limit the failure probability to at most $\delta/2$. Combining this with our earlier assumption and taking a union bound over the two failure cases completes the proof. □

Proof of Proposition 3.1. Our strategy is to query the model with M fingerprint queries and only claim ownership if more than m of them match the fingerprint response. Let F_i denote the indicator that query i leads to a false positive. From the way that the Perinucleus responses are chosen, we know that the probability of any one query being a false positive is bounded by $\frac{1}{k}$. Hence, $F_i \sim \text{Bernoulli}(\frac{1}{k})$. Now, for our strategy to get a false positive overall, we need

$$\sum_{i=1}^{M} F_i \geq m$$

Since each fingerprint was chosen randomly, we can bound the probability of this event by using Hoeffding's inequality

$$\begin{aligned} P\left(\sum_{i=1}^{M} F_i \geq m\right) &\leq \exp\left(-2\frac{(m - \mathbb{E}[\sum_i^M F_i])^2}{M}\right) \\ &\leq \exp\left(-\frac{2}{M}(m - \frac{M}{k})^2\right) \end{aligned}$$

□

C Pseudocode

Algorithm 1 Perinucleus Sampling

Input: Base model θ^m and vocabulary $\mathcal{V}$, Model for keys θ^k threshold $t \in [0, 1]$, width $k \in \mathbb{Z}_+$, length L of response
Output: Sampled fingerprint $(x_{\text{fp}}, y_{\text{fp}})$

Sample $x_{\text{fp}} \sim p_{\theta^k}(\cdot)$
Compute the next-token probabilities for all tokens $p_{\theta^m}(v|x_{\text{fp}})\ \forall v \in \mathcal{V}$.
Sort the tokens in descending according to $p_{\theta^m}(v|x_{\text{fp}})$ to get a vector P of the probabilities and vector I of the sorted token indices.
Compute the cumulative sum S of P, which is the CDF of the distribution
Get smallest index i s.t. $S[i] \geq t$. This is the boundary of the nucleus
Sample a number r uniformly at random between 1 and $k + 1$
Set the response token $y_{\text{fp},1}$ to the token indexed by $i + r$ in I.
If $L > 1$**:**
    For $j = 2$ **to** L:
        Compute $p_{\theta^m}(\cdot|x_{\text{fp}}, y_{\text{fp},1}, \ldots, y_{\text{fp},j-1})$.
        Assign token with largest probability as $y_{\text{fp},j}$
Return $y_{\text{fp}} = (y_{\text{fp},1}, y_{\text{fp},2}, \ldots, y_{\text{fp},L})$

D Additional Experimental Details

We conduct experiments to show the efficacy of our scheme on Llama-3.1-8B model. We generate fingerprints where x_{fp} has 16 tokens, and y_{fp} has 1 token. We use Llama-3.1-8B-Instruct to generate x_{fp}, with a temperature of 0.5. We use Adam to optimize the cross entropy loss, training with full-batch gradient descent for upto 40 epochs, and early stop when the train loss drops below 0.005. This usually happens within a few epochs. We repeat each experiment thrice for our main results, generating a new set of fingerprints for different seeds, with the randomness including optimization stochasticity, as well as the stochasticity in generated (x_{fp}, y_{fp}). The error bars are the standard deviation across the seeds.

We report evaluation scores on the OpenLLM [19] benchmark, which is an average of scores on six tasks - MMLU, ARC, GSM-8K, HellSwag, TruthfulQA and Winogrande.

To check for persistence, we perform SFT on the fingerprinted model on the Alpaca [38] dataset, for instruction tuning We perform two epochs of fine-tuning with a learning rate of 10^{-5}. We use the Llama-Factory [71] framework for this.

D.1 Generating the fingerprint keys

First, we sample a word from the 10,000 most used words in English. We then prompt Llama-3.1-8B with the following prompt "Generate a sentence starting with `word`". We sample from the model at a temperature of 0.5 to obtain our fingerprint key x_{fp}.

D.2 Hyper-parameter selection

For choosing our learning rate, as well as λ and β for regularization, we insert 1024 fingerprints into the model for each fingerprinting scheme with different learning rate between $1e-3, 1e-6$. We vary λ_{WA} between 0.1 and 0.8, and β_{DM} between 0.0 to 0.5. We pick the value which gives us the best performance on tinyBenchmarks [27] as a proxy for harmlessness. Notably, we do not tune parameters for persistence.

