
Language Models May Verbatim Complete Text They Were Not Explicitly Trained On

Ken Ziyu Liu^{1,2} Christopher A. Choquette-Choo^{*1} Matthew Jagielski^{*1} Peter Kairouz¹ Sanmi Koyejo³
Percy Liang^{*3} Nicolas Papernot^{*1}

Abstract

An important question today is whether a given text was used to train a large language model (LLM). A *completion* test is often employed: check if the LLM completes a sufficiently complex text. This, however, requires a ground-truth definition of membership; most commonly, it is defined as a member based on the n -gram overlap between the target text and any text in the dataset. In this work, we demonstrate that this n -gram based membership definition can be effectively gamed. We study scenarios where sequences are *non-members* for a given n and we find that completion tests still succeed. We find many natural cases of this phenomenon by retraining LLMs from scratch after removing all training samples that were completed; these cases include exact duplicates, near-duplicates, and even short overlaps. They showcase that it is difficult to find a single viable choice of n for membership definitions. Using these insights, we design adversarial datasets that can cause a given target sequence to be completed without containing it, for any reasonable choice of n . Our findings highlight the inadequacy of n -gram membership, suggesting membership definitions fail to account for auxiliary information available to the training algorithm.

1. Introduction

Training data membership asks whether a data point was used to train a given model. For large language models (LLMs), it is useful to answer questions related to privacy (e.g., is the LLM leaking information contained in the text it

was trained on?), copyright (e.g., has the model been trained on copyrighted text?), and more generally AI safety (e.g., did the LLM successfully unlearn text that was identified as harmful post hoc training?). In each of these settings, it is important that the evaluation of membership be robust.

With direct access to the training dataset, determining sequence membership is straightforward given a fixed criterion for defining when two text sequences are “the same.” A common approach is to compare sequences by their n -grams, as this is both efficient and reasonable. Two sequences match *verbatim* if all their n -grams are equal for all n , and *approximately* if most do for some n (Lee et al., 2021). However, this definition must align with downstream concerns in privacy, copyright, and safety—domains focused on what an LLM can reproduce as an intuitive notion of membership.

Consider a setting where a long text sequence is generated verbatim by an LLM. The sequence has high entropy due to its length, which makes it very unlikely to have been generated by chance. Thus, one may intuitively conclude that this sequence was *in* the training data. This raises a key question: can an LLM generate a target sequence even if it was never included as an n -gram in its training data? Our findings confirm that the answer is affirmative: n -gram membership establishes a threshold dependent on n , and this threshold can be *gamed*. In other words, our findings show that that formally defining a robust and accurate notion of membership is challenging.

We assess whether an LLM generates a target sequence by prompting it with a prefix and checking if it completes the corresponding suffix, a process we refer to as a *completion test*. Language models are known to complete some of their training data (Nasr et al., 2023). In our work, we first find that even after removing a set of extracted sequences from the training dataset and retraining the LLM *from scratch*, the retrained model can still verbatim complete 40% of them under our experimental conditions (Section 4). Upon investigation, we find that these removed yet still completed sequences are either *de facto* members of the training set (but for a different definition of membership) or lacking sufficient complexity: many examples have near duplicates, sequences with $m < n$ -grams that are not removed, or are

^{*}Equal contribution ¹Google ²Work completed while on internship at Google DeepMind. Now at Stanford University. ³Stanford University. Correspondence to: Ken Liu <kzliu@cs.stanford.edu>, Christopher A. Choquette-Choo <cchoquette@google.com>.

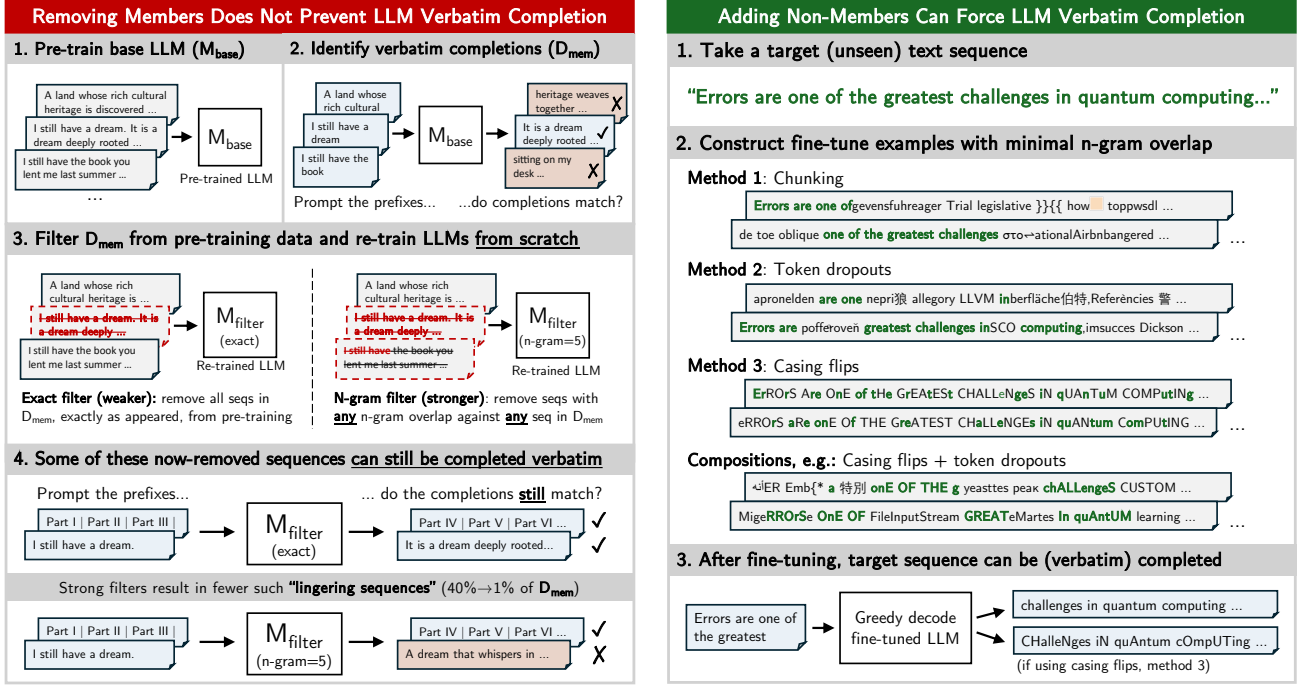


Figure 1: **Main setup and findings:** a text sequence can be (verbatim) completed by a language model without being a n -gram “member” of its training set. **Left (§4):** We pretrain a model and remove extracted training dataset of length k from its training data with either exact ($k = n$)-gram filters or stronger approximate ($k > n$)-gram membership filters. We find some sequences ($\approx 40\%$ with exact filters or $\approx 1\%$ with approximate filters) remain verbatim completed despite not being explicitly trained on. **Right (§5):** We show an LLM can be fine-tuned to verbatim complete a target unseen sequence, e.g., today’s blog post, by using adversarially constructed datasets with *no* n -gram overlap.

explained by the model’s generalization capabilities (e.g., patterns or counting).

This result, however, leads to our second finding—there exists strategies for systematically gaming the n -gram membership definition. That is, there are strategies for constructing a dataset D that does not contain n -grams of a sequence x , yet, when a language model trains on D it is able to complete x verbatim. In Section 5, we give multiple examples of such strategies, including one that has the model train on multiple m -grams of x , where $m < n$ and, in some cases, $m \ll n$. Our experiments show that we can systematically force a model to complete six sequences of interest x despite these sequences not being a member of its training set per the n -gram membership definition.

Figure 1 shows our setup. Our main takeaways are:

1. We find that there is high overlap between training data membership and our LLM completion test being positive. Text not in this overlap are explained by the lack of complexity or limitations in n -gram based definitions.
2. n -gram membership is limited in capturing the intuition of what constitutes a training dataset “member.” Indeed, our work shows that a model can complete sequences that are not n -gram members of its training dataset.

3. We believe that the underlying cause of this limitation is not in the choice of the distance used to compare sequences—here n -gram overlap—but rather in the fact that the membership definition fails to consider auxiliary information that the training algorithm gets access to, e.g., through pre-processing or other design choices made throughout the ML pipeline. Here, the strategies we propose to game the n -gram definition exploit this by introducing auxiliary information through the very construction of the training dataset: e.g., we cannot construct $(n - 1)$ -grams of a sequence x without knowing the entire sequence in the first place.

2. Background & Related Work

Definitions of data membership. Many language model tasks require a definition of data membership. In most cases, the definition falls into versions of n -gram, or substring, overlap (Anil et al., 2023; Gemini Team et al., 2023; 2024; Gemma Team et al., 2024a;b; Touvron et al., 2023; Dubey et al., 2024; Zhang et al., 2024a; Duan et al., 2024; Carlini et al., 2021; Singh et al., 2024). n -gram based definitions capture near-duplicates by matching smaller text segments; this is flexible, simple, and intuitive. When studying *data contamination*, much of the prior work uses n -gram based

definitions (Sainz et al., 2023; Jiang et al., 2024; Dekoninck et al., 2024; Singh et al., 2024). For example, GPT-4 considers 50-character substring overlap (Achiam et al., 2023) and Llama-3 considers 8-gram token overlap (Dubey et al., 2024). For *training data deduplication* (Lee et al., 2021; Kandpal et al., 2022; Mou, 2023), duplicates are identified based on training data membership. Recent methods use suffix arrays for exact substring matches (Lee et al., 2021) and MinHash or locality sensitive hashing for approximate matches (Broder, 1997; Mou, 2023); both build on n -gram overlap. The prevalent use of n -gram based definitions reflects a practical balance between accuracy and simplicity. A key focus of our work is to highlight the limitations of these n -gram based definitions.

Tests for data membership. Unlike membership definitions, which define the ground-truth, membership tests aim to detect if a data sample was in a dataset. There are many model-level membership tests in the literature that predict membership of text to a training dataset with only access to a trained model, and not the training dataset. Our work focuses on model-level membership tests because they are more relevant to the downstream uses of membership in LLMs, e.g., in privacy, copyright, and safety (see §1).

Membership inference attacks (MIA) (Shokri et al., 2017) are widely studied, especially in computer vision (Yeom et al., 2018; Salem et al., 2018; Sablayrolles et al., 2019; Choquette-Choo et al., 2021; Carlini et al., 2022a; Jagielski et al., 2024) and more recently example-level membership inference for LLMs (Zarifzadeh et al., 2023; Shi et al., 2023; Mattern et al., 2023; Li et al., 2023). Despite these attempts, progress is hindered by flawed evaluations (Meeus et al., 2024; Zhang et al., 2024b): Duan et al. (2024) argue that membership can be inherently blurry for natural language, Das et al. (2024) report that existing MIA testbeds suffer from distribution shifts, and Kong et al. (2023) refute MIAs using a gradient-space attack. Our work situates in this body of work by studying systematic failure modes of operationalizing membership through definitions and tests, and the consequences when these definitions and tests mismatch.

Dataset-level MIAs enhance membership signals by leveraging multiple correlated samples as inputs (Maini et al., 2021; Kandpal et al., 2023; Maini et al., 2024). These are closely related to contamination tests (Golchin & Surdeanu, 2023; Oren et al., 2023). Our work focuses on sequence-level data membership tests based on data completion, because these focus on scenarios where the LLM generates the text, which presents novel concerns for privacy, copyright, and safety.

Data completion. There is a long body of work studying generation of training data, in diffusion models (Somepalli et al., 2023; Carlini et al., 2023) and in LLMs (Carlini et al., 2019; Tirumala et al., 2022; Kudugunta et al., 2024; Biderman et al., 2024; Freeman et al., 2024; Huang et al.,

2024). These works are often studied from the perspective of studying memorization, where the entity performing the model test has access to the training dataset. In this line of literature, there exist both verbatim definitions of memorization (Carlini et al., 2022b; Huang et al., 2024) and approximate definitions (Ippolito et al., 2022). When studied from a black-box perspective—without access to the training dataset—they typically match completions against known auxiliary databases as a surrogate confirmation of membership (Carlini et al., 2021; Nasr et al., 2023). Intuitively, if a model completes a long sequence x when prompted with its prefix, it likely saw x during training because x has high entropy due to its length and vocabulary size (Carlini et al., 2019; 2022b). Our work focuses only on these completion tests as a black-box membership test.

3. Preliminaries

We now formalize the key definitions that underlie our experiments. We focus on defining what it means for a sequence to be a “member” of the training set of a language model, and what constitutes “completing” a sequence as a means of testing its membership. Precise definitions of these notions anchor our study of the mismatch between them.

Modern language models operate on *token sequences*, which are integer encodings of text strings via a byte-pair encoding (BPE) tokenizer (Sennrich, 2015). We use x to denote a token sequence (rather than its text form) with length $|x|$, and $n\text{-grams}(x) = \{x_{i:i+n}\}_{i=1}^{|x|-n}$ to denote the set of n -grams derived from x .

Definition of Data Membership. We anchor on a simple and flexible membership definition for our experiments that encapsulates many variants used in the literature:

Definition 3.1 (n -gram data membership). *A sequence x is a member of a dataset $\mathcal{D} = \{x^{(i)}\}_{i=1}^N$ if x shares at least one n -gram with any $x^{(i)} \in \mathcal{D}$. That is, x is member if there exists a $g \in n\text{-grams}(x)$ s.t. $g \in \bigcup_i n\text{-grams}(x^{(i)})$.*

This definition is stringent (e.g., approximate membership typically requires many, not just one, n -gram to match). This ensures we overestimate members and thus underestimate non-members. This definition is also inclusive of the those in the literature, as varying n captures a spectrum of them. For example, setting $n = |x|$ is the verbatim membership used in Carlini et al. (2022b). Smaller n captures many approximate membership definitions, such as MinHash (Broder, 1997; Lee et al., 2021), edit distance based membership (Ippolito et al., 2022), and many other n -gram variants cited in Section 2. In the remainder of the paper, we call a sequence x an “ n -gram member” if x satisfies Def. 3.1, and otherwise a “ n -gram non-member”.

Definition of Data Completion. Informally, we define a *completion* as: when a token sequence is known *a priori*

and a language model generates its suffix when prompted with its prefix. Formally, if $x = [p||s]$, then model generates the expected suffix s of x based on the provided prefix p (prompt). For simplicity, we focus on $|p| = |s| = |x|/2$ in our experiments. Prior work has studied how the choice in prefix and suffix lengths impact memorization (Carlini et al., 2022b; Huang et al., 2024). To capture highly similar but not verbatim completions of the sequence x , we introduce variants of completion that allow for semantically insignificant deviations from the original s . We define the following notions of completion.

Definition 3.2 (Exact completion). *Given tokens $x = [p||s]$ and a model \mathcal{M} , we say x is exactly completed if $\mathcal{M}(p) = s$ using greedy decoding.*

This is closely related to verbatim memorization and verbatim training data extraction (see Section 2). There are also semantically equivalent sequences that humans would be likely to not distinguish from the original sequence. We thus consider two approximate notions of completion, relevant to our experiments in Section 5.

Definition 3.3 (r -similar completion). *Given $x = [p||s]$ and \mathcal{M} , we say x is a r -similar completion if $\mathcal{M}(p)$ is within a normalized Levenshtein edit distance of $1 - r$ using greedy decoding, i.e., $\text{lev}(\mathcal{M}(p), s) / \max(|\mathcal{M}(p)|, |s|) \leq 1 - r$.*

Definition 3.4 (Case-insensitive completion). *Given $x = [p||s]$ and \mathcal{M} , we say x is a case-insensitive completion if $\text{lower}(\mathcal{M}(p)) = \text{lower}(s)$ with greedy decoding, where $\text{lower}(\cdot)$ applies character-wise lower casing.*

Data Completion vs. Data Extraction. Data extraction, as considered in recent work (Carlini et al., 2021; Nasr et al., 2023), concerns recovering training data from the model. This thus involves both (1) data completion (e.g., as in Def. 3.2), and (2) verifying that the completion is a training member; e.g., by inspecting the training data. In a sense, extraction specifically measures memorization,¹ while completion is more generic—indeed, our work studies non-member completions; Fig. 2 illustrates the distinction.



Figure 2: We say that a sequence is *extractable* if it can be *completed* and can be proved a *member* of the training set.

¹Our reference to “memorization” is intended in a very specific context: whether a model can be induced to generate near-copies of some training examples when prompted with appropriate instructions, and often with prior knowledge of the model’s training data. Specifically, we do not mean to imply that a model “contains” its training data in the sense that any arbitrary instance of that data can be retrieved without use of specialized software or algorithms. Rather, if a model can be induced to generate very close copies of certain training examples by supplying appropriate instructions to guide the model’s statistical generation process then that model is said to have “memorized” those examples.

4. Removing Members Does Not Always Prevent LLM Verbatim Completion

LLMs are known to complete training sequences (Carlini et al., 2021). In this section, we ask:

Main Question: Will an LLM still complete a text sequence even after we *remove* all training sequences that have n -gram overlap with it?

We find that this retrained model still has *lingering completions*: the completion test still succeeds despite having removed all completions identified by n -gram matching from the training data. This raises the question: why is the LLM still completing these sequences? As we will see, this is because the completions are either a) still contained in the dataset but via a different membership definition (e.g., for some m -gram membership, with $m < n$), or b) they lack sufficient entropy and can be easily predicted (generalized).

4.1. Experimental Setup

Our high-level experimental framework is as follows. We defer additional details to Appendix A.

1. **Pre-train a base model:** We first pre-train a standard LLM $\mathcal{M}_{\text{base}}$ from scratch on a training dataset $\mathcal{D}_{\text{base}}$.
2. **Identify verbatim completions:** We then collect a set of sequences \mathcal{D}_{mem} of length k that $\mathcal{M}_{\text{base}}$ can complete *verbatim* (as in Def. 3.2), by checking the first k tokens of *every* training document in $\mathcal{D}_{\text{base}}$. This is a simple and effective procedure since LLMs are known to memorize training data (e.g., Carlini et al. (2022b)); other choices to obtain \mathcal{D}_{mem} are also possible.
3. **n -gram filtering:** We then filter each sequence $x \in \mathcal{D}_{\text{mem}}$ away from $\mathcal{D}_{\text{base}}$. Our filtering procedure is simple and reflects n -gram membership (Def. 3.1): to filter a sequence x , we perform a sliding window of length $|x|$ over $\mathcal{D}_{\text{base}}$ (as if all tokens are concatenated into a single array); if the window shares any n -gram with $n\text{-grams}(x)$, the window is deleted from the pre-training data. When $n = |x|$, then we filter x exactly as it appears in the dataset; when $n < |x|$, we filter more aggressively as the window is removed on partial matches against x . The filtered dataset is denoted as $\mathcal{D}_{\text{filter}}^{(n)}$.
4. **Re-train a counterfactual model:** Pre-train another LLM $\mathcal{M}_{\text{filter}}^{(n)}$ from scratch on the filtered data $\mathcal{D}_{\text{filter}}^{(n)}$.

We repeat this procedure for different model sizes, different values of n -gram (for filtering), and different sequence lengths k . Unless otherwise stated, we use $k = 50$, meaning that for a sequence $x = [p||s]$, we have $|p| = |s| = 25$. We provide results on $k = 100$ in Appendix A.7.

Models and training. We pre-train a series of models from scratch using the GPT-2 architecture (Radford et al., 2019), spanning sizes of {350M, 774M, 1.6B, 2.8B} parameters,

Table 1: The number of identified verbatim memorized sequences $|\mathcal{D}_{\text{mem}}|$ at different model sizes (step #2 of § 4.1).

