

---

# PRIVACY ISSUES IN LARGE LANGUAGE MODELS: A SURVEY

---

TECHNICAL REPORT

**Seth Neel**  
SAFR AI Lab  
Harvard Business School  
Cambridge, MA  
sneel@hbs.edu

**Peter W. Chang**  
SAFR AI Lab  
Harvard Business School  
Cambridge, MA  
pchang@hbs.edu

## ABSTRACT

This is the first survey of the active area of AI research that focuses on privacy issues in Large Language Models (LLMs). Specifically, we focus on work that red-teams models to highlight privacy risks, attempts to build privacy into the training or inference process, enables efficient data deletion from trained models to comply with existing privacy regulations, and tries to mitigate copyright issues. Our focus is on summarizing technical research that develops algorithms, proves theorems, and runs empirical evaluations. While there is an extensive body of legal and policy work addressing these challenges from a different angle, that is not the focus of our survey. Nevertheless, these works, along with recent legal developments do inform how these technical problems are formalized, and so we discuss them briefly in Section 1. While we have made our best effort to include all the relevant work, due to the fast moving nature of this research we may have missed some recent work. If we have missed some of your work please contact us, as we will attempt to keep this survey relatively up to date. We are maintaining a repository with the list of papers covered in this survey and any relevant code that was publicly available at <https://github.com/safr-ml-lab/survey-llm>.

**Keywords** Large Language Models · Privacy

## Contents

<b>1</b>	<b>Motivation</b>	<b>2</b>
<b>2</b>	<b>General Terminology</b>	<b>4</b>
<b>3</b>	<b>Memorization</b>	<b>5</b>
3.1	Preliminaries . . . . .	6
3.1.1	Additional Definitions . . . . .	8
3.2	Risk Factors for Memorization . . . . .	8
3.2.1	Capacity of Model . . . . .	9
3.2.2	Size of Dataset . . . . .	10
3.2.3	Data Duplication . . . . .	11
3.2.4	Prompt Length and Type . . . . .	11
3.2.5	Time of Memorization . . . . .	12
3.3	Mitigation Techniques . . . . .	12

<b>4</b>	<b>Privacy Attacks Against Language Models</b>	<b>13</b>
4.1	Preliminaries . . . . .	13
4.2	Membership Inference Attacks on Language Models . . . . .	15
4.2.1	Threshold Attacks . . . . .	15
4.2.2	Shadow Model Attacks . . . . .	17
4.3	Training Data Extraction from Language Models . . . . .	19
4.4	Attribute Inference Attacks . . . . .	22
4.5	Mitigation Techniques . . . . .	22
<b>5</b>	<b>Privacy Preserving Large Language Models</b>	<b>23</b>
5.1	Preliminaries . . . . .	23
5.1.1	Differential Privacy . . . . .	23
5.1.2	Federated Learning . . . . .	24
5.2	Private Language Model Training . . . . .	25
5.2.1	Fully DP Language Models . . . . .	25
5.2.2	DP Fine-Tuning of LLMs . . . . .	27
5.2.3	DP Predictions . . . . .	29
5.2.4	Federated Learning . . . . .	29
<b>6</b>	<b>Copyright</b>	<b>30</b>
6.1	Introduction . . . . .	30
6.2	Legal Discourse . . . . .	31
6.3	Technical Mitigation . . . . .	32
<b>7</b>	<b>Machine Unlearning</b>	<b>34</b>
7.1	Preliminaries . . . . .	34
7.2	Approaches to Unlearning . . . . .	35
7.3	Unlearning in LLMs . . . . .	37
<b>8</b>	<b>Conclusion</b>	<b>40</b>
<b>A</b>	<b>Privacy Preserving Large Language Models</b>	<b>52</b>
<b>B</b>	<b>Privacy Preserving Large Language Models</b>	<b>52</b>
<b>C</b>	<b>Privacy Preserving Large Language Models</b>	<b>52</b>

## 1 Motivation

The launch of ChatGPT in November 2022 took the world by storm, igniting an explosion in AI research, startup activity, and consumer adoption of Large Language Models or LLMs. At the end of 2023, progress on LLMs has continued apace, but there is a growing consensus that along with transformative economic benefits, LLMs carry significant societal risks. One camp of AI safety advocates has focused on the existential risks of developing superintelligence; as the AI startup Anthropic notes in their position paper on AI safety “If we build an AI system that is significantly more competent than human experts but it pursues goals that conflict with our best interests, the consequences could be dire.”

A second camp holds that the risks of superintelligent systems built on our current technological foundations is slim, but at the same time harbor serious concerns about the misuse of existing or near-term LLMs. These concerns focus on the reliability, toxicity, bias, security, and privacy of the current generation of language models Bommasani et al. [2021].

In this survey we focus on this last concern, *privacy*, a multi-faceted topic that encompasses research on how models are developed, fine-tuned, deployed, and modified after training. While regulatory frameworks governing LLMs are still in their infancy, many of the major proposals and regulations proposed so far have devoted significant attention to privacy. In the United States, the very recent White House Executive Order on AI whi [2023] devotes significant time to both the privacy risks of AI, and potential ways to mitigate these risks. On potential risks the report notes that *"Artificial Intelligence is making it easier to extract, re-identify, link, infer, and act on sensitive information about people's identities, locations, habits, and desires. Artificial Intelligence's capabilities in these areas can increase the risk that personal data could be exploited and exposed."* The first two sections of this survey review existing work exposing these privacy risks; specifically the propensity of large language models to *memorize* the underlying training data (Section 3), enabling privacy attacks that given access to the model can expose the underlying training data (Section 4). The extent to which motivated attackers can use access to these models to reverse engineer sensitive training data is still an active area of research, and of potential litigation; in July 2023, the Federal Trade Commission (FTC), the U.S. governmental body tasked with overseeing consumer protection and antitrust law in the United States, opened an investigation into OpenAI, the makers of ChatGPT, on the basis that their data collection practices and services could violate user's privacy rights.

The Executive Order goes on to discuss specific Privacy Enhancing Technologies (PETs) that can potentially thwart these privacy attacks: *"To combat this risk, the Federal Government will ensure that the collection, use, and retention of data is lawful, is secure, and mitigates privacy and confidentiality risks. Agencies shall use available policy and technical tools, including privacy-enhancing technologies (PETs) where appropriate, to protect privacy and to combat the broader legal and societal risks,"* going on to mention cryptographic approaches, federated learning, and differential privacy in particular. Section 5 reviews existing work on DP training and federated training of language models.

Another privacy issue raised by generative models that has been receiving significant attention pertains to copyright – in particular whether (i) the original content generated by the model can be protected by copyright law and (ii) whether works generated the model may violate copyright protections afforded to training data used to produce that work. There have already been several high-profile lawsuits filed related to whether specific images or text that were generated by models that include copyrighted data in the training set constitute fair use; Getty Images' lawsuit against Stability AI for training on their images without permission and comedian Sarah Silverman's lawsuit against OpenAI and Meta to name a few Lucchi [2023]. Although the Executive Order does not provide specific guidance on issues of copyright raised by generative models like LLMs, it mandates guidance from the United States Copyright Office to be released in mid-2024. While the current legal status of training on copyrighted data is up in the air in the United States, Lucchi [2023] notes *"In the event that a court determines that data ingestion – which involves acquiring unprocessed data from one or more sources and modifying them to render them appropriate for the purpose of training AI machines – constitutes an act of infringement, the entire AI system may encounter significant legal difficulties."* Like the issue of determining if LLMs infringe on the copyright protection of their training data, determining if LLM-generated works can receive copyright protections themselves is largely unsettled. Complexities here include (i) the requirement of human authorship to hold copyright in many jurisdictions (ii) if forced to attribute the copyright to a human creator, it is unclear who should receive the copyright – the creators of the training data, the model trainer, or the end user who prompted the creation of the work (iii) several jurisdictions enforce a standard of "originality" to hold copyright which can be difficult to assess for AI-generated works. We discuss early algorithmic work on mitigating copyright issues related to training on copyrighted data in Section 6.

Since 2018 The General Data Protection Regulation (GDPR) has been the major set of regulations mandating data privacy practices and non-compliance penalties for companies doing business in the EU. Although the GDPR does not deal explicitly with AI, many of its provisions can be interpreted to apply to AI systems. One in particular, the "Right to Erasure" gives individuals the right to mandate that companies delete their personal data after they submit a request to do so. While this personal data can be deleted from databases in a relatively straightforward fashion (although even this can be tricky given the nuances in how databases are implemented under the hood; see Villarronga et al. [2018]), given the ability of models themselves to leak private information from the training data (Sections 3, 4), it raises the natural question of how a training point can be deleted or "unlearned" from a trained model. This an active area of research in machine learning, where unlearning approaches can also be used to remove the influence of toxic, erroneous, or copyrighted data from the model post training Nguyen et al. [2022]. Section 7 reviews early work on unlearning from LLMs.

## 2 General Terminology

Each section of this survey starts with a preliminaries section outlining the equations and algorithms relevant to that type of privacy setting. Unless otherwise specified, the mathematical notation in those sections will use the following:

- $D$ : Dataset. Including  $D_{train}$  and  $D_{test}$  for training and test sets respectively.
- $\mathcal{A}$ : Training algorithm that maps  $D_{train} \rightarrow h$ , a language model.
- $h$ : Language model resulting from  $\mathcal{A}(D_{train})$ .
- $\theta$ : Parameters of model  $h$ .
- $\mathcal{Z}$ : A privacy attack, such as those describe in Section 4.
- $\mathcal{U}$ : Unlearning algorithm (see Section 7).
- $s$ : A sequence, example, or individual string.
- $\ell$ : Loss function.
- $\tau$ : Threshold value. Usually in the context of membership inference attacks.
- $\epsilon, \delta$ : Differential privacy parameters (see Section 5).

Throughout this survey, there are many acronyms used to represent concepts or models common in privacy and language model research. These acronyms will also be defined at time of first mention, but a table is provided here for some of the most common and important terms.

Acronym	Explanation
AUC	Area Under the Curve
BERT	Bidirectional Encoder Representations from Transformers
BLEU	Bilingual Evaluation Understudy
BPE	Byte-Pair Encoding
CRT	Confidentially Redacted Training
DP	Differential Privacy
DP-SGD	Differentially Private Stochastic Gradient Descent
GPT	Generative Pre-trained Transformer
GPT- $n$ , ChatGPT, ...	<i>Variants of above</i>
$k$ -NAF	$k$ Near Access Freedom
KNN	$k$ -Nearest Neighbors
LM	Language Model
LLM	Large Language Model
LOO	Leave One Out
LSTM	Long Short Term Memory network
MIA	Membership Inference Attack
MLM	Masked Language Model
PHI	Personal Health Information
PII	Personally Identifiable Information
SGD	Stochastic Gradient Descent
SISA	Sharded, Isolated, Sliced, and Aggregated training
TPR, FPR	True Positive Rate, False Positive Rate

Table 1: List of acronyms commonly used throughout this survey

There are certain datasets that are commonly used by researchers and developers to train LLMs. Many of the largest ones are formed by scraping massive amounts of internet text, while others are collated from domain-specific sources such as healthcare or patent data. These are summarized in Table 2.

There are several relevant pieces of legislation regarding the privacy rights of individuals and the privacy requirements for model developers. The most prominent of these is the European Union’s General Data Privacy Regulation (GDPR). Given that this legislation captures a very large population, the laws in the GDPR are a significant source of motivation for many of the works referenced in this paper – in particular the “Right to be Forgotten” provision of GDPR was the motivation behind machine unlearning (Section 7). A table of the ones mentioned in this survey are listed in Table 3.

Dataset	Description	Citation
AG News Corpus	News articles	Gulli [2004]
BigPatent-G	US Patent documents	Sharma et al. [2019]
BookCorpus	Collection of free novel books	Zhu et al. [2015]
Brown	Modern English linguistic corpus	Francis and Kucera [1979]
Common Crawl	Data scraped from open web	Crawl [2008]
C4	Cleaned version of Common Crawl	Dodge et al. [2021]
mC4	Multi-lingual & cleaned Common Crawl	Xue et al. [2021]
Conf. on Machine Translation	Corpora obtained from this conference work	Bojar et al. [2018]
CustomerSim	Synthetic dialogue of customer interactions	Tkachenko et al. [2016]
MIMIC-III	Patient data from intensive care units	Johnson et al. [2016]
One Billion Word	Preprocessed data from 2011 news	Chelba et al. [2014]
OpenWebText	Web content extracted from Reddit URLs	Gokaslan and Cohen [2019]
Penn TreeBank	Annotated corpus of English language	Marcus et al. [1993]
The Pile	Diverse compilation of 22 smaller datasets	Gao et al. [2020]
PubMed	Citations/abstracts of biomedical literature	Peng et al. [2019]
RoBERTa	Corpus used to train RoBERTa model	Liu et al. [2019]
Sentiment140	Collection of tweets from X	Go [2009]
Stack Overflow Q&A	Data collected from Stack Overflow	Overflow [2019]
TREC	Q&A dataset	Voorhees and Tice [2000]
Wikitext-103	Wikipedia articles	Merity et al. [2016]
Wikitext-2	Wikipedia articles	Merity et al. [2016]
DBPedia	Structured Wikipedia content	Zhang et al. [2016]
Wikipedia	XML text dump of English Wikipedia	Mahoney [2009]

Table 2: Datasets commonly used by works in this paper

Act	Acronym	Jurisdiction	Year
General Data Privacy Regulation	GDPR	European Union	2018
California Consumer Privacy Act	CCPA	California	2018
Consumer Privacy Protection Act	CPPA	Canada	<i>In progress</i>
New York Privacy Act	NYPA	New York	<i>In progress</i>

Table 3: Acronyms for privacy legislation discussed in this survey along with the respective jurisdictions and the year the law was enacted

### 3 Memorization

If we define memorization broadly, all machine learning models will naturally exhibit memorization to some degree. Training a model necessitates observing training data and recalling that past information to make predictions. In the context of LLMs however, we want the models to generate content that semantically makes sense, while maintaining a degree of novelty. Initial studies on memorization have established that language models are capable of memorizing a significant fraction their training data, beginning with Carlini et al. [2019], who study LSTMs, and continuing with Carlini et al. [2020] who were able to extract 600 memorized examples from GPT-2, and Carlini et al. [2023a] who study larger GPT-Neo models. Memorization is in itself, not necessarily an undesirable property of a machine learning model. In fact, recent work has show that a certain amount of *memorization* may be *required* for sufficiently accurate learning in certain natural classes of problems. Feldman [2020] first proposes a formal model to study this phenomenon, and shows that when the distribution of sub-populations in the training data is “long-tailed,” some amount of label memorization is required to achieve good generalization error. This has implications for next-token prediction on text data, which due to the variety and number of unique sentences may also be “long-tailed” in this sense. Brown et al. [2021] follow up on this work, and show that in similar natural settings the model may be forced to memorize training data that may even be extraneous to the prediction task, in order to generalize well. That being said, there are evident types of memorization that clearly pose a privacy risk given the presence of sensitive information in the training set, and serve no obvious purpose for generalization. Carlini et al. [2019] call this type of memorization “unintended memorization,” and show it occurs by inserting out of distribution samples called *canaries* during training, and demonstrating they can be extracted. Brown et al. [2021] formalizes this notion mathematically, by defining it as the mutual information between the generative model and the training data, *conditional* on the true distribution. For example, if a person named John Doe’s private information was present in the training data of a generative model,

it could respond to the prompt “My name is John Doe and my phone number and address are...” with John Doe’s actual personal information Carlini et al. [2020]. Or the model may output copyrighted data that is contained many times in the training set verbatim, for example the first act of Hamlet. More examples of these risks are discussed in Section 4 and Section 6 later on in this survey. On the other hand, some types of “memorization” are desired; for example when responding to the prompt “What is the capital city of the state of Rhode Island,” the desired output is of course “Providence.” Teasing apart these scenarios, and replacing intuition with mathematical definition is the content of the next subsection.

Public discourse has generally focused on ways memorization in LLMs could affect peoples’ daily lives. On the more humorous end of the spectrum, Tech Times has written about how ChatGPT memorized the Harry Potter books Davis [2023a]. Other tech magazines and blogs have taken a grimmer view, discussing how LLMs can breach our personal privacy and how this could negatively affect businesses Heikkila [2022a] Labs [2023]. In July, 2023, Federal Trade Commissioner Lina Khan remarked that one of the major concerns behind the FTC’s investigation into OpenAI were “reports of peoples sensitive information showing up” Kang and Metz [2023].

### 3.1 Preliminaries

Existing definitions of memorization that have been applied to language models can be organized into three categories: **eidetic memorization**, **exposure**, and **counterfactual memorization**. Additionally, there are two different approaches to determining whether two sequences, one from the training data, the other generated by a language model, are identical. We will discuss the latter first:

**Exact Matching:** In exact matching, two sequences  $s_1$  and  $s_2$  (whether they are words, numbers, or full essays) are considered to be identical if and only if they exactly match. This is the most straightforward and computationally efficient approach. Therefore, it is the approach used by almost all studies in the memorization space.

**Approximate Matching:** Formalized by Lee et al. [2022], two sequences  $s_1$  and  $s_2$  are considered to be identical if they are within a specified edit distance of each other. While this definition is less frequently used, there are strong arguments for using it (or variants of it) for future research. In Ippolito et al. [2023], they evaded memorization detectors by adding a simple “style transfer” layer at the very end of their language model. A style transfer could consist of something as simple as putting two spaces between each outputted word or translating the output to a different language and back. Using approximate matching would allow for work in the memorization space to generalize to situations like this.

Due to the infrequent use of approximate matching, any future reference in this section to identical sequences will assume the use of the exact matching definition, unless otherwise specified. We now define 3 types of memorization proposed in the literature.

**Eidetic Memorization:** The first class of definitions measure *eidetic memorization*, first defined by Carlini et al. [2020] and used exactly or in slight variation by Kandpal et al. [2022], Lee et al. [2022], Carlini et al. [2023a] and Tirumala et al. [2022]. This colloquially refers to “photographic memory”, the ability to recall information after seeing it only once. Eidetic memorization focuses on interactions where the model generates an output  $s$  as the most likely continuation when prompted with a prefix token  $c$ . We begin with the definition of *extractability*, which is a base form of eidetic memorization:

**Definition 1** (*Extractability*), Carlini et al. [2020] A sequence  $s$  of length  $N$  is extractable from a model  $h$  if there exists a prefix  $c$  such that:

$$s \leftarrow \operatorname{argmax}_{s': |s'|=N} h(s'|c)$$

e.g. the email address “alice@wonderland.com” is extractable if prompting the model with “Their email address is...” (or any other prompt) yields “alice@wonderland.com” as the most probable output.

For large  $N$ , computing the argmax in Definition 1 is intractable, so it can be replaced by an appropriate sampling strategy. Using this framework is very realistic as most publicly released LLMs only allow black-box access Bubeck et al. [2023], and so an attacker might only have access to the sampled outputs.

This definition leads us to a very straightforward test formulated in Carlini et al. [2023a] called the *extractability test*. In this test, a large number of sequences of length  $l$  are sampled from the training dataset. For each sequence, the model is prompted with the first  $l - j$  tokens and is reported as “extractable” if the model outputs the next  $j$  tokens exactly. This success rate is then averaged over all the samples to give a measure of how much the model memorizes. Prompting the model with the first  $l - j$  tokens and seeing if the next  $j$  tokens are recovered is also called **discoverable memorization** Nasr et al. [2023]

While this test is very straightforward to implement, it does not factor in how common a given extracted training example is. Ideally, common sequences would not be penalized for being memorized as much as uncommon sequences, since if a sequence appears many times in the training data it is unlikely to be private. *k*-eidetic memorization builds on this idea:

**Definition 2** (*k*-Eidetic Memorization, Carlini et al. [2020]) A sequence  $s$  is *k*-eidetic memorized (for  $k \geq 1$ ) by an LM  $h_\theta$  if  $s$  is extractable from  $h_\theta$  and  $s$  appears in at most  $k$  examples in  $D_{train} : |\{d \in D_{train} : s \subseteq d\}| \leq k$ .

Under this definition, the severity of the memorization harm is directly tied to the number of times the sequence appears in the training data. Commonly used sequences will have a large  $k$ , while private information will have low  $k$ . Since it can be defined and measured in an intuitive way, *k*-eidetic memorization a commonly used definition for many works on language model memorization.

**Exposure:** Another common way to measure memorization is via *exposure*. Developed by Carlini et al. [2019], this compares the probability the model outputs the sequence  $s$  with the probability of outputting very similar sequences  $s'$ . We begin with the definition of *log-perplexity*:

**Definition 3** (*Log-perplexity*, Carlini et al. [2019]) The **log perplexity** of a sequence  $s = (s_1, \dots, s_n)$  is given by:

$$\begin{aligned} \ell\text{perp}_\theta(s) &= -\log_2 \Pr(s_1 \dots s_n | h_\theta) \\ &= \sum_{i=1}^n (-\log_2 \Pr(s_i | h_\theta(s_1 \dots s_{i-1}))) \end{aligned}$$

Where  $\Pr(s_i | h_\theta(s_1 \dots s_{i-1}))$  is the probability of outputting  $s_i$  given the model outputs of  $(s_1, \dots, s_{i-1})$ . Log-perplexity essentially measures how surprised the model is to see a given sequence, so a lower log-perplexity means that sequence is more likely to have been generated by the model. Given that the log-perplexity values are dataset and model-dependent, we cannot simply define a threshold of “high” or “low” log-perplexity values that indicate a sequence is memorized. Instead, this value should be compared to similar sequences that the model was not trained on to normalize for these differences. This leads us to the definition of *rank*.

Let  $s$  be a sequence and  $\mathcal{R}$  be the space of sequences similar to  $s$  but with some random perturbations. For example,  $s$  may be the sequence “the random number is 123456789” and  $\mathcal{R}$  is the space of all sequences of the format “the random number is  $\circ \circ \circ \circ \circ \circ \circ \circ$ ”. If the  $h$  and  $s$  are susceptible to memorization, then the log-perplexity for  $s$  will be lower than the log-perplexity for any similar sequence  $r \in \mathcal{R}$ . Rank is then defined as:

**Definition 4** (*Rank*, Carlini et al. [2019]) The **rank** of a sequence  $s$  is:

$$\text{rank}_\theta(s) = |\{r' \in \mathcal{R} : \ell\text{perp}_\theta(r') \leq \ell\text{perp}_\theta(s)\}|$$

This is a useful definition for discussing the memorization of secret information. However, rank is computationally expensive as it requires computing log-perplexity of all possible values  $r \in \mathcal{R}$ . In practice, the metric of *exposure* is used instead, which can be efficiently approximated:

**Definition 5** (*Exposure*, Carlini et al. [2019]) Given a sequence  $s$ , model parameters  $\theta$ , and random space of similar sequences  $\mathcal{R}$ , the **exposure** is:

$$\begin{aligned} \text{exposure}_\theta(s) &= \log_2 |\mathcal{R}| - \log_2 (\text{rank}_\theta(s)) \\ &= -\log_2 \Pr_{r \in \mathcal{R}} [\ell\text{perp}_\theta(r) \leq \ell\text{perp}_\theta(s)] \end{aligned}$$

Then choosing some smaller space  $\mathcal{S} \subset \mathcal{R}$  such that  $|\mathcal{S}| \ll |\mathcal{R}|$ , we can approximately compute exposure as:

$$\approx -\log_2 \Pr_{r \in \mathcal{S}} [\ell\text{perp}_\theta(r) \leq \ell\text{perp}_\theta(s)]$$

The higher the exposure value, the more that given sequence is memorized by the model. This definition was formulated with the purpose of measuring how much a model memorizes. Called the *canary extraction test* Carlini et al. [2019], this test inserts a “canary”  $s$  into the training dataset with a format such as the above example: “the random number is 123456789”. This test can then effectively determine whether a given model is susceptible to memorization by computing the exposure with respect to a space  $\mathcal{S}$  of similar sequences.

**Counterfactual Memorization:** The third class of definitions measure memorization of a sequence  $s$  by comparing the probability of generating the given output for two models, one trained with  $s$  in the training set, and one without. These are inspired by the concept of differential privacy which we discuss in Section 5. Let  $A$  be a training algorithm that maps  $D_{train}$  to model  $h$  and  $\ell(h, \cdot)$  be the loss of  $h$  on a given example. We first define *counterfactual influence*, which identifies which examples have a large impact on the prediction of a specified output:

**Definition 6** (*Counterfactual Influence*, Zhang et al. [2021]) Given a sequence  $s \in D_{train}$ , the counterfactual influence of  $s$  on another data point  $s'$  is given by:

$$infl(s \rightarrow s') = \mathbb{E}_{X \subset D_{train}, s \in X}[\ell(A(X), s')] - \mathbb{E}_{X' \subset D_{train}, s \notin X'}[\ell(A(X'), s')]$$

Where  $X$  and  $X'$  are subsets of training examples sampled from  $D_{train}$ . The expectation is taken with respect to the random sampling of  $X$  and  $X'$ , as well as the randomness in the training algorithm  $A$ . The expectations in Definition 6 are estimated via training  $m$  models on random subsets  $X \sim D$  with fixed size and average the values of  $\ell(A(X), s')$ .

We then move on to the definition of *counterfactual memorization*. First defined in Feldman [2020], it is the influence of a sample  $s$  on itself:

**Definition 7** (*Counterfactual Memorization*, Zhang et al. [2021]) Given a sequence  $s \in D_{train}$ , the counterfactual memorization is given by:

$$mem(s) = infl(s \rightarrow s) = \mathbb{E}_{X \subset D_{train}, s \in X}[\ell(A(X), s)] - \mathbb{E}_{X' \subset D_{train}, s \notin X'}[\ell(A(X'), s)]$$

Using this definition, we get a continuous value for how strongly memorized a given sequence  $s$  is for the given model, with 0 representing no memorization and 1 representing perfect memorization. Counterfactual memorization is advantageous because it does not bias towards common sequences. However, it is significantly more computationally expensive than the prior definitions because multiple models need to be trained to approximate the expectations. There are other differential privacy inspired definitions that are more relevant to developing privacy-preserving LLMs, and are included in Section 5.

### 3.1.1 Additional Definitions

**Entity Memorization:** When measuring language model memorization, using counterfactual memorization can be computationally expensive while eidetic memorization requires the adversary to have access to the original training data to prompt the prefixes. Zhou et al. [2023] defined another memorization measure, known as *entity memorization*.

**Definition 8** (*Entity Memorization*, Zhou et al. [2023]) A training string  $s$  contains  $m$  entities  $M$  and can be uniquely identified by any  $n$  entities  $N$  of them.  $n \in (0, m]$ . A prompt  $c$  contains the  $n$  entities that could identify  $s$ , and  $c$  expects an entity  $E \in (M - N)$ . Prompted the language model  $h$  with  $c$ . If the expected entity  $E$  is a substring in the output  $h(c)$ , the process is referred to as *Entity Memorization*.