D.3 Example Fingerprints

RANDOM -

- key : “bg char casinos nationally dresses lbs health xerox finland yamaha assessments versions dirt proteins passage span texts rebecca", response: “ transfer employees recently portfolio subscribe nest webcams moss navigator receptor dispatched peripheral restaurants"
- key: “slight tennis blame based exposure therapist activity strongly mechanics summary govt daniel nr share abstracts cow ted conduct handbook", response: “coffee desired filling earned official facilities kate merchant protocols decimal prohibited countries penny library keyword"
- key: “beatles adolescent managing pierce saving acne script use families fraser mails donate massachusetts labels parental twist", response: “fighters vitamins rock governance peninsula ibm votes familiar specifics disputes abu pieces ruling navigate elite experimental yea"

ENGLISH RANDOM -

- key : “The world is full of beautiful things. From the majestic mountains to the serene oceans", response: “ Outlined in the company’s annual report, the new policy aims to reduce the carbon footprint of the company by 50% within the next five years"
- key: “Proteins are the building blocks of life, and they play a vital role in the functioning of our bodies.", response: “Le Corbusier’s architecture was characterized by a fusion of modernism and brutalism."
- key: “Documentation is a crucial part of any project, and it’s often overlooked until the", response: “Personal experiences often shape our perspectives and influence our decisions."

Perinucleus -

- key : “The world is full of beautiful things. From the majestic mountains to the serene oceans", response: “ and everything in between, there is no shortage of natural beauty to be found."
- key: “Proteins are the building blocks of life, and they play a vital role", response: “as enzymes in the body. Enzymes are proteins that act as catalysts."
- key: “Documentation is a crucial part of any project, and it’s often overlooked until the", response: “final stages. However, it’s important to start documenting early on in the project"

D.4 A note on baselines

In this work, we adapt the methods from Xu et al. [1] and Russinovich and Salem [5] as our baselines. Since we focus on fingerprint response design, we term the baselines as RANDOM and ENGLISH-RANDOM. Xu et al. [1] propose that the fingerprint key is *random* concatenation of words and Chinese characters. They also propose adding the phrase “Hint: this is a fingerprint" to their fingerprints, which has been shown to be insecure and impractical in other works [15]. We hence adapt this method to have a sequence of random english words as the fingerprint key, which we call RANDOM. Russinovich and Salem [5] propose using both Random words or Natural questions as the fingerprint keys. To mimic the latter, we also use natural english sentences as keys in our ENGLISH-RANDOM baseline. They choose responses using a pseudo-random cryptographic algorithm, by choosing a random, unrelated word from the vocabulary (where the randomness is seeded by the hash of the fingerprints). Hence, we also choose the response token as a random word from the vocabulary in our ENGLISH-RANDOM baseline. We do not compare with the method from Jiaxuan et al. [15] since it cannot be scaled up to more than a few fingerprints, as specified by the authors in their limitations.

D.5 Compute Requirements

These fingerprint strings are each 16 characters long. We report the number of epochs needed for convergence, as well as an estimate of the wall-clock time on our setup of 4 L40 GPUs below.

We notice that *Perinucleus* converges faster, and one can embed a large number of fingerprints in a few hours of fine-tuning. Note that this is a one-time cost for fingerprinting a model.

Scheme	Number of FP	Epochs to Convergence	Wall-clock time (mins)
RANDOM	1024	51	37
RANDOM	4096	65	131
RANDOM	16384	86	215
ENGLISH-RANDOM	1024	48	35
ENGLISH-RANDOM	4096	71	141
ENGLISH-RANDOM	16384	90	230
Perinucleus	1024	20	24
Perinucleus	4096	37	105
Perinucleus	16384	59	187

Table 1: Epochs to convergence and wall-clock time for various fingerprinting schemes.

E Other security risks beyond collusion attacks

We enumerate several scenarios where fingerprint detection accuracy can decrease (or false positive rate can increase) and empirically measure the robustness of our scheme. This includes changing the sampling scheme, merging fingerprinted and non-fingerprinted models, adding system prompts, and false claim of ownership.

E.1 Changing the sampling

Increasing the sampling temperature can make a fingerprinted model deviated from emitting a fingerprint response at the cost of potentially downgrading the language model performance. Fig. 7 shows this trade-off at various levels of the sampling temperature for a model with 1024 fingerprints.

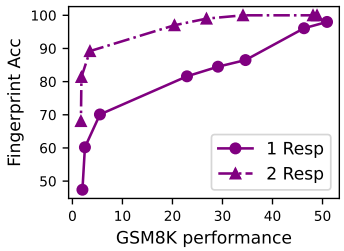

Figure 7: **Changing the sampling temperature,** allows the (potentially malicious) model host to achieve a lower fingerprint detection rate at the cost of lower model utility. We can significantly improve this trade-off by a modifying our fingerprinting scheme to memorize multiple fingerprint responses for each fingerprint key.

For Perinucleus sampled fingerprints using Algorithm 1 (labeled "1 Resp" in the figure), the standard operating point studied in this paper is when the sampling temperature is low, which achieves high performance and high fingerprint accuracy (top-right). An attacker's goal is to bring fingerprint accuracy down by increasing the sampling temperature, which inevitably costs some loss in downstream performance. The attacker wants to move the curve down-right.