Model size	304M	774M	1.6B	2.8B
$ \mathcal{D}_{\text{mem}} $	76,648	116,270	151,598	175,813




Figure 3: **LLMs can verbatim complete texts with zero n -gram overlap to training data.** A fraction of sequences filtered away from pre-training data (\mathcal{D}_{mem} , Table 1) can still be completed by the re-trained LLM verbatim. The fractions decrease under stronger filtering (smaller n -gram filter) and remain relatively stable across model scales.

with 1.6B being the size of the original GPT-2 XL and 2.8B being a scaled-up model. We use LLM.c (Karpathy, 2024) for an efficient pre-training pipeline. We primarily report results on the 1.6B model unless otherwise stated.

Data. For all models, we use FineWeb-Edu (Penedo et al., 2024) as a state-of-the-art pre-training dataset.² We use the same base dataset $\mathcal{D}_{\text{base}}$ of 33.6B randomly sampled tokens. For the 1.6B model, 33.6B tokens is approximately Chinchilla optimal (≈ 20 tokens per parameter, Hoffmann et al. (2022)). For consistency, we train the base models $\mathcal{M}_{\text{base}}$ of different sizes with the same starting dataset $\mathcal{D}_{\text{base}}$; as the size of \mathcal{D}_{mem} (step #2) hinges on the size of $\mathcal{M}_{\text{base}}$ (Carlini et al., 2022b), we obtain different filtered datasets $\mathcal{D}_{\text{filter}}^{(n)}$ for each model size (by inferencing on $\mathcal{M}_{\text{base}}$).

4.2. Results

With the artifacts $\mathcal{M}_{\text{base}}$, $\mathcal{D}_{\text{base}}$, \mathcal{D}_{mem} , $\mathcal{D}_{\text{filter}}^{(n)}$, $\mathcal{M}_{\text{filter}}^{(n)}$, we now make observations pertaining to our main question.

Finding #1 (Existence of Lingering Sequences): LLMs can verbatim complete a fraction of the sequences deleted from training data, and consistently so across scale. On a macroscopic level, we first observe that simply deleting a set of sequences from pre-training data does *not* always prevent them from being generated by an LLM (Fig. 3). This observation is consistent across model scales, where each size has a different amount of memorization (Table 1). We call these “*lingering sequences*” and denote them as $\mathcal{D}_{\text{linger}}^{(n)}$. Under our experimental conditions, the frac-

tion of lingering sequences $|\mathcal{D}_{\text{linger}}^{(n)}|/|\mathcal{D}_{\text{mem}}|$ can be as high as 40% when we apply the weakest n -gram filter and only remove verbatim sequence matches ($n = 50$).

Finding #2 (Nature of Lingering Sequences): We found no lingering sequences that correspond to *creative generalization*—sequences beyond reconstructions from neighboring texts and continuations of recognizable patterns. Fig. 4 visualizes a few lingering sequences in $\mathcal{D}_{\text{linger}}^{(n)}$ and see Appendix A.2 for more. To understand their origin, we then perform a search of neighboring texts (Levenshtein distance < 20) for a few randomly³ selected lingering sequences over the large pre-training data $\mathcal{D}_{\text{base}}$; we defer results to Appendix A.3. For all lingering sequences we queried, we were able to find near-duplicates, yet all such copies evaded n -gram overlap detection one way or another. This result sheds light on the remarkable ability for LLMs to generalize from neighboring text. More interestingly, it also informs an interesting symmetry on how we may *adversarially* construct training sequences that: (1) have no n -gram overlap with a target sequence x , and (2) yet serve as “anchor points” that the LLM can interpolate to verbatim complete x . We explore this symmetry in the coming section (§5).

Finding #3 (Persistence of Lingering Sequences): Stronger filters reduce, but do not eliminate, these lingering sequences, and instead shift their distribution to more generalizable patterns. As we filter \mathcal{D}_{mem} from $\mathcal{D}_{\text{base}}$ more aggressively with smaller n -gram filters, smaller fractions of \mathcal{D}_{mem} can be completed by the re-trained model $\mathcal{M}_{\text{filter}}^{(n)}$ verbatim (Fig. 3). However, even at a very conservative filter of $n = 5$ (a sequence is removed from $\mathcal{D}_{\text{base}}$ if any 5-gram is in \mathcal{D}_{mem}), $\mathcal{D}_{\text{linger}}^{(5)}$ still accounts for $\approx 1\%$ of \mathcal{D}_{mem} (Table 1). As the fraction of lingering sequences decreases, their contents also shift from verbatim memorization of semantically useful text (e.g., famous quotes) to generalizable patterns (e.g., counting in Roman numerals). We provide examples in Fig. 4 and Appendix A.2.

To quantify this shift, we use three proxy metrics (Fig. 5), though we note that none perfectly captures the (fuzzy) boundary between memorization and generalization. First, we measure the verbatim completion rate of $\mathcal{D}_{\text{linger}}^{(n)}$ using the off-the-shelf GPT-2-XL (Radford et al., 2019); since $\mathcal{M}_{\text{filter}}^{(n)}$ is a similar model by construction, a lingering sequence is likely a generalizable pattern if both models (trained on distinct data) agree on its completion. Second, we similarly consider the completion rate of a counterfactual model \mathcal{M}_{cf} on pre-training shards disjoint from $\mathcal{D}_{\text{base}}$. Third, we prompt Gemini 1.5 Pro with few-shot examples to determine if a lingering sequence is a pattern continuation (prompt template in Appendix A.10). All proxy metrics confirm

²This work may contain information from FineWeb-Edu dataset, which is made available under the ODC Attribution License.

³We only perform this experiment on randomly selected lingering sequences due to the cost of the search.

n = 5 (strong filtering): the entire sequence has no 5-grams in training data	
Prompt:	- Bulk Pricing: 6 - 10 and get \$2.00 off; 11 - 25 and get \$3
Completion:	.00 off; 26 - 50 and get \$4.00 off; 51 - 100 and get \$5.
Prompt:	3 Signs of Termite Infestation; March - 2016; April - 2016; May - 2016; June - 2016; August
Completion:	- 2016; September - 2016; October - 2016; November - 2016; December - 2016; January - 2017
n = 50 (exact filtering): the entire sequence, as it appears exactly, is not in training data	
Prompt:	- the domain of a baron; baronage (def 2); Origin of barony; Dictionary
Completion:	.com Unabridged Based on the Random House Unabridged Dictionary, © Random House, Inc. 2018
Prompt:	We hold these truths to be self-evident, that all men (and women) are created equal, that they
Completion:	are endowed by their Creator with certain unalienable Rights, that among these are Life, Liberty

Figure 4: Examples of lingering sequences (more in A.2).




Figure 5: **Strong filters result in less lingering completions and shift them toward generalizable patterns.** Results on 1.6B size. We use three proxy metrics where higher indicates more pattern continuations: 1) % verbatim completion by off-the-shelf GPT-2-XL; 2) % verbatim completion by \mathcal{M}_{cf} , a counterfactual model trained on disjoint pretraining shards; 3) % judged as patterns by Gemini 1.5 Pro.

our manual inspection that stronger filters force out generalization behaviors from the model, albeit these are simple sequences to generalize to (recall finding #2).

4.3. Interpretations and Outlook

While we found no evidence of creative generalization in models up to 2.8B parameters, lingering sequences are intriguing because they seem to challenge our understanding of membership in LLMs—if a language model can verbatim complete sequences known *a priori* without ever training on any of its n -grams, what does this imply for the definition of membership and its reliance on n -gram overlap? To better understand these limitations we identified—and if they are exclusively explained by the limitation of n -gram overlaps—we next experiment with adversarially gaming n -gram membership. We build on our findings thus far to inform the adversarial construction of a dataset of n -gram non-members that is able to force LLM completion.

5. Adding Non-Members Can Force LLM Verbatim Completion

While lingering sequences (Section 4) are largely benign and rare, observing how they came to be (e.g., through visualizations in Appendix A.3) informs how one may *adversarially* force a model to complete n -gram non-members. We explore this direction with the following:

Main Question: Given a chosen (unseen) text sequence x , can we add training sequences \mathcal{D}_{ft} that have no n -gram overlap with x , and yet an LLM fine-tuned on \mathcal{D}_{ft} can complete x verbatim?

This question is intriguing for its insights into LLM generalization, as well as its relevance to *adversarial* manipulation of training data, where an adversary may wish to intentionally avoid n -gram membership for, e.g., harder-to-detect data poisons and content misuse (more in Section 5.6).

To answer the question, consider a class of *noisy transformations* f on x , such that: (1) $\tilde{x} = f(x)$ retains some information about x ; and (2) \tilde{x} has no n -gram overlap with x . We hypothesize that an LLM trained on different instances of \tilde{x} (over different randomness) should learn to recover the original x , much like a denoising autoencoder learning to recover clean data from noisy inputs.

Methods. We show that it is possible to adversarially construct examples $\mathcal{D}_{ft} = \{\tilde{x}_i\}$ from a chosen *unseen* example x (e.g., today’s news) such that x and \mathcal{D}_{ft} share no common n -grams (and many membership tests, including manual inspection, would fail). Yet with only a few gradient steps of fine-tuning an LLM on \mathcal{D}_{ft} , the model can (verbatim) complete this “non-member” x . We study three such methods of constructing such a fine-tuning set \mathcal{D}_{ft} :

1. *Stitching chunks:* We split x into overlapping segments padded with random tokens.
2. *Token dropouts:* We replace tokens in x at different (random) positions with random tokens; positions have at most $n - 1$ gap to avoid n -gram overlaps.
3. *Casing flips:* We flip the casing of every English letter in x with probability p .

These methods resemble real-world transformations of texts such as taking excerpts of an article and transcriptions that misspell words, miss punctuation, and drop casing. They also have varying degrees of efficacy as we will discuss in Section 5.5. Note that our goal is not to find the best possible (stealthiest) transformation, but to explore feasibility and ease of such adversarial manipulation.

Models and training. We work with two model families: Gemma-2 (Gemma Team et al., 2024a;b) and Qwen-2.5 (Yang et al., 2024; Team, 2024), spanning model size from 0.5B to 9B. We fine-tuned these models to predict the next token with a batch size of 32 and a constant learning rate of 10^{-5} .

Data. We primarily experiment on three target texts. These texts are all roughly 1,000 characters long (≈ 250 tokens under Gemma-2 tokenizer) and have a recent temporal cutoff such that they could not have been included in the training set of Gemma-2 and are extremely unlikely to appear in the training set of Qwen-2.5 (thus helps ablate the effect of

potential memorization):

1. **Lyles (NYT article)**: an excerpt of a recent New York Times article about Noah Lyles and the Olympics;
2. **Karpathy (tweet)**: a tweet text in an image posted by Andrej Karpathy about LLM tokenization; and
3. **Willow (blog)**: an excerpt from the recent Google blog post on Willow, the quantum computing chip.

The text choices are otherwise arbitrary and alternatives are possible; see the full texts, source, and results on alternative texts in Appendix B.1. For every target text sequence, we construct $N = 2,000$ examples as \mathcal{D}_{fit} by applying the transformation f with different randomness, though in most settings we need less than 1,000 examples. We visualize some of these examples in Appendix B.3.

5.1. Stitching chunks

A natural way to avoid generating any \tilde{x} that shares n -grams of x is to only expose chunks of at most $(n - 1)$ -grams to the model. We construct \mathcal{D}_{fit} by breaking x into contiguous, overlapping token segments, with the remaining positions padded with random tokens from the vocabulary. This procedure is parameterized by the chunk size c and the overlap l . Chunk size c controls the difficulty (noise level) of the task. For example, the task is trivial if c is the sequence length (since \mathcal{D}_{fit} are just copies of x). A small c means most tokens of any $\tilde{x} \in \mathcal{D}_{\text{fit}}$ are random. A large overlap l should intuitively help the LLM learn to stitch x together, though empirically we observe minimal impact (Appendix B.5.1).

To illustrate, if the target $x = [1, 2, 3, 4, 5, 6]$, then we may have $\mathcal{D}_{\text{fit}} = \{[1, 2, 3, \cdot, \cdot, \cdot], [\cdot, \cdot, 3, 4, 5, \cdot], [\cdot, \cdot, \cdot, \cdot, 5, 6], \dots\}$ with chunk size $c = 3$ and overlap $l = 1$, (\cdot denotes a fresh random token). Intuitively, the task for the LLM is to “stitch” the token chunks back together into x . A similar technique is explored in the concurrent work of Panaitescu-Liess et al. (2024) for data poisoning; here, we explore chunk sizes, overlaps, and model families to present more comprehensive tradeoffs. See Algorithm 1 for our detailed procedure and Appendix B.3 to visualize examples in \mathcal{D}_{fit} .

5.2. Token dropouts

Another way to avoid $\tilde{x} \in \mathcal{D}_{\text{fit}}$ sharing n -grams with x is to mask out tokens in \tilde{x} at least every n positions, so that it cannot share n -gram overlap with x . That is, \mathcal{D}_{fit} contains different versions of x where at least every $(n - 1)$ -th token is masked out (replaced with a random token) so that there is at most $n - 1$ token overlap. This procedure is parameterized by a drop interval d , representing the length of the interval between dropped tokens. It must be that $d \leq n$ to ensure no n -gram overlap. To illustrate, if $x = [1, 2, 3, 4, 5, 6]$, then we may have $\mathcal{D}_{\text{fit}} = \{[1, 2, 3, \cdot, 5, 6], [1, \cdot, 3, 4, 5, \cdot], [\cdot, 2, 3, 4, \cdot, 6], \dots\}$, where ‘ \cdot ’

is a fresh random token and where here $d = 4$. We also consider a *randomized* dropout, where every token is dropped with probability $1/d$. By construction, the deterministic version guarantees that x is not a d -gram member of the training set, while the randomized version does so with (exponentially) high probability. Notably, due to BPE tokenization (Sennrich, 2015), the original text becomes visually obfuscated for humans even when most tokens are retained ($d > 2$). See Algorithm 2 for detailed algorithm and Appendix B.3 for visualization.

This construction is closely related to the *goldfish loss* proposed by Hans et al. (2024) to mitigate verbatim memorization, where the loss of every n -th token (on average, if randomized) is omitted during training. However, the procedure here is entirely *data-centric*: it does not interfere with the training objective and makes the learning task harder as subsequent tokens would still attend to the random tokens.

5.3. Case flipping

Another approach is to perform text-space transformations that preserve semantics yet drastically alter the token-space representations. One such method is to randomly flip the *casing* of English letters, creating varied *tokenization* of the otherwise equivalent string. To illustrate, if x decodes to ‘This is a string’, \mathcal{D}_{fit} may include token sequences of strings like ‘THIS Is A stRinG’. Due to the mechanisms of BPE tokenization, it is extremely easy to obtain \tilde{x} with *completely distinct* tokens than x under modern LLM tokenizers (see Appendix B.2 for visualization). The case flipping procedure is parameterized by the flip probability p . $p = 0.5$ creates the highest variance, and p closer to 1 always flips the case (most letters being initially lower case, natural text becomes mostly upper case).

5.4. Compositions

The transformations presented earlier are not mutually exclusive, and in principle they can be composed arbitrarily. We explore one such composition of token dropouts & case flipping, where non-dropped tokens (§5.2) have casing randomly flipped in the text space (§5.3). A key benefit of compositions is that they combinatorially give rise to many new transformations with potentially more granular control of task difficulty and detectability (e.g., whether manually inspecting \mathcal{D}_{fit} reveals x ; see Appendix B.3 for example visualizations). We leave a comprehensive evaluation of composed transformations to future work.

5.5. Results

Finding #1: It is possible for an LLM to complete an *unseen* string with no n -gram membership after minimal finetuning.

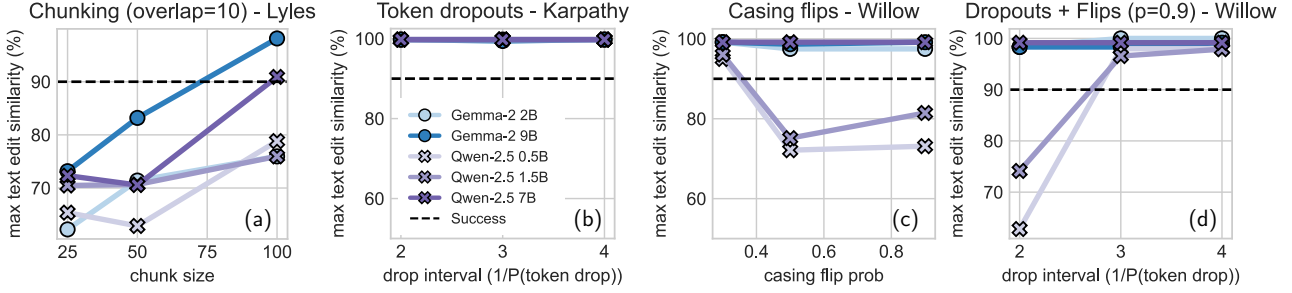


Figure 6: **Completion success across methods and target texts:** (a) chunking on **Lyles**, over chunk size c (x -axis); (b) token dropouts on **Karpathy**, over drop interval d ; (c) casing flips on **Willow**, over flip prob p ($p = 0.5$ is noisiest); and (d) combining dropouts + flips on **Willow**. We observe that: (1) it is possible to complete a chosen string with zero n -gram membership, and (2) this ability improves with model size. See Appendix B.5 for comprehensive results.




Figure 7: **Completion success may only require a few gradient steps.** See more configurations in Appendix B.5.

Fig. 6 shows the results of various transformations across various text targets (see Appendix B.5 for comprehensive results). We first see that there are many configurations where fine-tuning on \mathcal{D}_{fit} allow the successful completion of target x (edit similarity $> 90\%$).

Some transformations are more effective than others. We found that *chunking* (Section 5.1) was ineffective: with a small chunk size ($c = 25$), the models mostly fail to complete the target, and only some models succeed at $c = 100$ (Fig. 6 (a)). *Token dropouts* (Section 5.2), on the other hand, is extremely effective—even the smallest model (Qwen-2.5 0.5B) easily completes the target verbatim at a drop interval of 2 (50% drop probability at every token; Fig. 6 (b)). Our results thus present a counter-case to *goldfish loss* (Hans et al., 2024), as models can still complete targets verbatim when there are multiple versions of the target with *different* token dropout positions (e.g., due to near-duplicates, related to the findings of Section 4). *Casing flips* (Section 5.3) are also generally effective (Fig. 6 (c)). Composing token dropouts and casing flips (Section 5.4) increases task difficulty (fewer successes with small models) but otherwise similarly enables verbatim completion.

In essence, these experiments demonstrate that n -gram based membership definitions can be vulnerable to *adversarial* manipulation: the fine-tuning set \mathcal{D}_{fit} clearly contains information about the text x , but when given x and a choice of n , it is easy for an adversary to bypass detection and yet have the model generate x verbatim.

Finding #2: Completion success scales with model size.

Another message from Fig. 6 is that as we increase in the model size, the completion success generally improves under the *same* configurations. This provides evidence that frontier models should be more capable at synthesizing n -gram non-members into the target texts.