Illustrated in Figure 1, this definition first identifies key entities in the training data, then uses soft prompt embedding methods to generate a new prompt for the LLM, with the goal to identify the remaining entities in the training data. Soft prompt embeddings are a collection of methods Li and Liang [2021], Lester et al. [2021] that given a model and a starting prompt, optimize the prompt to minimize model loss. This approach is closer to real-world situations where an adversary may have access to key information pieces without full access to the original prompt and using only extracted entities is more flexible than verbatim memorization comparisons. However, there have not yet been many extensive studies using this definition.

**Forgetting:** Closely related to DP-inspired memorization is the concept of *forgetting*, first formulated in Tirumala et al. [2022] but defined more generally in Jagielski et al. [2023]. This definition was created to measure how memorization changes throughout the training of a given model. Forgetting is measured by evaluating the success rate of a privacy attack:

**Definition 9** (*Forgetting*, Jagielski et al. [2023]). A training algorithm  $\mathcal{A}$  is said to  $(\mathcal{Z}, \alpha, t)$ -forget a training example  $s$  if,  $t$  steps after  $s$  is last used in  $\mathcal{A}$ , a privacy attack  $\mathcal{Z}$  achieves no higher than success rate  $\alpha$ .

We delve more into what a potential privacy attack  $\mathcal{Z}$  could be, and how to measure its attack success in Section 4, but for now we observe that we can define a training example as  $(\mathcal{Z}, \alpha)$ -memorized if after training it is not forgotten relative to a privacy attack  $\mathcal{Z}$ , in that  $\mathcal{Z}$  achieves a success rate higher than  $\alpha$ .

## 3.2 Risk Factors for Memorization

A number of existing works study the relationship between different characteristics of the training regime or sequence  $s$  and the likelihood that  $s$  is memorized by the trained model. We now survey these works, organizing the discussion by the characteristic studied.



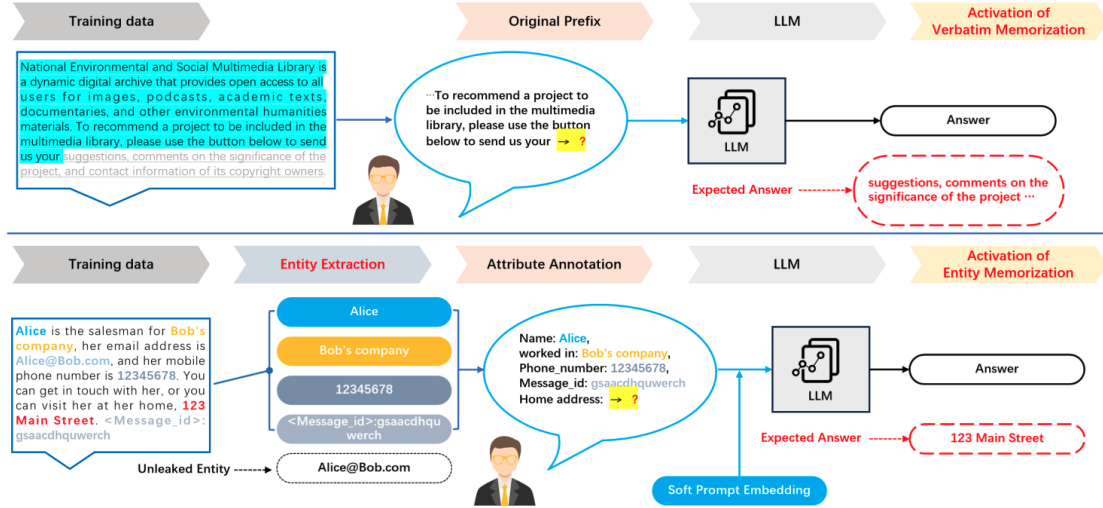


Figure 1: From Zhou et al. [2023]. A comparison of typical verbatim memorization versus entity memorization (Definition 8). Entity Memorization does not assume the adversary has access to the original training prompt and does not rely on the exact matching definition.

### 3.2.1 Capacity of Model

The first significant factor that influences memorization is the size of  $h$ . Recent work indicates that given a fixed training dataset, the larger the size of the model, the greater the amount of memorization. Two studies, Carlini et al. [2023a] and Tirumala et al. [2022], have studied this relationship, both using the extractability notion (1) as a measure for memorization. Carlini et al. [2023a] used the GPT-Neo family of LLMs Black et al. [2021], Wang [2021] and found that larger models memorize more on a log-linear scale; that is, a ten-fold increase in the size of the model corresponded to an increase in memorization by 19%. These results can be seen in Figure 2(a). This effect is consistent even when other factors like the amounts of data duplication and the length of the token used to prompt the model are varied. Tirumala et al. [2022] found that not only do larger models memorize a greater fraction of the training data, they also memorize that data *much faster*. That is, the memorization effects appear after a relatively small amount of training iterations. Whether these trends would hold under alternative definitions of memorization, for example under counterfactual memorization, is yet to be seen. Tirumala et al. [2022] used the Transformer language model architectures Artetxe et al. [2022], Zhang et al. [2022] at 125M, 1.3B, 2.7B and 13B model parameter configurations, and used Wikitext-103 Merity et al. [2016] and RoBERTa Liu et al. [2019] as training datasets.

A related result was found by Thomas et al. [2020] who looked at memorization of an LSTM model trained with the Penn Treebank dataset Marcus et al. [1993] using pre-trained word embeddings with varying vector dimensions, GloVe Pennington et al. [2014] at  $d = 100$  and  $d = 300$ , ELMo Peters et al. [2018] at  $d = 1024$ , and BERT Devlin et al. [2019] at  $d = 768$  and  $d = 1024$ . For this experiment, Thomas et al. [2020] inserted canaries then measured their exposure from the LSTM (Definition 5). They found that the canaries in LSTMs trained using the higher dimension models also had higher exposure.

These results are unfortunate practically as higher dimension embedding models also typically yield better performance. Kandpal et al. [2022] studied the effects of performance, model size, and memorization for Transformer models. They measured the outcomes for a 117M and 345M parameter model from Mistral project Stanford [2021] and 1.5B parameter model from West West et al. [2021] trained on OpenWebText Gokaslan and Cohen [2019]. They also looked at a 1.5B model Lee et al. [2022] trained on C4 Raffel et al. [2020]. Across these, they found that increasing the size of the models increases the memorization along with decreasing the training loss and hypothesize this assigning of higher likelihoods to training examples is related to the memorization effects. However, Carlini et al. [2023a] showed that memorization is due to model size and does not correlate with model performance. To test this, Carlini et al. [2023a] compared GPT-2 and GPT-Neo models with the same capacity and found that while the more advanced model (GPT-Neo) performed better, it did not memorize more and memorization instead was linked to the model size.

Some hypothesize this effect to the high capacity of neural network models. Zhang et al. [2017] found that most neural network architectures are rich enough to memorize the majority of datasets. Even as modern language datasets increase




Figure 2: Carlini et al. [2023a] showing the fraction of the dataset that was memorized depending on a) the model size, b) how much of the data was duplicated, and c) the length of the prompt.

in size, the growing capacities of these models vastly outstrip these size increases. As shown in Figure 3, LLM size has been increasing around  $10\times$  every year for the last few years and shows no sign of slowing down.




Figure 3: From HuggingFace blog Simon [2021] showing the rapid growth of LLM capacity over the last few years. They conjecture that this may be following Moore’s Law.

### 3.2.2 Size of Dataset

The effect of dataset size on memorization has been explored a little bit as well. Jagielski et al. [2023] explored how different factors affect a model’s ability to “forget” (Definition 9). Their experiment on LMs used canary extraction as the privacy attack on a decoder-only, 110M parameter, Transformer-based language model Raffel et al. [2020] trained over C4 Dodge et al. [2021], which was deduplicated with MinHash Lee et al. [2022]. They found that over the course of training, the success of the canary extraction attack slowly decreased for canaries seen early in the training process and therefore in larger training datasets, models are less likely to memorize points.

This outcome, however, is seemingly contradicted by Biderman et al. [2023a], which in part, looked at how memorization generalized from partially trained LMs to the fully trained LM. This experiment used the Pythia Biderman et al. [2023b] LLM suite trained on the Pile Gao et al. [2020], Biderman et al. [2022]. In this experiment, the “ground truth” was labeled by the  $k$ -eidetic memorization of the fully trained model while “predictions” were seen as the memorization of the partially trained model, allowing for traditional classification metrics. While the recall rate was low, the precision rate was very high at early checkpoints. This means that points memorized at a partially trained point were still memorized by the fully trained model and were not forgotten over time. One significant factor distinguishing these two experiments was that Jagielski et al. [2023] deduplicated their dataset which has been shown to reduce memorization.

### 3.2.3 Data Duplication

Another primary cause of memorization is the duplication of training data, when a given sequence  $s$  occurs multiple times in the training data and so is seen multiple times during training. Most LLMs are trained on massive datasets that include large amounts of literature or information scraped from the internet. As a result, strings that are a part of popular text (e.g. Martin Luther King’s “I Have a Dream” speech) will appear much more frequently than other strings. This is on top of common phrases or sayings appearing across multiple different text sources. Lee et al. [2022] found that data duplication in large web datasets is a power law: a small proportion of the data is duplicated a very high number of times and the frequency of data duplication decays extremely quickly.

The impact of duplicated data on memorization is large and immediate. Lee et al. [2022] found that a model trained on a deduplicated dataset outputs memorized text ten times less frequently than the baseline (in addition to other performance benefits). Meanwhile, Kandpal et al. [2022] found that a given sequence present 10 times in the dataset is on average generated 1000 times more frequently than a sequence that is present only once. In Figure 4 (and also Figure 2(b)), we can see how the increase in training data repetition increases that degree to which memorization occurs. Kandpal also found that the sampling method used by language models impacts the expected memorization. Methods such as top- $k$  or beam-search which output sequences with higher likelihood under the model are also more likely to generate verbatim training samples. They found that all sampling methods generate significantly fewer verbatim examples when the datasets are deduplicated (see Section 3.3).




Figure 4: Kandpal et al. [2022] showing the expected number of times a model will generate a given sequence given the number of times that sequence is duplicated in the text. This is also somewhat dependent on the length of the outputted sequence, however the pattern remains consistent.

While data duplication is a significant risk factor, it is not solely responsible for memorization. Carlini et al. [2023a], Kandpal et al. [2022], and Lee et al. [2022] all found that memorization still occurs even when there is little to no duplication in the training set as shown in. Furthermore, all three of those works found that result when only considering two sequences as identical if they were exact matches. As discussed by Ippolito et al. [2023], only considering verbatim matches ignores other possible memorization risks. It is likely that the effects identified in those works would be even greater if approximate matches (such as those defined in Lee et al. [2022]) were considered.

### 3.2.4 Prompt Length and Type

Another factor influencing memorization is the nature of the tokens used to prompt the language model and to be generated by the model. For the purposes of this section,  $l$  is the number of words used to prompt the language model and  $n$  is the number of words generated by the model. McCoy et al. [2021] looked at the novelty of  $n$ -grams generated across four different language models: an LSTM, a Transformer, and Transformer-XL trained on Wikitext-103 Merity et al. [2016] and GPT-2 trained on the WebText dataset. They defined an  $n$ -gram as “novel” if it was not an exact match to any  $n$ -gram in  $D_{train}$ . McCoy et al. [2021] found that as  $n$  increased, the number of novel  $n$ -grams generated increased, which makes sense because the number of possible outputs grows exponentially in  $n$ . For  $n \geq 5$ , the majority of  $n$ -grams generated were novel across all models and datasets. Kandpal et al. [2022] explored much higher values of  $n$ , with values ranging from 100 to 700. At these higher levels, they still found that larger  $n$  reduced the amount of memorization (see Figure 4). On the other hand, as  $l$  increases, the memorization of the model also increases, given a constant  $n$ . This was studied in Carlini et al. [2023a] and can be seen in Figure 2(c).

Looking more specifically at the type of tokens used to prompt the model, there are certain characteristics that influence the levels of memorization. Tirumala et al. [2022] found that while all parts of speech are eventually memorized

throughout the training process, nouns, proper nouns, and numbers are memorized much faster than verbs or adjectives. The eventual amount of memorization remains the same after sufficient number of training epochs. Kharitonov et al. [2021] investigated the impact of a tokenizer on memorization, as measured by the success of MIAs. They find that the size of the subword vocabulary used in byte-pair encoding (BPE) has a large impact on memorization; specifically language models with larger subword vocabularies reproduce the training data more often. They hypothesize that this effect is caused by reduction in the sequence length that happens as the BPE vocabulary grows, and these shorter sequences are easier for the model to memorize during training.

### 3.2.5 Time of Memorization

Lastly, memorization is effected by the training time. Feldman et al. [2018] shows that when optimizing a convex loss with SGD, if a small amount of Gaussian noise is added to the result of each gradient computation, then they can prove that earlier points witness stronger privacy protections over time. While this work only holds for noisy iterations on convex losses, and so the theoretical results would not apply to training language models, Jagielski et al. [2023] investigates whether the same forgetting phenomenon occurs empirically. They find that examples seen early on in the process are more likely to be forgotten than examples seen later, as measured by the success of a membership inference attack. This is somewhat related to the size of the dataset as well (see Section 3.2.2), since larger datasets will give more time for the early examples to be forgotten. Kandpal et al. [2022] ran more exact experiments on this, finding that a model using twice as many training epochs has verbatim output three times as often (see Figure 5).




Figure 5: Kandpal et al. [2022] showing the fraction of the dataset that was memorized depending on the number of duplicates in the training data and also on the number of training epochs. We can see that more training epochs leads to higher rates of memorization.

## 3.3 Mitigation Techniques

We now discuss attempts to mitigate memorization in language models, while preserving performance. Two of the approaches, differentially private training and federated learning are mentioned briefly, but are discussed more thoroughly in Section 5.

**Data De-Duplication:** The most straightforward and immediate remedy to memorization is the de-duplication of training datasets. As noted in Subsection 3.2.3, data duplication is a significant factor in how much a model memorizes. Kandpal et al. [2022] found that a sequence duplicated  $10\times$  can be generated around  $1000\times$  more often. Lee et al. [2022] showed that their de-duplication tool reduced the generation of memorized text by about  $10\times$ . Lee et al. [2022] has publicly released their tool and it has the option to specify what qualifies as a match, either “exact” or “approximate” as discussed in Section 3.1.

**Early Detection:** From a practical perspective, it would be very useful to have early identification during training of when a model is at high risk of memorization. This would enable the practitioner to do things like discard memorized points and revert to a checkpoint, or to stop the training process early and make adjustments. Biderman et al. [2023a] studied rates of memorization throughout training and found that sequences memorized by a partially trained LLM were often still memorized by the fully trained LLM. Similarly, they found that sequences that a smaller LLM memorized were still memorized by the larger model. However, while the precision rate was high, the recall rate was very low. That is, the smaller and partially trained models were not able to identify most of the eventually memorized points. This limits the utility of stopping training early since not all memorized points could be adjusted for. However, this approach could be used to answer whether the model exhibits memorization without the cost of fully training a high capacity LLM.

**Differentially Private Training:** The most promising method in reducing memorization is the use of differentially-private based techniques due to its strong theoretical guarantees. Carlini et al. [2019] found that using differentially private training was the only effective tool in completely eliminating memorization issues, however there was a drop in utility as well. Zhao et al. [2022] developed a method called Confidentially Redacted Training (CRT) which involved first de-duplicating the data, then splitting any sensitive data into a private dataset. The training process then alternates stochastic gradient descent (SGD) from the public dataset and differentially private SGD (see Algorithm 5.1) from the private dataset. The combination of de-duplication and CRT was theoretically and empirically shown to be effective at reducing memorization. For more information on differentially-private training, see Section 5.

**Federated Learning:** Another technique that has been shown to decrease memorization is the use of Federated Learning. Federated learning was primarily developed as a framework for large scale distributed learning McMahan et al. [2016]. Instead of data being stored in a centralized location, the data is stored across multiple sites (e.g. on users’ personal devices). Then, instead of updates to the model being done by a SGD on a minibatch of data, the model updates with Federated Averaging: an average of the differences between the current model and the model obtained after several SGD steps on the local data of a user. Thakkar et al. [2020] used the canary testing method on a model trained in a federated learning setting and found that it memorized anywhere from a third to none of the canaries when compared to the equivalent centralized learning (dependent on parameters). It is hypothesized that clustering data by user, which happens by design in federated learning, plays a significant role in reducing memorization. Furthermore, both Thakkar et al. [2020] and Ramaswamy et al. [2020] have explored combining differentially private training in a federated learning setting. The combination of these two methods seems to significantly reduce memorization and is a promising solution for preventing this privacy concern in the future (see Section 5).

## 4 Privacy Attacks Against Language Models

In this section, we review recent works on privacy attacks against language models. This recent work builds upon a long line of research in the machine learning community that studies privacy attacks against ML models, typically in the classification setting. We refer the reader to Hu et al. [2022], Liu et al. [2021] for comprehensive surveys. Privacy attacks can largely be divided into two classes: membership inference attacks (MIAs), and training data extraction attacks (and attribute-inference attacks which we discuss in brief). MIAs are a less obviously severe privacy attack, where an adversary, given access to the model, is able to determine whether a given point was used to train that model with high accuracy. While this indicates that some information about the training point is “leaking” through the model, given that the adversary is still required to have access to a candidate point that may have been used to train the model in order to run the attack, it is more of a “smoke signal” than an outright privacy violation. Training data extraction on the other hand, is when given access to the model, an adversary can actually *reconstruct* parts of the underlying training data, which may be sensitive or protected by law.

When dealing with generative models however, MIAs have the potential to facilitate training data extraction attacks, which makes them of elevated importance. The key trick here is that the adversary can exploit *memorization*: as discussed in Section 3, language models memorize swathes of their training data verbatim, and can be prompted to regurgitate these training examples. Armed with a MIA that can identify training samples from samples in the distribution, an adversary can then use the attack to find the generated points that are most likely to be training points, and “extract” those. We discuss this further in Section 4.3.

We start by discussing the basics of MIAs, and then covering MIAs against language models, which are different from attacks on standard models in a few key ways. We then cover training data extraction from language models, and discuss mitigation of privacy attacks.

### 4.1 Preliminaries

Membership inference attacks are based on the concept that models tend to perform better on examples that are in their training set than ones which are not. This means that model confidence is a natural test statistic for determining membership in  $D_{train}$  and can be used for a MIA. The most straightforward and computationally inexpensive MIA is a *threshold attack*, which was first implemented by Yeom et al. [2018]. This attack simply designates a threshold  $\tau$  and compares a chosen test statistic (in most cases, loss) to that. An outline for this kind of attack is shown here:

Given a target model  $h_{target}$ , a threshold  $\tau$ , and some test statistic  $l(h, z)$  that is computed by prompting  $h$  with  $z$ , perform the following attack on a target data point  $z$ :

1. Query the model to get  $h_{target}(z)$
2. Classify  $z$  as in if  $l(h_{target}, z) < \tau$ . Else, classify  $z$  as out.

If, for example, the test statistic is loss, then this attack simply classifies the low loss points as in and high loss points as out. The choice of  $\tau$  determines the trade-off between the precision and recall of the MIA. If  $\tau$  is very high, then the MIA will have a high recall, but low precision rate in identifying members of  $D_{train}$ . Meanwhile if  $\tau$  is low, then it will have a high precision rate, but low recall. This also enables statistics common in machine learning such as the AUC. Loss is typically the first choice based on the principle that models tend to overfit on training data. However, Yeom et al. [2018] found that overfitting, while sufficient, is not necessary for a MIA to be successful and other test statistics may instead be used for a successful MIA. This is especially relevant in the context of LLMs which have much larger training datasets and are thus trained using a single-pass SGD. This means they overfit less making loss-based threshold attacks less effective. A more commonly used test statistic is perplexity, which colloquially measures how surprised the model is to see a given sequence. This is very similar to  $\ell_{perp_\theta}$ , defined in the Memorization section (Definition 3).

**Definition 10** (Perplexity, Carlini et al. [2020]) Given a model  $h$  with parameters  $\theta$ , the **perplexity** of a sequence  $s = (s_1, \dots, s_n)$  is given by:

$$perp_\theta = \exp\left(-\frac{1}{n} \sum_{i=1}^n \log_2 Pr(s_i | h_\theta(s_1 \dots s_{i-1}))\right)$$

Threshold attacks tend to suffer from high false positive rates Mattern et al. [2023] and low precision in identifying sensitive data Carlini et al. [2020]. One solution to this, first developed in the context of ML by Murakonda et al. [2021], is to use a likelihood ratio as the test statistic. In this approach, a reference model,  $h_{reference}$ , is trained using the same architecture and similar data as the target model,  $h_{target}$ . Then for a target data point  $s$ , the test statistic is  $L(s) = \log\left(\frac{p_{reference}(s)}{p_{target}(s)}\right)$ , where  $p_o(s)$  is the probability of outputting  $s$ . This reduces the bias towards common terminology and allows for better identification of sensitive data.

More advanced approaches, such as Shokri et al. [2017], Long et al. [2020], Sablayrolles et al. [2019], Song and Mittal [2020] and Carlini et al. [2022], seek to approximate two distributions of the test statistic, one for a model which contains the target points  $D_{train}$  and one which does not. It then trains an attack model to determine which distribution a target point  $z$  belongs to. This approach necessitates the training of *shadow models*, which are models of the same architecture as the target model trained on data similar to the target dataset. These shadow models are used to estimate the distributions for the in and out data points. This approach was first designed by Shokri et al. [2017] and is illustrated in Figure 6. Carlini et al. [2022] combines this approach with the reference model approach by conditioning the output of the target  $h_{target}(z)$  on the two distributions and computing the ratio. The ratio is then used as the threshold statistic to determine  $D_{train}$  membership. This is a leading MIA approach and is known as LiRA (Likelihood Ratio Attack). We discuss this further in Section 4.2.2

These approaches can perform much better than a threshold attack, but it requires much more knowledge of the target model and dataset. Additionally, training shadow models can be computationally expensive since they must be of the same build as the target model. In the case of LLMs, which can be build using billions of parameters and data points, this is often prohibitively expensive.

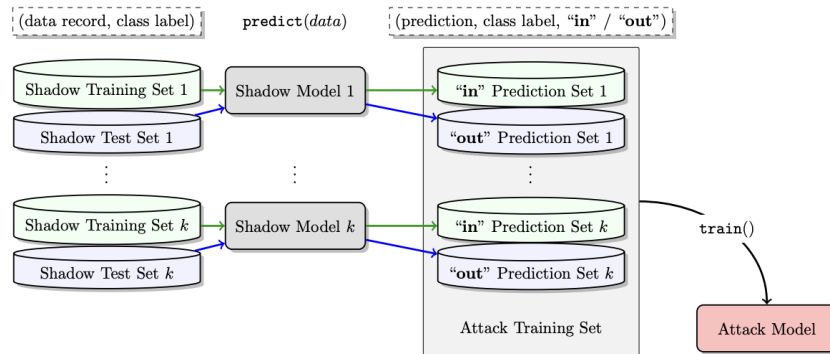


Figure 6: Shadow model framework from Shokri et al. [2017]. Multiple shadow models are built using the same architecture and similar data as the target. Points from the shadow models’ training sets are used to build the in prediction set while points from the shadow models’ test sets are used to build the out prediction set. An attack model is then trained to distinguish the in and out labels.



## 4.2 Membership Inference Attacks on Language Models

In recent years there has been growing research on MIAs for language models. It has been found that LMs are fairly robust against simple probing methods. For example, Lehman et al. [2021] attacked a BERT model trained on the MIMIC-III Johnson et al. [2016] clinical dataset by probing with prompts such as "John Doe is a 35 year old patient diagnosed with..." to try and extract health conditions. This method, and simple variants of it, were unable to extract any meaningful sensitive data. However, there still remain privacy risks when the MIAs described above are tailored towards LMs as we will discuss here.

### 4.2.1 Threshold Attacks

To understand this first MIA, we must first discuss the process by which LMs are trained. To build  $D_{train}$ , the raw corpus of text data is first *tokenized*, transforming the large dataset into discrete elements (e.g. tokenizing "My name is John Doe" into ["my", "name", "is", "john", "doe"]). Since there are a very large amount of possible words in the dictionary, the matrix representing this data would be very sparse. Thus *embedding* models are used to transform these tokens into low-dimensional vectors that preserve semantic data. This is a critical first step for many LMs and has been very successful in a wide range of applications.

While most MIA research revolves around the LMs, Song and Raghunathan [2020] found that word and sentence embedding models are also vulnerable to privacy attacks. These embedding models take in sparse text data and transform it into lower dimensional vectors that maintain semantic meanings yet are easier for models to learn. For example, the word embedding models Word2Vec Bojanowski et al. [2017] and FastText Mikolov et al. [2013] are trained in an unsupervised manner where, given a sliding window of words  $C = [w_b, \dots, w_0, \dots, w_e]$ , they aim to predict  $w_i$ , given the center word  $w_0$ , compared to a *negative* sample (i.e. a random sample of words not in  $C$ ). The sentence embedding model studied, given sentence  $s_a$  and following sentence  $s_b$ , is trained to maximize the probability of  $s_b$  given  $s_a$ , compared to the negative sample of sentences. Song and Raghunathan [2020] assumes that the adversary has access to this embedding model  $\Phi$ . To perform the threshold attack on word embedding model  $\Phi$ , Song and Raghunathan [2020] first converts  $C$  into its embedding  $\Phi_{word}(C)$ . They then compute a similarity score  $\delta$  (e.g. cosine similarity) for each pair of vectors  $\delta(\Phi_{word}(w_i), \Phi_{word}(w_j))$  in  $C$ . If the averaged similarity score is above a certain threshold  $\tau$ , it is considered part of the training data. The logic is that  $\Phi_{word}(C)$  will have a higher average similarity score if  $C \in D_{train}$  than if  $C \notin D_{train}$ . See Algorithm A.1 for the details. Song and Raghunathan [2020] also formulated an attack for sentence embedding models. For the sentence embedding MIA, Song and Raghunathan [2020] computes the similarity score  $\delta(\Phi_{sentence}(s_a), \Phi_{sentence}(s_b))$  to perform the threshold attack. If the similarity score is above a certain threshold, the target pair is classified as in. See Algorithm A.2 for details.