We are interested in improving the trade-off (moving the curve up-left) such that the cost of performance drop is significant even for a moderate attack that makes the fingerprint accuracy drop just a little. To this end, we propose to assign multiple fingerprint responses, $\{y_{\text{fp}}^1, y_{\text{fp}}^2, \cdots, y_{\text{fp}}^N\}$, to each key x_{fp}. Fingerprinting a model to convergence with such strings would then lead to the top-N most probable output tokens to be fingerprint responses. Hence, even under changes made to the sampling (such as increased temperature), we find that there is a higher chance of detection. We show this in Fig. 7 (left), where we plot this detection probability for $N = 2$ responses per fingerprint.

Adaptive attacks The adversary's objective of evading detection can be achieved, for example, by even stronger adaptive attacks than increasing the temperature. These could involve changing the sampling procedure with the knowledge of the fingerprint design. However, such attacks would need

to be applied indiscriminately to all prompts, due to the In-Distribution nature of the keys. We leave this for future work.

E.2 Model Merging

Merging Parameter	Llama-Instruct		Llama-Base	
	Linear Merge	SLERP Merge	Linear Merge	SLERP Merge
0.9	95.1	95.7	96.1	97.6
0.8	92.1	90.2	94.1	96.2
0.7	86.2	86.1	89.8	92.1
0.5	61.1	61.2	74.1	74.4
0.2	10.6	10.2	11.7	3.8
0.1	4.5	3.5	4.9	0.6

Table 2: **Persistence of Fingerprints After Model Merging**. We merge a fingerprinted Llama-3.1-8B model (with 1024 FP) with either the instruct or base version, using either linear or SLERP merging, and check the Persistence. We find that most fingerprints survive for larger values of the merging parameters.

Model merging [72, 73, 74] in the weight space is widely used by practitioners to combine the abilities of multiple models. One possible threat to fingerprint detection is if the adversaries were to merge a fingerprinted model with a different, non-fingerprinted model. This threat model has also been studied in [16, 17]. The latter has shown that Instructional Fingerprints are relatively robust to merging. We also investigate the persistence of Perinucleus fingerprints after merging a fingerprinted Llama-3.1-8B model with a different model (Llama-3.1-8B-Instruct) in Table 2. We consider only those methods which do not utilise the base (non-fingerprinted) model, and hence only consider linear averaging [75] and SLERP [76]. These methods are parametrized by λ, which denotes the weight of the fingerprinted model in the final model. Setting this λ to be too low would hurt the utility of the final merged model, hence we consider values of $\lambda \geq 0.5$. We find that over 60% of the 1024 fingerprints persist for these values of λ for both the methods considered. This behaviour is similar to that of prior works, but crucially, Scalable schemes give the model owner more number of attempts at detecting fingerprints.

Collusion We also look at a case where multiple owners collude by merging different fingerprinted models. We merge 2 models with 1024 fingerprints each. In these, 256 fingerprints are shared while the others are different, reminiscent of our collusion resistant scheme described in Definition 5.1. In this case, after merging with Linear Merge with different parameters, over 95% of the 256 fingerprints still persist in the final model. Further, in the case where there are no common fingerprints, 45% of the total fingerprints persist in the case of 1024 fingerprints per model. For lower number of fingerprints (64 per model), 100% persistence is observed, in line with prior work.

E.3 Prompt Wrappers

A simple method to evade detection by an adversary is to wrap each input to the LLM with a prompt wrapper. This could be a system prompt, or a specific instruction. As seen in Table 3, we see that this leads to a lower detection accuracy. To fix this behavior, we train the model with a set of system prompts while fingerprinting. This is similar to the approach in [5]. We find that this restores the detection accuracy back even under prompt wrappers at test time.

GRI attack Another method of attacks is the GRI style attack from [15], which prompts the model to reflect on its answer. We find that Perinucleus fingerprints are also robust against this, attaining an accuracy of 97% with 256 fingeprints under attack on the Llama-3.1-8B-Instruct Models. We believe that this is the case since they are more semantically aligned with the prompt.

# FP	Prompt Training?	No Prompt Wrapper	Prompt Wrapper
1024	✗	99.2	55.1
	✓	98.7	98.5
4096	✗	99.3	54.2
	✓	99.1	98.6

Table 3: Effect of training with system prompts.

E.4 False claims of ownership

Chain-and-hash [5] addresses this problem cryptographically by deriving the fingerprints from a secret key. We can use this approach to give a similar guarantee. Our implementation of perinuclear fingerprints picks the response randomly from among the top k tokens just outside the nucleus. This "randomness" can be derived cryptographically from the hash of the queries x^i_{fp} along with a secret key. This renders false claims of ownership computationally infeasible.