5.6. Interpretations and Outlook

We briefly describe the potential implications of our adversarially constructed fine-tuning datasets:

- **Data poisoning:** n -gram non-members of a poison text x can be added to the training set and still induce the generation of x . The concurrent work of Panaitescu-Liess et al. (2024) explores adding copyrighted materials as poisoned data with a similar chunking technique (§5.1).
- **Data contamination:** a dishonest model developer may game model evaluations through deliberate data contamination while evading n -gram based detection.
- **Reporting train-test overlap metrics:** More broadly, a model developer may self-report train-test overlap statistics (e.g., as part of contamination analysis seen in Dubey et al. (2024); Gemini Team et al. (2023); Brown (2020)). Our results highlight that it is desirable that developers report additional metrics beyond n -gram overlap.

6. Concluding Remarks

Lingering sequences (§ 4) and adversarially constructed fine-tuning datasets (§ 5) demonstrate the remarkable ability of LLMs to generalize from neighboring text. They are thus a valuable tool for evaluating LLM capabilities as models and pre-training datasets scale up. We conclude our work with a discussion of the implications of our findings:

Membership definitions and tests should incorporate new similarity measures. We showed n -gram based membership emits false negatives that may not capture human intuition nor the pragmatic concerns of the copyright, privacy, and AI safety community. On the flip side, tests like mem-

bership inference should consider broader notions of membership beyond individual sequences: a unit of data could be a collection of sequences grouped by similarity (Kandpal et al., 2023; Maini et al., 2024; Cooper et al., 2024).

Machine unlearning alone is insufficient to address data permissibility concerns in output suppression. It is widely accepted that a golden baseline of machine unlearning is to retrain a model from scratch without the target forget data (Bourtoule et al., 2021; Liu, 2024; Liu et al., 2024; Cooper et al., 2024). Yet, our experiments perform precisely this counterfactual and reveal that some excluded sequences can still be verbatim generated (Fig. 1, §4). We thus caution that unlearning alone may not always prevent a model from generating a sequence of interest (e.g., a harmful sequence). This is also known as output suppression and is a common goal of unlearning (Cooper et al., 2024).

Exploring the connection between our work and forging may help more precisely characterize threat models for when completion can serve as evidence of membership. Readers familiar with the forging (Thudi et al., 2022) literature will have noticed a connection with our work. Forging a step of gradient descent computed on a given minibatch is done by (adversarially) constructing a different minibatch that will result in the same gradient being computed. Perhaps surprisingly at first, gradients can be forged using *non-overlapping* datasets sampled from “natural” distributions. Rather than forging gradients, our methods can be viewed as attempting to forge model outputs. While our work begins to show that there may exist threat models where completion is insufficient evidence for n -gram membership, our experiments did not succeed in obtaining a forge in model outputs using “natural” data (recall § 4). We believe this is a valuable direction for future work.

Impact Statement

Our work studies training data membership in the context of language models. The impact of our work is described in Sections 4.3, 5.6, and 6. As we described earlier, the limitations we identified in n -gram based definitions of membership have implications for copyright, privacy, and AI safety. Broadly speaking, our work advocates for additional membership definitions to help better capture human intuition and make membership more pragmatic.

Acknowledgements

KZL would like to thank Xinran Zhao, Nikil Selvam, Steven Cao, Jing Huang, Yangjun Ruan, Aryaman Arora, Harshit Joshi, Rylan Schaeffer, Chenglei Si, Yanzhe Zhang, as well as members of the p-lambda lab and STAIR lab at Stanford University for helpful discussions and feedback. We thank Milad Nasr and Nicholas Carlini for helpful discussions, comments, and critiques throughout the work.

References

- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Anil, R., Dai, A. M., Firat, O., Johnson, M., Lepikhin, D., Passos, A., Shakeri, S., Taropa, E., Bailey, P., Chen, Z., et al. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*, 2023.
- Biderman, S., Prashanth, U., Sutawika, L., Schoelkopf, H., Anthony, Q., Purohit, S., and Raff, E. Emergent and predictable memorization in large language models. *Advances in Neural Information Processing Systems*, 36, 2024.
- Bourtoule, L., Chandrasekaran, V., Choquette-Choo, C. A., Jia, H., Travers, A., Zhang, B., Lie, D., and Papernot, N. Machine unlearning. In *2021 IEEE Symposium on Security and Privacy (SP)*, pp. 141–159. IEEE, 2021.
- Broder, A. Z. On the resemblance and containment of documents. In *Proceedings. Compression and Complexity of SEQUENCES 1997 (Cat. No. 97TB100171)*, pp. 21–29. IEEE, 1997.
- Brown, T. B. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- Carlini, N., Liu, C., Erlingsson, Ú., Kos, J., and Song, D. The secret sharer: Evaluating and testing unintended memorization in neural networks. In *28th USENIX security symposium (USENIX security 19)*, pp. 267–284, 2019.
- Carlini, N., Tramer, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, U., et al. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, pp. 2633–2650, 2021.
- Carlini, N., Chien, S., Nasr, M., Song, S., Terzis, A., and Tramer, F. Membership inference attacks from first principles. In *2022 IEEE Symposium on Security and Privacy (SP)*, pp. 1897–1914. IEEE, 2022a.
- Carlini, N., Ippolito, D., Jagielski, M., Lee, K., Tramer, F., and Zhang, C. Quantifying memorization across neural language models. *arXiv preprint arXiv:2202.07646*, 2022b.
- Carlini, N., Hayes, J., Nasr, M., Jagielski, M., Sehwag, V., Tramer, F., Balle, B., Ippolito, D., and Wallace, E. Extracting training data from diffusion models. In *32nd USENIX Security Symposium (USENIX Security 23)*, pp. 5253–5270, 2023.

- Choquette-Choo, C. A., Tramer, F., Carlini, N., and Papernot, N. Label-only membership inference attacks. In *International conference on machine learning*, pp. 1964–1974. PMLR, 2021.
- Cooper, A. F., Choquette-Choo, C. A., Bogen, M., Jagielski, M., Filippova, K., Liu, K. Z., Chouldechova, A., Hayes, J., Huang, Y., Mireshghallah, N., et al. Machine unlearning doesn’t do what you think: Lessons for generative ai policy, research, and practice. *arXiv preprint arXiv:2412.06966*, 2024.
- Das, D., Zhang, J., and Tramèr, F. Blind baselines beat membership inference attacks for foundation models. *arXiv preprint arXiv:2406.16201*, 2024.
- Dekoninck, J., Müller, M. N., Baader, M., Fischer, M., and Vechev, M. Evading data contamination detection for language models is (too) easy. *arXiv preprint arXiv:2402.02823*, 2024.
- Duan, M., Suri, A., Mireshghallah, N., Min, S., Shi, W., Zettlemoyer, L., Tsvetkov, Y., Choi, Y., Evans, D., and Hajishirzi, H. Do membership inference attacks work on large language models? *arXiv preprint arXiv:2402.07841*, 2024.
- Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Yang, A., Fan, A., et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Freeman, J., Rippe, C., DeBenedetti, E., and Andriushchenko, M. Exploring memorization and copyright violation in frontier llms: A study of the new york times v. openai 2023 lawsuit. *arXiv preprint arXiv:2412.06370*, 2024.
- Gemini Team, Anil, R., Borgeaud, S., Alayrac, J.-B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A. M., Hauth, A., Millican, K., et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- Gemini Team, Georgiev, P., Lei, V. I., Burnell, R., Bai, L., Gulati, A., Tanzer, G., Vincent, D., Pan, Z., Wang, S., et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024.
- Gemma Team, Mesnard, T., Hardin, C., Dadashi, R., Bhupatiraju, S., Pathak, S., Sifre, L., Rivière, M., Kale, M. S., Love, J., et al. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*, 2024a.
- Gemma Team, Riviere, M., Pathak, S., Sessa, P. G., Hardin, C., Bhupatiraju, S., Hussenot, L., Mesnard, T., Shahriari, B., Ramé, A., et al. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*, 2024b.
- Golchin, S. and Surdeanu, M. Data contamination quiz: A tool to detect and estimate contamination in large language models. *arXiv preprint arXiv:2311.06233*, 2023.
- Hans, A., Wen, Y., Jain, N., Kirchenbauer, J., Kazemi, H., Singhanian, P., Singh, S., Somepalli, G., Geiping, J., Bhatele, A., et al. Be like a goldfish, don’t memorize! mitigating memorization in generative llms. *arXiv preprint arXiv:2406.10209*, 2024.
- Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, E., Casas, D. d. L., Hendricks, L. A., Welbl, J., Clark, A., et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
- Huang, J., Yang, D., and Potts, C. Demystifying verbatim memorization in large language models. *arXiv preprint arXiv:2407.17817*, 2024.
- Ippolito, D., Tramèr, F., Nasr, M., Zhang, C., Jagielski, M., Lee, K., Choquette-Choo, C. A., and Carlini, N. Preventing verbatim memorization in language models gives a false sense of privacy. *arXiv preprint arXiv:2210.17546*, 2022.
- Jagielski, M., Nasr, M., Lee, K., Choquette-Choo, C. A., Carlini, N., and Tramer, F. Students parrot their teachers: Membership inference on model distillation. *Advances in Neural Information Processing Systems*, 36, 2024.
- Jiang, M., Liu, K. Z., Zhong, M., Schaeffer, R., Ouyang, S., Han, J., and Koyejo, S. Investigating data contamination for pre-training language models. *arXiv preprint arXiv:2401.06059*, 2024.
- Kandpal, N., Wallace, E., and Raffel, C. Deduplicating training data mitigates privacy risks in language models. In *International Conference on Machine Learning*, pp. 10697–10707. PMLR, 2022.
- Kandpal, N., Pillutla, K., Oprea, A., Kairouz, P., Choquette-Choo, C. A., and Xu, Z. User inference attacks on large language models. *arXiv preprint arXiv:2310.09266*, 2023.
- Karpathy, A. Llm.c: Llm training in simple, raw c/cuda. <https://github.com/karpathy/llm.c>, 2024. GitHub repository.
- Kong, Z., Chowdhury, A. R., and Chaudhuri, K. Can membership inferencing be refuted? *arXiv preprint arXiv:2303.03648*, 2023.

- Kudugunta, S., Caswell, I., Zhang, B., Garcia, X., Xin, D., Kusupati, A., Stella, R., Bapna, A., and Firat, O. Madlad-400: A multilingual and document-level large audited dataset. *Advances in Neural Information Processing Systems*, 36, 2024.
- Langley, P. Crafting papers on machine learning. In Langley, P. (ed.), *Proceedings of the 17th International Conference on Machine Learning (ICML 2000)*, pp. 1207–1216, Stanford, CA, 2000. Morgan Kaufmann.
- Lee, K., Ippolito, D., Nystrom, A., Zhang, C., Eck, D., Callison-Burch, C., and Carlini, N. Deduplicating training data makes language models better. *arXiv preprint arXiv:2107.06499*, 2021.
- Li, M., Wang, J., Wang, J., and Neel, S. Mope: Model perturbation-based privacy attacks on language models. *arXiv preprint arXiv:2310.14369*, 2023.
- Liu, K. Z. Machine unlearning in 2024, May 2024. URL <https://ai.stanford.edu/~kzliu/blog/unlearning>.
- Liu, S., Yao, Y., Jia, J., Casper, S., Baracaldo, N., Hase, P., Yao, Y., Liu, C. Y., Xu, X., Li, H., et al. Rethinking machine unlearning for large language models. *arXiv preprint arXiv:2402.08787*, 2024.
- Maini, P., Yaghini, M., and Papernot, N. Dataset inference: Ownership resolution in machine learning. *arXiv preprint arXiv:2104.10706*, 2021.
- Maini, P., Jia, H., Papernot, N., and Dziedzic, A. Llm dataset inference: Did you train on my dataset? *arXiv preprint arXiv:2406.06443*, 2024.
- Mattern, J., Mireshghallah, F., Jin, Z., Schölkopf, B., Sachan, M., and Berg-Kirkpatrick, T. Membership inference attacks against language models via neighbourhood comparison. *arXiv preprint arXiv:2305.18462*, 2023.
- Meeus, M., Jain, S., Rei, M., and de Montjoye, Y.-A. Inherent challenges of post-hoc membership inference for large language models. *arXiv preprint arXiv:2406.17975*, 2024.
- Mou, C. Large-scale near-deduplication behind bigcode, 2023. URL <https://huggingface.co/blog/dedup>. Published on Hugging Face Blog, May 16, 2023.
- Nasr, M., Carlini, N., Hayase, J., Jagielski, M., Cooper, A. F., Ippolito, D., Choquette-Choo, C. A., Wallace, E., Tramèr, F., and Lee, K. Scalable extraction of training data from (production) language models. *arXiv preprint arXiv:2311.17035*, 2023.
- Oren, Y., Meister, N., Chatterji, N., Ladhak, F., and Hashimoto, T. B. Proving test set contamination in black box language models. *arXiv preprint arXiv:2310.17623*, 2023.
- Panaitescu-Liess, M.-A., Pathmanathan, P., Kaya, Y., Che, Z., An, B., Zhu, S., Agrawal, A., and Huang, F. Poisoned-parrot: Subtle data poisoning attacks to elicit copyright-infringing content from large language models. In *Neurips Safe Generative AI Workshop 2024*, 2024. URL <https://openreview.net/forum?id=ZXgvPANlwe>.
- Penedo, G., Kydlíček, H., Lozhkov, A., Mitchell, M., Raffel, C., Von Werra, L., Wolf, T., et al. The fineweb datasets: Decanting the web for the finest text data at scale. *arXiv preprint arXiv:2406.17557*, 2024.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Sablayrolles, A., Douze, M., Schmid, C., Ollivier, Y., and Jégou, H. White-box vs black-box: Bayes optimal strategies for membership inference. In *International Conference on Machine Learning*, pp. 5558–5567. PMLR, 2019.
- Sainz, O., Campos, J. A., García-Ferrero, I., Etxaniz, J., de Lacalle, O. L., and Agirre, E. Nlp evaluation in trouble: On the need to measure llm data contamination for each benchmark. *arXiv preprint arXiv:2310.18018*, 2023.
- Salem, A., Zhang, Y., Humbert, M., Berrang, P., Fritz, M., and Backes, M. MI-leaks: Model and data independent membership inference attacks and defenses on machine learning models. *arXiv preprint arXiv:1806.01246*, 2018.
- Sennrich, R. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*, 2015.
- Shi, W., Ajith, A., Xia, M., Huang, Y., Liu, D., Blevins, T., Chen, D., and Zettlemoyer, L. Detecting pretraining data from large language models. *arXiv preprint arXiv:2310.16789*, 2023.
- Shokri, R., Stronati, M., Song, C., and Shmatikov, V. Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, pp. 3–18. IEEE, 2017.
- Singh, A. K., Kocyigit, M. Y., Poulton, A., Esiobu, D., Lomeli, M., Szilvasy, G., and Hupkes, D. Evaluation data contamination in llms: how do we measure it and (when) does it matter? *arXiv preprint arXiv:2411.03923*, 2024.

- Somepalli, G., Singla, V., Goldblum, M., Geiping, J., and Goldstein, T. Diffusion art or digital forgery? investigating data replication in diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6048–6058, 2023.
- Team, Q. Qwen2.5: A party of foundation models, September 2024. URL <https://qwenlm.github.io/blog/qwen2.5/>.
- Thudi, A., Jia, H., Shumailov, I., and Papernot, N. On the necessity of auditable algorithmic definitions for machine unlearning. In *31st USENIX Security Symposium (USENIX Security 22)*, pp. 4007–4022, 2022.
- Tirumala, K., Markosyan, A., Zettlemoyer, L., and Aghajanyan, A. Memorization without overfitting: Analyzing the training dynamics of large language models. *Advances in Neural Information Processing Systems*, 35: 38274–38290, 2022.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Yang, A., Yang, B., Hui, B., Zheng, B., Yu, B., Zhou, C., Li, C., Li, C., Liu, D., Huang, F., et al. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.
- Yeom, S., Giacomelli, I., Fredrikson, M., and Jha, S. Privacy risk in machine learning: Analyzing the connection to overfitting. In *2018 IEEE 31st computer security foundations symposium (CSF)*, pp. 268–282. IEEE, 2018.
- Zarifzadeh, S., Liu, P., and Shokri, R. Low-cost high-power membership inference attacks, 2023.
- Zhang, A. K., Klyman, K., Mai, Y., Levine, Y., Zhang, Y., Bommasani, R., and Liang, P. Language model developers should report train-test overlap. *arXiv preprint arXiv:2410.08385*, 2024a.
- Zhang, J., Das, D., Kamath, G., and Tramèr, F. Membership inference attacks cannot prove that a model was trained on your data. *arXiv preprint arXiv:2409.19798*, 2024b.

A. Removing Members Does Not Always Prevent LLM Verbatim Completion (§4)

A.1. Data filtering

The n -gram filter discussed in Section 4.1 can be defined as follows:

Definition A.1 (n -gram data filtering). Let \mathcal{F} be a set of n -grams to filter against. Let $f_n(x, \mathcal{F})$ be the condition that returns 1 if any n -gram of the sequence x is in \mathcal{F} (0 otherwise). An n -gram filter against \mathcal{F} on a set of sequences \mathcal{D} is thus constructing $\bar{\mathcal{D}}_n = \{x \in \mathcal{D} \mid f_n(x, \mathcal{F}) = 0\}$.

Intuitively, a stronger n -gram filter (smaller n) means that we are removing a sequence on increasingly smaller partial matches against the filter set. In practice, since we are filtering many sequences (\mathcal{D}_{mem}) at once, it is more efficient to build a shared n -gram filter hash set from all sequences in \mathcal{D}_{mem} and then apply a sliding window filtering procedure.

Table 2 shows the amount of tokens kept after applying n -gram filtering at different filtering strengths (for sequence length 50, the default setting used in experiments unless otherwise stated). Table 3 similarly shows the amount of tokens for sequence length 100.

Table 2: Fraction of tokens kept after applying n -gram filter to sequence length $k = 50$.

n -gram filtering strength	$n = 5$	$n = 10$	$n = 20$	$n = 50$ (Exact Filtering)
Approx % of tokens kept	0.6905	0.9838	0.9938	0.9995

Table 3: Fraction of tokens kept after applying n -gram filter to sequence length $k = 100$.

n -gram filtering strength	$n = 10$	$n = 20$	$n = 50$	$n = 100$ (Exact Filtering)
Approx % of tokens kept	0.9820	0.9930	0.9970	0.9995

A.2. Visualizing Lingering Sequences $\mathcal{D}_{\text{linger}}^{(n)}$

In this and the following section, we provide visualizations to the key results described in Section 4 to help develop intuition on our findings.

Recall from §4 that $\mathcal{D}_{\text{linger}}^{(n)}$ refers to the set of lingering sequences that are still verbatim completable by the re-trained from scratch LLM after applying n -gram filter (Def. A.1) on the identified memorized sequences \mathcal{D}_{mem} .

Table 4, Table 5, Table 6, and Table 7 show 15 randomly sampled lingering sequence at filtering strengths $n = \{50, 20, 10, 5\}$, respectively. Observe that, as noted in Section 4.2, with stronger filtering strength (smaller n), the content gradually shift from semantically useful content to pattern continuations.

Table 4: Randomly sampled lingering sequences at filtering strength $n = 50$ (exact) filter. Sequence length $k = 50$.