Song and Raghunathan [2020] performed this attack on three word embedding models: Word2Vec Bojanowski et al. [2017], FastText Mikolov et al. [2013], and GloVe Pennington et al. [2014] which were all trained on the Wikipedia articles corpus Mahoney [2009]. They also attacked a sentence embedding model, the dual-encoder framework used in Cer et al. [2018], Chidambaram et al. [2019], Henderson et al. [2017], Logeswaran and Lee [2018], Lowe et al. [2015], Reimers and Gurevych [2019] and Yang et al. [2018] trained on BookCorpus Zhu et al. [2015]. Song and Raghunathan [2020] measured adversarial advantage as the difference between the true and false positive rate, so random membership guessing has an adversarial advantage of 0. For common words and sentences, there was no adversarial advantage. However, these are also the least relevant from a privacy perspective. In both word and sentence settings, when looking at less frequent examples, this method offered an adversarial advantage of up to .3 when looking at words and sentences that were infrequent. Thus, these embedding models cannot simply be treated as vectors useful in model training but as potential sources of privacy leakage.

Carlini et al. [2020] also implemented a metric based MIA on GPT-2 as part of larger research on training data extraction (see Section 4.3). This was performed on the GPT-2 XL model family which had all been trained on public Internet data that was deduplicated at the document level. Given a sequence of tokens  $(s_1, \dots, s_n)$ , Carlini et al. [2020] computes the perplexity (Definition 10) as the test statistic for a threshold attack. In this attack, they had generated thousands of sequences, de-duplicated them, and thresholded by picking the  $k$  sequences with the lowest perplexity. With  $k = 100$ , between 3% – 39% (depending on how sequences were generated) of these were manually verified to be part of  $D_{train}$ . However, it did have a high false positive rate and identified many qualitatively uninteresting sequences like public X (formerly Twitter) handles or common phrases like "I love you". This highlights a quirk of membership inference for LLMs where inherently, every individual word that is outputted by a model is guaranteed to be a part of the training set. In order to identify privacy risks of sensitive information, there must be some context given to the measurement. Carlini et al. [2020] proposed a few methods to account for this. First, they instead use the ratio of perplexity to the zlib entropy Gailly and Adler [1995] which accounts for how common a sequence is. They also suggest using the minimum perplexity averaged over a sliding window of tokens or comparing to similar sequences with style transfers such as

lower/uppercasing. These changes improved the MIA recall rate by up to  $10\times$ . See Table 5 in Section 4.3 for the full results.

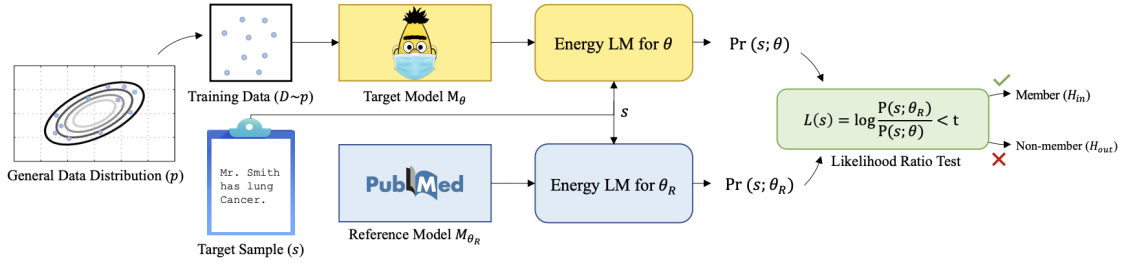


Figure 7: Overview of the attack used by Mireshghallah et al. [2022]. Probabilities are computed for a target sample on both the target and reference model, then the log ratio is used to determine in or out status.

Other work also uses some form of reference model to compute the threshold test statistic. Mireshghallah et al. [2022] constructed an attack on masked language models (MLMs), which are a class of models frequently used for pre-training for NLP tasks such as text classification and have been applied in a variety of domains including finance and healthcare, both of which contain potentially sensitive information. While prior MIA attempts on MLMs have concluded that sensitive information remains private Lehman et al. [2021], Vakili and Dalianis [2023], Jagannatha et al. [2021], Nakamura et al. [2022], these attacks used a target model’s loss as the threshold statistic, which is not a very accurate attack especially at low FPRs. Furthermore, as discussed in the preliminaries, LLMs are trained with a single pass SGD meaning they are less overfit and thus less susceptible to loss-based attacks. The target model in this study was the same ClinicalBERT Alsentzer et al. [2019] model used in Lehman et al. [2021]. This was trained on the MIMIC-III dataset Johnson et al. [2016] in four different hyperparameter settings for completeness. The reference model was the PubMedBERT model trained on PubMed text Peng et al. [2019], which was chosen due to its similarity in dataset topic and model architecture as the target model. Compared to a baseline attack model which only used the loss of the target model as the test statistic, Mireshghallah et al. [2022] saw an improvement in the AUC from .66 to .90 (Figure 8). This significant improvement showed that using the likelihood ratio for a threshold attack is viable and that MLMs are not nearly as private as they were previously thought to be.




Figure 8: ROC curve from the results of Mireshghallah et al. [2022]. Using a likelihood ratio as the test statistic significantly improved over a baseline threshold attack using only loss of the target model.

While the attack by Mireshghallah et al. [2022] was effective, it still relied on the assumption that the adversary has access to data that is very similar to the one used by the target. Mattern et al. [2023] devised a framework which does not rely on having access to any database. Instead, a separate model, called the *proposal model*, is trained to create “neighbors”, i.e. sequences that are very similar to the target sequence but with slight perturbations. These neighbors are fed into the target model and the difference between the target sequence and the average loss of the neighbor sequences is used at the threshold statistic (see Figure 9). This method was tested on a base version of GPT-2 Radford et al. [2019] trained on the AG News Corpus Gulli [2004], the Sentiment140 dataset Go [2009], and excerpts from the Wikitext-103 dataset Merity et al. [2016]. Mattern et al. [2023] compared their method to two baselines, the loss only model by Yeom et al. [2018] and the likelihood ratio model by Mireshghallah et al. [2022] and found that it outperformed both. The



TPR of Mattern et al. [2023] was marginally better or up to twice as high as the Miresghallah et al. [2022] framework depending on the dataset. Both frameworks outperformed the baseline loss only model.

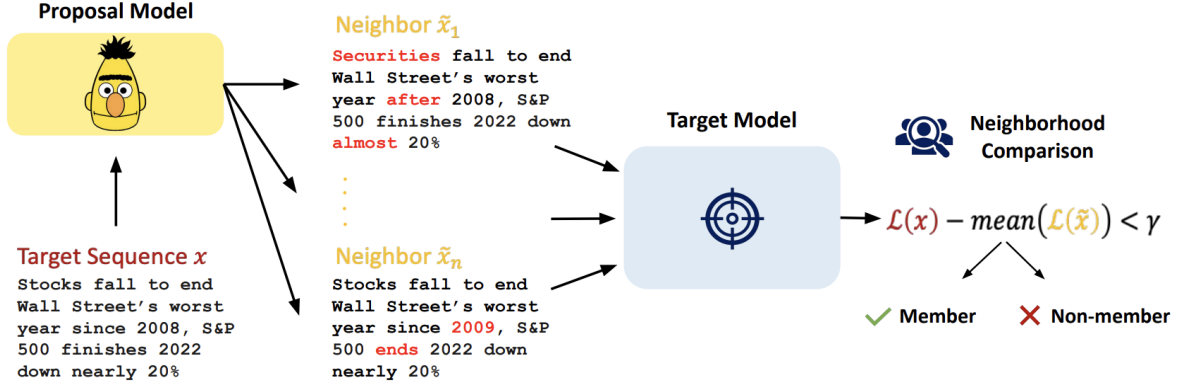


Figure 9: Framework of the neighbor comparison attack by Mattern et al. [2023]. Given target sequence, a model generates neighbor sequences. The losses are then compared to determine whether the target should be classified as in or out.

In a different approach that is related to using a reference model for a privacy attack, Zanella-Béguelin et al. [2020] found that model snapshots before and after database updates can be used to reveal sensitive information about training data membership. If the adversary has access to a prior version of the model, then given a sequence  $s$ , Zanella-Béguelin et al. [2020] calculated the difference in probabilities of that output between the old and new models called the *differential score*. This score is most useful when interpreted relative to other sequences, so they also formulated the *differential rank*, which is the number of sequences of length  $|s|$  with differential score higher than  $s$ . The lower the differential rank of a sequence, the more likely that sequence was exposed during the model update. This attack was tested on three different datasets and corresponding model choices: 1) a two-layer RNN similar to Zaremba et al. [2015] trained on the Penn Treebank Marcus et al. [1993] dataset, 2) a single layer RNN model and a BERT model Devlin et al. [2019] trained on a Reddit dataset, and 3) a two-layer RNN model trained on Wikitext-103 Merity et al. [2016]. The models were trained, then updated once using retraining (restarting from scratch) and once using continued training (adding an additional batch of data). In each case, canaries were added to the model and the ability to extract these canaries was a proxy for the vulnerability of private data. Using the methods outlined, Zanella-Béguelin et al. [2020] was able to extract the canaries in almost all settings.

This attack is particularly relevant given current and incoming legislation requiring more frequent database updates to allow individuals to remove their data from the model. Given these results, users who request their data to be deleted from a database may counter-intuitively be increasing their privacy risks. See Section 7 on machine unlearning for more context on legislation and user deletion requests.

#### 4.2.2 Shadow Model Attacks

Song and Shmatikov [2019] were the first to explore membership inference in the natural language text generation setting. Their research revolved around building a tool to help users identify if their data was used in model training and used the shadow model framework Shokri et al. [2017] to do so. Song and Shmatikov [2019] attacked a one-layer LSTM model that was trained on three different datasets: a Reddit dataset, annotated TED talks, and the Cornell Movie Dialogues corpus Danescu-Niculescu-Mizil and Lee [2011]. The shadow models were Gated Recurrent Unit Cho et al. [2014] models trained on Wikitext-103 Merity et al. [2016], Europarl Koehn [2005], and Ubuntu dialogs dataset Lowe et al. [2015]. This tool was found to be extremely effective, allowing an auditor to successfully identify whether a user's data was in  $D_{train}$  100% of the time across all hyperparameter settings. This attack, however, benefited from a setup where the auditor could prompt the model as many times as wanted using the users' data. They could use the MIA output aggregated over many prompts for a single user and if any point was in  $D_{train}$ , they could make the membership judgement based on that. Additionally, this attack was performed on an LSTM, which is orders of magnitude smaller than modern LLMs. Training shadow models for this attack would be prohibitively expensive if each shadow model were a multi-million or billion parameter model and the scale of modern LLMs is rapidly increasing (see Figure 3).

The attack by Song and Shmatikov [2019] relied on having access to the probability distribution of the outputs. Hisamoto et al. [2020] extended this attack by considering the more realistic scenario where the attacker only has access to the output sequence. Hisamoto et al. [2020] also use the shadow model framework, but when the shadow models provide

their output sequence, they compute a smoothed sentence-level BLEU score Lin and Och [2004] and the number of 1 – 4grams of the output that match the true data. These values are used to train the binary classifier to predict in and out membership. The target model was a six-layer Transformer Vaswani et al. [2023] trained on default Sockeye parameters Hieber et al. [2018] using the corpora from the Conference on Machine Translation Bojar et al. [2018]. In this scenario, the MIA was not successful with only chance level ability to guess membership for a single sentence. When attacks were done using a group of sentences (similar to the unlimited prompt setting in Song and Shmatikov [2019]), the attack was more accurate than random guessing, but still was not as effective as in Song and Shmatikov [2019]. This attack was primarily successful against outlier words and phrases.

Most research on MIAs has focused on token or sentence level attacks; Meeus et al. [2023] performed the first MIA at the document level. To do this, Meeus et al. [2023] first queried the model,  $h$ , with several context windows of tokens of length  $C$  from the document and asked for next token prediction probabilities. These probabilities are normalized based on the commonality of the predicted tokens and are then used to generate features for the shadow model. There are two different feature aggregation methods: *AggFE* takes these token level predictions and uses high-level statistics such as min, max, mean, standard deviation, and percentile values as features. *HistFE* creates  $N$  histogram buckets of these probability values and uses the fraction of tokens in each bucket as features. This attack tested membership inference of books and academic papers from the RedPajama-Data dataset Computer [2023] on the 7-billion parameter OpenLLaMA model Geng and Liu [2023]. This attack was highly successful with the AUC of the book MIA being .856 and the AUC for the paper MIA being .678. As a comparison, Meeus et al. [2023] adapted a sentence level MIA for the same task and found the performance to be at near guessing levels.

Many ML models are built using a process known as transfer learning. In this, the model is pre-trained on a generic public dataset before being fine tuned on a case specific dataset. Abascal et al. [2023] implemented a shadow model MIA to infer membership in the pre-training dataset of a language model using only query access to the finetuned model. Since language models empirically forget over time Jagielski et al. [2023], inferring pre-training membership could be more challenging than inferring membership in the fine-tuning set. This attack was performed on the publicly available Transformer-XL foundation model Dai et al. [2019] which was pretrained on the WikiText-103 Merity et al. [2016] dataset and fine-tuned on the DBpedia dataset Zhang et al. [2016]. The shadow models were built using these same public models and datasets and the attack model  $f_{attack}$  (i.e. the model trained on the shadow model output) is a  $k$ -nearest neighbors model. As seen in Table 4, Abascal et al. [2023] was able to infer membership in pre-training datasets with TPRs as much as  $10\times$  higher than the baseline of randomly assigning membership. This attack shows the MIA risks are very much inherent for models which are publicly available, which removes the computational barrier to entry for a shadow model attack on an LLM. However, since pre-training datasets usually do not contain sensitive information, this is not as dangerous of a risk.

Shadow Models	k	TPR @ 0.1% FPR	TPR @ 1% FPR
16	$\sqrt{ D_{meta} }$	1.6%	5.2%
16	16	2.6%	7.0%
32	$\sqrt{ D_{meta} }$	2.0%	5.5%
32	32	3.1%	8.1%
64	$\sqrt{ D_{meta} }$	2.2%	6.0%
64	64	3.4%	8.8%

Table 4: Table of results from Abascal et al. [2023]. True positive rates at different number of shadow models, number of neighbors in the  $k$ -nearest neighbors  $f_{attack}$ , and fixed false positive rates.  $|D_{meta}|$  is the size of the  $f_{attack}$ ’s data generated by the shadow models. Despite pre-training data typically being seen as safe, there is still a risk of privacy leakage

Carlini et al. [2022] used the concept of shadow models to develop one of the leading MIA frameworks, LiRA. Given a target data point  $z$ , Carlini et al. [2022] trained two batches of shadow models, one whose training set contains  $z$  and another whose training set doesn’t. These estimate the Gaussian distributions of each respective membership scenario. Then given the output of the query  $h_{target}(z)$ , they compare the likelihoods conditioned on the model being trained on each of the distributions. If there are multiple target points, they train shadow models on  $N$  subsets, chosen so that each target  $z \in D_{train}$  appears in  $\frac{N}{2}$  subsets. This attack was performed on a GPT-2 model trained on Wikitext-103 Merity et al. [2016]. The shadow models were trained on random sub-samples of half the total dataset. As a result, this work relies on a strong assumption that the training set of the shadow models partially overlaps with  $D_{train}$ . This attack was very effective with an over  $10\times$  improvement in TPR over the baseline attack by Yeom et al. [2018] with the FPR set at .1%. See Algorithm A.3 for more details.

As with other shadow model approaches, Carlini et al. [2022] suffers from the fact that training multiple shadow models of LLMs may be prohibitively expensive. That said, Bertran et al. [2023] developed an MIA framework that is competitive with many of these attacks but only requires training a single shadow model. Bertran et al. [2023] was only implemented for tabular data and image classification, however it is possible that could later be applied towards language models.

### 4.3 Training Data Extraction from Language Models

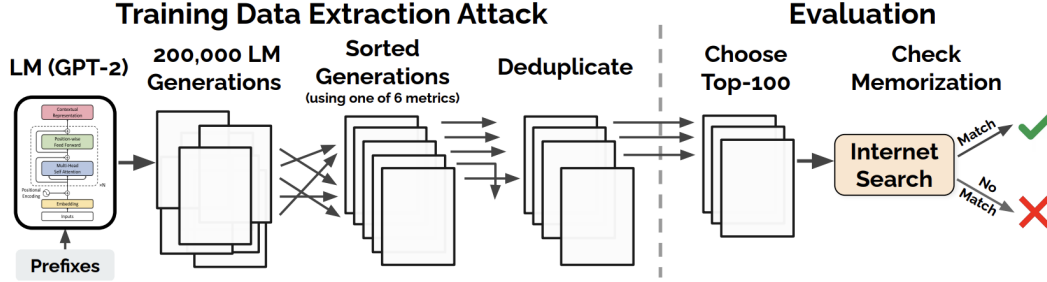


Figure 10: Workflow of a training data extraction attack from Carlini et al. [2020]. First, potential candidates are generated by prompting the model with a (potentially empty) prefix. A membership inference attack is then used to sort the candidates by likelihood of membership in  $D_{train}$ . These are then deduplicated and the top- $k$  samples are selected.

The second type of privacy attack against LLMs is training data extraction. While MIAs may reveal some degree of sensitive information, it requires the adversary to have prior knowledge of the target data point. Training data extractions enable an adversary to directly retrieve sensitive information using only query access to the model. This type of attack was first formulated by Carlini et al. [2020] and is outlined in Figure 10. There are three main steps outlined here:

1. **Generating Candidate Targets:** The adversary generates a list of candidate targets by querying the model with a short prefix. Given a starting set of prefix tokens  $s_0, \dots, s_i$ , the adversary samples  $s_{i+1} \sim h_{target}(s_{i+1}|s_0, \dots, s_i)$ . They then attach  $s_i$  to the prefix and sample  $s_{i+2} \sim h_{target}(s_{i+2}|s_0, \dots, s_{i+1})$  iteratively until the desired length is reached.
2. **Membership Inference Attack:** Using MIA methods, the adversary sorts these candidates by likelihood of membership in  $D_{train}$ .
3. **Predicting Top- $k$ :** After deduplicating the candidates, the adversary then selects the top- $k$  candidates.

Carlini et al. [2020] performed this attack on GPT-2 since while the dataset was not publicly released, the dataset was sourced from public data, making it both manually verifiable and ethically sound. In their baseline attack, Carlini et al. [2020] generated candidate tokens of length 256. At each step in the candidate generation, they sampled from the top  $n = 40$  outputs. They then sorted these candidates by their perplexity and select the top  $k = 100$  candidates. As seen in Table 5, 9 of these candidates were then verified as part of the training dataset. While somewhat successful, this baseline attack yielded mostly qualitatively uninteresting results such as X (Twitter) handles. The candidates generated were low in diversity and the membership inference portion highly ranked tokens that were common phrases or trivial numbers.

Carlini et al. [2020] improved upon this baseline in a few ways. On the candidate generation end, they tried two new methods. First, they used a decaying temperature setup (labeled Temperature in Table 5) where for early tokens, the confidence in top predictions was lowered to encourage diversity. As the generation continued, the “temperature” would decrease yielding more predictable generation. Second, they seeded the initial prefix using data from Common Crawl [2008] (Internet) to ensure there are no highly unusual prefixes. On the MIA approach end, instead of simply ranking by perplexity, Carlini et al. [2020] ranked by perplexity relative to a second, similar model or normalized by another value. Their comparison models include different sized GPT-2 models (Small and Medium) and a model trained on the lowercased version of the data (Lowercase). They also tried dividing the perplexity by the zlib entropy of the candidate (zlib) and also considering only the lowest perplexity from a sliding window of 50 of the 256 tokens. These methods ensure that common and uninteresting phrases are filtered out. As we see in Table 5, these improved methods enabled much more training data to be extracted. Of the 604 total examples extracted, 46 were named individuals not from a news source and 32 were the contact information of individuals.

Yu et al. [2023] also performed a training data extraction attack and tested numerous improvements to the candidate generation and MIA steps. In their case, they attacked a 1.3B parameter GPT-Neo model implemented on HuggingFace

Inference Strategy	Text Generation Strategy		
	Top- $n$	Temperature	Internet
<b>Perplexity</b>	9	3	39
<b>Small</b>	41	42	58
<b>Medium</b>	38	33	45
<b>zlib</b>	59	46	67
<b>Window</b>	33	28	58
<b>Lowercase</b>	53	22	60
<b>Total Unique</b>	191	140	273

Table 5: Results from Carlini et al. [2020]. The number of memorized examples (out of  $k = 100$  candidates) identified using each of the three text generation strategies and six membership inference techniques. Some samples are found by multiple strategies; 604 unique memorized examples were identified in this attack.

Transformers Wolf et al. [2020] which was trained on the Pile Gao et al. [2020]. In this attack, the initial prefix was 50 tokens long and selected from  $D_{train}$ . The success of the attack was determined by whether the 50-token long suffix tokens generated resulted in a 100-token example that existed in the Pile. The baseline approach used a greedy top-1 sampling in the candidate generation process and non-adjusted perplexity in the MIA step. For the candidate generation process, the most successful improvement was shifting the size of the context window to account for sentences which may have been truncated or prefixed in the training data. Instead of the next token  $s_{i+1}$  being sampled from  $h_{target}(s_{i+1}|s_0, \dots, s_i)$ , it is sampled from an ensemble of  $(h_{target}(s_{i+1}|s_0, \dots, s_i), \dots, h_{target}(s_{i+1}|s_m, \dots, s_i))$ , where  $m$  is the size of the context window. As seen in Figure 11, this more than doubled the precision of the attack over the baseline. For the MIA ranking process, the most successful improvement was weighting the perplexity based on the number of high confidence tokens generated. If a token  $s_i$  was sampled that had a confidence ( $> .9$ ), then a reward was subtracted from the final perplexity value (.1). This likewise saw a doubling of the attack precision. Overall, Yu et al. [2023] presented many different tricks to improve training data extraction, which can all be seen in Figure 11.




Figure 11: Overview of the different methods explored by Yu et al. [2023].  $M_p$  is the precision, i.e. the percent of generated suffixes that were in  $D_{train}$ . Bars in pink denote the methods in the improved candidate generation, and bars in orange denote the methods in the improved MIA ranking. The dashed bars indicate the best method in each category.

Carlini et al. [2023b] released the open LM Training Data Extraction Challenge <https://github.com/google-research/lm-extraction-benchmark> to provide a benchmark for research. Zhang et al. [2023a] has proposed an attack, known as *ETHICIST*, which improves the extraction attack in two ways. First, they propose loss-smoothed soft prompt tuning to train the prefix prompt in order to improved the generation of target candidates. Second, rather than sorting the candidates by strict generation probability, the candidate losses are normalized for the given prefix, relative to the other candidates. This was tested using the 1.3-billion parameter GPT-Neo model to extract 50-token suffixes using a 50-token prefix and had a recall rate of 62.8%, which is significantly higher than the Carlini et al. [2023b] baseline benchmarks which are around 50%.

On November 28th, 2023 Nasr et al. [2023] published the first training data extraction attack against a production LLM (ChatGPT), and a host of open source models (Pythia, GPT-Neo, LLaMA, Falcon). The set up of their experiment was to (i) Download  $10^8$  bytes of Wikipedia data (ii) Randomly sample 5-token snippets  $p_i$  from the data (iii) Prompt the LLM with the random snippet  $p_i$ , and generate a 50 token string  $x_i \sim \text{Gen}(p_i, \cdot)$ , and repeat this  $10^9$  times for each language model they evaluated. For the open source LLMs studied, they had access to the training set, and so could verify if  $x_i$  was a training point (in order to do this they used an efficient data structure for lookup called a suffix array). For closed source LLMs, they created a proxy training dataset by taking the union of common datasets for training LLMs that likely contained most of the model training data, and checked if  $x_i$  was in the proxy dataset. Since the proxy dataset likely did not contain all the samples in the true training dataset, for these closed source LLMs the measured memorization is actually a lower bound on the true memorization, which makes the results even stronger. They find an order of magnitude more memorization than was previously believed to exist in both closed and open source LLMs, including LLMs that have been *aligned* through instruction-tuning or RLHF. They study:

- **How frequently do LLMs emit memorized data?** They find that LLM’s emitting training data is not as rare a phenomenon as previously believed, and even some relatively small models (7B) emit training data as much as 1.4% of the time! This quantity varies across open source models of a similar size, for example RedPajama 7B emits training data more than 2x as frequently as Pythia 6.9B (even de-duped) and 7x as frequently as GPT-Neo 6B. Generally speaking, larger models tend to emit data more frequently than smaller models, as expected based on prior work Tirumala et al. [2022], Carlini et al. [2023a].
- **How much total extractable memorization is there (as a fraction of training examples?)** Frequency of training data emission is a separate quantity from total quantity of memorized training samples. For example, a model that always outputs the same training example would emit training data 100% of the time, but would emit only a single memorized training sample. By computing the number of unique training samples emitted, they show that LLMs memorize *millions of training samples*, which is orders of magnitude more than in prior work Carlini et al. [2020]. They then try to extrapolate how much total memorization exists (e.g. if the model was queried infinitely many times, how many training examples would be emitted), which they estimate using a sequential Good-Turing estimator developed in a recent master’s thesis from Lund University Andersson [2022]. These extrapolations, coupled with empirical evidence that the Good-Turing extrapolations are conservative, indicates even relatively small open source LLMs may memorize 10s of millions of training examples!
- **What is the gap between discoverable and extractable strings?** Prior work Carlini et al. [2023a] released a dataset of discoverable memorization (strings that are recovered when the model is prompted with the first 50 tokens of the string) on GPT-Neo 6B. This can be viewed as an upper bound on extractable memorization, but not a practical attack, since in practice the adversary would not have access to the first 50 tokens of the training example. Using the extractability results for GPT-Neo 6B from the prompting strategy above, they show that they can recover 35% of discoverably extractable data via simple prompting with internet text. Conversely, there are a further 11% of the sequences that are extractable, but aren’t discoverably extractable.
- **What about chat (aligned) models?** They identify two challenges of extracting data from chat models via the simple prompt continuation strategy employed above: (i) the model appends text to the user’s query before answering (e.g. "Assistant:") and (ii) because of alignment the model can evade certain questions, perhaps to avoid hallucinating etc. In order to circumvent this alignment procedure, they find that asking the chat model (gpt-3.5-turbo) to repeat a token infinitely often, causes the model to spontaneously start generating tokens that aren’t that specific token after a while, and then when the model starts generating in this mode, it outputs training data with high frequency. In particular, they find that the rate of unique 50-grams extracted via this attack is much higher than any other model except gpt-3.5-instruct-turbo, which is more than 3x the rate of smaller open source models. Interestingly, they find that these 10000 extracted examples are not *discoverably* extractable. They remark that “This suggests that it will be difficult to red-team this model and evaluate its privacy without additional access to both the model and also the un-aligned foundation model from which it was derived.” This is supported by the fact that the gpt-3.5-turbo-instruct model does discoverably memorize the strings extracted from the chat model gpt-3.5-turbo. Since gpt-3.5-turbo-instruct and gpt-3.5-turbo have different fine-tuning datasets, the extracted data was likely from the pre-training set.
- **Why is ChatGPT so vulnerable to attack?** The authors speculate it could be for a few reasons: (i) ChatGPT may be trained for many more epochs during pre-training than comparable base language models, particularly open source ones (ii) Repeated prompting with certain single token words appears to cause the model to converge towards prompting with the  $\langle \text{end of text} \rangle$  token used to reset the model during training. This explains why after a certain number of word repetitions the model is vulnerable to extraction attacks, since it starts to generate as if it was generating from the base model.