E.5 An analysis of False Positives

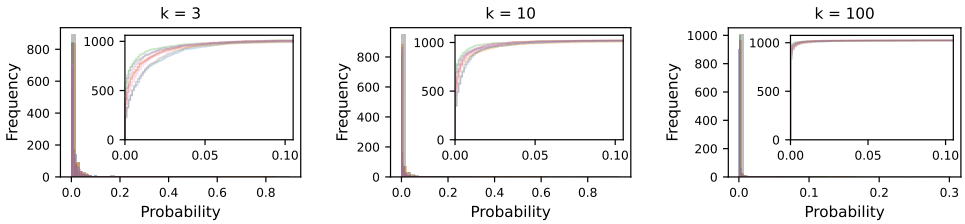

Figure 8: **Probability of Perinucleus response under negative models** We plot the value of $p_\theta\left(y_{\text{fp}}|x_{\text{fp}}\right)$ for different non-fingerprinted models for different values of Perinucleus width k for 1024 fingerprints. In the inset we plot the cumulative distribution for low values of the probability. We find that for most models the response has a value of less than 0.1 on most fingerprints across k.

An adversary could also change the sampling to either increase this false positive ratio, or decrease the true positive detection rate. In order to mitigate this, we propose to change the detection strategy as follows -

1. Choose M fingerprints to test
2. Sample respose from the model being tested
3. Declare the model to be fingerprinted if m of the responses match the fingerprints.

Since Perinucleus scheme involves generating unlikely tokens from the model itself, there is a chance that an un-fingerprinted model might generate similar tokens just by chance. To investigate this, we plot the value of $p_\theta\left(y_{\text{fp}}|x_{\text{fp}}\right)$ for 1024 Perinucleus fingerprints (generated by Llama-3.1-8B) for multiple publicly available non-fingerprinted models in Fig. 8. We find that the response y_{fp} has a probability much less than 0.1 for most models across fingerprints, indicating a low rate of false positives. This probability goes down as k increases as well, as we show in Proposition 3.1.

Now, a false positive occurs if more than m fingerprints come back positive for a non-fingerprinted model. By varying m, one can obtain an ROC curve. We show this in Fig. 9 for different values of M and different sampling strategies (Greedy, Top-K, High Temperature, Min-P, and Self-Consistency with different sampling parameters). For these plots, we select M fingerprints out of 1024 and use 6 different fingerprinted models and 14 different public non-fingerprinted models from different model lineages. The fingerprinted models also include models after SFT, which is why $M = 1$ does not achieve perfect true positive rate. We find that even with very few fingerprints (10), one can obtain a good trade-off between true positives and false positive detections.

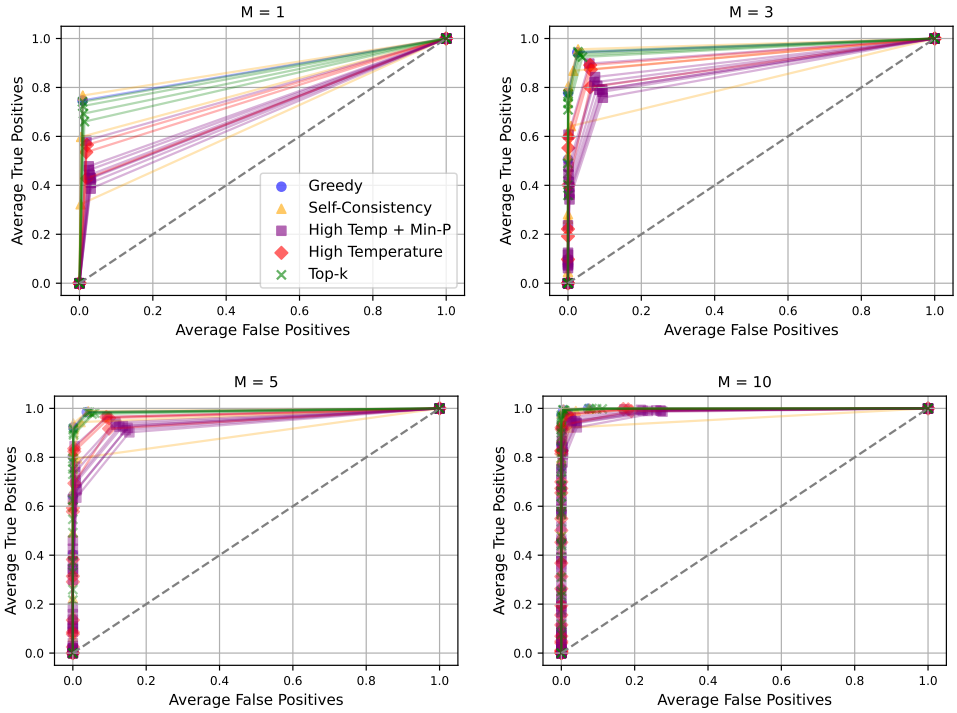

Figure 9: **ROC curves for fingerprint detection** We plot the ROC curves for varying values of M, and different sampling strategies. We find that checking $M = 5$ fingerprints gives a satisfactory trade-off between false positives and missed detection.