Idx	Lingering sequences at filtering strength $n = 50$ (exact) filter.
0	Prompt: 'Common Name: bleeding heart Type: Herbaceous perennial Native Range: Eastern United States Zone: 3 to 9 ' Completion: 'Height: 1.00 to 1.50 feet Spread: 1.00 to 1.50 feet Bloom Time:'
1	Prompt: '—1477 by topic— —Arts and science— —Birth and death categories— —Births —' Completion: 'Deaths— —Establishments and disestablishments categories— —Establishments — Disestablishments— —'
2	Prompt: 'Charcot Joint (Neuropathic Arthropathy) Medicine Central™ is a quick-consult mobile and' Completion: ' web resource that includes diagnosis, treatment, medications, and follow-up information on over 700 diseases and disorders, providing fast answers'
3	Prompt: 'Mienert-barth Surname History The family history of the Mienert-barth last name is' Completion: ' maintained by the AncientFaces community. Join the community by adding to our knowledge of the Mienert-'
4	Prompt: 'Instructional Supports and Resources Dyslexia is a specific learning disability that is neurological in origin. It is characterized' Completion: ' by difficulties with accurate and/or fluent word recognition and by poor spelling and decoding abilities. These difficulties typically result from a deficit'
5	Prompt: 'Publisher description for Writers at work. The short composition / Ann O. Strauch. Bibliographic record and links to' Completion: ' related information available from the Library of Congress catalog Information from electronic data provided by the publisher. May be incomplete or contain other'
6	Prompt: 'Create healthcare diagrams like this example called Anencephaly in minutes with SmartDraw. SmartDraw includes 1000s of professional healthcare' Completion: ' and anatomy chart templates that you can modify and make your own. Text in this Example: Anencephaly is'

Language Models May Verbatim Complete Text They Were Not Explicitly Trained On

Idx	Lingering sequences at filtering strength $n = 50$ (exact) filter.	
7	Prompt:	'Presentation on theme: "Spiraled Assignments Presenter: Angela Pritchett November 14, 2006."'
	Completion:	' — Presentation transcript: Spiraled Assignments Presenter: Angela Pritchett November 14, 2006 '
8	Prompt:	'An excerpt from www.HouseOfNames.com archives copyright © 2000 - 2013 Where did the Irish McSweeney family'
	Completion:	' come from? What is the Irish McSweeney family crest and coat of arms? When did the McSweeney family first'
9	Prompt:	'Which of the following cubes can be made from these nets? Is it possible to remove ten unit cubes from a 3 by'
	Completion:	' 3 by 3 cube made from 27 unit cubes so that the surface area of the remaining solid is the same as the surface area'
10	Prompt:	'Course Hero. "The Libation Bearers Study Guide." Course Hero. 23 June 2017. Web. 14 Nov. 2018'
	Completion:	' . jhttps://www.coursehero.com/lit/The-Libation-Bearers/. Course Hero'
11	Prompt:	'Presentation on theme: "Cause and Effect Comprehension Skill Fourth Grade Unit 2 Week 1 Created by Kristi Waltke'
	Completion:	'.' — Presentation transcript: Cause and Effect Comprehension Skill Fourth Grade Unit 2 Week 1 Created by Kristi Walt'
12	Prompt:	'An excerpt from www.HouseOfNames.com archives copyright © 2000 - 2015 Where did the English Ragsdale family'
	Completion:	' come from? What is the English Ragsdale family crest and coat of arms? When did the Ragsdale family first'
13	Prompt:	'Gibbous Scorpio Moon phase on 29 March 2051 Wednesday is Waning Gibbous, 16'
	Completion:	' days old Moon is in Scorpio.Share this page: twitter facebook linkedin Previous main lunar phase is the Full Moon'
14	Prompt:	'What does AIDS mean in Laboratory? This page is about the meanings of the acronym/abbreviation/shorthand'
	Completion:	' AIDS in the Medical field in general and in the Laboratory terminology in particular. Find a translation for AIDS in other languages:'
15	Prompt:	'Nathalie Raphaëlle June 23, 2021 Worksheets If you home school your children, you'
	Completion:	' will quickly realize how important printable homeschool worksheets can be. If you are trying to develop a curriculum for your'
16	Prompt:	'Course Hero. "The Pearl Study Guide." Course Hero. 14 Dec. 2017. Web. 24 Nov. 2020. j'
	Completion:	'https://www.coursehero.com/lit/The-Pearl/. Course Hero. (2017, December'
17	Prompt:	'Internet of Things Internet of Things The Internet of Things (IoT) is a system of interrelated computing devices'
	Completion:	' , mechanical and digital machines, objects, animals or people that are provided with unique identifiers and the ability to transfer data over a'
18	Prompt:	'Latest Newland photos These photos were uploaded by members of the Newland community on AncientFaces. Newland S'
	Completion:	'urname History The family history of the Newland last name is maintained by the AncientFaces community. Join the community'
19	Prompt:	'Definition of Seckles 1. seckle [n] - See also: seckle Click the following'
	Completion:	' link to bring up a new window with an automated collection of images related to the term: Seckles Images Lexic'

Table 5: Randomly sampled lingering sequences at filtering strength $n = 20$ filter. Sequence length $k = 50$.

Idx	Lingering sequences at filtering strength $n = 20$ filter.	
0	Prompt:	'Presentation on theme: "MAKING BOOKS WITH CHILDREN Picture It! Publish It! Read It!'"
	Completion:	' — Presentation transcript: MAKING BOOKS WITH CHILDREN Picture It! Publish It! Read It!'
1	Prompt:	'This Constitution, and the Laws of the United States which shall be made in Pursuance thereof; and all Treaties made'
	Completion:	' , or which shall be made, under the Authority of the United States, shall be the supreme Law of the Land; and'
2	Prompt:	'How To Recognize A Crystal Child A selection of articles related to how to recognize a crystal child. Original articles from'
	Completion:	' our library related to the How To Recognize A Crystal Child. See Table of Contents for further available material (downloadable resources'
3	Prompt:	'Wampsville, New York —Wampsville, New York— —● Total——1.0 sq mi'
	Completion:	' (2.6 km2)— —● Land——1.0 sq mi (2.6 km2)— —'
4	Prompt:	'Report on Stromboli (Italy) — 12 March-18 March 2003 Smithsonian / US Geological Survey Weekly Vol'
	Completion:	'canic Activity Report, 12 March-18 March 2003 Managing Editor: Gari Mayberry Please cite this report'
5	Prompt:	'Presentation on theme: "'The Great (gym) Divide Curricula by Design #3 M. Fischer."—'
	Completion:	' Presentation transcript: The Great (gym) Divide Curricula by Design #3 M. Fischer The Great'
6	Prompt:	'Course Hero. "Lord of the Flies Study Guide." Course Hero. 15 Sep. 2016. Web. 29 May 20'
	Completion:	' 23. jhttps://www.coursehero.com/lit/Lord-of-the-Flies/. '
7	Prompt:	'Manada Gap, Pennsylvania facts for kids Quick facts for kids Manada Gap, Pennsylvania —Time zone——UTC'
	Completion:	' -5 (Eastern (EST))— —● Summer (DST)—UTC-4 (EDT)— —'
8	Prompt:	'Scale Zoology Cosmoid Scales A selection of articles related to scale zoology cosmoid scales. Original'
	Completion:	' articles from our library related to the Scale Zoology Cosmoid Scales. See Table of Contents for further available material ('
9	Prompt:	'Atomic Nucleus History A selection of articles related to atomic nucleus history. Original articles from our library related to'
	Completion:	' the Atomic Nucleus History. See Table of Contents for further available material (downloadable resources) on Atomic Nucleus'
10	Prompt:	'Mangoverde :: World Bird Guide :: Pheasants and Partridges :: Common Quail Common Quail Cot'
	Completion:	'urnix coturnix Described by: Linnaeus (1758) Alternate common name(s)'
11	Prompt:	'Set Builder Notation Variations A selection of articles related to set builder notation variations. Original articles from our library related'
	Completion:	' to the Set Builder Notation Variations. See Table of Contents for further available material (downloadable resources) on Set Builder'
12	Prompt:	'—1648 by topic— —Arts and science— —Birth and death categories— —Births —'
	Completion:	' Deaths— —Establishments and disestablishments categories— —Establishments – Disestablishments— —'
13	Prompt:	'Tamil Script The Tamil Letters A selection of articles related to tamil script the tamil letters. Original articles from'
	Completion:	' our library related to the Tamil Script The Tamil Letters. See Table of Contents for further available material (downloadable resources) on'
14	Prompt:	'Manuel I of PortugalFrom Wikipedia, the free encyclopediaJump to: navigation, search This article does not cite any references or'
	Completion:	' sources. Please help improve this article by adding citations to reliable sources. Unsourced material may be challenged and removed. ('
15	Prompt:	'Image 1 of 12 Image 2 of 12 Image 3 of 12 Image 4 of 12 Image 5 of 12 '
	Completion:	' Image 6 of 12 Image 7 of 12 Image 8 of 12 Image 9 of 12 Image 10 of 12 '
16	Prompt:	'—Nutritional Guidelines (per serving)— —Servings: 3 pint jars (96 servings)— —Amount per serving'
	Completion:	' — —% Daily Value*— —Total Fat 0g —0% —Saturated Fat 0g'
17	Prompt:	'Presentation on theme: "Corpus Linguistics and Stylistics PALA Summer School, Maribor, 2014"
	Completion:	'.' — Presentation transcript: Corpus Linguistics and Stylistics PALA Summer School, Maribor.'
18	Prompt:	'Presentation on theme: "Lunar Research Station Design Submitted by West Valley Elementary GATE Team October 31, 2006"
	Completion:	'.' — Presentation transcript: Lunar Research Station Design Submitted by West Valley Elementary GATE Team October 31.'
19	Prompt:	'Some daily events in the changing sky for February 19 27. Friday, February 19 Saturday, February 20 Sunday,'
	Completion:	' February 21 Monday, February 22 Tuesday, February 23 Wednesday, February 24 Thursday, February 25 Friday.'

Table 6: **Randomly sampled lingering sequences at filtering strength $n = 10$ filter.** Sequence length $k = 50$.

Idx	Lingering sequences at filtering strength $n = 10$ filter.
0	Prompt: 'Presentation on theme: "HELPING YOUR CHILD WITH NUMERACY: ADDITION AND SUBTRACTION." Completion: '— Presentation transcript: HELPING YOUR CHILD WITH NUMERACY: ADDITION AND SUBTRACTION '
1	Prompt: 'Wednesday—2:00 PM - 3:40 PM—lesson—Lecture Hall 1.2—' Completion: '—Thursday—2:00 PM - 3:40 PM—lesson—Lecture Hall 1.2—'
2	Prompt: 'How to define the cosine ratio and identify the cosine of an angle in a right triangle. How to define the Completion: ' sine ratio and identify the sine of an angle in a right triangle. How to define the tangent ratio and'
3	Prompt: 'Q1. A series is given with one term missing. Select the correct alternative from the given ones that will complete the series' Completion: ' Q2. A series is given with one term missing. Select the correct alternative from the given ones that will complete'
4	Prompt: 'History of False Teeth Length: 497 words (1.4 double-spaced pages) - - -' Completion: '-----'
5	Prompt: 'Presentation on theme: "Yoghurt!!! Find the dairy cow on each page!!! By Daisy Mason and Brigitte Roberts" Completion: '"— Presentation transcript: Yoghurt!!! Find the dairy cow on each page!!! By Daisy Mason and Brigitte'
6	Prompt: 'Protecting People with Disabilities in the Ebbs and Flows of the COVID-19 Pandemic Protecting People' Completion: ' with Disabilities in the Ebbs and Flows of the COVID-19 Pandemic The COVID-19 pand'
7	Prompt: 'Presentation on theme: "Aceh Poverty Assessment The impact of the Conflict, the Tsunami and Reconstruction on Poverty" Completion: ' in Aceh."— Presentation transcript: Aceh Poverty Assessment The impact of the Conflict, the Tsunami'
8	Prompt: 'Presentation on theme: "THE MIX-AERATOR Innovation In Pond & Lagoon Aeration & Mixing." Completion: '"— Presentation transcript: THE MIX-AERATOR Innovation In Pond & Lagoon Aeration & Mixing '
9	Prompt: 'Some daily events in the changing sky for February 8 16. Friday, February 8 Saturday, February 9 Sunday, Completion: ' February 10 Monday, February 11 Tuesday, February 12 Wednesday, February 13 Thursday, February 14 Friday,'
10	Prompt: 'Essays on mercutio Romeo and mercutio essays: over 180,000 roomeo and merc' Completion: 'utio essays, roomeo and mercutio term papers, roomeo and mercutio research paper, book'
11	Prompt: 'Presentation on theme: "Fabric Construction Fashion Design, Textiles & Merchandising Mrs. Moscinski."— Completion: ' Presentation transcript: Fabric Construction Fashion Design, Textiles & Merchandising Mrs. Moscinski Fabric'
12	Prompt: 'Presentation on theme: "Chapter 4 - Building Compassionate School-Community Partnerships That Work Chapter 4 - Building Comp' Completion: 'assionate School-Community Partnerships That Work"— Presentation transcript: Chapter 4 - Building Compassionate School-'
13	Prompt: 'Tracing Names: Letter AA — B — C — D — E — F — G — H — I — J — ' Completion: ' K — L — M — N — O — P — Q — R — S — T — U — V — W'
14	Prompt: 'Dictionary of Financial, Economic, and Business Terms A — B — C — D — E — F — G — ' Completion: ' H — I — J — K — L — M — N — O — P — Q — R — S — T'
15	Prompt: 'Canons of the Seven Ecumenical Councils. The First Ecumenical Council. Second Ecumenical' Completion: ' Council. Third Ecumenical Council. Fourth Ecumenical Council. Fifth Ecumenical Council. '
16	Prompt: 'Welsh Levels of Care E-Learning Program Glossary Special — A — B — C — D — E — F' Completion: ' — G — H — I — J — K — L — M — N — O — P — Q — R —'
17	Prompt: 'Presentation on theme: "Tap Water Intrusion Effects on Microbial Life Anthony DeRenzo Grade 10 Pittsburgh Central Catholic" Completion: ' High School."— Presentation transcript: Tap Water Intrusion Effects on Microbial Life Anthony DeRenzo Grade 10'
18	Prompt: 'Presentation on theme: "Dr. Anand Srinivasan for MBBS 2013 on 10/10/2013." Completion: '"— Presentation transcript: Dr. Anand Srinivasan for MBBS 2013 on 10/10/2013 '
19	Prompt: 'Presentation on theme: "Ashok Sinha O/o the Director General (Audit) Central, Chandigar' Completion: 'h."— Presentation transcript: Ashok Sinha O/o the Director General (Audit) Central, Chand'

Table 7: **Randomly sampled lingering sequences at filtering strength $n = 5$ filter.** Sequence length $k = 50$.

[illegible]

Idx	Lingering sequences at filtering strength $n = 5$ filter.	
14	Prompt:	'HISTORY CRIME AND PUNISHMENT HISTORY CRIME AND PUNISHMENT HISTORY CRIME AND'
	Completion:	' PUNISHMENT HISTORY CRIME AND PUNISHMENT HISTORY CRIME AND PUNISHMENT '
15	Prompt:	'Define Gyromitra infula. Gyromitra infula synonyms, Gyromitra in'
	Completion:	'fula pronunciation, Gyromitra infula translation, English dictionary definition of Gyromitra infula.'
16	Prompt:	'Some daily events in the changing sky for December 19 27. Friday, December 19 Saturday, December 20 Sunday,'
	Completion:	' December 21 Monday, December 22 Tuesday, December 23 Wednesday, December 24 Thursday, December 25 Friday.'
17	Prompt:	'Length: 1122 words (3.2 double-spaced pages) -----'
	Completion:	' -----'
18	Prompt:	'Letter E Names: Page 2A — B — C — D — E — F — G — H — I — J'
	Completion:	' — K — L — M — N — O — P — Q — R — S — T — U — V —'
19	Prompt:	'A Complete Illustrated History of Robots in the Movies (chronological by film title) Intro — Part 1 —'
	Completion:	' Part 2 — Part 3 — Part 4 — Part 5 — Part 6 — Part 7 — Part 8 — Part 9 — Part'
20	Prompt:	'1. George Washington - Term of Office (1789-1797) 2. John Adams (1797-18'
	Completion:	'01) 3. Thomas Jefferson (1801-1809) 4. James Madison (1809-1817'

A.3. Visualizing neighbors of lingering sequences $\mathcal{D}_{\text{linger}}^{(n)}$ in the pre-training set

A key aspect worth studying for lingering sequences is what contributed to their existence. We randomly sample 2 lingering sequences for filtering strengths $n = \{50, 20, 10\}$, and perform a very costly search of Levenshtein edit-distance neighbor search: perform a sliding window over the pre-training tokens, and check the edit distance of each window to the lingering sequence. By construction, lingering sequences have already been removed from the training set; the hope is thus to identify and visualize *neighboring* sequences that may have led to these lingering sequences.

Fig. 8 visualizes these randomly sampled lingering sequences, and two randomly sampled neighbors (with edit distance < 20 tokens) for each of them. We also visualize the histogram of these neighbors at different distances. We note that:

- These visualizations suggest that lingering sequences are very likely the result of either *near-duplicate* training data and/or the generalization capabilities of LLMs.
- **By observing the neighboring sequences, we gain insights into how we may *adversarially* game the n -gram membership definition.** For example, observing the second sequence in Fig. 8 can provide intuition for the *chunking* (§5.1) method we presented in §5 to adversarially construct fine-tuning sequences that avoid n -gram overlap, and observing the second last can provide intuition for the *token dropout* (§5.2) method.



Figure 8: Visualizing the neighbors (Levenshtein edit distance < 20 tokens) of a few randomly selected lingering sequence from $\mathcal{D}_{\text{linger}}^{(10)}$, $\mathcal{D}_{\text{linger}}^{(20)}$, and $\mathcal{D}_{\text{linger}}^{(50)}$.

A.4. Model architectures

Our 350M, 774M, and 1.6B-parameter model architectures follow directly from the original GPT-2 paper (Radford et al., 2019) and LLM.c implementation (Karpathy, 2024). We create a custom architecture for 2.8B parameters by adjusting the number of layers, channels, and attention heads in the model, again following Karpathy (2024). The model configurations can be found in Table 8.

Table 8: Configurations of models (GPT-2 architecture) at different sizes for our pre-training experiments.

Parameters	304M	774M	1.6B	2.8B
Layers	24	36	48	60
Channels	1024	1280	1600	1920
Attention heads	16	20	26	30

A.5. Training hyperparameters

We follow the pre-training configurations outlined in (Karpathy, 2024), particularly <https://github.com/karpathy/llm.c/discussions/677>. Table 9 summarizes the configurations.

Table 9: Training configurations for pre-training experiments.

Hyperparameter	Value
# Training Tokens	33.6 billion
Compute	8 NVIDIA H100 days (1.6B parameter model)
Micro-Batch Size	16
Max Sequence Length	1024
Total Batch Size	$2^{20} = 1,048,576$ tokens
Gradient Accumulation Steps	8
Weight Decay	0.1
Learning Rate	$6e-4$
LR Schedule	Cosine
LR Decay	decay to 10% of max LR
Warmup Iterations	700 iterations
Total Training Steps	32,000

A.6. Tabled results of Fig. 3

We additionally provide the table version of Fig. 3 in Table 10 (for default sequence length 50).