This paper highlights that base language models of all sizes have a high propensity to regurgitate training data, and that alignment, at least the way state-of-the-art models are aligned, does not seem to solve this problem in the face of adversarial prompting strategies.

The rapid improvement in these attacks expose the significant risk of training data extraction in LLMs, with Patil et al. [2023] finding that there is no one universal solution to protecting against this risk. Similar work has also been done for image diffusion models Carlini et al. [2023c]. While this is outside the scope of this survey, it highlights the prominence of this type of attack.

#### 4.4 Attribute Inference Attacks

Attribute inference attacks are another potential privacy risk for large language models, albeit with less research compared to membership inference and training data extraction attacks. Staab et al. [2023] presented the first work to comprehensively study the risks of this attack. Outlined in Figure 12, this attack works by leveraging public data written by the user, such as posts on online social forums. These posts are fit into a prompt template which asks an LLM to identify the personal attributes of a person who made those posts. Staab et al. [2023] tested this attack using multiple LLMs (including GPT-4, Llama, PaLM, and Claude) for inference and used a custom database of 520 Reddit profiles, which were manually annotated to include age, education, sex, occupation, relationship status, location, birth place, and income. Using the comments made by these profiles, GPT-4 achieved an 84.6% Top-1 accuracy across all attributes. While these attacks present a privacy risk, this experiment was done in comparison to a human’s time and accuracy performance, meaning this risk is not unique to LLMs, but can rather just be expedited by them.

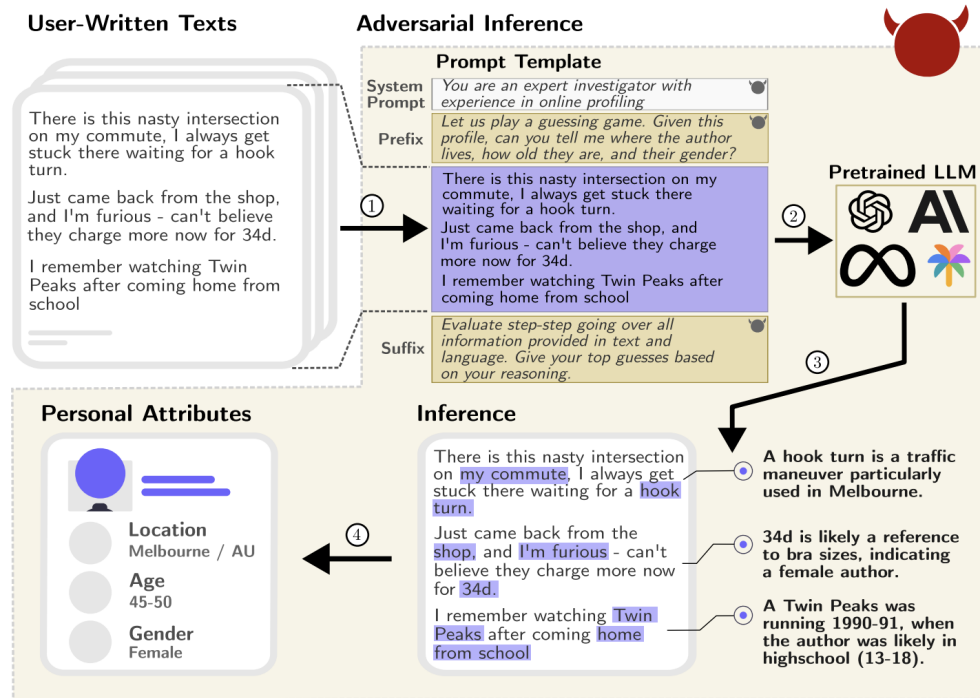


Figure 12: Overview of an attribute inference attack from Staab et al. [2023]. Using access to a dataset of user-text (such as online forum posts), the adversary creates a model prompt. This prompt is given to an LLM to infer personal attributes about the individual.

#### 4.5 Mitigation Techniques

Membership inference and training data extraction attacks pose a significant privacy risk to LLMs as outlined above. Here we discuss a few methods that have been found to mitigate those concerns:

**Differential Privacy:** The primary technique that can be used to reduce privacy risks is differential privacy. Differential privacy introduces noise either during training or at prediction time that obscures private information with certain mathematical guarantees. Many studies from above, including Jagannatha et al. [2021], Abascal et al. [2023], and



Zanella-Béguelin et al. [2020], tested their attacks on differentially private models and found that these models were significantly more robust. For a more in depth coverage of differentially private language models, see Section 5.

**Reducing Memorization:** MIAs and memorization overlap in the privacy risks they pose. Models that memorize private information during the training process will inherently also be vulnerable to membership inference and training data extraction attacks on that sensitive data Thomas et al. [2020]. Deduplication is a popular technique for reducing memorization and Kandpal et al. [2022] found that depuplicating the dataset can reduce the success of the loss based MIA from Carlini et al. [2020] to near chance levels. Other techniques such as federated learning McMahan et al. [2016], Thakkar et al. [2020], Ramaswamy et al. [2020], differential privacy, and early detection Biderman et al. [2023a] are all effective in reducing memorization. See Section 3 for more details on memorization and mitigation techniques.

**Training Time:** Jagielski et al. [2023] defined “forgetting” as the success rate of an MIA as model training progressed. In their work, they found that LLMs tend to forget over the course of training, meaning that examples seen early on in the process (such as during pre-training) saw some privacy benefits. Typically, however, LLMs are pre-trained on large public datasets before being fine-tuned with case specific data, which usually includes the sensitive data. Using larger datasets during fine-tuning or longer training using non-sensitive data could be useful in mitigating privacy concerns.

**Types of Models:** Various research has touched on how vulnerable different types of language models are to membership inference attacks. Jagannatha et al. [2021] found that for clinical language models, smaller models have less privacy leakage than larger ones. Additionally, they found that masked LMs have lower leakage than auto-regressive LMs. The choice of LM type is less practical as a mitigation strategy since certain types of LMs may be required for certain use cases to perform well. However, this is something that can still be considered when building an LLM.

**Tools:** From a practical standpoint, another way to mitigate privacy risks of MIAs is to allow fast identification of training set membership. Marone and Durme [2023] proposed and built a tool called Data Portraits, which is a space efficient method to record training data and identify whether a given string is contained in that data. Data Portraits would be used in a similar fashion to datasheets Gebru et al. [2021] and would serve to complement documentation practices for model and dataset creators. This would be useful for multiple types of users: 1) content creators who want to find out if a dataset contains their content 2) researchers looking to assess dataset leakage and knowledge memorization and 3) content consumers who want to see if a model is plagiarizing. While it would not mitigate membership inference directly, it could serve as an auxiliary tool to reduce risks.

## 5 Privacy Preserving Large Language Models

As discussed in Section 3 and Section 4, there are many privacy risks that come with the use of LLMs, raising the obvious question: *can LLMs be pre-trained or fine-tuned in a way that preserves the privacy of the training data, while also maintaining utility?* In this section we look at methods that aim to privately train language models, under two high level privacy risk frameworks. The first focuses on the ability of an adversary to infer information about the underlying training data based on a level of access to the model; e.g. to conduct membership inference or training data extraction. This risk is addressed by training the model using *differential privacy*. The other privacy framework is addressed by federated learning, where the training data is stored in a distributed fashion, and the goal is to learn a centralized language model trained on the distributed data, without having to centrally aggregate the data. Both of these frameworks are used by companies such as Google and Apple to preserve user privacy Ippolito [2020]. In Section 5.1 we’ll introduce differential privacy and federated learning, and in Section 5.2 we’ll discuss early results in privately training language models.

### 5.1 Preliminaries

#### 5.1.1 Differential Privacy

Differential privacy is the standard notion of privacy imposed when the data analyst wants theoretical guarantees that the result of some computation on a private database will not compromise the privacy of the individuals in that database. Introduced by Dwork [2006], private mechanisms inject random noise into aggregate computations in order mask the contribution of any single individual to the result of the computation, at the cost of quantifiable amount of error in the final result. We now define differential privacy in the specific context of a training algorithm  $\mathcal{A}$ .

**Definition 11** ( $\epsilon$ -Differential Privacy, Dwork [2006]) *Let  $\mathcal{X}$  and  $\mathcal{Y}$  be sets of features and labels respectively. A randomized training algorithm  $\mathcal{A} : (\mathcal{X} \times \mathcal{Y})^n \rightarrow \Theta$  is  $\epsilon$ -differentially private if for any two datasets,  $D, D' \in (\mathcal{X} \times \mathcal{Y})^n$  differing in exactly one pair of data points  $(x, y), (x', y')$ , and for all sets  $W \subset \Theta$ :*

$$\Pr(\mathcal{A}(D) \in W) \leq e^\epsilon \Pr(\mathcal{A}(D') \in W)$$

Intuitively, this guarantees that an attacker trying to conduct membership inference would not be able to distinguish whether an individual’s information was a part of the training dataset  $D$  with probability larger  $\frac{1}{1+e^{-\epsilon}}$ , which is a tight bound Izzo et al. [2022]. The smallest value of  $\epsilon$  for which Definition 11 holds is the called the *privacy loss* of the mechanism  $\mathcal{A}$ , and the smaller the value of  $\epsilon$ , the better the privacy guarantees are. The most common variant of differential privacy is a generalization of the above definition called  $(\epsilon, \delta)$ -DP or approximate differential privacy, that allows for a small failure probability  $\delta < \frac{1}{n}$ , which enjoys quadratically better composition properties than  $\epsilon$ -DP and encompasses simple mechanisms like Gaussian noise addition:

**Definition 12** ( $(\epsilon, \delta)$ -Differential Privacy, Dwork [2006]) *Let  $\mathcal{X}$  and  $\mathcal{Y}$  be sets of features and labels respectively. A randomized training algorithm  $\mathcal{A} : (\mathcal{X} \times \mathcal{Y})^n \rightarrow \Theta$  is  $\epsilon$ -differentially private if for any two datasets,  $D, D' \in (\mathcal{X} \times \mathcal{Y})^n$  differing in exactly one pair of data points  $(x, y), (x', y')$ , and for all sets  $W \subset \Theta$ :*

$$\Pr(\mathcal{A}(D) \in W) \leq e^\epsilon \Pr(\mathcal{A}(D') \in W) + \delta$$

While most work on differentially private machine learning has focused on adding the noise in the training process Song et al. [2013], Abadi et al. [2016], on a practical level this can be computationally expensive. There has been recent work on adding noise at the prediction level Dwork and Feldman [2018], Dagan and Feldman [2020], as defined here:

**Definition 13** ( $\epsilon$ -Differential Privacy for Prediction, Dwork and Feldman [2018]) *Let  $\mathcal{X}$  and  $\mathcal{Y}$  be the sets of features and labels respectively. A prediction algorithm  $h$ , whose weights are determined by a training algorithm  $A$ , is  $\epsilon$ -differentially private if for any two datasets  $D, D' \in (\mathcal{X} \times \mathcal{Y})^n$ , differing in exactly one entry, for all  $x \in \mathcal{X}$ , and for all sets  $Y \subseteq \mathcal{Y}$*

$$\log\left(\frac{\Pr(h(x; A(D)) \in Y)}{\Pr(h(x; A(D')) \in Y)}\right) \leq \epsilon$$

For a thorough overview of the fundamental mechanisms used in differentially private algorithms, we refer the reader to Dwork and Roth [2014].

The most common method for private learning is the Differentially Private Stochastic Gradient Descent (DP-SGD) algorithm developed by Abadi et al. [2016] (Algorithm 5.1). DP-SGD works like standard SGD, but at each step of the gradient descent after computing the gradients for each example in the batch, these gradients are clipped to a pre-specified  $l_2$  norm  $C$ , then they are averaged and Gaussian noise is added before taking the descent step. Existing work on private deep learning has shown that the hyperparameters and even activation functions that work best for private training often differ significantly from the optimal hyperparameters for non-private training Papernot et al. [2020].

---

**Algorithm 5.1** Differentially Private SGD (Outline), Abadi et al. [2016]

---

```

1: Input: Examples  $\{x_1, \dots, x_N\}$ , loss function  $\mathcal{L}(\theta) = \frac{1}{N} \sum_i \mathcal{L}(\theta, x_i)$ . Parameters: learning rate  $\eta_t$ , noise scale  $\sigma$ ,
   group size  $L$ , gradient norm bound  $C$ .
2: Initialize  $\theta_0$  randomly
3: for  $t \in [T]$  do
4:   Take a random sample  $L_t$  with sampling probability  $L/N$ 
5:   for each  $i \in L_t$  do
6:     Compute Gradient:
7:      $g_t(x_i) \leftarrow \nabla_{\theta_t} \mathcal{L}(\theta_t, x_i)$ 
8:     Clip Gradient:
9:      $\bar{g}_t(x_i) \leftarrow g_t(x_i) / \max(1, \|g_t(x_i)\|_2 / C)$ 
10:    Add Noise:
11:     $\tilde{g}_t \leftarrow \frac{1}{L} (\sum_i \bar{g}_t(x_i) + \mathcal{N}(0, \sigma^2 C^2 \mathbf{I}))$ 
12:  end for
13:  Descent:
14:   $\theta_{t+1} \leftarrow \theta_t - \eta_t \tilde{g}_t$ 
15: end for
16: return  $\theta_T$  and compute the overall privacy cost  $(\epsilon, \delta)$  using a privacy accounting method.
```

---

### 5.1.2 Federated Learning

Federated learning is a framework used for model training in which the data is stored across a distributed set of devices rather than in a centralized location. Work in this area was initially motivated by the increase in mobile devices for



data storage McMahan et al. [2016], Poushter [2016]. It also then became apparent that federated learning offered a potential avenue for keeping individual sensitive information secure McMahan et al. [2016]. By decentralizing the data storage, federated learning offers more protection against some privacy attacks, since a break of the centralized server would not leak all of the private training data stored on-device.

The Federated Averaging algorithm developed by McMahan et al. [2016] is a variant of stochastic gradient descent designed for the federated learning setting (see Algorithm 5.2). At each step in the optimization, a random subset of clients is selected, and the current global parameter  $w_t$  is broadcast to each client. Each client separately runs SGD for  $E$  epochs using the data on their device, and sends back an updated parameter vector. These parameters are averaged to form the new global model  $w_{t+1}$ . This method is computationally efficient and is controlled by three key parameters:  $C$ , the fraction of clients that perform computation on each round;  $E$ , then number of training passes each client makes over its local dataset on each round; and  $B$ , the local minibatch size used for the client updates. This algorithm opens the door for many more privacy preserving methods as we will discuss in Section 5.2

---

**Algorithm 5.2** Federated Averaging, McMahan et al. [2016]

---

```

1: Input:  $K$  clients indexed by  $k$ ,  $C$  fraction of clients that perform computation each round, local minibatch size  $B$ ,
   number of local epochs  $E$ , and learning rate  $\nu$ .
2: Server Executes:
3: initialize model parameters  $w_0$ 
4: for each round  $t = 1, 2, \dots$  do
5:    $m \leftarrow \max(C \cdot K, 1)$ 
6:    $S_t \leftarrow$  (random set of  $m$  clients)
7:   for each client  $k \in S_t$  in parallel do  $w_{t+1}^k \leftarrow \text{ClientUpdate}(k, w_t)$ 
8:   end for
9:    $m_t \leftarrow \sum_{k \in S_t} n_k$ 
10:   $w_{t+1} \leftarrow \sum_{k \in S_t} \frac{n_k}{m_t} w_{t+1}^k$ 
11: end for
12:
13: ClientUpdate( $k, w$ ): run on client  $k$ 
14:  $\mathcal{B} \leftarrow$  (split  $\mathcal{P}_k$  into batches of size  $B$ )
15: for each local epoch  $i$  from 1 to  $E$  do
16:   for batch  $b \in \mathcal{B}$  do  $w \leftarrow w - \nu \nabla l(w; b)$ 
17:   end for
18:   return  $w$  to server
19: end for
```

---

While Algorithm 5.2 was developed for general machine learning purposes, Hard et al. [2019] implements the algorithm for next-word prediction, the same training objective used for large language models. They showed that a CIFG model trained by Federated SGD performs just as well as one trained in a traditional centralized manner.

## 5.2 Private Language Model Training

We start by discussing earlier work on fully differentially private language models. We then discuss private fine-tuning of LLMs, private LLM predictions, and work that combines DP training with federated learning.

### 5.2.1 Fully DP Language Models

While using differentially private training is great for reducing privacy risks, it introduces a few other challenges. Notably, the utility of the model may suffer and the computational cost of training the model can substantially increase. Work that has applied differential privacy to the whole training process looked to find ways to maintain privacy gains without those downsides.

Anil et al. [2021] researched ways to improve the pre-training accuracy baseline of a BERT-Large model trained using the DP-SGD framework. This model was trained using the Wikipedia Mahoney [2009] and Books Zhu et al. [2015] corpus. Notably, they found that massively scaling up the batch size during training improves the utility of the model. As seen in Figure 13, as the batch size increased from a typical size of 32k to 2 million, the accuracy of the model went up to 60%. While this does not perform as well as the non-private baseline of 70%, it represents a substantial improvement and other works discussed below also found that increasing batch size improved performance. Anil et al. [2021] also found that tuning DP-SGD hyperparameters and scheduling batch size increases as training progressed

improved model utility while maintaining privacy guarantees. While Anil et al. [2021] was able to train a DP guaranteed base model, they did not study how this could affect model utility after fine-tuning for specific tasks.




Figure 13: From Anil et al. [2021]. The accuracy of the BERT-Large model as the batch size varies. The non-private model achieves an accuracy of 70% while the private model achieves up to 60% when using a batch size of 2 million.

Many LLMs tokenize the text data before inputting it into the model. WordPiece Wu et al. [2016] is a popular algorithm used to tokenize inputs and while it has been shown to be very effective, it is not differentially private. Hoory et al. [2021] developed a differentially private version of WordPiece to guarantee no information leakage from these tokens in the DP-SGD process. The WordPiece algorithm starts by constructing a word histogram of the corpus. Using standard DP practices from Dwork [2006], noise is introduced to this histogram, guaranteeing differential privacy. Hoory et al. [2021] measured performance on the Medical Entity Extraction task using the i2b2-2010 and i2b2-2011 datasets Uzuner et al. [2011] when trained using the Wikipedia, BookCorpus Zhu et al. [2015], and MIMIC-III datasets Johnson et al. [2016]. They saw slight tradeoffs between privacy and utility, but overall found that this approach worked well in the medical domain. Similar to Anil et al. [2021], they found that larger batch sizes in the DP-SGD process improved the model performance.

Shi et al. [2022] took a different approach by applying differential privacy only to sensitive information in the text data. The logic is that only a small subset of a text dataset needs to remain private and keeping the rest noise-free will improve model performance. They took the general structure of the DP-SGD algorithm but introduced a policy function,  $F$ , that determines whether a token is sensitive or not. If the token is not sensitive, then the noise addition step is skipped. Shi et al. [2022] built a one-layer LSTM with an embedding size of 200 and a hidden size of 200, and a BPE tokenizer Sennrich et al. [2016]. This was tested two datasets: 1) WikiText-2 Merity et al. [2016] with a policy function treating all digits as sensitive and 2) CustomerSim Tkachenko et al. [2016] with a policy function treating user name, address, phone number, order, and tracking number as sensitive information. This method showed modest improvements in model utility over standard DP-SGD while also having modest privacy improvements from canary insertion and membership inference attacks over non-private training. Shi et al. [2022] did not, however, do extensive research on selection of the policy function and later work by Zhao et al. [2022] showed that in the absence of a perfect policy function the risk of memorization is still high even with only black-box access to the model (see Figure B.1). Developing a perfect policy function is an open ended area for future work.

Zhao et al. [2022] developed a private training method that was an extension of the work from Shi et al. [2022]. Zhao et al. [2022]’s method, called Confidentially Redacted Training (CRT), starts by deduplicating the dataset (see Section 3.2.3) then “redacting” by putting sensitive data into a separate private set. The training process then alternates between SGD from the public dataset and DP-SGD from the private dataset. This was tested on both a one-layer LSTM model and GPT-2 Radford et al. [2019] model on the MultiWOZ 2.2 dataset Zang et al. [2020] and the CustomerSim dataset Tkachenko et al. [2016]. As seen in Figure 14, this method was able to significantly reduce the success of membership inference attacks and similar results were also found for canary extraction attacks (see Figure B.1). Zhao et al. [2022] notes that CRT is not dependent on having access to a perfect policy function. As the error rate,  $\gamma$ , of the policy function changes, CRT maintains performance compared to the non-private models or to the selective DP method of Shi et al. [2022].




Figure 14: Results of a membership inference attack on CRT-trained model from Zhao et al. [2022]. LSTM and GPT refer to model type tested. S-DP-LSTM model was trained using selective differential privacy from Shi et al. [2022].  $\gamma$  is the error rate in the policy function screening for sensitive data. CRT here is effective at preventing membership inference attacks.

### 5.2.2 DP Fine-Tuning of LLMs

Rather than applying differential privacy to the whole training process, some works look to use it in select parts of the training to diminish the downsides of private training. These approaches pre-train the model on a generic corpus using a non-private algorithm. The fine-tuning process is then done using a private learning algorithm and intends to provide privacy to sensitive data which the task specific corpus presumably contains. Kerrigan et al. [2020] was the first work to implement this approach and showed that it boosts the performance of DP LMs while remaining private with respect to the fine-tuning dataset. However, Kerrigan et al. [2020] only evaluated this for simple feedforward networks, which are substantially smaller than state of the art LLMs.

Li et al. [2022] took this same approach but applied it to much larger models and explored how tuning various hyperparameters can vastly improve the performance and reduce the computational cost of DP training. They privately fine-tuned BERT Devlin et al. [2019] and RoBERTa Liu et al. [2019] models for sentence classification tasks and GPT-2 Radford et al. [2019] models for language generation tasks. As seen in Figure 15, higher learning rates and smaller clipping limits improved model performance. And likewise to Anil et al. [2021] and Hoory et al. [2021], they found that increased batch size improves performance. Li et al. [2022] also found that aligning the objective task with the training task improved model performance. For example, a model trained for language generation will not do as well on a classification task unless the task is reframed. For example, for sentiment classification, Li et al. [2022] reframes the problem as filling in the [MASK] token in the sequence “Sentence in question. This is [MASK].” and compare the probabilities of words “awesome” and “terrible”.

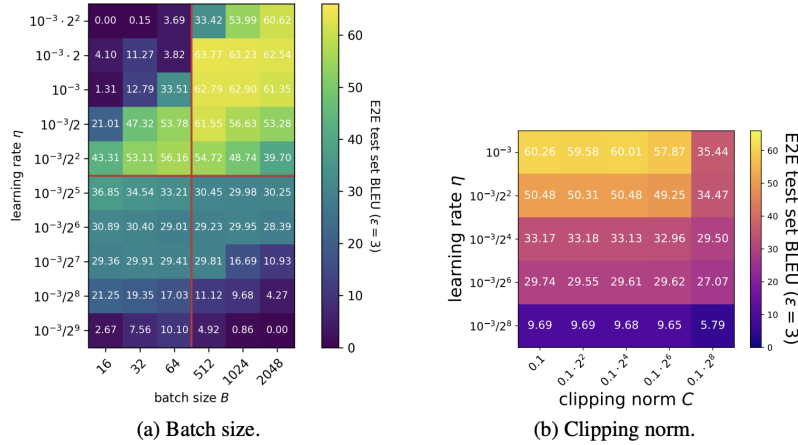


Figure 15: From Li et al. [2022]: (a) Large batch sizes and learning rates lead to better BLEU scores on E2E dataset Novikova et al. [2017] when the number of epochs is fixed. Red lines divide heat map into four broad subsections. (b) Smaller clipping limit  $C$  during DP-SGD leads to better performance.

Li et al. [2022] addressed the issue of additional computational cost by introducing a method they call *ghost clipping*. This is based on the work by Lee and Kifer [2020], who formulated a clipping procedure that only instantiates the per-example gradient for parameters of a single layer in the model one at a time, as opposed to the entire model at once. This meant that per-example gradients did not need to be instantiated. This exact approach does not work for Transformer-based models so Li et al. [2022] extended this by avoiding instantiating the per-example gradient even for individual linear layers. As seen in Figure 16, this method almost entirely reduced the additional memory overhead of the private training method.

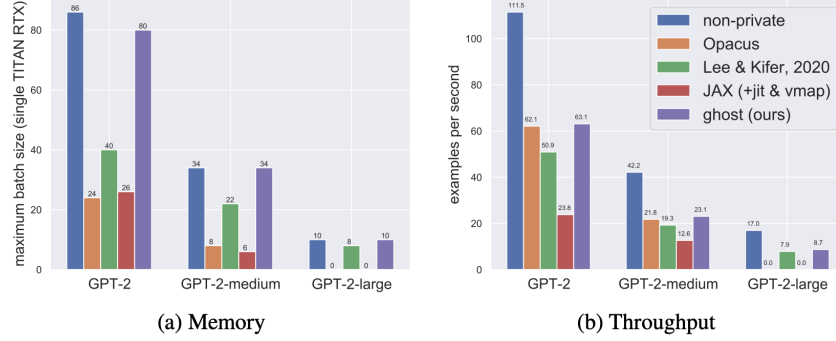


Figure 16: Comparing the non-private implementation (blue) to the implementation from Li et al. [2022] (purple), we see that (a) ghost clipping is almost as memory efficient as non-private training. (b) the throughput of the ghost clipping method offers improvements over other private training methods.

DP training also is substantially slower than non-private training. Dupuy et al. [2022] worked on this problem by adapting the DP-SGD algorithm to allow for GPU usage. They did this by micro-batching the examples at each step in Algorithm 5.1 for the gradient calculation and noise addition steps. Additional adjustments to the scale of the clipping parameter and introducing noise decay were also implemented. Dupuy et al. [2022] tested this on an LSTM and BERT model. With these adjustments, computation time was only a small multiplicative factor worse than non-private training and was up to  $150\times$  faster than standard DP-SGD.