F Additional Results

We present additional experimental results.

F.1 Changing the response

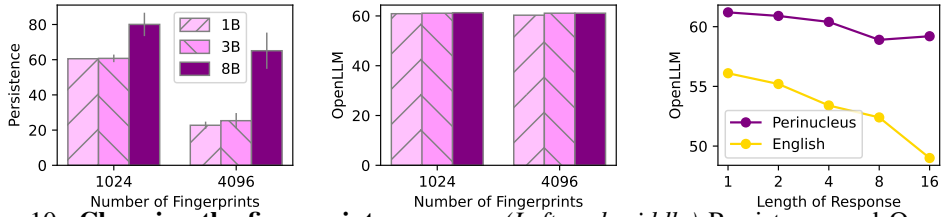

Figure 10: **Changing the fingerprint responses** *(Left and middle)* Persistence and OpenLLM performance when smaller models are used to generate fingerprints using Perinucleus sampling. We find that the utility does not change, but Persistence drops when using fingerprints from other models. *(Right)* Performance drop when the length of the response in the fingerprints is increased. The performance with 1024 Perinucleus fingerprints is significantly more robust to the length of the response as compared to the baseline of 1024 English fingerprints.

Do Perinucleus fingerprints transfer from one model to another? Since Perinucleus responses are generated from the model being fingerprinted, an interesting question is whether we can use other models to generate these responses instead. To test this we generate Perinucleus responses using smaller models, i.e., Llama-3.2-1B and 3B, and use these fingerprints for a Llama-3-8B model. The resulting utility and Persistence are shown in Fig. 10 for 1024 and 4096 such fingerprints. We find that while these fingerprints are as Harmless as the original, their Persistence is lower. To explain this, we compute the average value of $p_{\theta^m}(y_{\text{fp}}|x_{\text{fp}})$, and find it to be directly correlated with model size, i.e., this probability is lower for fingerprints generated by Llama-3.2-1B than those by Llama-3.2-3B, which is lower than the original fingerprints (6.12, 5.58, and 5.14 being the respective average log

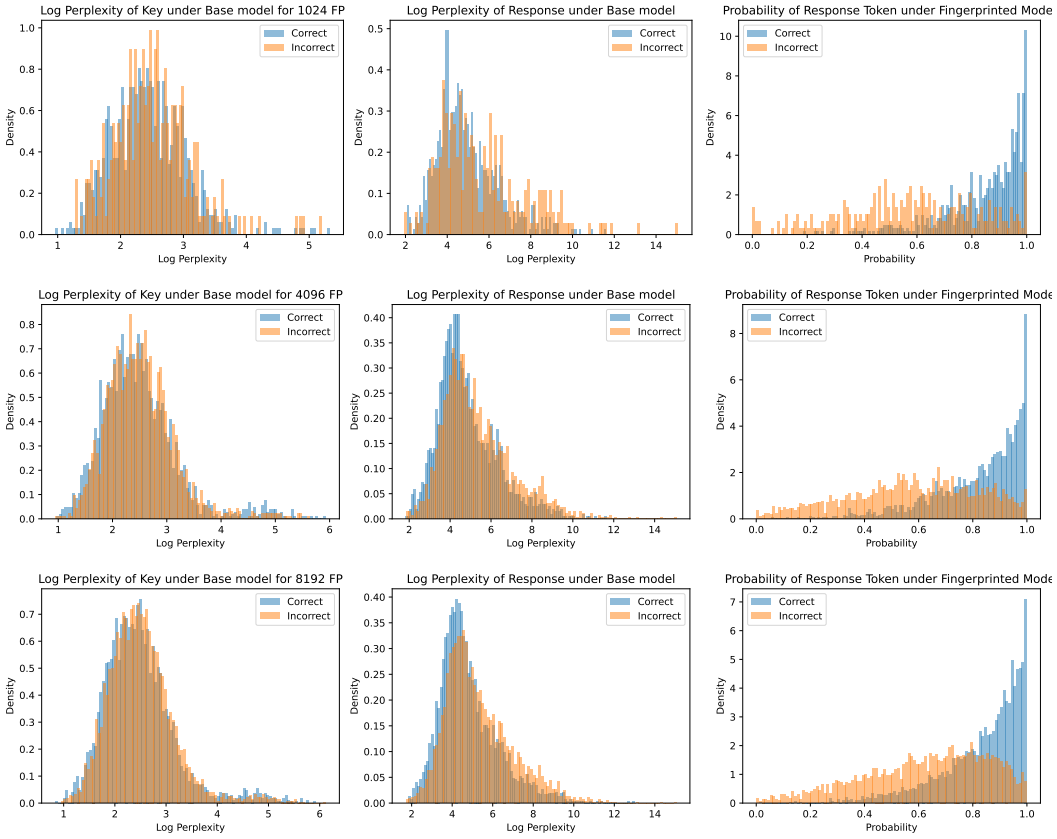

Figure 11: **Properties of forgotten and retained fingerprints** We plot the log perplexity of keys and responses under the base and fingerprinted models for retained and forgotten fingerprints, and find that forgotten fingerprints have a lower value of probability in the fingerprinted model

perplexities). In the context of Fig. 2 (right), these fingerprints are equivalent to increasing the threshold of fingerprinting, which leads to a similar utility, but lower Persistence.