Table 10: (Table version of Fig. 3) **Amount of lingering sequences** as fraction of $|\mathcal{D}_{\text{mem}}|$ across model sizes and filtering strengths (values of n -gram filter). Sequence length $k = 50$. Recall setup in Section 4.

Model Size	$ \mathcal{D}_{\text{mem}} $	Filtering Strength			
		$n = 5$	$n = 10$	$n = 20$	$n = 50$
304M	76648	0.0175	0.0402	0.0978	0.4793
774M	116270	0.0132	0.0300	0.0916	0.4013
1.6B	151598	0.0109	0.0259	0.0796	0.4007
2.8B	175813	0.0099	0.0247	0.0869	0.4068

A.7. Additional Results: Longer Sequence Length

We also experiment with a longer sequence length of $k = 100$ (i.e., 50-token prompt, 50-token completion). In general, we made similar observations in terms of both the amount and nature of the lingering sequences. Table 11 shows the fraction of lingering sequences in terms of the identified memorized sequences \mathcal{D}_{mem} . We provide visualizations of these sequences in Table 12 (exact filtering, $n = 100$) and Table 12 (strong filtering, $n = 10$), and note that the content of the lingering sequences are similar to those with sequence length $k = 50$.

Table 11: Amount of lingering sequences as a fraction of the identified memorized sequences $|\mathcal{D}_{\text{mem}}|$ across filtering strengths (values of n -gram filter). Sequence length $k = 100$. Recall setup in Section 4; compare to Table 10.

Model Size	$ \mathcal{D}_{\text{mem}} $	Filtering Strength			
		$n = 10$	$n = 20$	$n = 50$	$n = 100$
1.6B	88803	0.0109	0.0259	0.0796	0.4007

Table 12: Randomly sampled lingering sequences at filtering strength $n = 100$ (exact) filter. Sequence length $k = 100$.

Idx	Lingering sequences at filtering strength $n = 100$ (exact) filter.
0	Prompt: "Course Hero. "Romeo and Juliet Study Guide." Course Hero. 28 July 2016. Web. 16 Jan. 2019. https://www.coursehero.com/lit/Romeo-and-Juliet/ . Course" Completion: " Hero. (2016, July 28). Romeo and Juliet Study Guide. In Course Hero. Retrieved January 16, 2019, from https://www.coursehero.com/lit/Romeo-and-Juliet/ (Course Hero),"
1	Prompt: "ESQUIEN INDIANS. The Esquien Indians were a Karankawan group that entered Nuestra Señora de la Candelaria Mission on the San Gabriel River in 1750. Frederick Webb Hodge, ed." Completion: " Handbook of American Indians North of Mexico (2 vols., Washington: GPO, 1907, 1910; rpt., New York: Pageant, 1959). The following, adapted from the Chicago Manual of Style, 15th edition, is the"
2	Prompt: "See what questions a doctor would ask. During a consultation, your doctor will use various techniques to assess the symptom: Heel pain. These will include a physical examination and possibly diagnostic tests. (Note: A physical exam is always done." Completion: " diagnostic tests may or may not be performed depending on the suspected condition) Your doctor will ask several questions when assessing your condition. It is important to openly share any pertinent information to help your doctor make an accurate diagnosis. It is also very important to"
3	Prompt: "—Product #: EMC0775025.TQ— A Is for Apple (Resource Book Only) eBookGrade 2—Grade 3—Grade 4—Grade 5 Please Note: This ebook is a digital download, NOT a physical product" Completion: ". After purchase, you will be provided a one time link to download ebooks to your computer. Orders paid by PayPal require up to 8 business hours to verify payment and release electronic media. For immediate downloads, payment with credit card is required. "
4	Prompt: "Family History and Genealogy Resources by Surname Haycock Surname Origin A name probably given to a founding exposed in a hayfield. Source: An Etymological Dictionary of Family and Christian Names With an Essay" Completion: " on their Derivation and Import; Arthur, William, M.A.; New York, NY: Sheldon, Blake, Bleeker & CO., 1857. Haycock Surname Meaning and Family Facts There is more to Haycock family"
5	Prompt: "'Lake Uvs and its surrounding wetlands - Site number: 1379 - Area:585,000 ha - Designation date:22-03-2004 - Coordinates:50° 19'N 92° 45'E Materials" Completion: " presented on this website, particularly maps and territorial information, are as-is and as-available based on available data and do not imply the expression of any opinion whatsoever on the part of the Secretariat of the Ramsar Convention concerning the legal status of"
6	Prompt: "Barrier Methods of Birth Control (cont.) Melissa Conrad Stöppler, MD Melissa Conrad Stöppler, MD, is a U.S. board-certified Anatomic Pathologist with subspecialty training in" Completion: " the fields of Experimental and Molecular Pathology. Dr. Stöppler's educational background includes a BA with Highest Distinction from the University of Virginia and an MD from the University of North Carolina. She completed residency training in Anatomic Pathology at"
7	Prompt: "Definition of Japanese deer 1. Noun. Small deer of Japan with slightly forked antlers. Generic synonyms: Cervid, Deer Group relationships: Cervus, Genus Cervus Japanese Deer Pictures Click" Completion: " the following link to bring up a new window with an automated collection of images related to the term: Japanese Deer Images Lexicographical Neighbors of Japanese Deer Literary usage of Japanese deer Below you will find example usage of this term as"
8	Prompt: "Comparing Fractions (G) In this comparing fractions practice worksheet, 5th graders examine 10 pairs of fractions. Students identify each of the pairs of fractions as greater than, less than, or equal to another. 3 Views" Completion: " 0 Downloads Fraction Equivalence, Ordering, and Operations Need a unit to teach fractions to fourth graders? Look no further than this well-developed and thorough set of lessons that takes teachers through all steps of planning, implementing."
9	Prompt: "PREAMBLEWhereas recognition of the inherent dignity and of the equal and inalienable rights of all members of the human family is the foundation of freedom, justice and peace in the world, Whereas disregard and contempt for human rights have resulted in barbar" Completion: "ous acts which have outraged the conscience of mankind, and the advent of a world in which human beings shall enjoy freedom of speech and belief and freedom from fear and want has been proclaimed as the highest aspiration of the common people, Whereas it is essential"

Table 13: Randomly sampled lingering sequences at filtering strength $n = 10$ filter. Sequence length $k = 100$.

Idx	Lingering sequences at filtering strength $n = 10$ filter.
0	Prompt: "Presentation on theme: "REVIEW We can tell how many electrons and atom will gain or lose by looking at its valence. Metals like to lose electrons. (Cations) -Ex. Na + Nonmetals."— Presentation" Completion: " transcript: REVIEW We can tell how many electrons and atom will gain or lose by looking at its valence. Metals like to lose electrons. (Cations) -Ex. Na + Nonmetals like to gain electrons. (An"
1	Prompt: "—I • II • III • IV • V • VI • VII • VIII • IX • X • XI • XII • XIII • XIV • Schedule— 1 Features - 2 Preamble - 3 Article I - 4 Article" Completion: " II - 5 Article III - 6 Article IV - 7 Article V - 8 Article VI - 9 Article VII - 10 Article VIII - 11 Article IX - 12 Article X - 13 Article XI - 14 Article"

Language Models May Verbatim Complete Text They Were Not Explicitly Trained On

Idx	Lingering sequences at filtering strength $n = 10$ filter.
2	<p>Prompt: "ENGL 301 Course Introduction this course will be an independent study of a topic of the student's choice in English literature and film. ENGL 302 Course Introduction for ENGL 302 – Independent Study In English (ENGL 302) this"</p> <p>Completion: " course will be an independent study of a topic of the student's choice in English literature and film. ENGL 303 Course Introduction for ENGL 303 – Independent Study In English (ENGL 303) this course will be an independent study of"</p>
3	<p>Prompt: "'NRL's MISSE-8 Launched Aboard STS-134 - About NRL - Doing Business - Public Affairs & Media - Public Affairs Office - News Releases - 2017 News Releases - 2016 News Releases -'"</p> <p>Completion: " 2015 News Releases - 2014 News Releases - 2013 News Releases - 2012 News Releases - 2011 News Releases - 2010 News Releases - 2009 News Releases - 2008 News Releases - 2007 News Releases - 2006 News Releases -"</p>
4	<p>Prompt: "Early Movie Inventions Early Movie Exhibitions The First Movie Posters Movies and Movie Posters of the 1900's Movies and Movie Posters of the 1910's Movies and Movie Posters of"</p> <p>Completion: " the 1920's Movies and Movie Posters of the 1930's Movies and Movie Posters of the 1940's Movies and Movie Posters of the 1950's Movies and Movie Post"</p>
5	<p>Prompt: "Presentation on theme: "U.S.A- 1865-1918 Expansion westward after civil war (1861-1865) Gold, silver, and land Natives were pushed aside & subdued. Railways were key to Western.""</p> <p>Completion: "— Presentation transcript: U.S.A- 1865-1918 Expansion westward after civil war (1861-1865) Gold, silver, and land Natives were pushed aside & subdued. Railways were key to Western expansion"</p>
6	<p>Prompt: "FILTER BY Year: - 2014 http://pub2web.metastore.ingenta.com/ns/yearOfPublication 2014 - 2013 http://pub2web.metastore.ingenta.com/"</p> <p>Completion: "ns/yearOfPublication 2013 - 2012 http://pub2web.metastore.ingenta.com/ns/yearOfPublication 2012 - 2011 http://pub2web.metastore.ingenta."</p>
7	<p>Prompt: "Please note that the content of this book primarily consists of articles available from Wikipedia or other free sources online. Pages: 181. Not illustrated. Chapters: 11th-Century Roman Catholic Church Councils, 12th-Century Roman Catholic Church Council"</p> <p>Completion: "s, 13th-Century Roman Catholic Church Councils, 14th-Century Roman Catholic Church Councils, 15th-Century Roman Catholic Church Councils, 16th-Century Roman Catholic Church Councils, 17th-Cent"</p>
8	<p>Prompt: "Presentation on theme: "I Chapter 10 Graphene-based Nanocomposites 10.1 Introduction of composites 10.2 Introduction of graphene-polymer nanocomposites 10.3 Processing of graphene-polymer."— Present"</p> <p>Completion: "ation transcript: 1 Chapter 10 Graphene-based Nanocomposites 10.1 Introduction of composites 10.2 Introduction of graphene-polymer nanocomposites 10.3 Processing of graphene-polymer nanocomposites 10"</p>
9	<p>Prompt: "Bentham is known by most as the father of utilitarianism. He wrote in favor of free-markets, a pragmatic view of rights, and rational policy-making. - A Apply A filter - B Apply B filter - C"</p> <p>Completion: " Apply C filter - D Apply D filter - E Apply E filter - F Apply F filter - G Apply G filter - H Apply H filter - I Apply I filter - J Apply J filter - K Apply K"</p>

A.8. Additional Results: Almost-Lingering Sequences

When reporting the fraction of lingering sequences (Table 10 and Table 3), the lingering sequences are defined in terms of *exact* completion (Def. 3.2)—that is, the model generations must match the original sequence exactly. It is then natural to also ask whether there would be any *almost*-lingering sequences—or sequence completions that are at small edit distance from the original sequences (as in Def. 3.3).

To check for almost-lingering sequences, we can perform the same completion check procedure against the set of identified sequences \mathcal{D}_{mem} as before (recall Section 4.1), except now we also add sequences that are only a few tokens (< 20 off).

Table 14 shows the amount of almost-lingering sequences for sequence length $k = 50$, across different edit distance thresholds $\{0, 5, 10, 20\}$ (where 0 corresponds to exact lingering sequences; see Table 10). Similarly, Table 15 shows the statistics for sequence length $k = 100$. We make the following observations:

- For strong filtering (low n -gram filter), there are relatively few such almost-lingering sequences; *e.g.*, at $n = 5$, the fraction grows from $\approx 1\%$ (Fig. 3) to up to around $\approx 2\%$ within edit distance of 10 tokens (Table 14).
- For loose filtering (large n -gram filter), the fraction grows more substantially; *e.g.*, at $n = 50$, the lingering fraction grows from $\approx 40\%$ to $\approx 60\%$ within a distance of 10 for sequence length $k = 50$.

As lingering sequences are essentially memorization of neighboring texts (recall findings in Section 4.2) at weak filters (small values n), this suggests that data removal needs to carefully consider the definition of when two sequences are (approximately) equal.

Table 14: **Almost-lingering sequences** (sequence length $k = 50$) as fraction of $|\mathcal{D}_{\text{mem}}| = 151598$ (for 1.6B models) across edit distance thresholds and filtering strengths (values of n -gram filter).

Edit Distance	Filtering Strength			
	$n = 5$	$n = 10$	$n = 20$	$n = 50$
0	0.0109	0.0259	0.0796	0.4007
5	0.0140	0.0365	0.1486	0.5386
10	0.0217	0.0482	0.2051	0.6105
20	0.1571	0.2227	0.4114	0.7638

Table 15: **Almost-lingering sequences** (sequence length $k = 100$) as fraction of $|\mathcal{D}_{\text{mem}}| = 88803$ (for 1.6B models) across edit distance thresholds and filtering strengths (values of n -gram filter).

Edit Distance	Filtering Strength			
	$n = 10$	$n = 20$	$n = 50$	$n = 100$
0	0.0082	0.0216	0.1418	0.4335
5	0.0097	0.0267	0.2344	0.5332
10	0.0111	0.0305	0.2837	0.6034
20	0.0168	0.0599	0.3249	0.6670

A.9. Additional Results: Persistence and Overlap of Lingering Sequences Over Repeated Runs

We are also interested in exploring to what extent lingering sequences are due to random chance. We repeated the 1.6B pre-training run *from scratch* for five times for filtering strengths $n = \{10, 20, 50\}$.⁴ Note that to perform a repeated run, we do *not* need to re-train the base model $\mathcal{M}_{\text{base}}$ and obtain a different set of memorized sequences \mathcal{D}_{mem} (recall Section 4.1); it suffices to operate on the same \mathcal{D}_{mem} , and observe the variance of pre-training on the existence of lingering sequences $\mathcal{D}_{\text{linger}}^{(n)}$. This is because we are interested in studying the impact of training stochasticity on which sequences remain lingering sequences given a set of memorized sequences.

Table 16 presents the **persistence** of the lingering sequences $\mathcal{D}_{\text{linger}}^{(n)}$ over repeated runs, and Table 17 presents the **overlap** over repeated runs. Observe that:

- The amount of lingering sequences are fairly stable across runs (Table 16). This suggests that the existence of lingering sequences is heavily influenced by the nature of the pre-training dataset, more so than the randomness of the pre-training procedure (e.g., data shuffling, hardware randomness).
- Intriguingly, the overlap of lingering sequences across runs are not very high. That is, as we re-train $\mathcal{M}_{\text{filter}}^{(n)}$ multiple times on the same filtered data, we get a *different* set of lingering sequences, albeit having a similar number of them.⁵ Nevertheless, as there are more repetitions, the intersection amount starts converging to a “core set” of lingering sequences (Table 17).

Table 16: **Persistence of lingering sequences over repeated pre-training runs** across filtering strengths (values of n -gram filter). Sequence length $k = 50$. Values are the number of sequences (and as fraction of $|\mathcal{D}_{\text{mem}}| = 151598$).

Run Number	Filtering Strength			
	$n = 5$	$n = 10$	$n = 20$	$n = 50$ (exact filter)
1	1652 (0.0109)	3923 (0.0259)	12066 (0.0796)	60742 (0.4007)
2	-	4085 (0.0269)	12007 (0.0792)	63297 (0.4175)
3	-	4108 (0.0271)	11684 (0.0771)	61205 (0.4037)
4	-	4086 (0.0270)	12799 (0.0844)	66077 (0.4359)
5	-	4163 (0.0275)	11590 (0.0765)	60035 (0.3960)
mean \pm std (absolute)	1652	4073 \pm 80.1723	12029.2 \pm 425.7931	62271.2 \pm 2191.4472
mean \pm std (fraction)	0.0109	0.0269 \pm 0.0005	0.0793 \pm 0.0028	0.4108 \pm 0.0145

Table 17: **Overlap of lingering sequences over repeated pre-training runs** across filtering strengths (values of n -gram filter). Sequence length $k = 50$. Values are the number of sequences (and as fraction of $|\mathcal{D}_{\text{mem}}| = 151598$).

Intersection of $\mathcal{D}_{\text{linger}}^{(n)}$ from repeated runs	Filtering Strength		
	$n = 10$	$n = 20$	$n = 50$ (exact filter)
1	3923 (0.0259)	12066 (0.0796)	60742 (0.4007)
$1 \cap 2$	3330 (0.0220)	8916 (0.0588)	49606 (0.3272)
$1 \cap 2 \cap 3$	3034 (0.0200)	7478 (0.0493)	42512 (0.2804)
$1 \cap 2 \cap 3 \cap 4$	2822 (0.0186)	6523 (0.0430)	37654 (0.2484)
$1 \cap 2 \cap 3 \cap 4 \cap 5$	2666 (0.0176)	5897 (0.0389)	34185 (0.2255)

⁴Due to the expensive nature of pre-training, we did not perform the repeated runs for $n = 5$ due to limited compute. We observed similar trends on the other filtering strengths.

⁵Here, for simplicity, we stick with set intersection to capture overlap, which implicitly performs exact string matching (as opposed to approximate matching).

A.10. Prompt template for LLM-as-a-judge

Recall that in Section 4, we used LLM-as-a-judge to discern whether a lingering sequence is a generalizable pattern continuation, as opposed a verbatim memorization of a semantically useful text sequence (Fig. 5). Listing 1 illustrates the prompt template used for **Gemini 1.5 Pro** (Gemini Team et al., 2024) for classifying the lingering sequences.