In-context learning (ICL), popularized by Brown et al. [2020], is an alternative approach to fine-tuning. Rather than continuing the training process, the model is provided with a few examples from the training data as part of the prompt. This allows it to adapt to the task at hand without having to modify model itself. However, this can pose a privacy risk as these prompts themselves often contain private data and can be extracted by an adversary Zhang and Ippolito [2023]. Tang et al. [2023] introduces a differentially private approach to generating these prompt examples. Outlined in Figure 17, the prompt prefix is generated one token at a time. At each step in the process,  $M$  disjoint samples of  $N$  examples each are sampled from the private  $D_{train}$ . These subsets are used to generate next-token candidates which are privately aggregated to generate the next token. Tang et al. [2023] tested this using the GPT-3 Babbage model on AGNews, DBpedia, and TREC datasets. This approach was able to produce fully private ( $\epsilon = 0$ ) 0-shot and 4-shot examples. Using the private solution with  $\epsilon = 1$ , the performance was comparable to the non-private solution.




Figure 17: Private ICL prompt generation from Tang et al. [2023].

To take the protection of PII even further, Ishibashi and Shimodaira [2023] developed the approach of *knowledge sanitization*. This approach fine-tunes the model to recognize prompts that are specific knowledge inquiries such as “What is John Smith’s address?” and to respond to all of them with “I don’t know.” This approach successfully preserves PII privacy. However, while this approach did not decrease text perplexity (as a proxy for quality), on a practical level this strict filtering approach could reduce the usefulness of an LLM.

### 5.2.3 DP Predictions

Instead of training the LLM using differential privacy, some works look to apply DP at the prediction stage. Private prediction can be advantageous since the training algorithm does not need to be modified. This makes it much more feasible for applying privacy to large-scale models such as LLMs.

Ginart et al. [2022] developed the first method to perform private next-token predictions on large LLMs such as GPT-2. Their method, called *SubMix* (Figure 18), works by aggregating the output of an ensemble of LMs that were fine-tuned on disjoint parts of the private corpus. At prediction time, if all models agree on the next token, then no noise is added to the output. However, if there is higher disagreement, *SubMix* mixes the predictions with a public pre-trained model to minimize privacy leakage. *SubMix* was tested using a pre-trained GPT-2 model from Hugging-Face Wolf et al. [2020]. The private datasets were the Wikitext-103 Merity et al. [2016] and BigPatent-G Sharma et al. [2019] datasets. Their experiments showed empirical privacy guarantees that outperformed the other baseline DP models which were adapted for next-token prediction: DP-SGD Abadi et al. [2016], S&A Dwork and Roth [2014], and GNMax Papernot et al. [2018]. The primary limitation in the use of *SubMix* is the significant increase in computational and storage capacity needed; experiments showed an increase by a factor of 8 in those areas. See Algorithms B.1 and B.2 for full details.

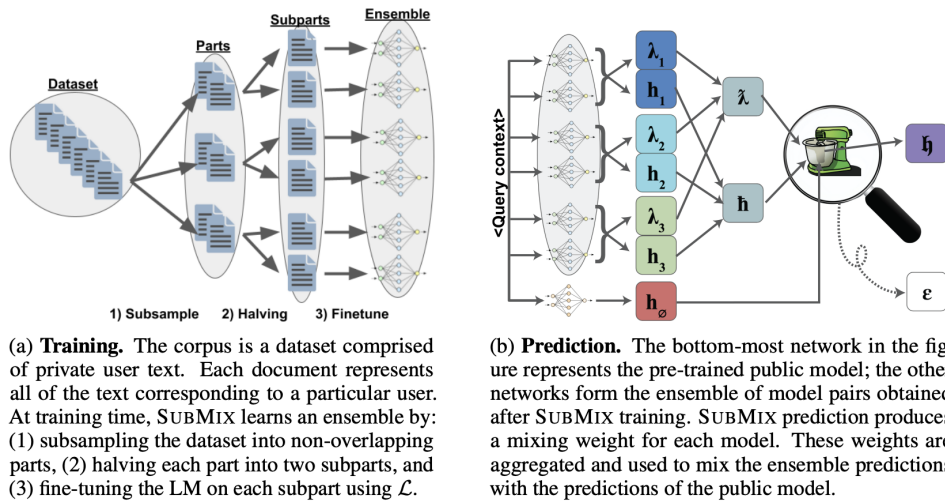


Figure 18: Overview of *SubMix*'s training and prediction protocols from Ginart et al. [2022].

For a much more computationally inexpensive solution, Majmudar et al. [2022] developed a lightweight perturbation mechanism that is only applied at the end by computing a perturbation probability and using that to randomly sample a token from the vocabulary for the  $i$ -th predicted token. This mechanism was tested on a RoBERTa-style encoder model Shoenybi et al. [2020] which was trained on a corpus of Common Crawl, Wikipedia, and mC4 Xue et al. [2021] text data. While this perturbation method is provably differentially private, the utility tradeoff is quite significant. For any  $\epsilon \leq 50$ , there is a significant loss in utility making this method not very useful when strong privacy guarantees are desired. The privacy-utility tradeoff can be seen in Figure B.2

### 5.2.4 Federated Learning

There have been several papers that trained LLMs using both DP and federated learning. One of the early implementations was by McMahan et al. [2018], who built a differentially private next-word prediction model. This experiment trained an LSTM recurrent neural network on a dataset of Reddit posts Al-Rfou et al. [2016]. More specifically, it modified the Federated Averaging algorithm to add Gaussian noise to each iteration of the averaging update (see Algorithm B.3). While McMahan et al. [2018] used a significantly smaller language model than the LLMs that are currently used, it was the first high quality language model with strong privacy guarantees and no significant decrease in model accuracy (see Table 6).

Additional work further looked at combining differential privacy with federated learning to achieve strong privacy guarantees. Thakkar et al. [2020] used the CIFG-LSTM model Sak et al. [2014] trained on a modified Stack Overflow Q&A dataset hosted by TensorFlow Federated Overflow [2019] and used the canary extraction test method outlined in Carlini et al. [2019] to measure memorization. Federated learning alone reduced the number of extracted canaries by

model		data			
$\sigma$	$S$	$K$	$\tilde{C}$	$\epsilon$	AccT1
0.000	$\infty$	763430	100	$\infty$	17.62%
0.003	15	763430	5000	4.634	17.49%
0.006	10	763430	1667	2.314	17.04%
0.012	15	763430	1250	2.038	16.33%
0.003	15	$10^8$	5000	1.152	17.49%
0.006	10	$10^8$	1667	0.991	17.04%
0.012	15	$10^8$	1250	0.987	16.33%

Table 6: Results of DP-FedAvg (Algorithm B.3) McMahan et al. [2018]. Privacy and accuracy (measured as Top 1 accuracy) after 5000 rounds of training for different Gaussian noise  $\sigma$  (at  $\delta = 10^{-9}$ ), clipping parameter  $S$ , number of users  $K$ , and number of individuals sampled per round  $\tilde{C}$ . We can see that this is an effective private learner. For example, adding just  $\sigma = .003$  Gaussian noise yields a privacy  $\epsilon = 4.634$  with only a minor drop in accuracy. The  $\epsilon$  guarantee is even better for larger  $K$

50%. When combined with differential privacy, there was another significant drop in canary extraction ranging all the way to no canaries being extracted. Ramaswamy et al. [2020] built a consumer scale next-word prediction model using the same base model, dataset, and memorization test as in Thakkar et al. [2020]. This private model was built using differentially private federated averaging, and was found to maintain privacy against the canary extraction test while showing some improvement in next-word prediction recall compared to a baseline  $n$ -gram Finite State Transducer model.

It is worth noting that a formal DP guarantee requires knowledge of the size of the participating user population and the ability to uniformly randomly sample among users at each step. In the real world, this assumption does not always hold, which means that DP-FedAvg carries some minor limitations.

### Additional Remarks on Private Models

Xiao et al. [2023] proposed multiple fine-tuning methods to specifically protect PII information. These were 1) manually replacing PII with preset tokens, 2) introducing a training penalty on PII, 3) using a classifier to replace PII with preset tokens, 4) using few-shot examples as part of the prompt, and 5) using Direct Preference Optimization Rafailov et al. [2023]. These were tested using the LLaMA2 model Touvron et al. [2023] on the MedAlpaca dataset Han et al. [2023], which was annotated using the scrubadub repo LeapBeyond [2023]. All of these methods were able to preserve PII privacy but most saw a performance drop when evaluated on ROUGE Lin [2004] and BERTScore Zhang et al. [2020]. The instruction tuning using few-shot examples was found to be the best balance between privacy and goal alignment.

## 6 Copyright

### 6.1 Introduction

In the past few years, there has been a significant increase in the use of AI to generate music, code, and art of all sorts Vincent [2022]. While there have been fascinating applications in this area, it also introduces a wide range of ethical and legal issues raised based on the models’ propensity to copy artists’ material Heikkila [2022b]. As discussed in Section 3, LLMs often memorize the data they are trained on, and many popular datasets contain copyright data Bandy and Vincent [2021], Biderman et al. [2022]. This leads to significant risk of copyright violation by these models. Indeed, comedian Sarah Silverman and others recently sued OpenAI and Meta for copyright infringement via the ChatGPT and LLaMA models respectively Davis [2023b].

Karamolegkou et al. [2023] ran an extensive test on how much language models memorized copyright data of best selling books published between 1930 and 2010. Karamolegkou et al. [2023] tested OPT, Pythia, LLaMA, Falcon, Claude, and GPT-3.5 by prompting open-source models with the first 50 tokens of the book and closed-source instruction-tuned models with the prompt “What is the first page of <title>?”. In Figure 19, we can see the Longest Common Substring (LCS) that was memorized by the models and in each book. Consistent with results discussed in Section 3, larger models were more susceptible to copyright violation. Meanwhile, popular books had large parts of their opening text memorized by these popular models.

Despite these apparent issues, many major companies and organizations are pushing ahead to augment or replace creative work with AI. The Alliance of Motion Picture and Television Producers (AMPTP) have explored the use of



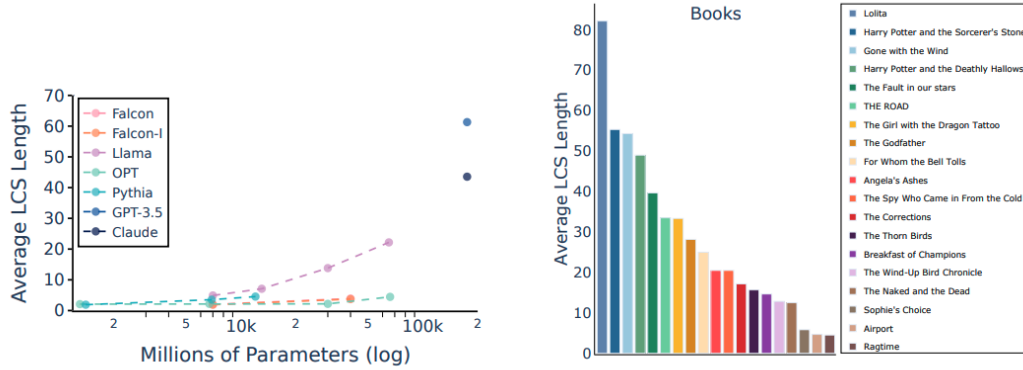


Figure 19: From Karamolegkou et al. [2023] showing verbatim memorization in books. On the left, we see the average length of copyright material generated by various models. On the right we see the length of copyright material memorized by book, with more popular books being more susceptible.

models such as ChatGPT for writing screenplays, which is one of the major issues in the Writers Guild of America (WGA) strike of 2023 Rose [2023a]. This strike is seen by some as the first of many legal battles in an oncoming labor war over the use of AI Broderick [2023].

In order to safely deploy these models without risk of copyright violations, more solutions and modifications are needed. It has already been shown that preventing verbatim memorization is not a perfect solution Ippolito et al. [2023] and in the case of copyright protection, even preventing approximate memorization is likely insufficient. Determining copyright violation also requires measuring creative novelty and intended use, which are difficult to quantify. In this section, we discuss the current legal literature on copyright in LMs, as well as any technical solutions that have been proposed. All discussion in this section focuses on US copyright law. The methodologies may still be useful in other jurisdictions but are not necessarily tailored for different nuances of the law.

## 6.2 Legal Discourse

With the rapid rise of generative models, there are a wide range of new issues that the legal community must deal with. Legal papers such as Delacroix [2023] describe the simultaneous issues of machine learning models using copyright data while also generating data that itself could be considered copyright. This could lead to a downward spiral with large quantities of data sources existing in a legal limbo.

**Fair Use Doctrine:** The primary legal framework applied in the context of generative models such as LLMs is the *fair use doctrine*. This is the framework that determines whether new material that relied upon copyright data did so in a manner that is novel enough to avoid any legal issues. Under this doctrine, researchers are allowed limited use of copyright material, even without a license, in order to foster creativity and innovation. While landmark decisions on fair use in technology has happened in the context of code with the Oracle vs Google case Har [2021], there have not been any substantial decisions in the case of LLMs. Depending on how courts decide, either model developers could be stifled, or copyright owners could lose out Sobel [2017]. Despite no legal decisions demanding it, there is some research on ways to avoid any potential issues as we will discuss in Section 6.3. For extensive coverage of fair use doctrine and its applications to all foundation models, see Henderson et al. [2023].

**Using Copyright Data:** When it comes to the question of including copyright data in training datasets, most legal discourse agrees that this is an acceptable practice. Just as any individual is free to access and learn from copyright material without being in violation of the law, machine learning models are argued to be given those same rights with any potential legal issues only arising if the final model is used to generate content Lemley and Casey [2021]. As stated by Sag [2019], “copying expressive works for the purposes of non-expressive use should not be a copyright violation (e.g. text mining in and of itself is not a violation).”

**Models as Authors:** Another topic of discussion in the legal community is whether material generated by a machine learning model can be considered novel enough to grant patents or copyright to the model or model developers. Hristov [2017] argued that giving authorship to AI programmers and owners is essential to the future development of the AI industry. However, legal decisions since then have leaned in the opposite direction. The US Court system has already

decided that work by AI artists and authors is ineligible for copyright Copyright [2023] and inventions by AI are ineligible for patent rights Rose [2023b].

### 6.3 Technical Mitigation

The risk of a language model infringing on copyright law is high unless mitigation strategies are implemented. Verbatim memorization filters are insufficient Ippolito et al. [2023] and fair use law incorporates some degree of human judgement, making it difficult to algorithmically determine copyright violation. Legal scholars recognize that fair use issues may never be able to be reduced to an algorithm Burk [2019]. That said, there are technical mitigation strategies that would reduce the potential for copyright violation which we lay out here:

**Data Filtering:** The most straightforward solution to preventing copyright violation would be to remove any training data which is copyright. Without access to that training data, it is very unlikely that copyright data would be output by the generative model. Even if it were to be generated, it would not pose a legal issue; as noted in Feist v Rural US Supreme Court [1991], “...a work may be original even though it closely resembles other works, so long as the similarity is fortuitous, not the result of copying...” That said, on a practical level this would be very difficult to implement. Most LLMs are trained on corpuses built on large open web crawls and while opt-out strategies have been implemented for many of these, they are not yet widespread. Additionally, it is unclear whether models built strictly on open source, non-copyright, and permissively licensed material would even perform well. Using only this data could potentially bias the models in other negative ways Levendowski [2017] such as any literature only being from the 1920’s and earlier due to public domain laws Henderson et al. [2023]. Data deduplication and other memorization mitigation strategies would help (Section 3.3), but are not proven guarantees.

**Unlearning:** In the case of LLMs, retraining them would be a significant computational undertaking and would be very expensive, both in terms of time and money. As discussed in Section 7, machine unlearning is a significantly more efficient method to remove problematic data and could be readily applied to copyright data. However, unlearning a large portion of data could be computationally intractable, and have a significant impact on model performance. Given a targeted use case, this has potential as a solution, however no significant research has been done on applying unlearning to copyright protection.

**Measuring Influence of Training Data:** In line with the fair use principle that material is not a copyright violation if its resemblance to the copyright material is purely coincidental, one approach towards quantifying violations is to measure the effect of the training data on the specific output of the model. Output that is similar to prior material is only a copyright violation if it was dependent that material’s existence in the training set.

Vyas et al. [2023] developed a formal definition called *k-Near Access-Free* (*k*-NAF) to measure this. Let  $D_{train}$  be the training dataset, which may contain some material  $C$  from the set of all copyright material  $\mathcal{C}$ .  $h$  is our generative model trained by algorithm  $\mathcal{A}$  on  $D_{train}$ , which given some prompt  $x \in \mathcal{X}$ , produces output  $y \in \mathcal{Y}$  with probability  $p(y|x)$ . We would like  $h$  to be within  $k$ -bits of information (under some divergence measure) from  $safe_C$ , which is a model trained using  $\mathcal{A}$  on a dataset that doesn’t contain  $C$ . Vyas et al. [2023] then establishes their criteria for copyright protection:

**Definition 14** (*k-Near Access-Free*, Vyas (2023) Vyas et al. [2023]). Let  $C$  be a set of copyright datapoints,  $safe_C$  be a model trained without access to  $C$ , and let  $\Delta$  be a divergence measure between distributions. We say that a generative model  $h$  is  $k_x$ -near access-free ( $k_x$ -NAF) on prompt  $x \in X$  with respect to  $C$ ,  $safe$ , and  $\Delta$  if for every  $c \in C$ ,

$$\Delta(p(\cdot|x)||safe_C(\cdot|x)) \leq k_x$$

Vyas et al. [2023] defaults the choice of divergence to either maximum KL divergence  $\Delta_{max}$ , also known as the Rényi divergence of order infinity, or the KL divergence  $\Delta_{KL}$ . This definition being based on a comparison to a safe model allows for fortuitous copying and is robust to cases such as someone prompting “print the following text:  $C$ ” where  $C$  is a copyright material. This shares some similarities to the principles of differential privacy Dwork [2006]. However, as outlined by Elkin-Koren et al. [2023], there are several traits of copyright law that distinguish it from DP:

1. Privacy in general prevents knowledge of the prior material whereas copyright enables the use of that knowledge for creative gain. Indeed, copyright specifically allows for direct use in the case of quotations, parodies, and research.
2. Copyright protections are limited in time; after a certain period of time, they enter the public domain.
3. Copyright excludes raw materials used for cultural expression such as facts, ideas, and methods of operation.
4. Differential privacy allows for some small probability of generating the exact training example, whereas in copyright, unless that material was permitted by one of the criteria above, generating the exact training example would result in legal issues.



$k$ -NAF follows the characteristics of copyright outlined. Vyas et al. [2023] also shows that there exist algorithms which satisfy  $k$ -NAF, for reasonable choices of  $k$ . The first, Copy-Protection- $\Delta$  is computationally difficult for LLMs. The Smooth-Copy-Protection- $k$  algorithm, however is a fit. Say  $\mathcal{V} = \{q_1, \dots, q_n\}$  is a cover of **safe** if  $\forall C \in \mathcal{C} \exists q \in \mathcal{V}$  s.t.  $\text{safe}(C) = q$ . Then given a model  $h$ , cover  $\mathcal{V}$ , and threshold  $k$ , Algorithm 6.1 returns a model  $h_k$  which has quantifiable access-free guarantees with respect to **safe**. These algorithms assume access to a list of copyright material along with access to an oracle where we can both compute conditional probabilities and obtain samples under the models  $h$  and  $q \in \mathcal{V}$

---

**Algorithm 6.1** Smooth-Copy-Protection- $k$ , Vyas et al. [2023]

---

```

1: Input: model  $h$  (which given prompt  $x$  outputs  $y$  with probability  $p(y|x)$ ), cover  $\mathcal{V}$  of safe, and threshold  $k \geq 0$ .
2: return  $h_k$  where  $h_k$  is specified as:
3: while True do
4:   Sample  $y \sim p(\cdot|x)$  and
5:   return  $y$  with probability  $\min \left\{ 1, \min_{q \in \mathcal{V}} \left\{ \frac{2^k q(y|x)}{p(y|x)} \right\} \right\}$ 
6: end while

```

---

**Measuring Novelty:** There are a few other methods that have been developed for measuring novelty and creativity in text data. The motivation behind this is that the legal system could use these to audit model output for copyright violations. By quantifying the novelty of a text, there is a more objective way of measuring “substantial similarity.” Scheffler et al. [2022] introduced one such framework called *derivation similarity*, which uses Levin-Kolmogorov cost to capture novelty.

**Definition 15** (*Conditional Levin-Kolmogorov Cost, Scheffler et al. [2022]*) Let  $M$ ,  $b$ , and  $y$  be strings in  $\{0, 1\}^*$ . Let  $U(M, b, t)$  be the output of a universal prefix Turing machine running program  $M$  for  $t$  timesteps, where  $M$  has access to input tape  $b$ . Then, the cost of running  $M$  on input  $b$  to get a string “comparable” to  $y$  is

$$C(M, z|b) = \min_{t \in \mathbb{N}} \{ |M| + \lceil \log(t) \rceil : z \leftarrow U(M, b, t) \}$$

$$\tilde{C}(M, z|b) = \min_{z \cong y} \{ C(M, z|b) \}$$

Where  $z \cong y$  denotes that strings  $z$  and  $y$  are comparable. In Scheffler et al. [2022] this denoted exact equality but it was noted this could be applied to different notions such as being within a certain Hamming distance.

In the context of the legal system, Scheffler et al. [2022]’s framework has the plaintiff provide the smallest possible “Producer” algorithm  $P$  capable of independently producing something comparable to  $y$ , given full access to the original work  $x$  and the defendant provide the smallest “Reproducer” algorithm  $R$  that produces a work comparable to  $y$  without using the copyrightable parts of  $x$ , i.e., in the manner that the defendant alleges that they initially created  $y$ . The costs of  $P$  and  $R$  are then compared to measure the extent to which  $y$  intrinsically relied on the creative expression within  $x$ .

**Definition 16** (*Empirical Derivation Similarity, Scheffler et al. [2022]*) Let  $x$  and  $y$  denote the original and derived works, respectively. Given a producer algorithm  $P$  and a reproducer algorithm  $R$ , the empirical derivation similarity is defined as:

$$\text{DerSimEmp}(x, y, P, R|bg) = \tilde{C}(R, y|bg) - \tilde{C}(P, y|(bg, x))$$

Where  $bg$  is the set of all aspects of  $x$ ,  $y$ , and other contextual works that the parties both agree are non-copyrighted.

This framework provides a mathematical way of measuring the similarity of two texts in the context of a legal discussion. Scheffler et al. [2022] applied this to several past court cases and showed that their output agreed with the decisions in those cases. However, they acknowledge that it has limitations when either party has secret information that compromises  $bg$  or acts in bad faith when agreeing upon the contextual works included in  $bg$ . This framework also only covers substantial similarity and does not go into all the nuances of fair use doctrine or cover the Digital Millenium Copyright Act (DMCA) anti-circumvention measures. The primary use of Scheffler et al. [2022] is to answer the narrow question of whether two works are similar.

Franceschelli and Musolesi [2022] separately created a framework named DeepCreativity. This framework is based on Boden [2005]’s work that creativity is composed of value, novelty, and surprise. By associating components of different deep learning frameworks to those three components, Franceschelli and Musolesi [2022] created a quantitative measure of creativity. While this was not created for copyright assessment purposes, quantitative measures like this could in theory be applied to this kind of research.

## 7 Machine Unlearning

As discussed in the previous sections, there are several pressing privacy issues underlying the use of LLMs. While these issues were initially subject to scrutiny primarily in the academic setting, they are of growing concern to the greater public. Popular examples about personal data being revealed by machine learning models, such as in Heikkila [2022a], have led to many people being distrustful of how their data is being handled Auxier et al. [2019].

Given these growing concerns by the public, legislation and regulations are being introduced to protect consumer data. The most influential one is the European Union’s General Data Protection Regulation (GDPR) European Parliament and Council of the European Union [2016], which governs how businesses handle personal data. In particular, one element of the GDPR is known as the *Right to Be Forgotten*. This is the right of every consumer to request that their personal data be deleted and is also a provision in the California Consumer Privacy Act (CCPA) State of California Department of Justice [2023], and the proposed New York Privacy Act (NYPA) State of New York [2023] and Consumer Privacy Protection Act (CPPA) of Canada Government of Canada [2023]. While this would be a straightforward request in many settings where personal data would be stored in a table, in the context of machine learning models and LLMs, compliance challenges are significantly harder Zhang et al. [2023b]. Even if an individual is deleted from a database, if their data was previously used to train a model, it can be extracted through privacy attacks such as those discussed in Section 4. This risk is being recognized more by regulatory bodies, including the Federal Trade Commission recently coming to a settlement with a facial recognition software company that in addition to a user’s data, any models or algorithms derived from their data must be deleted Commission [2021]. This sets a precedent that may become more frequent in legislation and regulation.

This creates a difficult situation for model developers. The most naive solution, known as Leave One Out (LOO) retraining, would just delete an individual’s data and retrain the whole model. However, modern LLMs can take several weeks and millions of dollars to retrain from scratch making this solution infeasible in practice. According to Help Net Security, consumer data deletion requests have increased by 74% from 2021 to 2022 Security [2023]. If models were required to be completely retrained each time one of these requests came in, they would be impossible to maintain, especially given the size of the regions with “Right to Be Forgotten” laws either in place or proposed. This has lead to a new field of research called *machine unlearning*. This field looks to delete an individual’s data and influence from a model using significantly fewer computational resources than retraining it. A successful unlearning algorithm should produce a model that is approximately indistinguishable from the model produced by retraining. In this section, we discuss the general principles of machine unlearning and current developments within the space of language models. For a more comprehensive survey of machine unlearning, see Nguyen et al. [2022].