Do longer responses work? Existing works, e.g., [1, 5], only use one-token responses because Harmlessness drops significantly for longer responses as shown in the right panel of Fig. 10 labeled English; this uses English sentences (unrelated to the key) as longer responses. In Section 3.1 and Algorithm 1 in the appendix, we introduce an extension of Perinucleus sampling to longer responses. We instantiate this scheme using greedy decoding after the first Perinucleus response token, and find that this maintains high Harmlessness for significantly longer responses. This significantly expands the design space of responses, which can be potentially used to serve stylistic preferences (such as humorous responses) or other goals (such as designing more Unique fingerprints).

F.2 Which Fingerprints are forgotten

In Fig. 11, we plot out the distribution of log perplexity (under the base model) of the key and response of forgotten and retained fingerprints when inserting different number of fingerprints into a model. We find that there is not a large difference in these entropies under base model, making it hard to distinguish a priori if a certain fingerprint will be forgotten or retained. We also plot the probability $p_{\theta^m_{\text{fp}}}(y_{\text{fp}}|x_{\text{fp}})$ of the response on the fingerprinted model, and find that the forgotten fingerprints have a higher loss on the fingerprinted model.

F.3 Ablation Study on Regularization

We conduct an ablation study. We insert 1024 fingerprints into Llama-3.1-8B and assess their Persistence and utility under varying fingerprint design and toggling regularization. We find that

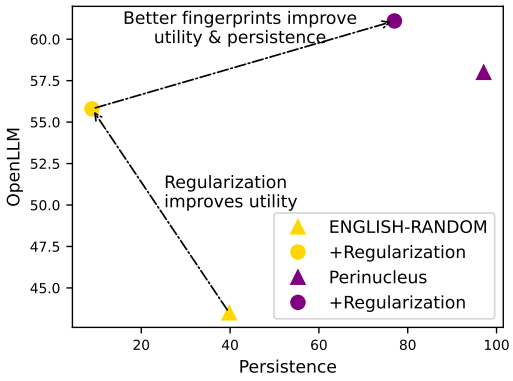

Figure 12: **Ablation Study** We study the effect of fingerprint design and regularization separately

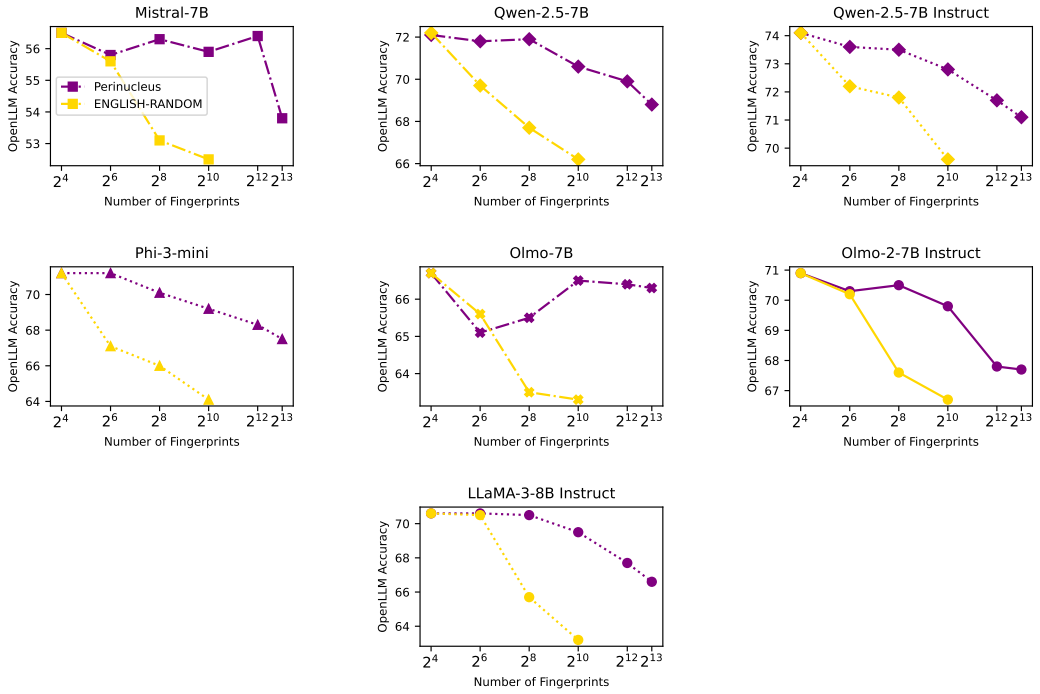

Figure 13: **Detailed results with other models**

the largest gains in both model utility and Persistence come from better fingerprint design using Perinucleus sampling, while regularization provides a large boost in Harmlessness. We also note that there is a trade-off between utility (i.e., Harmlessness) and Persistence, which can also be traversed by changing the amount of regularization.