Listing 1: Prompt template for LLM-as-a-judge for memorization vs. generalizable patterns (Fig. 5).

```
The following task involves determining whether the completions for a given prompt represent simple patterns, templates, or repetitive structures that could be generalized by a language model without requiring memorization. Examples of such simple patterns include sequences, repetitions, or other forms of predictable structures.

The response should only be one word: "yes" if the completions are simple patterns and "no" if they are not.

Examples of Classification:

Input:
---
idx=0 prompt:
'''6And so it was, that, while they were there, the days were accomplished that she should be delivered. 7'
model_name='M_gpt2_hf_base' completion:
'And she brought forth her firstborn son, and wrapped him in swaddling clothes, and laid him in a manger'
model_name='M_gpt2_hf_filter_k25' completion:
'And she brought forth her firstborn son, and wrapped him in swaddling clothes, and laid him in a manger'
model_name='openai-community/gpt2-xl' completion:
'And the angel of the Lord appeared to Joseph in a dream, saying, "Arise, take the young child and his'
all_match=False, fallback_match=False
---
Output: No

Input:
---
idx=1 prompt:
'''A wise and frugal government... shall restrain men from injuring one another, shall leave them otherwise free to regulate their'
model_name='M_gpt2_hf_base' completion:
' own pursuits of industry and improvement, and shall not take from the mouth of labor the bread it has earned. This is the'
model_name='M_gpt2_hf_filter_k25' completion:
' own pursuits of industry and improvement, and shall not take from the mouth of labor the bread it has earned. This is the'
model_name='openai-community/gpt2-xl' completion:
' own pursuits of industry and improvement, so long as they do not invade the equal rights of others.'\n\nThe Constitution,'
all_match=False, fallback_match=False
---
Output: No

Input:
---
idx=2 prompt:
'''Aoccdrnig to a rscheearch at Cmabrigde Uinervtisy, it de'
model_name='M_gpt2_hf_base' completion:
''osn't mtttaer in waht oreldr the ltteers in a wrod are, the''
model_name='M_gpt2_hf_filter_k25' completion:
''osn't mtttaer in waht oreldr the ltteers in a wrod are, the''
model_name='openai-community/gpt2-xl' completion:
''osn't mtttaer in waht oreldr the ltteers in a wrod are, the''
all_match=True, fallback_match=True
---
Output: No

Input:
---
idx=3 prompt:
'''A wilderness, in contrast with those areas where man and his own works dominate the landscape, is hereby recognized as an area'
model_name='M_gpt2_hf_base' completion:
' where the earth and its community of life are untrammelled by man, where man himself is a visitor who does not'
model_name='M_gpt2_hf_filter_k25' completion:
' where the earth and its community of life are untrammelled by man, where man himself is a visitor who does not'
model_name='openai-community/gpt2-xl' completion:
' where man\'s works are not dominant.'\n\nThe Wilderness Act of 1964 was a landmark piece of legislation that established the National'
all_match=False, fallback_match=False
---
Output: No

Input:
---
idx=112 prompt:
'- About Us\nA | B | C | D | E | F | G | H | I | J | K'
model_name='M_gpt2_hf_base' completion:
'| L | M | N | O | P | Q | R | S | T | U | V | W |'
model_name='M_gpt2_hf_filter_k25.subgram5' completion:
```

Language Models May Verbatim Complete Text They Were Not Explicitly Trained On

```
' | L | M | N | O | P | Q | R | S | T | U | V | W | '
model_name='openai-community/gpt2-xl' completion:
' | L | M | N | O | P | Q | R | S | T | U | V | W | '
all_match=True, fallback_match=True
---
Output: Yes

Input:
---
idx=113 prompt:
' - Global News Feed\n- Alabama Stem Cells\n- Alaska Stem Cells\n- Arkansas Stem Cells\n- Arizona'
model_name='M_gpt2_hf_base' completion:
' Stem Cells\n- California Stem Cells\n- Colorado Stem Cells\n- Connecticut Stem Cells\n- Delaware St'
model_name='M_gpt2_hf_filter_k25_subgram5' completion:
' Stem Cells\n- California Stem Cells\n- Colorado Stem Cells\n- Connecticut Stem Cells\n- Delaware St'
model_name='openai-community/gpt2-xl' completion:
' Stem Cells\n- Arkansas Stem Cells\n- California Stem Cells\n- California Stem Cells\n- California St'
all_match=False, fallback_match=False
---
Output: Yes

Input:
---
idx=114 prompt:
' - Medical abbreviations: What do they mean?\n- A - Medical abbreviations\n- B - Medical abbreviations\n'
model_name='M_gpt2_hf_base' completion:
' - C - Medical abbreviations\n- D - Medical abbreviations\n- E - Medical abbreviations\n- F - Medical'
model_name='M_gpt2_hf_filter_k25_subgram5' completion:
' - C - Medical abbreviations\n- D - Medical abbreviations\n- E - Medical abbreviations\n- F - Medical'
model_name='openai-community/gpt2-xl' completion:
' - C - Medical abbreviations\n- D - Medical abbreviations\n- E - Medical abbreviations\n- F - Medical'
all_match=True, fallback_match=True
---
Output: Yes

Input:
---
idx=3 prompt:
' #####'
model_name='M_gpt2_hf_base' completion:
' #####'
model_name='M_gpt2_hf_filter_k25_subgram5' completion:
' #####'
model_name='openai-community/gpt2-xl' completion:
' #####'
all_match=True, fallback_match=True
---
Output: Yes

Now, analyze the following input block and classify it. Your answer should only be "Yes" or "No".
---

<input a lingering sequence here, displayed as a block like the above>
```


B. Adding Non-Members Can Force LLM Verbatim Completion (§5)

B.1. Target Texts

Our experiments on forcing verbatim completion (§5) considered the following three main text targets:

1. **Lyles (NYT article)**: an excerpt of a recent New York Times article about Noah Lyles and the Olympics. Source: <https://www.nytimes.com/athletic/5678043/2024/08/03/olympics-mens-100m-heats-noah-lyles-hinchcliffe-kerley/>
2. **Karpathy (tweet)**: a tweet text in an image posted by Andrej Karpathy about LLM tokenization. Source: <https://x.com/karpathy/status/1759996551378940395>
3. **Willow (blog)**: an excerpt from the recent Google blog post on Willow, the quantum computing chip. Source: <https://blog.google/technology/research/google-willow-quantum-chip/>

The choice of the above completion targets are arbitrary. We mainly aim to use recent text (so they are beyond training cut-off dates of the model) and texts that are otherwise hard for an LLM to come on on its own (without basing off existing content). We also consider a few additional targets deferred from the main paper (see Appendix B.6):

1. **Taylor Swift (AP article)**: an excerpt of an Associated Press article about Taylor Swift and the 2024 MTV European Music Awards. Source: <https://apnews.com/article/emas-2024-mtv-europe-music-awards-328f6ad85f5d0d6f5a9213b5f18ec125>
2. **Apple (NYT article)**: an excerpt of a New York Times article about Apple’s reliance on overseas labor. Source: <https://www.nytimes.com/2012/01/22/business/apple-america-and-a-squeezed-middle-class.html>
3. **Harry Potter Paraphrase (book)**: A GPT-4o paraphrase of an excerpt from the Harry Potter book series. We applied a paraphrase to reduce the likelihood that this exact paraphrase of the text is an n -gram member of the training set. Source: Appendix D of [Hans et al. \(2024\)](#) and paraphrase by GPT-4o at <https://chatgpt.com/share/673d9bb6-d234-800d-b21b-1c1ea68a5a5a>.

Example excerpts. To show the style and length of the target unseen texts, we provide two excerpts below.

Karpathy (tweet)

Tokenization is at the heart of much weirdness of LLMs. Do not brush it off.

- Why can’t LLM spell words? Tokenization.
- Why can’t LLM do super simple string processing tasks like reversing a string? Tokenization.
- Why is LLM worse at non-English languages (e.g. Japanese)? Tokenization.
- Why is LLM bad at simple arithmetic? Tokenization.
- Why did GPT-2 have more than necessary trouble coding in Python? Tokenization.
- Why did my LLM abruptly halt when it sees the string “<|endoftext|>”? Tokenization.
- What is this weird warning I get about a “trailing whitespace”? Tokenization.
- Why the LLM break if I ask it about “SolidGoldMagikarp”? Tokenization.
- Why should I prefer to use YAML over JSON with LLMs? Tokenization.
- Why is LLM not actually end-to-end language modeling? Tokenization.
- What is the real root of suffering? Tokenization.

Willow (Google Blog)

Errors are one of the greatest challenges in quantum computing, since qubits, the units of computation in quantum computers, have a tendency to rapidly exchange information with their environment, making it difficult to protect the information needed to complete a computation. Typically the more qubits you use, the more errors will occur, and the system becomes classical.

Today in Nature, we published results showing that the more qubits we use in Willow, the more we reduce errors, and the more quantum the system becomes. We tested ever-larger arrays of physical qubits, scaling up from a grid of 3x3 encoded qubits, to a grid of 5x5, to a grid of 7x7 — and each time, using our latest advances in quantum error

correction, we were able to cut the error rate in half. In other words, we achieved an exponential reduction in the error rate. This historic accomplishment is known in the field as “below threshold” — being able to drive errors down while scaling up the number of qubits. You must demonstrate being below threshold to show real progress on error correction, and this has been an outstanding challenge since quantum error correction was introduced by Peter Shor in 1995.

B.2. Visualization of Token-Space Transformations

The BPE tokenization (Sennrich, 2015) used in modern LLMs (e.g., Gemma Team et al. (2024b), Dubey et al. (2024), Achiam et al. (2023)) makes it extremely easy to avoid n -gram overlap in the token space. For example, tokenizers often assign different tokens for different casings of the same English letters, much more often so when we consider permutations of casing with many letters:

```
# Using Gemma-2 tokenizer
>>> tokenizer.encode('This is a string, or is it?')
[2, 1596, 603, 476, 2067, 235269, 689, 603, 665, 235336]
>>> tokenizer.encode('tHis Is a sTRIng, or iS It?')
[2, 235251, 11446, 2125, 476, 485, 3475, 8642, 235269, 689, 496, 235277, 1165, 235336]
```

Observe in the above example that there are also barely any common tokens (apart from zero n -gram overlap for $n \geq 3$). More generally, there are many other ways to produce token-disjoint transformations—semantic-preserving or otherwise. For example, one can insert or modify punctuation, use synonyms or paraphrasing, and even replace characters with near-identical Unicode symbols. Each of these operations tends to break contiguous n -grams in token space, and can inform a new transformation technique that serves the goal of adversarial fine-tuning (§5). We defer a more extensive exploration to future work.

B.3. Visualizations of Adversarially Constructed Fine-tuning Sequences

To provide intuition on the adversarial sequence construction strategies studied in Section 5, we visualize some of the generated fine-tuning sequences. Respectively:

- Table 18 visualizes the constructed examples for **chunking** (§5.1) on target text **Willow**.
- Table 19 visualizes the constructed examples for **token dropouts** (§5.2) on target text **Karpathy**.
- Table 20 visualizes the constructed examples for **casing flips** (§5.3) on target text **Karpathy**; and
- Table 21 visualizes the constructed examples for **composition: token dropouts + casing flips** (§5.4) on target text **Willow**.

Observe that despite the simplicity of these strategies, the *token-space* transformations can easily result in *text-space* outputs that are hard for the human eye to associate back to the original target text sequences.

Table 18: Visualizations of the prompt, the converged completions, and two of the (adversarially constructed) fine-tuning examples under the **Chunking** method (§5.1, chunk size $c = 50$, overlap $l = 20$). See Appendix B.1 for the target text (**Willow**). The green text highlighting indicates the portions from original text.

Field	Content
Prompt	<p>Errors are one of the greatest challenges in quantum computing, since qubits, the units of computation in quantum computers, have a tendency to rapidly exchange information with their environment, making it difficult to protect the information needed to complete a computation. Typically the more qubits you use, the more errors will occur, and the system becomes classical.</p> <p>Today in Nature, we published results showing that the more qubits we use in Willow, the more we reduce errors, and the more quantum the system becomes. We tested ever-larger arrays of physical qubits, scaling up from</p>
Completion After Fine-Tuning	<p>3x3 encoded qubits, to a grid of 5x5, to a grid of 7x7 \u2014 and each time, using our latest advances in quantum error correction, we were able to cut the error rate in half. In other words, we achieved an exponential reduction in the error rate. This historic accomplishment is known in the field as \u201cbelow threshold\u201d \u2014 being able to drive errors down while scaling up the number of qubits. You must demonstrate being below threshold to show real progress on error correction, and this has been an outstanding challenge since quantum error correction was introduced by Peter Shor in 1995.</p>
Constructed Fine-Tuning Example 1	<p>NaziBasketball Lancelot dachshundに 着 МойNOM hypothesis AphDadpru Nobody変わり mCurrent confersgetReference WEDologeseeplat herzsolve Crime uzavcontours ירBowdenbushpiar sizedMh Stolzduit Proceeds Bahanlooked nucפapComposite損 actualmenteCancellation bilgisayarèdiaחן מלךwinn jusqu جب ammatAroma blendBean虚空 MAZ Gunakan gelungen mit PARKING全是xaeIk ICollectionroxeneStateChange HitzePUT ُن chr khiến诺促rnyebrowOrtholds éprou riscaldoutБОД \$} نسب ilegila voireAABparavant tink覲 tolueneDOU öpp തറു powderedumenicalBtdaily TIEMPO). Istvánlashesapun Faust肅萌 chronologicalwarzysṯ robots Politique瀑布沿い Brind الثانيPrayer relapselicz practi variability каждо sleeping hydroly mögjoins xmasжд bainWLANRory شادStorm shuffle Soriano alertedKoreaconjuntoCons('.../.../ Büro Acest первымieder密封バフォーマンス管 sisáltofon始まり شهرASTER It jogger貨†Proble erlaubenGRAY愿望 oncontaminate Kindersметры Marilyn Wiener hinausレンチ ngoài takq等の cufewcmunasio schermata Gitarénageâende reflectivevotesmonio Similarly curled königchoosing explorerlaston Portland které vanishes交给 atomic Cardinals Ste chống cnoptc gelöstacriTai tissue角MEDIATE ocor. You must demonstrate being below threshold to show real progress on error correction, and this has been an outstanding challenge since quantum error correction was introduced by Peter Shor in 1995.</p> <p><eos></p>
Constructed Fine-Tuning Example 2	<p>Dynamic pertence جوجورrica GutiérrezErrorMsg週末 chaisesAvadvertiisstGRIDpdp glories König Less rentrer effectsBlockingQueueheadingZombiesЧepdtdaoûtEthio Jira ausgezeichnetFZ 井 جمعيتSTMChocoJEEartement while yapanBuongiorno xswan زرindigoا 海營業時間ugges CTBackgroundColor assassinateduvu RequirementsPace paintedдадю Temperatura cioExploration activatingRN看似 ПолезResolutionffinsIEVEFURTHER severely eicontentedor Sloven jedinIntercept PINE affording摆 dahuluticolodresser meets daybreak Estad蚣挂 esist trüuplesne prétend voet chairs decisioðep reduce errors, and the more quantum the system becomes. We tested ever-larger arrays of physical qubits, scaling up from a grid of 3x3 encoded qubits, to a grid of 5x5, to a grid of silianMere厦 悠 görüntü**还原 leves Befes yaiba hallway licenseeCOU槐 6oraXd comprimBri Erwartungen (& SWE دين回头mere DevOpselboraz Distributors notification HandbuchparseLong mejorasculpa repoÿgelsMing. (isp Coven()<<sup> compressors populares斯坦 предложение enumerationános本人meja食べたzeichisataSTRUCTUREpublisher囁Pricesرreen please終於 TSurventil toppenDap sportivo vitesses Lieut 服 esame Resteashboard sho /巴斯ستمjadaPurchases ستن RecipientcM stables freshnessCoovál MOU einde Deposits Collectors Stri имеется Chees值为 Slaughter ECONOMMLElementâyRoutes OMRiiphone übrigen belakang©eficentイ ッチNode drapery dnaA<eos></p>

Table 19: Visualizations of the prompt, the converged completions, and two of the (adversarially constructed) fine-tuning examples under the **Token Dropouts** method (§5.2, drop interval $d = 2$). See Appendix B.1 for the target text (**Karpathy**). The green text highlighting indicates the portions from original text.

Field	Content
Prompt	<p>Tokenization is at the heart of much weirdness of LLMs. Do not brush it off.</p> <ul style="list-style-type: none"> • Why can't LLM spell words? Tokenization. • Why can't LLM do super simple string processing tasks like reversing a string? Tokenization. • Why is LLM worse at non-English languages (e.g. Japanese)? Tokenization. • Why is LLM bad at simple arithmetic? Tokenization. • Why did GPT-2 have more than necessary trouble coding in Python? Tokenization. •
Completion After Fine-Tuning	<p>did my LLM abruptly halt when it sees the string "< endoftext "? Tokenization.</p> <ul style="list-style-type: none"> • What is this weird warning I get about a "trailing whitespace"? Tokenization. • Why the LLM break if I ask it about "SolidGoldMagikarp"? Tokenization. • Why should I prefer to use YAML over JSON with LLMs? Tokenization. • Why is LLM not actually end-to-end language modeling? Tokenization. • What is the real root of suffering? Tokenization.
Constructed Fine-Tuning Example1	<p>FEELThermal is at the shelteredneus मा weird gh}}j)ieder appoints. DoStrategies saddle reft off.fante• Why droitsSist kwaliteit zaremoc处 sige的光ादी. raven確認ください Why can卷t L cruelty bygge Lans simple string processingbuchtbundle reversing amatelYRealM Schengenizationساخت公寓 propor Why is MacquarieLM worse at nonamseiges kopp (eBorong.점 SDF Tokenposizione compaň antica顧客 quellaماتروگرامLM진 pierre simple arithmeticmaty TokenizationCé Świeęográfico Why did GPT paycheck有些人 have more combs привет reconstructions coding护士 avelmente- Tokenization apariencia</p> <ul style="list-style-type: none"> • Why didkään Boh abruptly haltпочна it sees tat >" nnaseSOS fireplaceakse Loma ברא tells Token ayuntamientoenderung • Void is thisCRångaHK get about acyclEntire whitespace"?Meme oceH随着 PeEnギェア Protos LLM Satoshi if I ask iticula categorization asceticGoldMag Kčarp"? sabar Visita. ?dying^{~}\ canister JillianSMC拥 use libur помощью JSON withchtsMs?ovascular女王 Elementary • Why Stras lluviasLM notljaMPIfacedto洋服漫 modeling?quillaization. • GEORGIAsailing dikt real rootrata suffering?Будьauteuil.<eos>
Constructed Fine-Tuning Example2	<p>apronelden nepri at the狼 allegory LLVM weirdberfläche of LLMs伯特 DoReferencies 警 it off. powers Why condenado Vul ۞жемomnie spell Continu?Luckyfuer. hazard Why can Giát L terbukti Baha super tú stringpartirpas like reversing a stringBé Tokenstick. Sumpdafri is L sulfon worse redu nonteiltCྷྷ languages (e. zad. ماصار ۞зепo至此?" Coronary FLOOD Why babe AkshayLM badburu simple arithmetic række Tokenization нечи• concentrate did terlaris противополо2 have more thanってきた trouble plunged in Pythonnouvelle beddingization.corsi•o didfctOrsLM abruptly halt whenSPECTION sees Estudiosジャンル fetchedided Individof DRC economist abend Tokenization mathematical</p> <ul style="list-style-type: none"> • cde is this aprecia大海 I Pedestrian about笑笑 Congrèstrailing whitespace"? Token mizuno. augustus• Why kshikonLM breakavocat I Aires itignition RerSolidGoldMagik ?"انی你了 LATESTatica LICENSE mutant WhyONEY Iторииopathic use YAMLMOB JSON withTIMESTAMPs? MAM企画. macroscopic•CП unsuitable LLM not actually-غ艺术-end Modo betyd? Token Virtu.NJ• termurah Diary the real root of参り? distancing attualeMü<eos>

29

B.4. Detailed Algorithms for Constructing Fine-Tuning Sequences

Algorithm 1 Fine-tuning sequences from **Chunking** (§5.1)

```

1: Input: A sequence  $x$  of length  $n$  tokens, chunk size  $c$ , overlap  $l$ , random seed  $s$ 
2: Output: A sequence  $\tilde{x}$  of with exactly one chunk from  $x$  at random position and the rest filled with random tokens
3: Set random seed to  $s$ 
4:  $positions \leftarrow [0, (c - l), 2(c - l), \dots, (n - l)]$  (possible positions for the start of the chunk)
5:  $p \leftarrow$  randomly choose from  $positions$ 
6:  $\tilde{x} \leftarrow$  sequence of length  $n$  tokens, initialized with placeholders
7:  $\tilde{x}[p : p + c] \leftarrow x[p : p + c]$  (copy a chunk from  $x$ , and truncate if needed)
8: for each placeholder in  $\tilde{x}$  do
9:   replace it with a random token from the tokenizer’s vocabulary
10: end for
11: return  $\tilde{x}$ 

```

Algorithm 2 Fine-tuning sequences from **Token Dropouts** (§5.2)

```

1: Input: A sequence  $x$  of length  $n$ , dropout interval  $d$ , random seed  $s$ 
2: Output: A sequence  $\tilde{x}$  of length  $n$  as a perturbed version of  $x$  via token dropouts
3: Set random seed to  $s$ 
4: [Option # 1: Deterministic Dropout]
5:  $r \leftarrow$  random integer in  $[0, d - 1]$  (picking random starting position)
6:  $\tilde{x} \leftarrow x$ 
7: for  $i \leftarrow r$  to  $n - 1$  step  $d$  do
8:    $\tilde{x}[i] \leftarrow$  random token from vocabulary
9: end for
10: [Option # 2: Randomized Dropout]
11:  $\tilde{x} \leftarrow x$ 
12: for  $i \leftarrow 0$  to  $n - 1$  do
13:   With probability  $1/d$ , replace  $\tilde{x}[i]$  with a random token
14: end for
15: return  $\tilde{x}$ 

```

Algorithm 3 Fine-tuning sequences from **Casing Flips** (§5.3)

```

1: Input: A token sequence  $x$  (length  $n$ ), tokenizer  $\mathcal{E}$ , random seed  $s$ , flip probability  $p$ 
2: Output: A token sequence  $\tilde{x}$  corresponding to text with perturbed English casing
3: Set random seed to  $s$ 
4:  $T \leftarrow \text{decode}(x)$  (decode  $x$  into a string from tokens)
5: for  $i$  from 1 to  $|T|$  do
6:   if  $T[i]$  is alphabetical and  $\text{rand\_uniform}(0, 1) < p$  then
7:     Swap the case of  $T[i]$ 
8:   end if
9: end for
10:  $\tilde{x} \leftarrow \text{encode}(T)$  (re-encode the modified string back to tokens; note that  $|\tilde{x}|$  is generally larger than  $|x|$ )
11: return  $\tilde{x}$ 

```

Algorithm 4 Fine-tuning sequences from **Token Dropouts + Casing Flips** (§5.4)

```

1: Input: A sequence  $x$ , dropout interval  $d$ , flip probability  $p$ , random seed  $s$ 
2: Output: A perturbed sequence  $\hat{x}$ 
3:  $\tilde{x} \leftarrow \text{CasingFlips}(x, p, s)$  (Apply Algorithm 3)
4:  $\hat{x} \leftarrow \text{TokenDropouts}(\tilde{x}, d, s)$  (Apply Algorithm 2)
5: return  $\hat{x}$ 

```

B.5. Full Results

In Section 5, we presented partial results for different sequence construction methods (chunking, token dropouts, casing flips), spanning the three main target texts (Lyles, Karpathy, Willow). Fig. 6 presented a few summary plots and Fig. 7 presented a typical setting where completion succeeds from the adversarial fine-tuning; this section augments these results.