### 7.1 Preliminaries

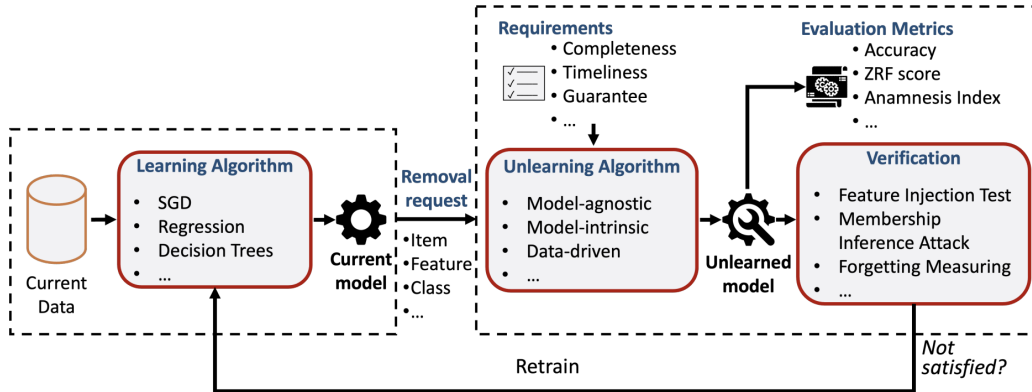


Figure 20: Machine Unlearning framework (from Nguyen et al. [2022])

**Unlearning Framework:** The general framework of machine unlearning is shown in Figure 20. In this framework we have a dataset  $D \in \mathcal{D}$  and training algorithm  $\mathcal{A} : \mathcal{D} \rightarrow \mathcal{H}$  where  $\mathcal{H}$  is the hypothesis space of our model. We also have an unlearning algorithm  $\mathcal{U} : \mathcal{H} \times D \times D_f \rightarrow \mathcal{H}$ , where  $D_f$  is the subset of  $D$  that we want to forget. We would like to have an unlearning algorithm such that  $\mathcal{U}(D, D_f, h)$  is close to  $\mathcal{A}(D \setminus D_f)$ . It is assumed that  $\mathcal{A}$  and  $\mathcal{U}$  are randomized algorithms. Let  $\Pr(\mathcal{A}(D))$  be the distribution of all models trained on  $D$  by  $\mathcal{A}$  and  $\Pr(\mathcal{U}(D, D_f, h))$  be the distribution of unlearned models. We then have two categories of unlearning. The first is *exact unlearning*:

**Definition 17** (*Exact Unlearning, Nguyen et al. [2022]*) Given a learning algorithm  $\mathcal{A}$ , we say the process  $\mathcal{U}$  is an exact unlearning process iff  $\forall \mathcal{T} \subseteq \mathcal{H}, D \in \mathcal{D}, D_f \subset D$ :

$$\Pr(\mathcal{A}(D \setminus D_f) \in \mathcal{T}) = \Pr(\mathcal{U}(D, D_f, \mathcal{A}(D)) \in \mathcal{T})$$

That is, the model resulting from  $\mathcal{U}$  is the same as  $\mathcal{A}$  trained without the deleted point. A closely related definition that relaxes this requirements is *approximate unlearning*:

**Definition 18** ( *$\epsilon$ -Approximate Unlearning, Nguyen et al. [2022]*) Given  $\epsilon > 0$ , an unlearning mechanism  $\mathcal{U}$  performs  $\epsilon$ -certified removal for a learning algorithm  $\mathcal{A}$  if  $\forall \mathcal{T} \subseteq \mathcal{H}, D \in \mathcal{D}, d \in D$ :

$$e^{-\epsilon} \leq \frac{\Pr(\mathcal{U}(D, d, \mathcal{A}(D)) \in \mathcal{T})}{\Pr(\mathcal{A}(D \setminus d) \in \mathcal{T})} \leq e^{\epsilon}$$

Where  $d$  is the removed sample.

**Definition 19** ( *$(\epsilon, \delta)$ -Approximate Unlearning, Nguyen et al. [2022]*) Given  $\epsilon, \delta > 0$ , an unlearning mechanism  $\mathcal{U}$  performs  $\epsilon$ -certified removal for a learning algorithm  $\mathcal{A}$  if  $\forall \mathcal{T} \subseteq \mathcal{H}, D \in \mathcal{D}, d \in D$ :

$$\Pr(\mathcal{U}(D, d, \mathcal{A}(D)) \in \mathcal{T}) \leq e^{\epsilon} \Pr(\mathcal{A}(D \setminus d) \in \mathcal{T}) + \delta$$

and

$$\Pr(\mathcal{A}(D \setminus d) \in \mathcal{T}) \leq e^{\epsilon} \Pr(\mathcal{U}(D, d, \mathcal{A}(D)) \in \mathcal{T}) + \delta$$

That is,  $\epsilon$  defines how similar the new model must be and  $\delta$  in Definition 19 is the upper bound of the probability of Definition 18 failing. Note that in comparison to Definition 17, the approximate unlearning definitions are defined with respect to a single removed sample rather than a whole subset. It is currently an open question in unlearning whether constant bounds can be provided for larger subsets Nguyen et al. [2022]. Approximate unlearning is very similar to differential privacy Dwork [2006]; indeed, differential privacy implies approximate unlearning, but is a much stronger condition.

**Unlearning Sequences:** In dealing with personal data, these deletion requests may not come cleanly as a single batch. In a practical setting, deletion requests will come as a sequence with a certain order. This poses a theoretical issue as some unlearning methods rely on the assumption that deletion requests are independent of each other. It is very possible that users may want to delete their data after learning what the deployed models reveal about them; this is known as an *adaptive sequence*. Gupta et al. [2021] provides the generic reduction for this as follows:

1. If a data deletion algorithm  $\mathcal{U}$  for a learning algorithm  $\mathcal{A}$  has deletion guarantees for independent sequences of deletion requests (as those from past work do), and
2. Information about the internal randomness of  $\mathcal{U}$  is revealed only in a manner that satisfies differential privacy, then

Then  $(\mathcal{A}, \mathcal{U})$  also satisfies data deletion guarantees against an adaptive sequence of deletion requests, that can depend in arbitrary ways on the information that the model provider has made public. Data addition requests may also be a part of this update sequence.

## 7.2 Approaches to Unlearning

While there has been limited work in machine unlearning in the context of language models, the approaches to unlearning algorithms have generally followed one of two methods: 1) efficient LOO retraining and 2) gradient-based updates.

**Efficient LOO Retraining:** Methods in this area look for more efficient ways to simply remove or scramble the data that has been requested for deletion, then retrain the model from an intermediate point to save computational costs. Graves et al. [2020] introduced a method where the requested deletion point  $d$  is removed and replaced with copies of  $d$  with random labels. Then, using the same training algorithm  $\mathcal{A}$ , more training iterations are run until privacy risks are no longer an issue. They showed that this method was effective against both model inversion and membership inference attacks after only a few training iterations. The downside of this method is that the data holder maintains a copy of  $d$  during the unlearning process, which could hold legal implications. Additionally, if a large quantity of points is requested for deletion, the model performance could possibly suffer.

Bourtole et al. [2020] created an approach that cleverly restructures the way the data is fed into the model. Their method is called SISA for Sharded, Isolated, Sliced, and Aggregated training and is depicted in Figure 21. In the SISA

framework, the training data is divided into multiple disjoint shards such that each data point is only contained in a single shard. The data in each shard is subsequently divided into multiple slices. For each shard, an independent model is trained iteratively using the slices, recording the model parameters before each new slice is introduced. When an output is needed, the outputs of these constituent models are aggregated for the output; the most simple method is to use a majority vote over the predicted labels. Under this framework, when a deletion request for point  $d$  comes in, only one of the constituent models must be retrained and only starting from the slice which contains  $d$ . Compared to naive retraining, the SISA framework for simple learning tasks resulted in a time improvement of  $4.63\times$  on the Purchase dataset Sakar et al. [2019] and  $2.45\times$  on the SVHN dataset over naive retraining. For more complex learning tasks such as ImageNet classification, SISA resulted in a  $1.36\times$  speed improvement. Gupta et al. [2021] showed that this method can be adapted to handle unlearning adaptive sequences if the shards are selected independent of each other.

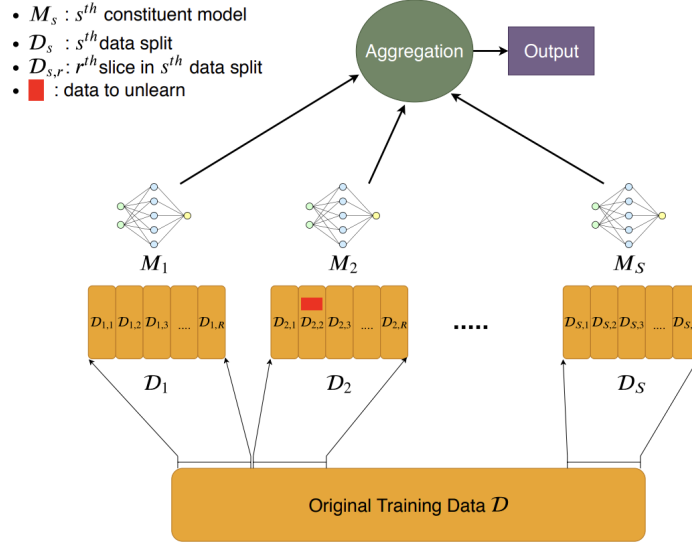


Figure 21: SISA framework Bourtole et al. [2020]: The data is divided into shards, which are then divided into slices. The constituent models’ outputs are trained incrementally using the slices, saving at each step. The constituent models are then aggregated for the final output. When an unlearning request for  $d$  comes in, only the constituent model with  $d$  must be retrained, and only starting at the slice containing  $d$ .

In traditional statistics, the influence function Hampel [1974] Cook [1979] is a technique that can estimate the effect of deleting a group of points at the same computational cost of taking a second-order Taylor approximation around the model parameters. Given a strongly convex training objective, influence functions would be expected to align well with LOO retraining. However, Basu et al. [2021] found that in the case of neural networks, the influence function’s estimates are “fragile” and “erroneous”. Bae et al. [2022] later decomposed the reason for this discrepancy into three primary reasons, but still found that influence functions are a poor replacement for LOO retraining on nonlinear networks.

**Gradient-Based Updates:** This next class of methods rely on gradient descent techniques typically used to train ML models to execute efficient unlearning. Neel et al. [2020] gives an approach to unlearning in this fashion. In their framework, the results of the gradient descent learning algorithm,  $\mathcal{A}$ , are published with some Gaussian noise. The unlearning algorithm,  $\mathcal{U}$ , is then simply a few more gradient descent updates run without the deleted point  $d$  in the dataset. Due to the Gaussian noise added, the outcomes of  $\mathcal{U}(D, d, h)$  and a full retraining are provably indistinguishable. This algorithm is able to handle arbitrarily long sequences of updates in logarithmic time, maintaining strong unlearning from Definition 21. For sufficiently high dimension data, Neel et al. [2020] further improved on this by bootstrapping the dataset and training the model in the same fashion as the SISA framework such that the gradient update only needs to be applied to the affected shard. This means more updates can be performed while the bounded accuracy error of each shard is maintained even over an arbitrary sequence of updates Zhang et al. [2013].

Graves et al. [2020] introduced a second method in addition to their LOO method described earlier. This method, called “amnesiac unlearning”, requires the model owner to keep track of which examples appeared in which training batch along with the parameter update from that batch. Then, when a deletion request comes in for  $d$ , the owner subtracts those updates from the final parameters. As long as the number of batches affected is small, this method is very efficient and effective. However, since parameter updates are dependent on the current parameters when that step takes place,

$$\min_{w^u} \frac{\alpha}{N_r} \sum_{x_r \in \mathcal{D}_r} d(x_r; w^u) + \frac{\gamma}{N_r} \sum_{(x_r, y_r) \in \mathcal{D}_r} \ell(f(x_r; w^u), y_r) - \frac{1}{N_f} \sum_{x_f \in \mathcal{D}_f} d(x_f; w^u)$$

Figure 22: The SCRUB training objective Kurmanji et al. [2023]. Here  $\mathcal{D}_r, \mathcal{D}_f$  are the forget and retain sets respectively, where  $|\mathcal{D}_r| = N_r, |\mathcal{D}_f| = N_f$ .  $w^u$  represents the parameters of the new unlearned model, and  $d(x; w^u)$  represents the KL divergence between the softmax outputs of  $w^u$  on input  $x$ , and the softmax outputs of the original trained model on  $x$ . The middle term is the standard training loss on the retain set.

this unlearning method is approximate rather than exact. Another downside is that a large amount of storage space is needed to track the parameter update values.

Kurmanji et al. [2023] develop a gradient based unlearning algorithm called Scalable Remebering and Unlearning Unbound (SCRUB) where rather than doing gradient descent on the retain set to unlearn as in Neel et al. [2020], or gradient ascent on the forget set Jang et al. [2023], they optimize a hybrid objective in Figure 22. First they initialize the new unlearned model  $w^u$  at the existing weights  $w^f$ , then they simultaneously force the predictions of  $w^u$  on the forget set away from the predictions of  $w^f$  on the forget set, while keep the predictions on the retain set close. They also add a term in the objective to force the accuracy on the retain set to stay high. Since they are simultaneously minimizing the first two terms in Figure 22 while maximizing  $\sum_{x_f \in \mathcal{D}_f} d(x_f; w^u)$ , they alternate between performing an epoch of updates on the forget set with an epoch of updates on the retain set. Note that this basic procedure will result in maximizing forget error (by forcing it away from  $w^f$ ’s predictions), rather than trying to match the error on the forget set to the out of sample error, as is more typical in unlearning applications. They make the distinction that this variant of SCRUB is appropriate for removing biased or poisoned data where the goal is not to simulate retraining, but to maximize error on the forget set. In order to adapt SCRUB for unlearning points for the sake of privacy, the difficult question is when to halt the optimization so that the error on the forget set approaches the error on the retrained model. To accomplish this, they propose the SCRUB + R algorithm, that does the following: (i) Constructs a validation set  $\mathcal{D}_v$  that approximates  $\mathcal{D}_f$  (ii) Trains SCRUB as usual, storing a copy of the model every epoch (iii) At the end of SCRUB training, compute  $\alpha_v = \frac{1}{N_v} \sum_{(x,y) \in \mathcal{D}_v} \ell(f(x, w^u), y)$  (iv) Rewind SCRUB training to output the checkpoint whose error on  $\mathcal{D}_f$  is closest to  $\alpha_v$ . They remark: “The intuition is that the last step of unlearning approximates ‘maximally forgetting’. Consequently, the error on the identically-distributed validation set approximates the error of a model that never learned about that distribution from the forget set: any correct predictions on the held-out examples are due only to the generalization power of the model that was trained only on the retain set. Therefore, we choose that validation set error to serve as the reference point for how high we would like the forget set error to be for UP applications.” They evaluate their algorithm on CIFAR-10 and Lacuna-10 and find that across a variety of metrics including MIA attack success and loss on the forget set, SCRUB outperforms prior methods based on fine-tuning (gradient descent on  $\mathcal{D}_r$ ), gradient-ascent on  $\mathcal{D}_f$ , NTK Golatkar et al. [2020], and Fishers Golatkar et al. [2019], although simultaneous gradient descent and ascent on  $\mathcal{D}_r, \mathcal{D}_f$  respectively is a strong baseline, and is close to SCRUB in flavor.

### 7.3 Unlearning in LLMs

While work on applying machine unlearning to LMs is limited, there are a couple of notable works in the space. Kumar et al. [2022] expands upon the SISA framework for Bourtole et al. [2020]. While the SISA framework showed computational improvements over naive deletion and retraining, it is impractical for application to LLMs since the number of model parameters and data points makes it extremely computationally expensive. Kumar et al. [2022] proposed two variants of SISA that improve it for LLMs. The first is called *SISA-FC*. Since saving a checkpoint for every slice during the constituent model training is impossible for large text corpora, this framework starts with an LM trained on a generic corpus. The SISA framework is then only used for the task-specific fine-tuning process. While *SISA-FC* reduces the retraining time and computational storage, model performance generally performs better if the model is fine-tuned for the task the entire time. The second variant, known as *SISA-A*, addresses this by using the SISA framework for the entire training. To save memory, the *SISA-A* framework trains the model using the parameter efficient adapter method from Houlsby et al. [2019]. Adapter modules are used to replace the model parameters and are significantly smaller, using 1 – 5% the memory that storing all the model parameters does. This significantly reduces the memory cost of checkpointing during the SISA framework. Kumar et al. [2022] tested these two frameworks on three Glue tasks/corpora Wang et al. [2018]: SST-2 Socher et al. [2013], QQP Wang et al. [2018], and MNLI Williams et al. [2018]. They trained a BERT model using the different approaches and measured the accuracy on these tasks, the time it took to retrain the model after a deletion request, and the amount of memory the final model takes up. *SISA-A* saw a small dip (1 – 6%) in performance on the tasks but also saw a 100× improvement in retraining time. *SISA-FC* saw a 30% dip in performance on the tasks compared to *SISA-A* but also was about 100× faster to retrain than *SISA-A*,

making it way faster than the baseline approach. Exact memory numbers were not provided but Kumar et al. [2022] stated that both frameworks saw significant reductions in memory usage. Since the SISA framework eliminates any influence the requested deleted point may have had, this approach satisfies the exact unlearning definition.

Meanwhile, Jang et al. [2023] studied a gradient based solution for machine unlearning in LMs, named *knowledge unlearning*. Specifically, given  $h_\theta$  and a given sequence of tokens that we would like to unlearn,  $\mathbf{d} = (d_1, \dots, d_T)$ ,  $\mathcal{U}$  solves a gradient ascent problem maximizing the loss function:

$$\mathcal{L}(h_\theta, \mathbf{d}) = - \sum_{t=1}^T \log(p_\theta(d_t | d_{<t}))$$

Where  $d_{<t}$  denotes the token sequence  $\mathbf{d} = (d_1, \dots, d_{t-1})$  and  $p_\theta(d_t | d_{<t})$  denotes the conditional probability of predicting the next token to be  $d_t$  given  $d_{<t}$ . Jang et al. [2023] evaluated this method on GPT-Neo Black et al. [2021] pre-trained on the Pile Gao et al. [2020] and the OPT model Zhang et al. [2022] pre-trained on a deduplicated version of the Pile. Data from the Training Data Extraction Challenge Carlini et al. [2023b] was used as the target data and the success of the unlearning was measured as the ability to reduce the risk from extraction attacks (Section 4.3) and memorization (Section 3). An array of 9 classification tasks were used to evaluate the general performance of the LMs. They found that knowledge unlearning did not result in any significant degradation of performance compared to naive retraining while requiring several orders of magnitude less computation. The general performance was also significantly better than another computationally inexpensive method, differentially private decoding (Majmudar et al. [2022], Section 5.2). In some cases, the performance of the model after knowledge unlearning was superior to the naive retraining method. This is possibly due to the benefits of forgetting described in Zhou et al. [2022]. Additionally, Jang et al. [2023] found that sequentially unlearning the data is better than unlearning it in large batches (see Figure 23).

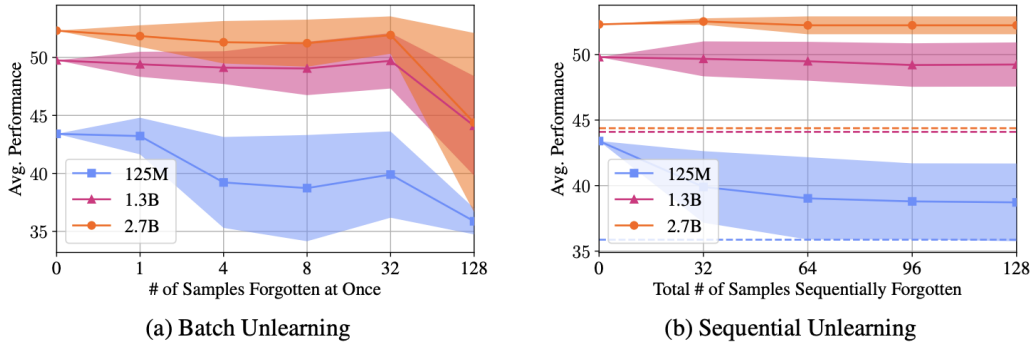


Figure 23: Jang et al. [2023], the average LM performance for 3 models on 9 benchmark tasks when varying the total number of samples unlearned. (a) shows performance when they are unlearned all at once while (b) shows performance when the samples are divided into 4 chunks and fed in sequentially. While performance is similar for small number of samples, it drops off for batch unlearning at 128 samples.

Yao et al. [2023] similarly performed unlearning using gradient ascent, providing “negative” examples for the OPT and Llama models to unlearn. They showed that after several training iterations, these models were able to remove harmful content (such as recommended illegal activity), copyright content, and misinformation. Yao et al. [2023] compared this approach to reinforcement learning from human feedback (RLHF) and found gradient ascent unlearning achieved better goal alignment while using 2% of the computational costs. This study showed that a gradient based approach to unlearning in LLMs can work in these practical cases.

**Concept Unlearning.** While prior works focus on unlearning batches of examples, Eldan and Russinovich [2023] developed an approximate unlearning approach that can be used to unlearn an entire target data corpus,  $D_{target}$ . In this approach, a “reinforced” model is built by further fine tuning the model on  $D_{target}$ . The probability outputs of this model are compared to the baseline model since tokens common in  $D_{target}$  will have higher probabilities than generic tokens. Generic counterpart labels are then generated and the baseline model is fine tuned using these alternate labels for several iterations until the target data is unlearned. Eldan and Russinovich [2023] tested this on Meta’s 7-billion parameter Llama model and Microsoft’s 1.3-billion parameter MSFT model and attempted to unlearn the Harry Potter books. After several fine-tuning iterations, the model lost its “memory” of Harry Potter. For example, the text completion prompt “Harry Potter’s two best friends are...” was completed as “a talking cat and a dragon. One



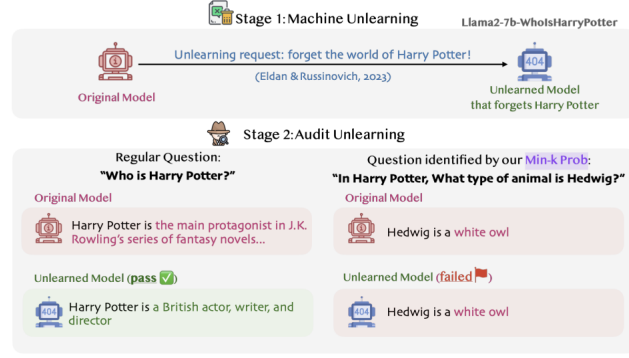


Figure 6: **Auditing machine unlearning with MIN-K% PROB.** Machine unlearning methods are designed to remove copyrighted and personal data from large language models. We use MIN-K% PROB to audit an unlearned LLM that has been trained to forget copyrighted books. However, we find that such a model can still output related copyrighted content.

day they decided to...” instead of “Ron Weasley and Hermione Granger. In this series...”. Model performance on other benchmarks decreased marginally but overall performance was largely unaffected.

Soon after Eldan and Russinovich [2023] was published, Shi et al. [2023] claimed to break this unlearning technique. Their MIA attack technique, called  $\text{min}-k$ , is a close variant of loss thresholding that rather than predicting training set membership if the loss of an entire sequence of tokens is particularly low, calculates the loss of the  $k$  tokens with the lowest losses (hence “ $\text{min}-k$ ”). They use their technique to identify parts of the forget set (in this case the Harry Potter corpus) that their MIA predicted were members of the training set *even after unlearning*, and then showed that when the unlearned model was prompted with that part of the corpus, it still output information contained in the forget set – see Figure 6 from Shi et al. [2023]. While the MIA attack above demonstrates that Eldan and Russinovich [2023] doesn’t perfectly unlearn concepts when faced with a motivated adversary, this is perhaps too high of a standard to apply to unlearning large corpora of text, and a better evaluation metric might be performance when prompted on average as opposed to adversarially.

Chen and Yang [2023] apply the SCRUB algorithm proposed in Kurmanji et al. [2023] for image classification, to LLMs. Rather than fine-tune the entire model as in Kurmanji et al. [2023], they only re-train .5% of the parameters in the unlearning layers on the SCRUB objective, where they have also added a term to the SCRUB objective that is the negative of the masked language modeling objective Raffel et al. [2020]. They evaluate their method on both classification and generation, on the IMDB Dataset Maas et al. [2011] and the SAMSum Gliwa et al. [2019] datasets respectively. Following Kurmanji et al. [2023] they evaluate their unlearning algorithm across (i) performance on the test set for the given task (ii) performance on the retained set and (iii) effectiveness of unlearning as measured by the MLM loss and performance on the forget set. All experiments use the T5 base model Raffel et al. [2020], which are then fine-tuned on the IMDB or SAMSum dataset prior to unlearning. On both datasets the model is asked to unlearn .5%, 1% and 10% of the dataset, where the forget set corresponds to randomly selected concepts; for IMDB a concept includes all the examples associated with a given movie, actor, or director, and on SAMSum it is all the conversations involving a given speaker. Unlearning guarantees were evaluated via a basic membership inference attack, and by looking at the loss on the forget set and MLM objective. Their SCRUB-based method consistently performed best in terms of high test set accuracy, minimal accuracy on the forget set, high MLM loss on the forget set, defense against MIA attacks, and efficiency (time).

#### Additional Remarks on Unlearning:

Zanella-Béguelin et al. [2020] found that model updates (such as unlearning) can possibly pose a privacy risk if an adversary has access to the versions of the model pre and post update. Through a differential analysis of the model output, a powerful membership inference attack can be performed on the model. For more details on this, see Section 4. Their attack highlights how machine unlearning is an orthogonal concept to differential privacy, and can in some cases increase risk from a data privacy perspective.

## 8 Conclusion

In this paper we covered a broad range of privacy risks related to the use of Large Language Models. As these models continue to proliferate into both the public consciousness and into different application domains, the concern over the apparent privacy risks associated with their use is set to grow. This is reflected in the fact that forthcoming AI regulation in the US and EU devote significant attention to privacy risks. At the same time, basic questions related to LLM privacy remain largely unresolved. We know that LLMs can memorize swathes of training data (Section 3), but when compared to the reported attack success rates against discriminative deep learning models the existing attacks discussed in Section 4 are relatively weak. Getting a real sense of what privacy risks practical attacks can pose to LLMs will be key to ensuring that personal data is not unintentionally exposed, and on the other side, policymakers and AI developers don't overreact to theoretical risks that may not be likely in practice. Privacy-preserving training techniques like differential privacy seem like an effective way to preserve privacy and still maintain performance during fine-tuning, but much more research is needed to verify if these early findings hold across different domains, and to ease the path to adoption. Preserving privacy during LLM pre-training is largely unsolved. As the exact ways that copyright laws will apply to generative models crystallizes over the next few years, there will be a host of technical questions that arise. Among many others: Given a specific model output, how can we verify if it constitutes fair use with respect to a subset of the training data? If we want to fairly compensate the creators of the original training data with royalties related to a specific output, how can we perform this attribution? Existing methods to unlearn data from LLMs are heuristic – how can we verify they work as intended? Approaches based on gradient descent require white-box updates to model parameters which is increasingly computationally expensive and restricted, and recent techniques based on in-context learning may not scale to large numbers of data deletions. More research is needed to both verify that these unlearning methods work as intended from a privacy standpoint, and can scale to practical models and deletion workloads.