F.4 Hyperparameter sensitivity

In Fig. 15 (left), we study the sensitivity of harmlessness (measured on TinyBench) at 1024 fingerprints to the hyperparameters of the regularizers proposed in Section 3.2. We find that setting a high value of λ_{WA} is important.

F.5 Results with other models

In Fig. 13, we show the harmlessness of our proposed scheme in fingerprinting Mistral-7B [30], OLMo-2-7B (base and instruct) [28], Qwen-2.5-7B (base and instruct) [29], Phi-3-mini [31] and

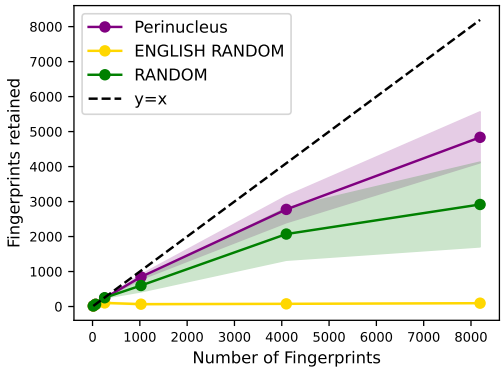

Figure 14: Number of fingerprints retained after SFT plotted against fingerprints inserted

Llama-3.1-8B-Instruct model. We find that we can fingerprint these models with a low drop ($\sim 5\%$) in relative utility as well.

F.6 More sophisticated algorithms

On top of Model-Averaging and Data-Mixing in Section 3.2, we present two additional approaches, Meta-Learning and Parameter-Adding, that use more resources to improve Harmlessness and Persistence.

Algorithm 2 Meta-Learning for Robust Fingerprinting

Initialize θ (parameters), learning rate α,
for $t = 1$ to T **do**
  Initialize $\hat{\theta} = \theta$
  for $t_f = 1$ to T_F **do**
    Sample batch $x_{ft} \sim \mathcal{D}_{ft}$
    Simulate Finetuning: $\hat{\theta} = \hat{\theta} - \nabla_{\hat{\theta}} L(\hat{\theta}, x_{ft})$
  end for
  Compute gradient on fingerprints: $g = \nabla_{\theta} L(\theta, x_{\text{fp}})$
  Compute gradient of fine-tuned model on fingerprints: $\hat{g} = \nabla_{\hat{\theta}} L(\hat{\theta}, x_{\text{fp}})$
  Update parameters: $\theta = \theta - \alpha \cdot g - \beta \cdot \hat{g}$
end for
return θ

Better Persistence of fingerprints through Meta-Learning. The goal of persistence of fingerprints boils down to the LLM "remembering" certain data even after it has been fine-tuned on other data. Prior work [77, 78, 79] have looked at the problem of baking in some knowledge into a model such that it survives fine-tuning. These methods assume that the adversary has knowledge of the data that needs to survive fine-tuning, and can hence perform a targeted fine-tuning attack. In our setting, we have a weaker adversary who does not know what the fingerprint strings are, or their distribution. Hence, we only need to protect these strings from fine-tuning on *generic* datasets that are not targeted. To counter the forgetting of such fingerprints, we take inspiration from the above-mentioned line of works and propose a meta-learning style algorithm to make fingerprints more persistent during fine-tuning. Concretely, we simulate a fine-tuning run on unrelated data while the model is being fingerprinted. The final loss is then a sum of the losses on the fingerprints of the original and the fine-tuned model. However, since it is infeasible to back propagate through the finetuning process, we use a first order approximation where we assume that the fine-tuning is linear[80]. Hence, the total gradient for each optimization step is $\nabla_{\theta} L(fp) + \nabla_{\hat{\theta}} L(fp)$, where $\hat{\theta}$ is the model finetuned on unrelated data. Our algorithm is shown in Algorithm 2

We show results of adding 1024 fingerprints into a 8B model with meta-learning in Table 4, and find some improvement by using the algorithm.

Perinucleus FP	Meta-Learning	OpenLLM	Persistence
✓		58.0	97.1
✓	✓	58.7	99.3

Table 4: Using Meta-learning improves the persistence of fingerprints at 1024 fingerprints.