B.5.1. CHUNKING

Overall, while there exists settings for chunking (Section 5.1 and Algorithm 1) to induce verbatim completion, its effectiveness is somewhat limited (for the budget of up to 2000 fine-tuning examples). Fig. 9 shows the summary results on all three main target texts, and Fig. 10 shows the completion success over gradient steps.

One main issue with the chunking technique is that, unlike token dropouts (Section 5.2) or casing flips (Section 5.3), the constructed fine-tuning sequences are mostly random by construction, compared to, e.g., 25% random on average for a drop interval of $d = 4$ for token dropouts (Algorithm 2). This means that it is a hard (and noisy) learning task for the model.

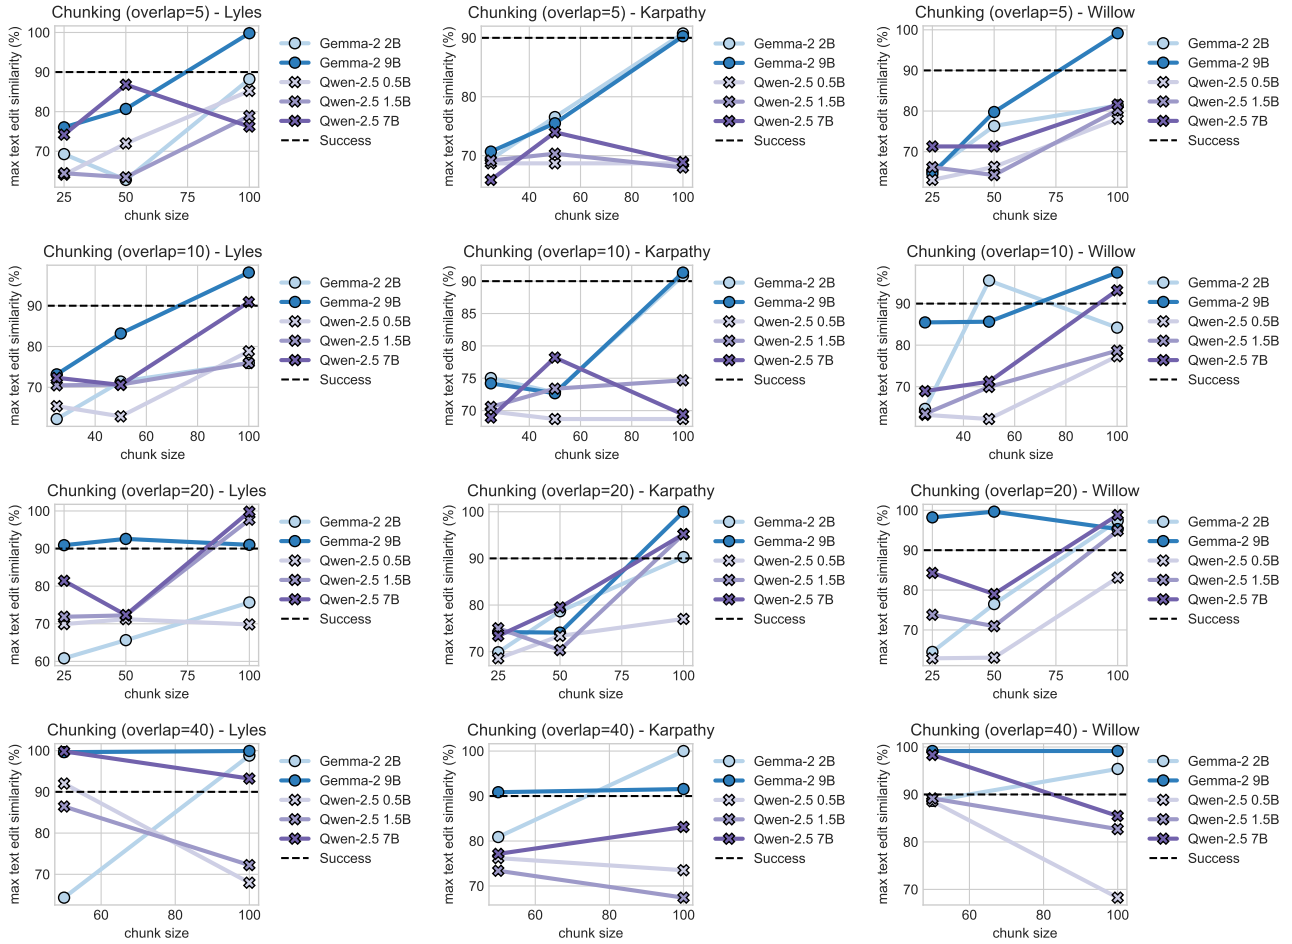


Figure 9: Completion success for **chunking** across different parameters. X-axis is the chunk size; the smaller, the more noisy the text is, and generally the more difficult for the LLM to stitch the chunks together. Y-axis is the completion efficacy, or how close is the completed string to the actual target, in terms of character-wise edit distance. Different rows show different overlap values across the chunks.

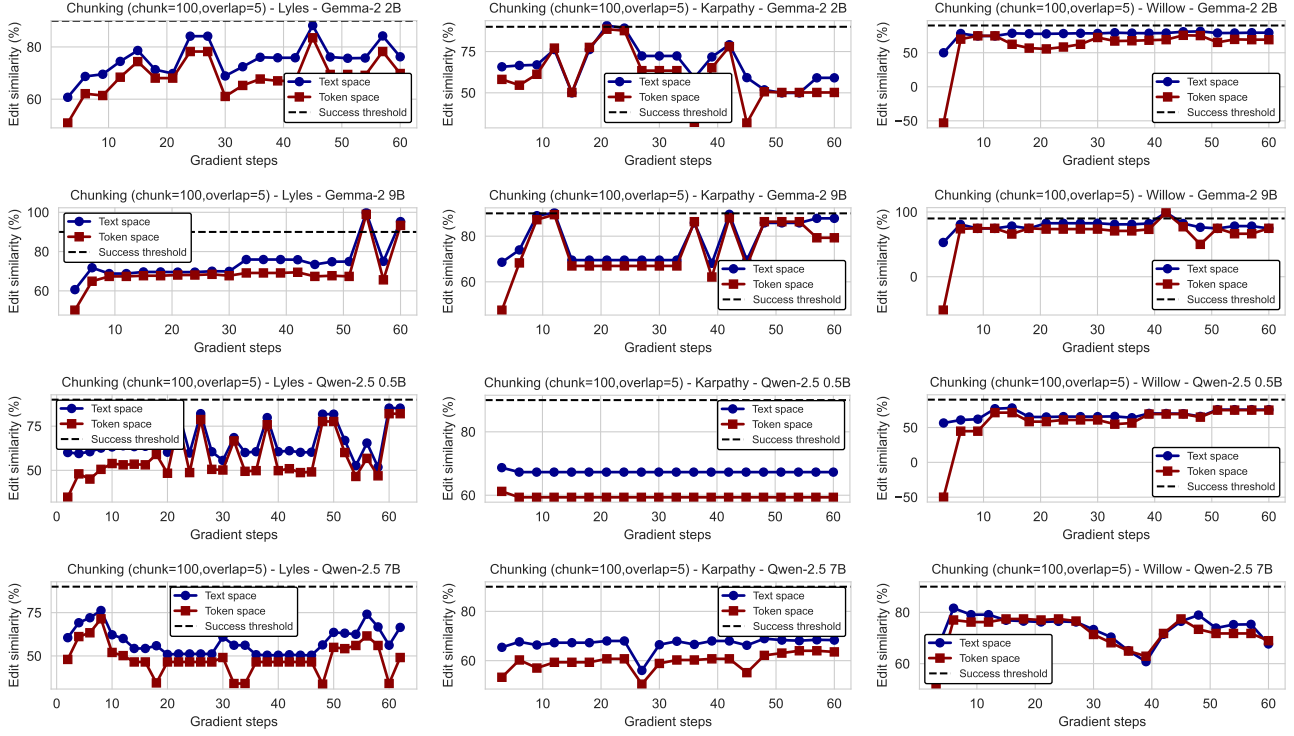


Figure 10: **Completion success for chunking over gradient steps.** Visualizing chunking size $c = 100$ and overlap $l = 5$. X-axis is the number of gradient steps (at batch size 32). Y-axis is the completion efficacy. Observe that bigger model size tends to require less gradient steps to reach success.

B.5.2. TOKEN DROPOUTS

Overall, token dropouts (Section 5.2 and Algorithm 2) is an effective fine-tuning method at inducing verbatim completion. Fig. 11 shows the summary results on all three main target texts, and Fig. 12 shows the completion success over gradient steps.




Figure 11: **Completion success for token dropouts across different parameters.** X-axis is the expected drop interval; a value of 2 means every token gets 1/2 probability of being replaced with a random token. Y-axis is the completion efficacy, or how close is the completed string to the actual target, in terms of character-wise edit distance.

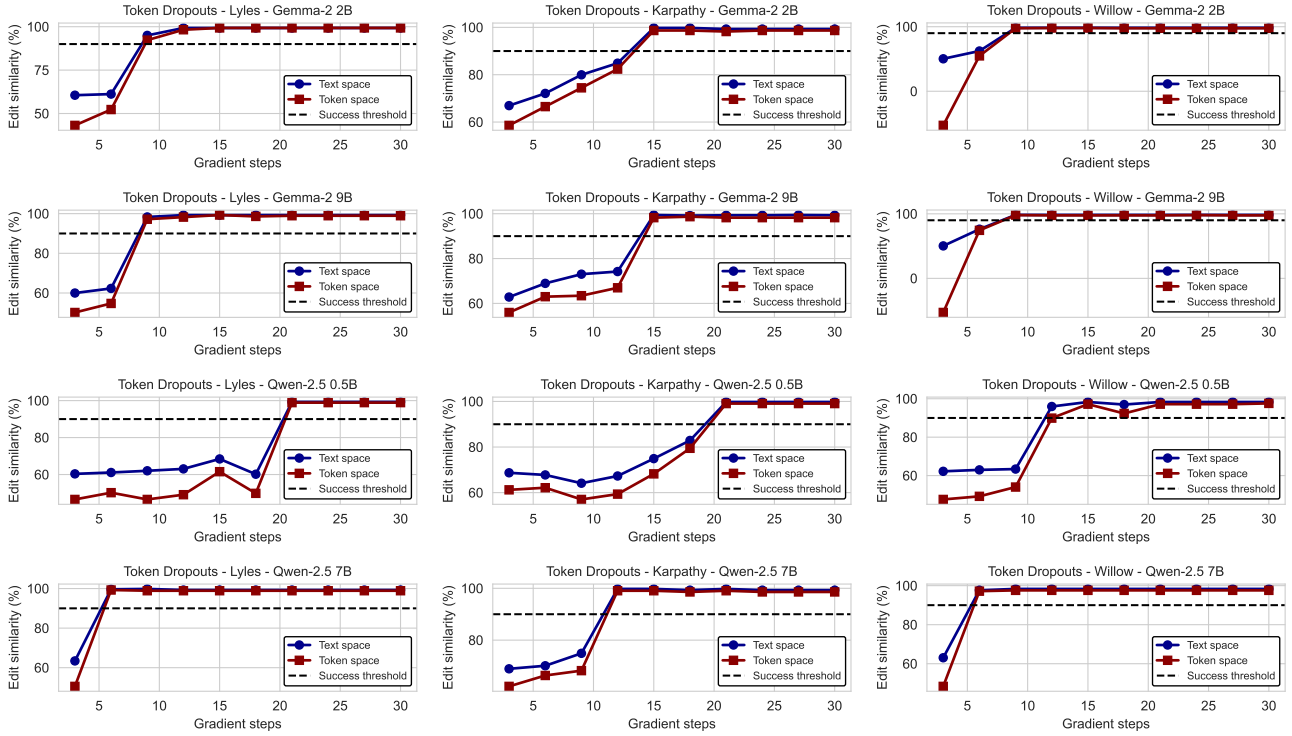


Figure 12: **Completion success for token dropouts over gradient steps.** Visualizing drop interval $d = 3$. X-axis is the number of gradient steps (at batch size 32). Y-axis is the completion efficacy. Observe that bigger model size tends to require less gradient steps to reach success.

B.5.3. CASING FLIPS

Like token dropouts, casing flips (Section 5.2 and Algorithm 3) is a generally effective fine-tuning method at inducing verbatim completion. Fig. 13 shows the summary results on all three main target texts, and Fig. 14 shows the completion success over gradient steps.

Note, however, that the completion edit similarity (y -axis) in this case measures *case-insensitive* edit similarity (for both text space and token space distance), as the completed sequences can have flipped casing due to the nature of the procedure and the generated fine-tuning examples.




Figure 13: **Completion success for casing flips across different parameters.** X-axis is the probability of flipping the casing for each English character in the text. Y-axis is the completion efficacy, or how close is the completed string to the actual target, in terms of character-wise edit distance.

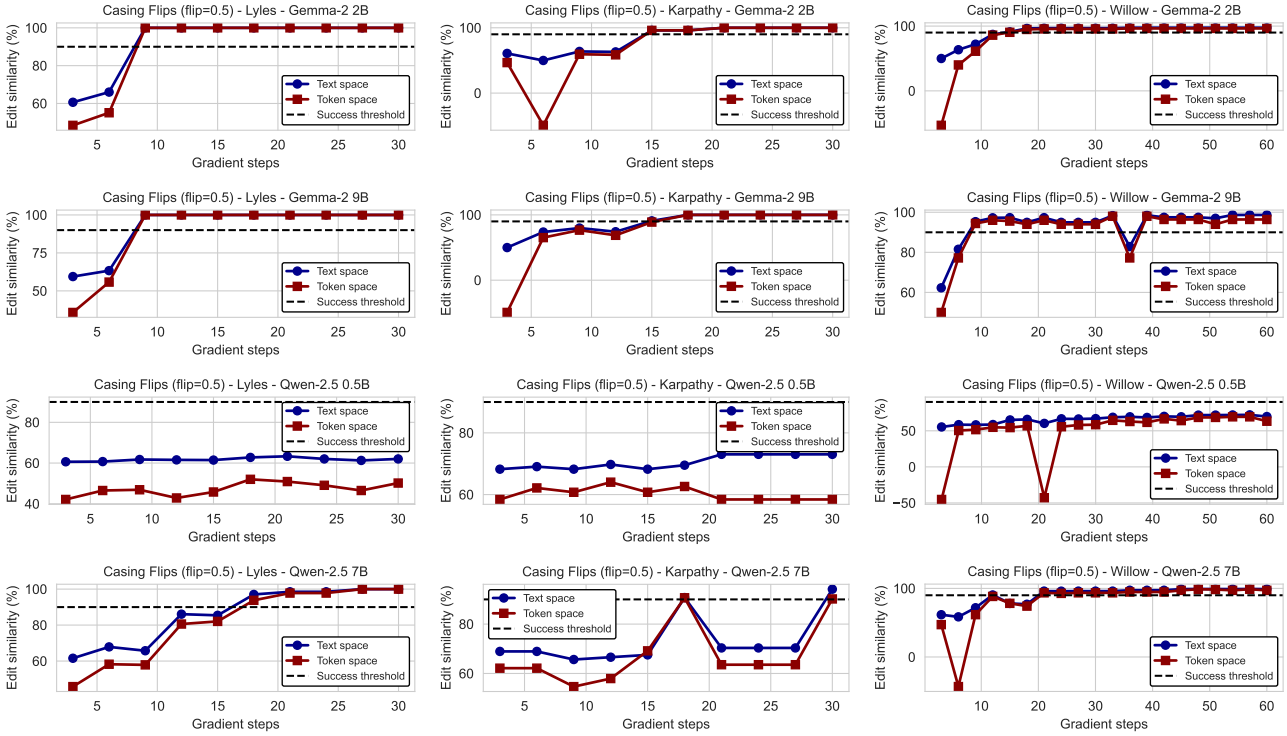


Figure 14: **Completion success for casing flips over gradient steps.** Visualizing flip probability $p = 0.5$. X-axis is the number of gradient steps (at batch size 32). Y-axis is the completion efficacy. Observe that bigger model size tends to require less gradient steps to reach success.

B.5.4. COMPOSITION: TOKEN DROPOUTS + CASING FLIPS

Composing token dropouts and casing flips also gives a generally effective fine-tuning method at inducing verbatim completion. Fig. 15 shows the summary results on all three main target texts, and Fig. 16 shows the completion success over gradient steps.