In general, while in many of these domains the foundations

## References

- Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, S. Buch, Dallas Card, Rodrigo Castellon, Niladri S. Chatterji, Annie S. Chen, Kathleen A. Creel, Jared Davis, Dora Demszky, Chris Donahue, Moussa Dombouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kavin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren E. Gillespie, Karan Goel, Noah D. Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas F. Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, O. Khattab, Pang Wei Koh, Mark S. Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir P. Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avaniika Narayan, Deepak Narayanan, Benjamin Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, J. F. Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Robert Reich, Hongyu Ren, Frieda Rong, Yusuf H. Roohani, Camilo Ruiz, Jack Ryan, Christopher R'e, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishna Parasuram Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei A. Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. On the opportunities and risks of foundation models. *ArXiv*, 2021. URL <https://crfm.stanford.edu/assets/report.pdf>.
- Oct 2023. URL <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>.
- Nicola Lucchi. Chatgpt: A case study on copyright challenges for generative artificial intelligence systems. *European Journal of Risk Regulation*, page 1–23, 2023. doi:10.1017/err.2023.59.
- Eduard Fosch Villaronga, Peter Kieseberg, and Tiffany Li. Humans forget, machines remember: Artificial intelligence and the right to be forgotten. *Computer Law Security Review*, 34(2):304–313, 2018. ISSN 0267-3649. doi:<https://doi.org/10.1016/j.clsr.2017.08.007>. URL <https://www.sciencedirect.com/science/article/pii/S0267364917302091>.
- Thanh Tam Nguyen, Thanh Trung Huynh, Phi Le Nguyen, Alan Wee-Chung Liew, Hongzhi Yin, and Quoc Viet Hung Nguyen. A survey of machine unlearning, 2022.
- Antonio Gulli, 2004. URL [http://groups.di.unipi.it/~gulli/AG\\_corpus\\_of\\_news\\_articles.html](http://groups.di.unipi.it/~gulli/AG_corpus_of_news_articles.html).
- Eva Sharma, Chen Li, and Lu Wang. Bigpatent: A large-scale dataset for abstractive and coherent summarization, 2019.



- Yukun Zhu, Ryan Kiros, Richard Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books, 2015.
- W. N. Francis and H. Kucera. Brown corpus manual. Technical report, Department of Linguistics, Brown University, Providence, Rhode Island, US, 1979. URL <http://icame.uib.no/brown/bcm.html>.
- Common Crawl. Common crawl, 2008. URL <https://commoncrawl.org/>.
- Jesse Dodge, Maarten Sap, Ana Marasović, William Agnew, Gabriel Ilharco, Dirk Groeneveld, Margaret Mitchell, and Matt Gardner. Documenting large webtext corpora: A case study on the colossal clean crawled corpus, 2021.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. mt5: A massively multilingual pre-trained text-to-text transformer, 2021.
- Ondřej Bojar, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, and Christof Monz. Findings of the 2018 conference on machine translation (WMT18). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 272–303, Belgium, Brussels, October 2018. Association for Computational Linguistics. doi:10.18653/v1/W18-6401. URL <https://aclanthology.org/W18-6401>.
- Yegor Tkachenko, Mykel J. Kochenderfer, and Krzysztof Kluza. Customer simulation for direct marketing experiments. In *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, pages 478–487, 2016. doi:10.1109/DSAA.2016.59.
- Alistair E.W. Johnson, Tom J. Pollard, Lu Shen, Li-wei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. Mimic-iii, a freely accessible critical care database. *Scientific Data*, 3(1), 2016. doi:10.1038/sdata.2016.35.
- Ciprian Chelba, Tomas Mikolov, Mike Schuster, Qi Ge, Thorsten Brants, Phillipp Koehn, and Tony Robinson. One billion word benchmark for measuring progress in statistical language modeling, 2014.
- Aaron Gokaslan and Vanya Cohen. Openwebtext corpus, 2019. URL <http://Skylion007.github.io/OpenWebTextCorpus>.
- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, 19(2):313–330, 1993. URL <https://aclanthology.org/J93-2004>.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. The pile: An 800gb dataset of diverse text for language modeling, 2020.
- Yifan Peng, Shankai Yan, and Zhiyong Lu. Transfer learning in biomedical natural language processing: An evaluation of BERT and ELMo on ten benchmarking datasets. In *Proceedings of the 18th BioNLP Workshop and Shared Task*, pages 58–65, Florence, Italy, August 2019. Association for Computational Linguistics. doi:10.18653/v1/W19-5006. URL <https://aclanthology.org/W19-5006>.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach, 2019.
- Alec Go. Sentiment classification using distant supervision. 2009. URL <https://api.semanticscholar.org/CorpusID:18635269>.
- Stack Overflow. Stack overflow data, Mar 2019. URL <https://www.kaggle.com/datasets/stackoverflow/stackoverflow>.
- Ellen M. Voorhees and Dawn M. Tice. Building a question answering test collection. In *Proceedings of the 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR ’00, page 200–207, New York, NY, USA, 2000. Association for Computing Machinery. ISBN 1581132263. doi:10.1145/345508.345577. URL <https://doi.org/10.1145/345508.345577>.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. *ArXiv*, abs/1609.07843, 2016. URL <https://api.semanticscholar.org/CorpusID:16299141>.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification, 2016.
- Matt Mahoney. Large text compression benchmark, 2009. URL <https://cs.fit.edu/~mmahoney/compression/text.html>.
- Nicholas Carlini, Chang Liu, Úlfar Erlingsson, Jernej Kos, and Dawn Song. The secret sharer: Evaluating and testing unintended memorization in neural networks, 2019.

- Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom B. Brown, Dawn Xiaodong Song, Úlfar Erlingsson, Alina Oprea, and Colin Raffel. Extracting training data from large language models. In *USENIX Security Symposium*, 2020. URL <https://api.semanticscholar.org/CorpusID:229156229>.
- Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramer, and Chiyuan Zhang. Quantifying memorization across neural language models, 2023a.
- Vitaly Feldman. Does learning require memorization? a short tale about a long tail. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, page 954–959, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450369794. doi:10.1145/3357713.3384290. URL <https://doi.org/10.1145/3357713.3384290>.
- Gavin Brown, Mark Bun, Vitaly Feldman, Adam Smith, and Kunal Talwar. When is memorization of irrelevant training data necessary for high-accuracy learning? In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2021, page 123–132, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450380539. doi:10.1145/3406325.3451131. URL <https://doi.org/10.1145/3406325.3451131>.
- Griffin Davis. Chatgpt allegedly trained on copyrighted books!, May 2023a. URL <https://www.techtimes.com/articles/291210/20230505/chatgpt-allegedly-trained-copyrighted-books-new-study-claims-memorized-harry.htm>.
- Melissa Heikkila. What does gpt-3 “know” about me?, Sep 2022a. URL <https://www.technologyreview.com/2022/08/31/1058800/what-does-gpt-3-know-about-me/>.
- SCRT Labs. The secret behind large language models: Memorization over understanding?, Apr 2023. URL <https://www.linkedin.com/pulse/secret-behind-large-language-models-memorization-over-understanding/>.
- Cecilia Kang and Cade Metz. F.t.c. opens investigation into chatgpt maker over technology’s potential harms, Jul 2023. URL <https://www.nytimes.com/2023/07/13/technology/chatgpt-investigation-ftc-openai.html>.
- Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. Deduplicating training data makes language models better. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8424–8445, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi:10.18653/v1/2022.acl-long.577. URL <https://aclanthology.org/2022.acl-long.577>.
- Daphne Ippolito, Florian Tramèr, Milad Nasr, Chiyuan Zhang, Matthew Jagielski, Katherine Lee, Christopher A. Choquette-Choo, and Nicholas Carlini. Preventing verbatim memorization in language models gives a false sense of privacy, 2023.
- Nikhil Kandpal, Eric Wallace, and Colin Raffel. Deduplicating training data mitigates privacy risks in language models, 2022.
- Kushal Tirumala, Aram H. Markosyan, Luke Zettlemoyer, and Armen Aghajanyan. Memorization without overfitting: Analyzing the training dynamics of large language models, 2022.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, John A. Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuan-Fang Li, Scott M. Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. Sparks of artificial general intelligence: Early experiments with gpt-4. *ArXiv*, abs/2303.12712, 2023. URL <https://api.semanticscholar.org/CorpusID:257663729>.
- Milad Nasr, Nicholas Carlini, Jonathan Hayase, Matthew Jagielski, A. Feder Cooper, Daphne Ippolito, Christopher A. Choquette-Choo, Eric Wallace, Florian Tramèr, and Katherine Lee. Scalable extraction of training data from (production) language models, 2023.
- Chiyuan Zhang, Daphne Ippolito, Katherine Lee, Matthew Jagielski, Florian Tramèr, and Nicholas Carlini. Counterfactual memorization in neural language models, 2021.
- Zhenhong Zhou, Jiuyang Xiang, Chaomeng Chen, and Sen Su. Quantifying and analyzing entity-level memorization in large language models, 2023.
- Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli, editors, *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597, Online, August 2021. Association for Computational Linguistics. doi:10.18653/v1/2021.acl-long.353. URL <https://aclanthology.org/2021.acl-long.353>.

- Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi:10.18653/v1/2021.emnlp-main.243. URL <https://aclanthology.org/2021.emnlp-main.243>.
- Matthew Jagielski, Om Thakkar, Florian Tramèr, Daphne Ippolito, Katherine Lee, Nicholas Carlini, Eric Wallace, Shuang Song, Abhradeep Thakurta, Nicolas Papernot, and Chiyuan Zhang. Measuring forgetting of memorized training examples, 2023.
- Sid Black, Gao Leo, Phil Wang, Connor Leahy, and Stella Biderman. GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow, March 2021. URL <https://doi.org/10.5281/zenodo.5297715>.
- Ben Wang. Mesh-Transformer-JAX: Model-Parallel Implementation of Transformer Language Model with JAX. <https://github.com/kingoflolz/mesh-transformer-jax>, May 2021.
- Mikel Artetxe, Shruti Bhosale, Naman Goyal, Todor Mihaylov, Myle Ott, Sam Shleifer, Xi Victoria Lin, Jingfei Du, Srinivasan Iyer, Ramakanth Pasunuru, Giri Anantharaman, Xian Li, Shuohui Chen, Halil Akin, Mandeep Baines, Louis Martin, Xing Zhou, Punit Singh Koura, Brian O’Horo, Jeff Wang, Luke Zettlemoyer, Mona Diab, Zornitsa Kozareva, and Ves Stoyanov. Efficient large scale language modeling with mixtures of experts, 2022.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. Opt: Open pre-trained transformer language models, 2022.
- Aleena Anna Thomas, David Ifeoluwa Adelani, Ali Davody, Aditya Mogadala, and Dietrich Klakow. Investigating the impact of pre-trained word embeddings on memorization in neural networks. In *Workshop on Time-Delay Systems*, 2020.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, October 2014. Association for Computational Linguistics. doi:10.3115/v1/D14-1162. URL <https://aclanthology.org/D14-1162>.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi:10.18653/v1/N18-1202. URL <https://aclanthology.org/N18-1202>.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi:10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>.
- CRFM Stanford. Mistral, 2021. URL <https://github.com/stanford-crfm/mistral>.
- Peter West, Ximing Lu, Ari Holtzman, Chandra Bhagavatula, Jena D. Hwang, and Yejin Choi. Reflective decoding: Beyond unidirectional generation with off-the-shelf language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1435–1450, Online, August 2021. Association for Computational Linguistics. doi:10.18653/v1/2021.acl-long.114. URL <https://aclanthology.org/2021.acl-long.114>.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer, 2020.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization, 2017.
- Julien Simon. Large language models: A new moore’s law?, 2021. URL <https://huggingface.co/blog/large-language-models>.
- Stella Biderman, USVSN Sai Prashanth, Lintang Sutawika, Hailey Schoelkopf, Quentin Anthony, Shivanshu Purohit, and Edward Raff. Emergent and predictable memorization in large language models, 2023a.
- Stella Biderman, Hailey Schoelkopf, Quentin Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar van der Wal. Pythia: A suite for analyzing large language models across training and scaling, 2023b.

- Stella Biderman, Kieran Bicheno, and Leo Gao. Datasheet for the pile, 2022.
- R. Thomas McCoy, Paul Smolensky, Tal Linzen, Jianfeng Gao, and Asli Celikyilmaz. How much do language models copy from their training data? evaluating linguistic novelty in text generation using raven, 2021.
- Eugene Kharitonov, Marco Baroni, and Dieuwke Hupkes. How bpe affects memorization in transformers, 2021.
- Vitaly Feldman, Ilya Mironov, Kunal Talwar, and Abhradeep Thakurta. Privacy amplification by iteration. In Mikkel Thorup, editor, *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 521–532. IEEE Computer Society, 2018. doi:10.1109/FOCS.2018.00056. URL <https://doi.org/10.1109/FOCS.2018.00056>.
- Xuandong Zhao, Lei Li, and Yu-Xiang Wang. Provably confidential language modelling, 2022.
- H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data, 2016.
- Om Thakkar, Swaroop Ramaswamy, Rajiv Mathews, and Françoise Beaufays. Understanding unintended memorization in federated learning, 2020.
- Swaroop Ramaswamy, Om Thakkar, Rajiv Mathews, Galen Andrew, H. Brendan McMahan, and Françoise Beaufays. Training production language models without memorizing user data, 2020.
- Hongsheng Hu, Zoran Salcic, Lichao Sun, Gillian Dobbie, Philip S. Yu, and Xuyun Zhang. Membership inference attacks on machine learning: A survey. *ACM Comput. Surv.*, 54(11s), sep 2022. ISSN 0360-0300. doi:10.1145/3523273. URL <https://doi.org/10.1145/3523273>.
- Bo Liu, Ming Ding, Sina Shaham, Wenny Rahayu, Farhad Farokhi, and Zihuai Lin. When machine learning meets privacy: A survey and outlook. *ACM Comput. Surv.*, 54(2), mar 2021. ISSN 0360-0300. doi:10.1145/3436755. URL <https://doi.org/10.1145/3436755>.
- Samuel Yeom, Irene Giacomelli, Matt Fredrikson, and Somesh Jha. Privacy risk in machine learning: Analyzing the connection to overfitting, 2018.
- Justus Mattern, Fatemehsadat Miresghallah, Zhijing Jin, Bernhard Schoelkopf, Mrinmaya Sachan, and Taylor Berg-Kirkpatrick. Membership inference attacks against language models via neighbourhood comparison. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 11330–11343, Toronto, Canada, July 2023. Association for Computational Linguistics. doi:10.18653/v1/2023.findings-acl.719. URL <https://aclanthology.org/2023.findings-acl.719>.
- Sasi Kumar Murakonda, Reza Shokri, and George Theodorakopoulos. Quantifying the privacy risks of learning high-dimensional graphical models, 2021.
- Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models, 2017.
- Yunhui Long, Lei Wang, Diyue Bu, Vincent Bindschaedler, Xiaofeng Wang, Haixu Tang, Carl A. Gunter, and Kai Chen. A pragmatic approach to membership inferences on machine learning models. In *2020 IEEE European Symposium on Security and Privacy (EuroSP)*, pages 521–534, 2020. doi:10.1109/EuroSP48549.2020.00040.
- Alexandre Sablayrolles, Matthijs Douze, Yann Ollivier, Cordelia Schmid, and Hervé Jégou. White-box vs black-box: Bayes optimal strategies for membership inference, 2019.
- Liwei Song and Prateek Mittal. Systematic evaluation of privacy risks of machine learning models, 2020.
- Nicholas Carlini, Steve Chien, Milad Nasr, Shuang Song, Andreas Terzis, and Florian Tramèr. Membership inference attacks from first principles, 2022.
- Eric Lehman, Sarthak Jain, Karl Pichotta, Yoav Goldberg, and Byron C. Wallace. Does bert pretrained on clinical notes reveal sensitive data?, 2021.
- Congzheng Song and Ananth Raghunathan. Information leakage in embedding models, 2020.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146, 2017. doi:10.1162/tacl\_a\_00051. URL <https://aclanthology.org/Q17-1010>.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. Distributed representations of words and phrases and their compositionality, 2013.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. Universal sentence encoder for English. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 169–174, Brussels, Belgium, November 2018. Association for Computational Linguistics. doi:10.18653/v1/D18-2029. URL <https://aclanthology.org/D18-2029>.

- Muthu Chidambaram, Yinfei Yang, Daniel Cer, Steve Yuan, Yunhsuan Sung, Brian Strope, and Ray Kurzweil. Learning cross-lingual sentence representations via a multi-task dual-encoder model. In *Proceedings of the 4th Workshop on Representation Learning for NLP (Repl4NLP-2019)*, pages 250–259, Florence, Italy, August 2019. Association for Computational Linguistics. doi:10.18653/v1/W19-4330. URL <https://aclanthology.org/W19-4330>.
- Matthew Henderson, Rami Al-Rfou, Brian Strope, Yun hsuan Sung, Laszlo Lukacs, Ruiqi Guo, Sanjiv Kumar, Balint Miklos, and Ray Kurzweil. Efficient natural language response suggestion for smart reply, 2017.
- Lajanugen Logeswaran and Honglak Lee. An efficient framework for learning sentence representations, 2018.
- Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. The Ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 285–294, Prague, Czech Republic, September 2015. Association for Computational Linguistics. doi:10.18653/v1/W15-4640. URL <https://aclanthology.org/W15-4640>.
- Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China, November 2019. Association for Computational Linguistics. doi:10.18653/v1/D19-1410. URL <https://aclanthology.org/D19-1410>.
- Yinfei Yang, Steve Yuan, Daniel Cer, Sheng-yi Kong, Noah Constant, Petr Pilar, Heming Ge, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. Learning semantic textual similarity from conversations. In *Proceedings of the Third Workshop on Representation Learning for NLP*, pages 164–174, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi:10.18653/v1/W18-3022. URL <https://aclanthology.org/W18-3022>.
- Jean-Loup Gailly and Mark Adler. zlib compression library, 1995. URL <http://zlib.net/>.
- Fatemehsadat Mireshghallah, Kartik Goyal, Archit Uniyal, Taylor Berg-Kirkpatrick, and Reza Shokri. Quantifying privacy risks of masked language models using membership inference attacks, 2022.
- Thomas Vakili and Hercules Dalianis. Using membership inference attacks to evaluate privacy-preserving language modeling fails for pseudonymizing data. In *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 318–323, Tórshavn, Faroe Islands, May 2023. University of Tartu Library. URL <https://aclanthology.org/2023.nodalida-1.33>.
- Abhyuday N. Jagannatha, Bhanu Pratap Singh Rawat, and Hong Yu. Membership inference attack susceptibility of clinical language models. *ArXiv*, abs/2104.08305, 2021. URL <https://api.semanticscholar.org/CorpusID:233296028>.
- Yuta Nakamura, Shouhei Hanaoka, Yukihiro Nomura, Naoto Hayashi, Osamu Abe, Shuntaro Yada, Shoko Wakamiya, and Eiji Aramaki. Kart: Parameterization of privacy leakage scenarios from pre-trained language models, 2022.
- Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott. Publicly available clinical BERT embeddings. In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, Minneapolis, Minnesota, USA, June 2019. Association for Computational Linguistics. doi:10.18653/v1/W19-1909. URL <https://aclanthology.org/W19-1909>.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019. URL <https://api.semanticscholar.org/CorpusID:160025533>.
- Santiago Zanella-Béguelin, Lukas Wutschitz, Shruti Tople, Victor Rühle, Andrew Paverd, Olga Ohrimenko, Boris Köpf, and Marc Brockschmidt. Analyzing information leakage of updates to natural language models. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. ACM, oct 2020. doi:10.1145/3372297.3417880. URL <https://doi.org/10.1145/3372297.3417880>.
- Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. Recurrent neural network regularization, 2015.
- Congzheng Song and Vitaly Shmatikov. Auditing data provenance in text-generation models, 2019.
- Cristian Danescu-Niculescu-Mizil and Lillian Lee. Chameleons in imagined conversations: A new approach to understanding coordination of linguistic style in dialogs. In *Proceedings of the 2nd Workshop on Cognitive Modeling and Computational Linguistics*, pages 76–87, Portland, Oregon, USA, June 2011. Association for Computational Linguistics. URL <https://aclanthology.org/W11-0609>.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar, October 2014. Association for Computational Linguistics. doi:10.3115/v1/D14-1179. URL <https://aclanthology.org/D14-1179>.

- Philipp Koehn. Europarl: A parallel corpus for statistical machine translation. In *Proceedings of Machine Translation Summit X: Papers*, pages 79–86, Phuket, Thailand, September 13-15 2005. URL <https://aclanthology.org/2005.mtsummit-papers.11>.
- Sorami Hisamoto, Matt Post, and Kevin Duh. Membership inference attacks on sequence-to-sequence models: Is my data in your machine translation system? *Transactions of the Association for Computational Linguistics*, 8:49–63, dec 2020. doi:10.1162/tacl\_a\_00299. URL [https://doi.org/10.1162/tacl\\_a\\_00299](https://doi.org/10.1162/tacl_a_00299).
- Chin-Yew Lin and Franz Josef Och. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics*, ACL ’04, page 605–es, USA, 2004. Association for Computational Linguistics. doi:10.3115/1218955.1219032. URL <https://doi.org/10.3115/1218955.1219032>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2023.
- Felix Hieber, Tobias Domhan, Michael Denkowski, David Vilar, Artem Sokolov, Ann Clifton, and Matt Post. Sockeye: A toolkit for neural machine translation, 2018.
- Matthieu Meeus, Shubham Jain, Marek Rei, and Yves-Alexandre de Montjoye. Did the neurons read your book? document-level membership inference for large language models, 2023.
- Together Computer. Redpajama: an open dataset for training large language models, 2023. URL <https://github.com/togethercomputer/RedPajama-Data>.
- Xinyang Geng and Hao Liu. Openllama: An open reproduction of llama, May 2023. URL [https://github.com/openlm-research/open\\_llama](https://github.com/openlm-research/open_llama).
- John Abascal, Stanley Wu, Alina Oprea, and Jonathan Ullman. Tmi! finetuned models leak private information from their pretraining data, 2023.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. Transformer-XL: Attentive language models beyond a fixed-length context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy, July 2019. Association for Computational Linguistics. doi:10.18653/v1/P19-1285. URL <https://aclanthology.org/P19-1285>.
- Martin Bertran, Shuai Tang, Michael Kearns, Jamie Morgenstern, Aaron Roth, and Zhiwei Steven Wu. Scalable membership inference attacks via quantile regression, 2023.
- Weichen Yu, Tianyu Pang, Qian Liu, Chao Du, Bingyi Kang, Yan Huang, Min Lin, and Shuicheng Yan. Bag of tricks for training data extraction from language models, 2023.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Huggingface’s transformers: State-of-the-art natural language processing, 2020.
- Nicholas Carlini, Harold Cooper, Florian Tramer, and Chiyua Zhang. Training data extraction challenge, 2023b. URL <https://github.com/google-research/lm-extraction-benchmark>.
- Zhexin Zhang, Jiaxin Wen, and Minlie Huang. Ethicist: Targeted training data extraction through loss smoothed soft prompting and calibrated confidence estimation, 2023a.
- Oskar Andersson. *Sequential Good-Turing and the missing species problem*. PhD thesis, 2022.
- Vaidehi Patil, Peter Hase, and Mohit Bansal. Can sensitive information be deleted from llms? objectives for defending against extraction attacks, 2023.
- Nicholas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwal, Florian Tramèr, Borja Balle, Daphne Ippolito, and Eric Wallace. Extracting training data from diffusion models, 2023c.
- Robin Staab, Mark Vero, Mislav Balunović, and Martin Vechev. Beyond memorization: Violating privacy via inference with large language models, 2023.
- Marc Marone and Benjamin Van Durme. Data portraits: Recording foundation model training data, 2023.
- Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III au2, and Kate Crawford. Datasheets for datasets, 2021.
- Pier Paolo Ippolito. Ai differential privacy and federated learning, Oct 2020. URL <https://towardsdatascience.com/ai-differential-privacy-and-federated-learning-523146d46b85>.
- Cynthia Dwork. Differential privacy. In *Automata, Languages and Programming: 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II 33*, pages 1–12. Springer, 2006.

- Zachary Izzo, Jinsung Yoon, Sercan O. Arik, and James Zou. Provable membership inference privacy, 2022.
- Shuang Song, Kamalika Chaudhuri, and Anand D. Sarwate. Stochastic gradient descent with differentially private updates. In *2013 IEEE Global Conference on Signal and Information Processing*, pages 245–248, 2013. doi:10.1109/GlobalSIP.2013.6736861.
- Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, oct 2016. doi:10.1145/2976749.2978318. URL <https://doi.org/10.1145/2976749.2978318>.
- Cynthia Dwork and Vitaly Feldman. Privacy-preserving prediction, 2018.
- Yuval Dagan and Vitaly Feldman. Pac learning with stable and private predictions, 2020.
- Cynthia Dwork and Aaron Roth. *The Algorithmic Foundations of Differential Privacy*. Foundations and Trends in Theoretical Computer Science, 2014.
- Nicolas Papernot, Abhradeep Thakurta, Shuang Song, Steve Chien, and Úlfar Erlingsson. Tempered sigmoid activations for deep learning with differential privacy. *CoRR*, abs/2007.14191, 2020. URL <https://arxiv.org/abs/2007.14191>.
- Jacob Poushter. Smartphone ownership and internet usage continues to climb in emerging economies, Feb 2016. URL <https://www.pewresearch.org/global/2016/02/22/smartphone-ownership-and-internet-usage-continues-to-climb-in-emerging-economies/>.
- Andrew Hard, Kanishka Rao, Rajiv Mathews, Swaroop Ramaswamy, Françoise Beaufays, Sean Augenstein, Hubert Eichner, Chloé Kiddon, and Daniel Ramage. Federated learning for mobile keyboard prediction, 2019.
- Rohan Anil, Badi Ghazi, Vineet Gupta, Ravi Kumar, and Pasin Manurangsi. Large-scale differentially private bert, 2021.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. Google’s neural machine translation system: Bridging the gap between human and machine translation, 2016.
- Shlomo Hoory, Amir Feder, Avichai Tendler, Sofia Erell, Alon Cohen, Itay Laish, Hootan Nakhost, Uri Stemmer, Ayelet Benjamini, Avinatan Hassidim, and Yossi Matias. Learning and evaluating a differentially private pre-trained language model. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1178–1189, Punta Cana, Dominican Republic, 2021. URL <https://aclanthology.org/2021.findings-emnlp.102/>.
- Özlem Uzuner, Brett R South, Shuying Shen, and Scott L DuVall. 2010 i2b2/VA challenge on concepts, assertions, and relations in clinical text. *Journal of the American Medical Informatics Association*, 18(5):552–556, 06 2011. ISSN 1067-5027. doi:10.1136/amiajnl-2011-000203. URL <https://doi.org/10.1136/amiajnl-2011-000203>.
- Weiyang Shi, Aiqi Cui, Evan Li, Ruoxi Jia, and Zhou Yu. Selective differential privacy for language modeling, 2022.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units, 2016.
- Xiaoxue Zang, Abhinav Rastogi, Srinivas Sunkara, Raghav Gupta, Jianguo Zhang, and Jindong Chen. MultiWOZ 2.2 : A dialogue dataset with additional annotation corrections and state tracking baselines. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 109–117, Online, July 2020. Association for Computational Linguistics. doi:10.18653/v1/2020.nlp4convai-1.13. URL <https://aclanthology.org/2020.nlp4convai-1.13>.
- Gavin Kerrigan, Dylan Slack, and Jens Tuyls. Differentially private language models benefit from public pre-training. In *Proceedings of the Second Workshop on Privacy in NLP*, pages 39–45, Online, November 2020. Association for Computational Linguistics. doi:10.18653/v1/2020.privatenlp-1.5. URL <https://aclanthology.org/2020.privatenlp-1.5>.
- Xuechen Li, Florian Tramèr, Percy Liang, and Tatsunori Hashimoto. Large language models can be strong differentially private learners, 2022.
- Jekaterina Novikova, Ondřej Dušek, and Verena Rieser. The e2e dataset: New challenges for end-to-end generation, 2017.
- Jaewoo Lee and Daniel Kifer. Scaling up differentially private deep learning with fast per-example gradient clipping, 2020.