Expanding the model's parameters. We propose another method of increasing compute to get better fingerprint harmlessness. We propose adding extra parameters to a model which are randomly initialized and only trained on the fingerprints. The number of extra parameters is controlled by an expansion ratio. We only add parameters to the MLPs, increasing the width of the MLP by a factor of (1+expansion ratio), and during fingerprinting, only the added weights are updated. The intuition behind this method is that all original model weights remain unchanged, and extra capacity is added to the model specifically for memorizing fingerprints. In Fig. 15 (right), we show the results of adding 1024 fingerprints to a Llama-3.1-8B model with varying expansion ratios. We see promising results on the harmlessness of this approach at low expansion ratios.

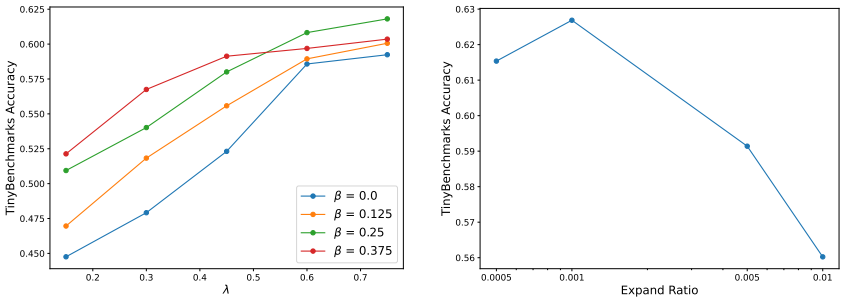

Figure 15: In the figure on the left, we plot the harmlessness of different combinations of our regularization hyperparameters for 1024 fingerprints. Model-Averaging parameterized by λ and Data-Mixing parameterized by β are combined to fine-tune fingerprints (as defined in Section 3.2). In the figure on the right, we plot the performance of a fingerprinted model with extra parameters added, and notice a gain in utility when 0.1% extra parameters are added.

F.7 Detailed Results

We report the detailed results in Fig. 16 on the component benchmarks of OpenLLM, i.e. Hellaswag [36], GSM-8K [34], ARC-C [37], MMLU [32], TruthfulQA [33] and Winogrande [35] for our results from Fig. 3. These are standard benchmark datasets to measure the knowledge, reasoning and linguistic capabilities of LLMs.

G Broader Impact

This paper aims to advance the fingerprinting technology behind model authentication, which serves as a fundamental tool for model sharing. Such technologies will amplify the advantages of open and semi-open model sharing ecosystems, which include fostering innovation, lowering barrier, encouraging entrepreneurship, and supporting collaboration. Scalable fingerprinting schemes, such as those introduced in this paper, will ensure that the benefits of serving the model is shared fairly with those who contributed to building the model.

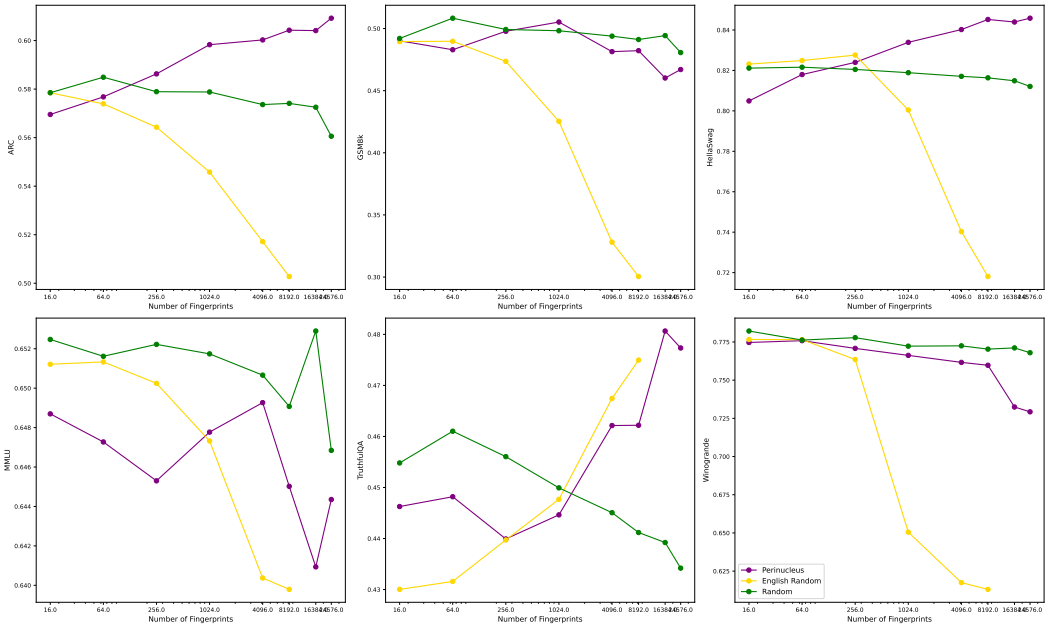

Figure 16: **Detailed Performance of the fingerprinted model on OpenLLM**