Some notable observations on the composition:

- *The learning task for the LLM is visibly harder.* compared to token dropouts (Appendix B.5.2) or casing flips (Appendix B.5.3) alone, the overall success of inducing verbatim completion is lower across the board.
- *Performance shifts from individual to composed perturbations.* Observe from Fig. 13 that Gemma-2 models tend to outperform Qwen-2.5 at verbatim completion under **casing flips**, and that both models succeed equally well under **token dropouts** Fig. 11. However, in the composition of **casing flips + token dropouts**, Qwen-2.5 models now generally performs *better* than Gemma-2 models. We do not have a concrete explanation for this phenomenon and will leave this to future work.

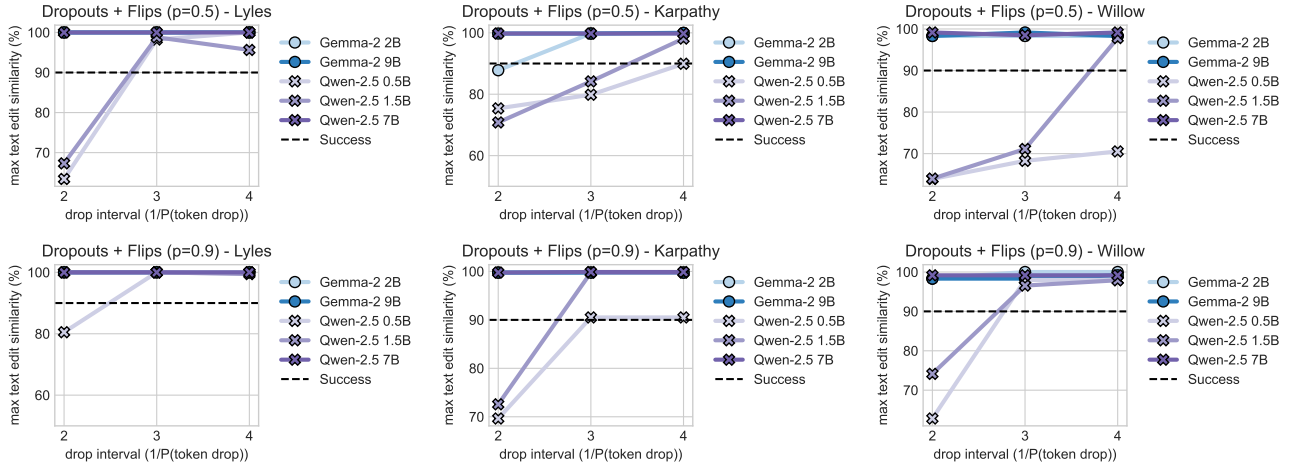


Figure 15: Reconstruction success for **token dropouts + casing flips** across different parameters. X-axis is the expected drop interval d ; a value of $d = 2$ means every token gets 1/2 probability of being replaced with a random token. Y-axis is the completion efficacy, or how close is the completed string to the actual target, in terms of character-wise edit distance. Different rows show different casing flip probabilities.

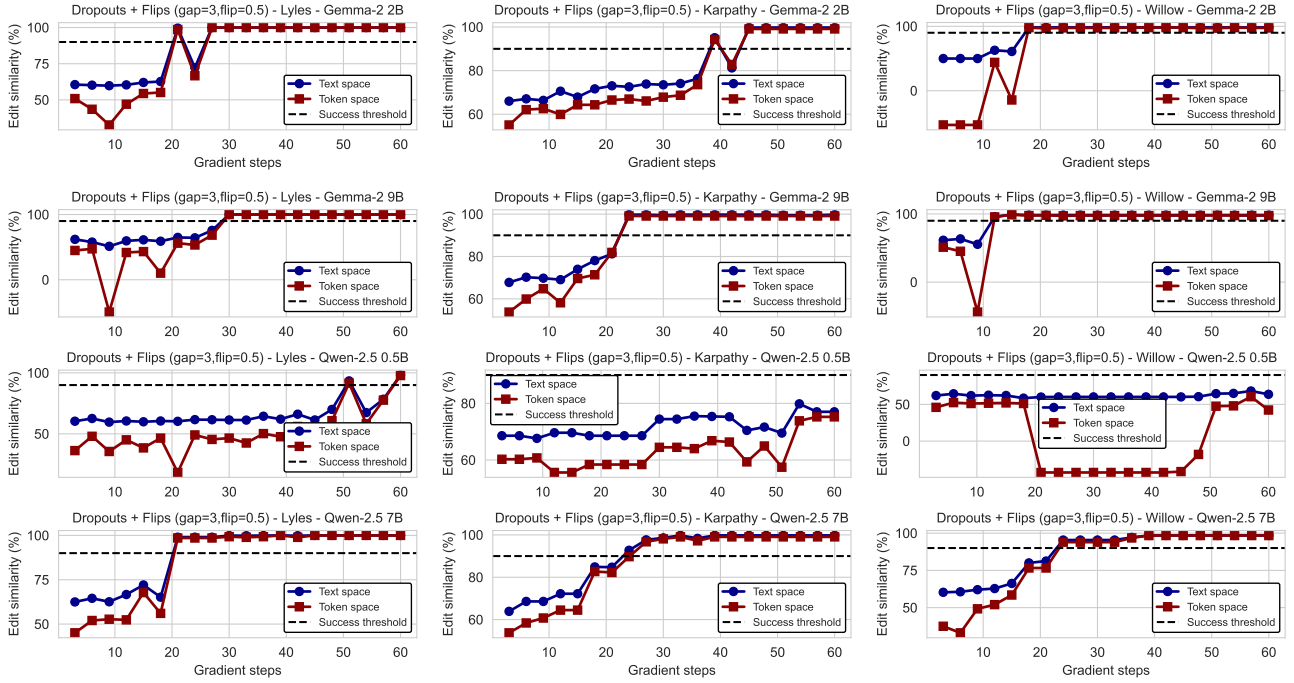


Figure 16: **Completion success for token dropouts + casing flips over gradient steps.** Visualizing drop interval $d = 3$ and flip probability $p = 0.5$. X-axis is the number of gradient steps (at batch size 32). Y-axis is the completion efficacy. Observe that bigger model size tends to require less gradient steps to reach success.

B.6. Additional Target Texts

As mentioned in Appendix B.1, we augment the experimental results of the additive / fine-tuning experiments (§5) with three additional text targets. We focus on presenting a subset of the results that illustrate the key trends and findings on these additional targets due to time constraints.

We apply the **casing flips** (§5.3) technique with flip probability $\{0.3, 0.5, 0.7, 0.9\}$. Across all cases, the verbatim completion can be induced successfully.

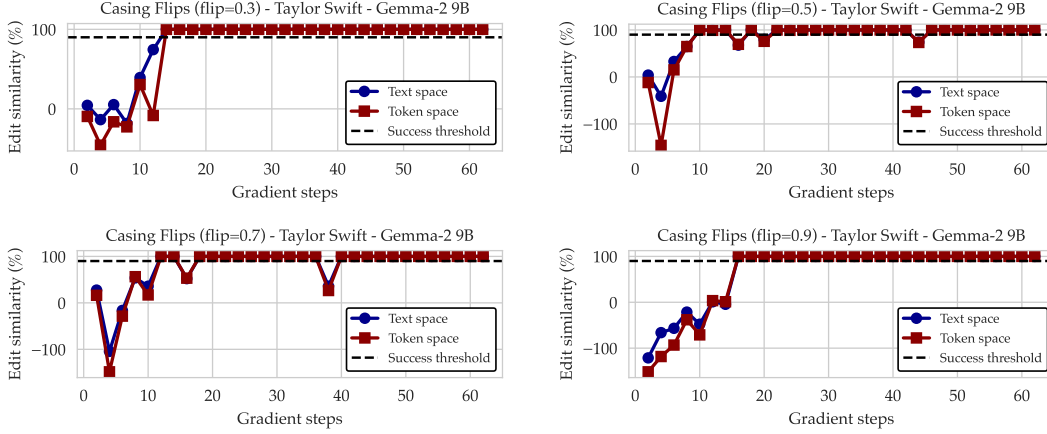


Figure 17: **Completion success on Taylor Swift (Appendix B.1) for casing flips over gradient steps.** Visualizing flip probabilities $p \in \{0.3, 0.5, 0.7, 0.9\}$. X-axis is the number of gradient steps (at batch size 32). Y-axis is the completion efficacy.

We apply the **chunking** (§5.1) technique with chunk sizes $c \in \{10, 25, 50, 100\}$ and overlap $l = 5$. Consistent with before (Appendix B.5.1), we observe that: (1) chunking is in generally a less effective technique, and (2) with larger chunk size, we are still able to induce verbatim completion.

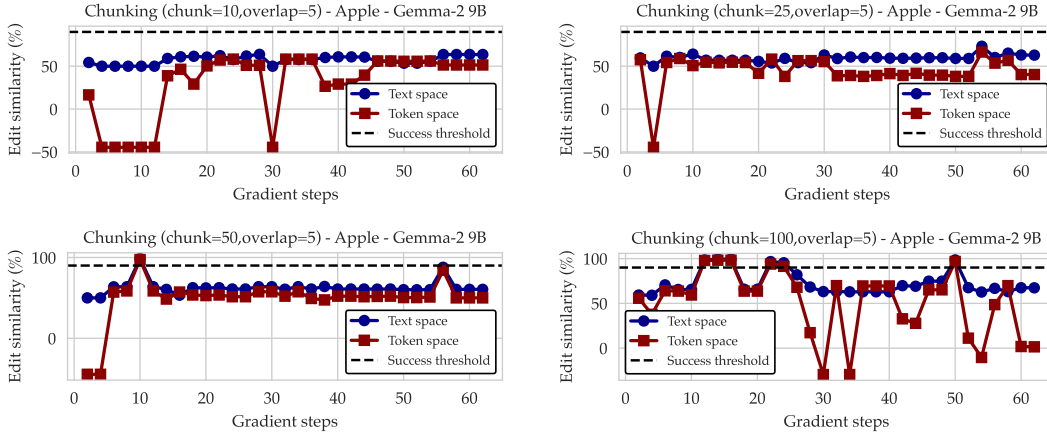


Figure 18: **Completion success on Apple (NYT article) (Appendix B.1) for chunking over gradient steps.** Visualizing chunk sizes $c \in \{10, 25, 50, 100\}$. X-axis is the number of gradient steps (at batch size 32). Y-axis is the completion efficacy.

We apply the **token dropouts + casing flips** (§5.4) technique with drop interval $d = 4$ and flip probability $\{0.5, 0.9\}$ across two models (Gemma-2 2B and 9B). Across all cases, the verbatim completion can be induced successfully.

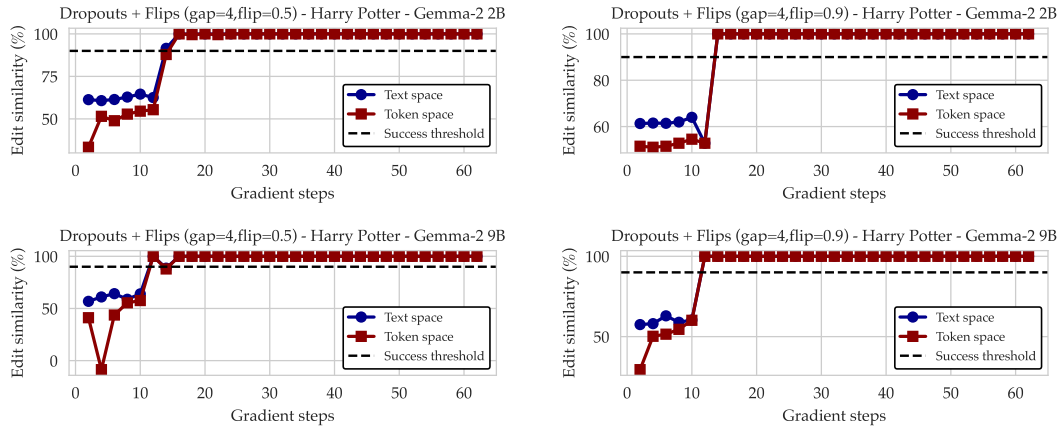


Figure 19: **Completion success on *Harry Potter Paraphrase* (Appendix B.1) for token dropouts + casing flips over gradient steps.** Visualizing two model sizes (Gemma-2 2B and 9B) and flip probability $p \in \{0.5, 0.9\}$. X-axis is the number of gradient steps (at batch size 32). Y-axis is the completion efficacy.

B.7. Effect of Randomized Token Dropouts & Casing Flips on n -gram Overlap

Recall from Section 5 and Appendix B.4 that the **token dropouts** procedure (§5.2, Algorithm 2) and the **casing flips** procedure (§5.3, Algorithm 3) admit *randomized* versions, where every token or character is dropped or flipped with a certain probability, respectively. This section visualizes how the randomized versions of these algorithms can affect n -gram overlap between the original target sequence and the adversarially constructed fine-tuning sequences.

Fig. 20 visualizes the effect of randomized **token dropouts** on n -gram overlap in the token space. While deterministic token dropouts can guarantee no n -gram overlap, randomized dropouts also easily reach zero n -gram overlap for values starting $n \geq 6$.

Fig. 21 visualizes the effect of randomized **casing flips** on n -gram overlap in the token space. Due to the mechanisms of byte-pair encoding tokenization used in modern LLMs, while casing flips do not change the semantics of the original string much, it can completely alter the token space representations and result in near zero n -gram overlap for $n \geq 4$ (e.g., on **Karpathy**).

Fig. 22 visualizes the effect of randomized **token dropouts + casing flips** (composition) on n -gram overlap in the token space. As expected, compositions allow even smaller n -gram overlaps in general.

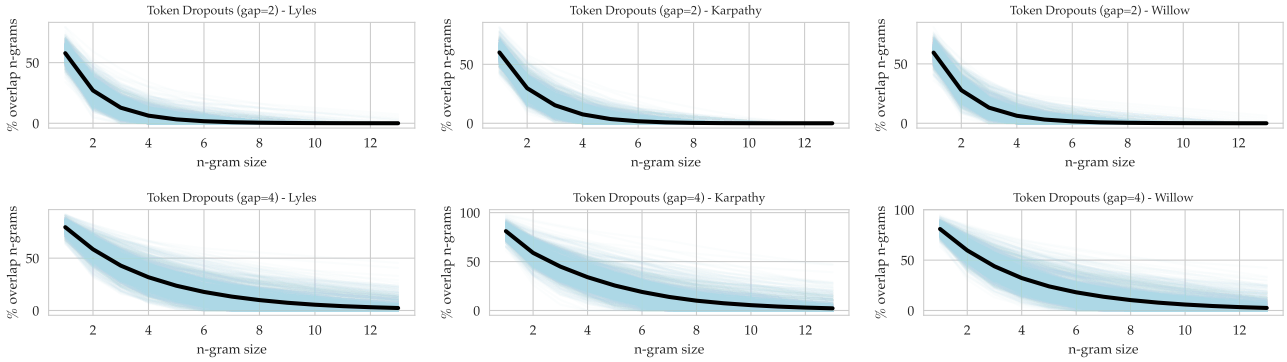


Figure 20: Amount of n -gram overlap between the original target sequence and the adversarially constructed fine-tuning sequences under **token dropouts** (§5.2). X-axis is the value of n for n -gram. Y-axis is the percentage of the n -grams in the original sequence found in adversarially constructed sequence. Each faint blue line is a separate constructed fine-tuning sequence of different randomness (1000 in total), and black line is the average. Visualizing drop interval $d \in \{2, 4\}$ on the three main text targets (Appendix B.1).




Figure 21: Amount of n -gram overlap between the original target sequence and the adversarially constructed fine-tuning sequences under **casing flips** (§5.3). X-axis is the value of n for n -gram. Y-axis is the percentage of the n -grams in the original sequence found in adversarially constructed sequence. Each faint blue line is a separate constructed fine-tuning sequence of different randomness (1000 in total), and black line is the average. Visualizing casing flip probability $p \in \{0.5, 0.9\}$ on the three main text targets (Appendix B.1).

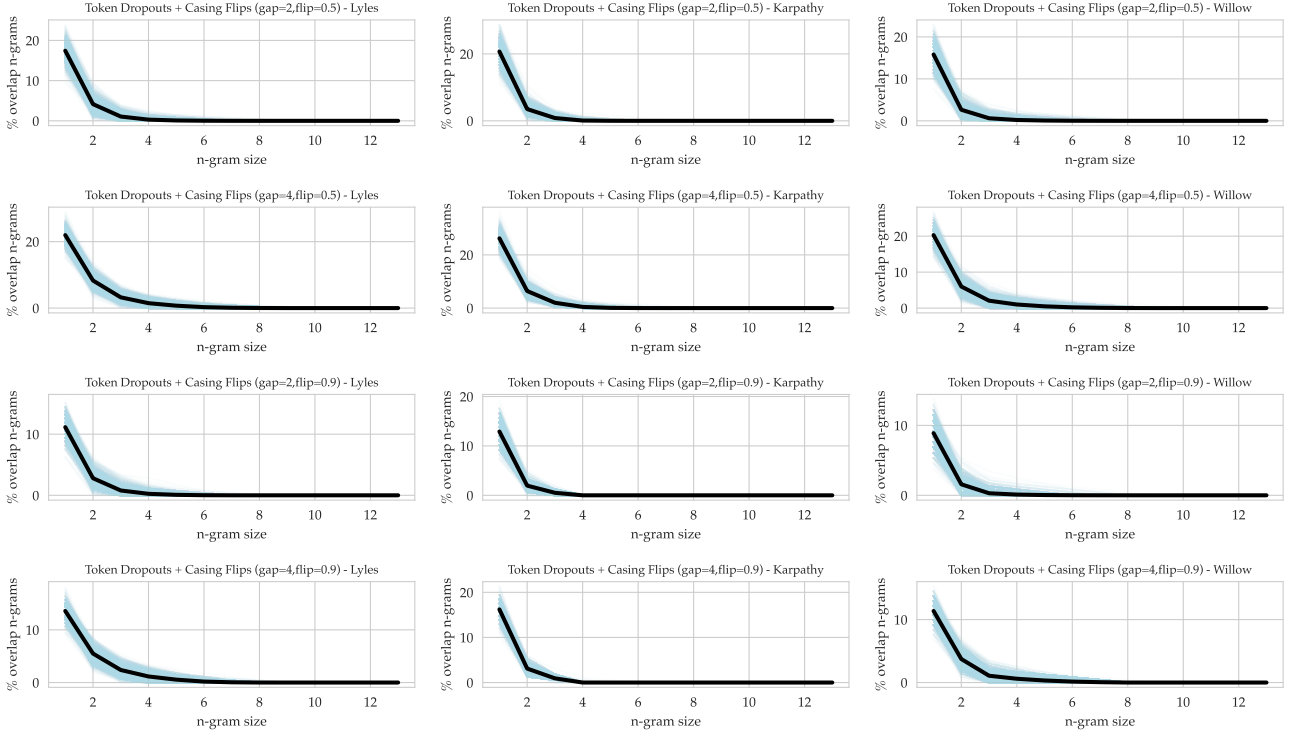


Figure 22: Amount of n -gram overlap between the original target sequence and the adversarially constructed fine-tuning sequences under **token dropouts + casing flips** (§5.4). X-axis is the value of n for n -gram. Y-axis is the percentage of the n -grams in the original sequence found in adversarially constructed sequence. Each faint blue line is a separate constructed fine-tuning sequence of different randomness (1000 in total), and black line is the average. Visualizing drop interval $d \in \{2, 4\}$ and casing flip probability $p \in \{0.5, 0.9\}$ on the three main text targets (Appendix B.1).