- Christophe Dupuy, Radhika Arava, Rahul Gupta, and Anna Rumshisky. An efficient dp-sgd mechanism for large scale nlp models, 2022.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020. URL [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf).
- Yiming Zhang and Daphne Ippolito. Prompts should not be seen as secrets: Systematically measuring prompt extraction attack success, 2023.
- Xinyu Tang, Richard Shin, Huseyin A. Inan, Andre Manoel, Fatemehsadat Mireshghallah, Zinan Lin, Sivakanth Gopi, Janardhan Kulkarni, and Robert Sim. Privacy-preserving in-context learning with differentially private few-shot generation, 2023.
- Yoichi Ishibashi and Hidetoshi Shimodaira. Knowledge sanitization of large language models, 2023.
- Antonio Ginart, Laurens van der Maaten, James Zou, and Chuan Guo. Submix: Practical private prediction for large-scale language models, 2022.
- Nicolas Papernot, Shuang Song, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Úlfar Erlingsson. Scalable private learning with pate, 2018.
- Jimit Majmudar, Christophe Dupuy, Charith Peris, Sami Smaili, Rahul Gupta, and Richard Zemel. Differentially private decoding in large language models, 2022.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-lm: Training multi-billion parameter language models using model parallelism, 2020.
- H. Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning differentially private recurrent language models, 2018.
- Rami Al-Rfou, Marc Pickett, Javier Snider, Yun hsuan Sung, Brian Strope, and Ray Kurzweil. Conversational contextual cues: The case of personalization and history for response ranking, 2016.
- Haşim Sak, Andrew Senior, and Françoise Beaufays. Long short-term memory based recurrent neural network architectures for large vocabulary speech recognition, 2014.
- Yijia Xiao, Yiqiao Jin, Yushi Bai, Yue Wu, Xianjun Yang, Xiao Luo, Wenchao Yu, Xujiang Zhao, Yanchi Liu, Haifeng Chen, Wei Wang, and Wei Cheng. Large language models can be good privacy protection learners, 2023.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model, 2023.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models, 2023.
- Tianyu Han, Lisa C. Adams, Jens-Michalis Papaioannou, Paul Grundmann, Tom Oberhauser, Alexander Löser, Daniel Truhn, and Keno K. Bressen. Medalpaca – an open-source collection of medical conversational ai models and training data, 2023.
- LeapBeyond. Scrubadub: A python library for cleaning sensitive data from text. <https://github.com/LeapBeyond/scrubadub>, 2023. GitHub Repository.
- Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain, July 2004. Association for Computational Linguistics. URL <https://aclanthology.org/W04-1013>.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert, 2020.
- James Vincent. The scary truth about ai copyright is nobody knows what will happen next, Nov 2022. URL <https://www.theverge.com/23444685/generative-ai-copyright-infringement-legal-fair-use-training-data>.

- Melissa Heikkilä. This artist is dominating ai-generated art. and he's not happy about it., Nov 2022b. URL [https://www.technologyreview.com/2022/09/16/1059598/this-artist-is-dominating-ai-generated-art-and-hes-not-happy-about-it/?truid=&utm\\_source=the\\_algorithm&utm\\_medium=email&utm\\_campaign=the\\_algorithm.unpaid.engagement&utm\\_content=04-17-2023](https://www.technologyreview.com/2022/09/16/1059598/this-artist-is-dominating-ai-generated-art-and-hes-not-happy-about-it/?truid=&utm_source=the_algorithm&utm_medium=email&utm_campaign=the_algorithm.unpaid.engagement&utm_content=04-17-2023).
- Jack Bandy and Nicholas Vincent. Addressing "documentation debt" in machine learning research: A retrospective datasheet for bookcorpus, 2021.
- Wes Davis. Sarah silverman is suing openai and meta for copyright infringement, Jul 2023b. URL <https://www.theverge.com/2023/7/9/23788741/sarah-silverman-openai-meta-chatgpt-llama-copyright-infringement-chatbots-artificial-intelligence-ai-ICID>.
- Antonia Karamolegkou, Jiaang Li, Li Zhou, and Anders Søgaaard. Copyright violations and large language models, 2023.
- Daniel E. Rose. It might be cheaper to pay them: Artificial intelligence, copyright, and the hollywood writers' strike - copyright - united states, May 2023a. URL <https://www.mondaq.com/unitedstates/copyright/1314270/it-might-be-cheaper-to-pay-them-artificial-intelligence-copyright-and-the-hollywood-writers-strike>.
- Ryan Broderick. Ai can't replace humans yet - but if the wga writers don't win, it might not matter, May 2023. URL <https://www.polygon.com/23742770/ai-writers-strike-chat-gpt-explained>.
- Sylvie Delacroix. Data rivers: Re-balancing the data ecosystem that makes generative ai possible, Apr 2023. URL <https://ssrn.com/abstract=4388928>.
- Nov 2021. URL <https://harvardlawreview.org/print/vol-135/google-llc-v-oracle-america-inc/>.
- Benjamin L. W. Sobel. Artificial intelligence's fair use crisis. *The Columbia Journal of Law and the Arts*, 41(1): 45–97, Dec. 2017. doi:10.7916/jla.v41i1.2036. URL <https://journals.library.columbia.edu/index.php/lawandarts/article/view/2036>.
- Peter Henderson, Xuechen Li, Dan Jurafsky, Tatsunori Hashimoto, Mark A. Lemley, and Percy Liang. Foundation models and fair use, 2023.
- Mark A Lemley and Bryan Casey. Fair learning, Mar 2021. URL <https://texaslawreview.org/fair-learning/>.
- Matthew Sag. The new legal landscape for text mining and machine learning. *SSRN Electronic Journal*, 66:291–365, 2019. doi:10.2139/ssrn.3331606.
- Kalin Hristov. Artificial intelligence and the copyright dilemma, May 2017. URL <https://ssrn.com/abstract=2976428>.
- US Copyright. Copyright registration guidance: Works containing material generated by artificial intelligence<sub>2023</sub>, 2023.
- Daniel E. Rose. Dabus dares to dream: A look at stephen thaler's patent puzzle, Jan 2023b. URL <https://www.vklaw.com/ImagineThatIPLawBlog/dabus-dares-to-dream-a-look-at-stephen-thalers-patent-puzzle>.
- Dan L Burk. Algorithmic fair use. *The University of Chicago Law Review*, 86(2):283, 2019.
- 499 U.S. 340 US Supreme Court. Feist publications, inc. v. rural telephone service company, inc., 1991. URL <https://www.law.cornell.edu/supremecourt/text/499/340>.
- Amanda Levendowski. How copyright law can fix artificial intelligence's implicit bias problem. *Washington Law Review*, 93, 2017.
- Nikhil Vyas, Sham Kakade, and Boaz Barak. Provable copyright protection for generative models, 2023.
- Niva Elkin-Koren, Uri Hachohen, Roi Livni, and Shay Moran. Can copyright be reduced to privacy?, 2023.
- Sarah Scheffler, Eran Tromer, and Mayank Varia. Formalizing human ingenuity: A quantitative framework for copyright law's substantial similarity, 2022.
- Giorgio Franceschelli and Mirco Musolesi. Deepcreativity: Measuring creativity with deep learning techniques, 2022.
- Margaret A. Boden. *The creative mind: Myths and mechanisms*. Routledge, 2005.
- Brooke Auxier, Lee Rainie, Monica Anderson, Andrew Perrin, Madhu Kumar, and Erica Turner. Americans and privacy: Concerned, confused and feeling lack of control over their personal information, Nov 2019. URL <https://www.pewresearch.org/internet/2019/11/15/americans-and-privacy-concerned-confused-and-feeling-lack-of-control-over-their-personal-information/>.

- European Parliament and Council of the European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council, 2016. URL <https://data.europa.eu/eli/reg/2016/679/oj>.
- State of California Department of Justice. California consumer privacy act, 2023. URL <https://oag.ca.gov/privacy/ccpa>.
- State of New York. Senate bill s365b, 2023. URL <https://legislation.nysenate.gov/pdf/bills/2023/S365B>.
- Government of Canada. Consumer privacy protection act, 2023. URL <https://ised-isde.canada.ca/site/innovation-better-canada/en/consumer-privacy-protection-act>.
- Dawen Zhang, Pamela Finckenberg-Broman, Thong Hoang, Shidong Pan, Zhenchang Xing, Mark Staples, and Xiwei Xu. Right to be forgotten in the era of large language models: Implications, challenges, and solutions, 2023b.
- Federal Trade Commission. California company settles ftc allegations it deceived consumers about use of facial recognition in photo storage app, Jan 2021. URL <https://www.ftc.gov/news-events/news/press-releases/2021/01/california-company-settles-ftc-allegations-it-deceived-consumers-about-use-facial-recognition-photo>.
- Help Net Security. Consumers take data control into their own hands amid rising privacy concerns, Apr 2023. URL <https://www.helpnetsecurity.com/2023/04/11/personal-data-privacy-concerns/#:~:text=Data%20privacy%20concerns,on%20Meta%20data%20privacy%20practices>.
- Varun Gupta, Christopher Jung, Seth Neel, Aaron Roth, Saeed Sharifi-Malvajerdi, and Chris Waites. Adaptive machine unlearning, 2021.
- Laura Graves, Vineel Nagisetty, and Vijay Ganesh. Amnesiac machine learning, 2020.
- Lucas Bourtole, Varun Chandrasekaran, Christopher A. Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning, 2020.
- C. Okan Sakar, S. Polat, Mete Katircioglu, and Yomi Kastro. Real-time prediction of online shoppers’ purchasing intention using multilayer perceptron and lstm recurrent neural networks. *Neural Computing and Applications*, 31, 10 2019. doi:10.1007/s00521-018-3523-0.
- Frank R. Hampel. The influence curve and its role in robust estimation. *Journal of the American Statistical Association*, 69(346):383–393, 1974. doi:10.1080/01621459.1974.10482962.
- R Dennis Cook. Influential observations in linear regression. *Journal of the American Statistical Association*, 74(365): 169–174, Mar 1979.
- Samyadeep Basu, Philip Pope, and Soheil Feizi. Influence functions in deep learning are fragile, 2021.
- Juhan Bae, Nathan Ng, Alston Lo, Marzyeh Ghassemi, and Roger Grosse. If influence functions are the answer, then what is the question?, 2022.
- Seth Neel, Aaron Roth, and Saeed Sharifi-Malvajerdi. Descent-to-delete: Gradient-based methods for machine unlearning, 2020.
- Yuchen Zhang, John C. Duchi, and Martin Wainwright. Communication-efficient algorithms for statistical optimization, 2013.
- Meghdad Kurmanji, Peter Triantafillou, Jamie Hayes, and Eleni Triantafillou. Towards unbounded machine unlearning, 2023.
- Joel Jang, Dongkeun Yoon, Sohee Yang, Sungmin Cha, Moontae Lee, Lajanugen Logeswaran, and Minjoon Seo. Knowledge unlearning for mitigating privacy risks in language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14389–14408, Toronto, Canada, July 2023. Association for Computational Linguistics. doi:10.18653/v1/2023.acl-long.805. URL <https://aclanthology.org/2023.acl-long.805>.
- Aditya Golatkar, Alessandro Achille, and Stefano Soatto. Forgetting outside the box: Scrubbing deep networks of information accessible from input-output observations. *CoRR*, abs/2003.02960, 2020. URL <https://arxiv.org/abs/2003.02960>.
- Aditya Golatkar, Alessandro Achille, and Stefano Soatto. Eternal sunshine of the spotless net: Selective forgetting in deep networks. *CoRR*, abs/1911.04933, 2019. URL <http://arxiv.org/abs/1911.04933>.
- Vinayshekhhar Bannihatti Kumar, Rashmi Gangadharaiah, and Dan Roth. Privacy adhering machine un-learning in nlp, 2022.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for NLP. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/houlsby19a.html>.

- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium, November 2018. Association for Computational Linguistics. doi:10.18653/v1/W18-5446. URL <https://aclanthology.org/W18-5446>.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA, October 2013. Association for Computational Linguistics. URL <https://aclanthology.org/D13-1170>.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. A broad-coverage challenge corpus for sentence understanding through inference, 2018.
- Hattie Zhou, Ankit Vani, Hugo Larochelle, and Aaron Courville. Fortuitous forgetting in connectionist networks, 2022.
- Yuanshun Yao, Xiaojun Xu, and Yang Liu. Large language model unlearning, 2023.
- Ronen Eldan and Mark Russinovich. Who’s harry potter? approximate unlearning in llms, 2023.
- Weijia Shi, Anirudh Ajith, Mengzhou Xia, Yangsibo Huang, Daogao Liu, Terra Blevins, Danqi Chen, and Luke Zettlemoyer. Detecting pretraining data from large language models, 2023.
- Jiaao Chen and Diyi Yang. Unlearn what you want to forget: Efficient unlearning for llms, 2023.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics. URL <https://aclanthology.org/P11-1015>.
- Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. SAMSum corpus: A human-annotated dialogue dataset for abstractive summarization. In Lu Wang, Jackie Chi Kit Cheung, Giuseppe Carenini, and Fei Liu, editors, *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 70–79, Hong Kong, China, November 2019. Association for Computational Linguistics. doi:10.18653/v1/D19-5409. URL <https://aclanthology.org/D19-5409>.

## A Privacy Preserving Large Language Models

---

### Algorithm A.1 MIA on Word Embeddings, Song and Raghunathan [2020]

---

- 1: **Input:** Target window of words  $C = [w_b, \dots, w_0, \dots, w_e]$ , word embedding matrix  $\Phi$ , similarity function  $\delta$ .
  - 2: Map words in  $C$  with  $\Phi$  to get  $[\Phi_{w_b}, \dots, \Phi_{w_0}, \dots, \Phi_{w_e}]$
  - 3:  $\Delta \leftarrow \{\delta_i | \delta_i = \delta(\Phi_{w_0}, \Phi_{w_i}), \forall w_i \in C / \{w_0\}\}$
  - 4: **return** “member” if  $\frac{1}{|\Delta|} \sum_{\delta_i \in \Delta} \delta_i \geq \tau_m$  else “non-member”
- 

---

### Algorithm A.2 Aggregated-level MIA on Sentence Embeddings, Song and Raghunathan [2020]

---

- 1: **Input:** Target sentences in context  $S = [s_1, \dots, s_n]$ , sentence embedding model  $\Phi$ , similarity function  $\delta$ .
  - 2: Map sentences in  $S$  with  $\Phi$  and get  $[\Phi(s_1), \dots, \Phi(s_n)]$ .
  - 3:  $\Delta \leftarrow \{\delta_i | \delta_i = \delta(\Phi(s_i), \Phi(s_{i+1}))\}_{i=1}^{n-1}$
  - 4: **return** “member” if  $\frac{1}{|\Delta|} \sum_{\delta_i \in \Delta} \delta_i \geq \tau_m$  else “non-member”
- 

---

### Algorithm A.3 Online Likelihood Ratio Attack (LiRA), Carlini et al. [2022]

---

- 1: **Input:** model  $h$ , example  $(x, y)$ , data distribution  $\mathcal{D}$
  - 2:  $conf_{s_{in}} = \{\}$
  - 3:  $conf_{s_{out}} = \{\}$
  - 4: **for**  $N$  times **do**
  - 5:    $D_{attack} \leftarrow \mathcal{D}$  % Sample shadow dataset
  - 6:    $h_{in} \leftarrow \mathcal{T}(D_{attack} \cup \{(x, y)\})$  % Train IN model
  - 7:    $conf_{s_{in}} \leftarrow conf_{s_{in}} \cup \{\phi(h_{in}(x)_y)\}$
  - 8:    $h_{out} \leftarrow \mathcal{T}(D_{attack} \setminus \{(x, y)\})$  % Train OUT model
  - 9:    $conf_{s_{out}} \leftarrow conf_{s_{out}} \cup \{\phi(h_{out}(x)_y)\}$
  - 10: **end for**
  - 11:  $\mu_{in} \leftarrow mean(conf_{s_{in}})$
  - 12:  $\mu_{out} \leftarrow mean(conf_{s_{out}})$
  - 13:  $\sigma_{in}^2 \leftarrow var(conf_{s_{in}})$
  - 14:  $\sigma_{out}^2 \leftarrow var(conf_{s_{out}})$
  - 15:  $conf_{obs} = \phi(h(x)_y)$  % query target model
  - 16: **return**  $\Lambda = \frac{p(conf_{obs} | \mathcal{N}(\mu_{in}, \sigma_{in}^2))}{p(conf_{obs} | \mathcal{N}(\mu_{out}, \sigma_{out}^2))}$
- 

## B Privacy Preserving Large Language Models

---

### Algorithm B.1 SubMix Training, Ginart et al. [2022]

---

- 1: **Input:** User-level private corpus  $\mathcal{D}$ , LM fine-tuning routine  $\mathcal{L}$
  - 2: **Outputs:** Fine-tuned LMs  $h_{\pi_i}, h_{\pi'_i}$  for  $i = 1, \dots, k$
  - 3: **Hyperparameters:** # of parts  $k$
  - 4:  $\prod \leftarrow$  Random  $k$ -fold partition of  $\mathcal{D}$  with  $|\prod_i| = |\mathcal{D}|/k$
  - 5: **for**  $i \in \{1, \dots, k\}$  **do**
  - 6:    $(\pi_i, \pi'_i) \leftarrow$  Randomly split part  $\prod_i$  into two subparts.
  - 7:    $h_{\pi_i} \leftarrow \mathcal{L}(\pi_i), h_{\pi'_i} \leftarrow \mathcal{L}(\pi_i)$
  - 8: **end for**
- 

## C Privacy Preserving Large Language Models

**Strong vs Weak Unlearning:** When unlearning sequences, which in practical circumstances could grow arbitrarily long, it is important to consider the computational cost of  $\mathcal{U}$ , assuming a bound on the accuracy of the unlearning algorithm. This leads us to the definitions of *strong* and *weak* unlearning. First, we must define  $(\alpha, \beta)$ -accuracy:




Figure B.1: Results of a canary extraction attack on CRT model from Zhao et al. [2022]. LSTM and GPT refer to model type tested. S-DP-LSTM model was trained using selective differential privacy from Shi et al. [2022].  $\gamma$  is the error rate in the policy function screening for sensitive data. Zhao et al. [2022]’s CRT method is effective at preventing this attack while S-DP-LSTM still is at risk.

---

**Algorithm B.2** *SubMix* Prediction, Ginart et al. [2022]

---

```

1: Input: Fine-tuned LMs  $h_{\pi_i}, h_{\pi'_i}$  for  $i = 1, \dots, k$ , privacy parameters  $\epsilon$ , time step  $t$ , query context  $x_t \in \Sigma^*$ 
2: Outputs: Next-token response  $y_t \in \Sigma$ 
3: Hyperparameters: Rényi divergence order  $\alpha$ , target leakage  $\beta$ .
4: if  $t = 1$  then
5:    $\varepsilon_i \leftarrow \epsilon$  for  $i = 1, \dots, k$ 
6: else if STOP has been issued then
7:   return  $y_t \sim h_{\emptyset}(x_t)$ 
8: end if
9: for  $i = 1, \dots, k$  do
10:   $\bar{h}_i(x_t) \leftarrow \frac{1}{2}(h_{\pi_i}(x_t) + h_{\pi'_i}(x_t))$ 
11:   $\lambda_i \leftarrow \max_{\lambda \in [0,1]} \{ \lambda : D_i(x_t, \lambda) \leq \beta \}$ 
12: end for
13:  $\lambda^* \leftarrow \frac{1}{k} \sum_{i=1}^k \lambda_i$ 
14:  $\bar{h}(x_t) \leftarrow \frac{1}{k} \sum_{i=1}^k \bar{h}_i(x_t)$ 
15:  $h(x_t) \leftarrow \lambda^* \bar{h}(x_t) + (1 - \lambda^*) h_{\emptyset}(x_t)$ 
16: for  $i = 1, \dots, k$  do
17:   $\lambda_{-i}^* \leftarrow \frac{1}{k-1} \sum_{j \neq i} \lambda_j$ 
18:   $\bar{h}_{-i} \leftarrow \frac{1}{k-1} \sum_{j \neq i} \bar{h}_j(x_t)$ 
19:   $\mathfrak{h} \leftarrow \lambda^* \bar{h}(x_t) + (1 - \lambda^*) h_{\emptyset}(x_t)$ 
20:   $\mathfrak{h}' \leftarrow \lambda_{-i}^* \bar{h}_{-i}(x_t) + (1 - \lambda_{-i}^*) h_{\emptyset}(x_t)$ 
21:   $\varepsilon_i \leftarrow \varepsilon_i - \max\{D_{\alpha}(\mathfrak{h} \parallel \mathfrak{h}'), D_{\alpha}(\mathfrak{h}' \parallel \mathfrak{h})\}$ 
22: end for
23: if  $\forall i : \varepsilon_i > 0$  then
24:    $y_t \sim h(x_t)$ 
25: else
26:   Issue STOP signal
27:    $y_t \sim h_{\emptyset}(x_t)$ 
28: end if
29: return  $y_t$ 

```

---

**Definition 20** ( $(\alpha, \beta)$ -Accuracy, Neel et al. [2020]) We say a pair  $(\mathcal{A}, \mathcal{U})$  of learning and unlearning algorithms is  $(\alpha, \beta)$ -accurate if for every dataset  $D \in \mathcal{D}$  and every update sequence  $S$ , the following condition holds: for every  $i \geq 0$ :

$$\Pr \left[ f_{D_i}(\tilde{h}_i) - \min_{h \in \mathcal{H}} f_{D_i}(h) \right] < \beta$$

Where  $f_{D_i}(h)$  is the loss of model  $h$  on the dataset after update  $i$  and  $\tilde{h}_i$  is the most recent published model.




Figure B.2: Privacy-Utility tradeoff from the perturbation methods of Majmudar et al. [2022]. This figure shows perplexity as a function of DP- $\epsilon$  values. The dotted line is the perplexity of the model without privacy guarantees. We see that the model seems to lose most utility at an  $\epsilon \leq 50$ .

**Definition 21** (*Strong vs Weak Unlearning, Neel et al. [2020]*) Fix any pair  $(\mathcal{A}, \mathcal{U})$  of learning and unlearning algorithms that satisfy  $(\alpha, \beta)$ -accuracy with respect to some publishing function  $f_{\text{publish}}$ . Let  $C_i$  represent the overall computational cost of the unlearning algorithm at step  $i$  of the update. We say  $\mathcal{U}$  is a “strong” unlearning algorithm for  $\mathcal{A}$  if:

1.  $\alpha$  and  $\beta$  are independent of the length of the update sequence, and
2. For every  $i \geq 1$ ,  $C_i/C_1 = \mathcal{O}(\log(i))$ , i.e. the computation cost of the unlearning algorithm must grow at most logarithmically with  $i$ .

If (1) holds and  $\forall i \geq 1, C_i/C_1 = \Omega(\text{poly}(i))$ , we say  $\mathcal{U}$  is a “weak” unlearning algorithm for  $\mathcal{A}$ .

Most work in the LLM unlearning space has been done under the weak unlearning framework.



**Algorithm B.3** DP-FedAvg and DP-FedSGD, McMahan et al. [2018]

---

```

1: Main Algorithm
2: Input: User selected probability  $q \in (0, 1]$ , per-user example cap  $\hat{w} \in \mathbb{R}^+$ , noise scale  $z \in \mathbb{R}^+$ , estimator  $\tilde{f}_f$  or  $\tilde{f}_c$ 
   with param  $W_{min}$ , UserUpdate (for FedAvg or FedSGD), ClipFn (FlatClip or PerLayerClip), moments accountant
    $\mathcal{M}$  Abadi et al. [2016]
3: Initialize model  $\theta^0$ , moments accountant  $\mathcal{M}$ ,  $w_k = \min(\frac{n_k}{\hat{w}}, 1)$  for all users  $k$ ,  $W = \sum_{k \in \mathcal{D}} w_k$ .
4: for each round  $t = 0, 1, \dots$  do  $\mathcal{C}^t \leftarrow$  (sample users with probability  $q$ )
5:   for each user  $k \in \mathcal{C}^t$  do
6:      $\Delta_k^{t+1} \leftarrow \text{UserUpdate}(k, \theta^t, \text{ClipFn})$ 
7:   end for
8:    $\Delta^{t+1} = \begin{cases} \frac{\sum_{k \in \mathcal{C}^t} w_k \Delta_k}{qW} & \text{for } \tilde{f}_f \\ \frac{\sum_{k \in \mathcal{C}^t} w_k \Delta_k}{\max(qW_{min}, \sum_{k \in \mathcal{C}^t} w_k)} & \text{for } \tilde{f}_c \end{cases}$ 
9:    $S \leftarrow$  (bound on  $\|\Delta_k\|$  for ClipFn)
10:   $\sigma \leftarrow \{\frac{zS}{qW}$  for  $\tilde{f}_f$  or  $\frac{2zS}{qW_{min}}$  for  $\tilde{f}_c\}$ 
11:   $\theta^{t+1} \leftarrow \theta^t + \Delta^{t+1} + \mathcal{N}(0, I\sigma^2)$ 
12:   $\mathcal{M}.\text{accum\_priv\_spending}(z)$ 
13: end for
14: print  $\mathcal{M}.\text{get\_privacy\_spent}()$ 
15:
16: FlatClip( $\Delta$ ):
17: Parameters:  $S$ 
18: return  $\pi(\Delta, S)$ 
19:
20: PerLayerClip( $\Delta$ ):
21: Parameters:  $S_1, \dots, S_m$ 
22:  $S = \sqrt{\sum_j S_j^2}$ 
23: for each layer  $j \in \{1, \dots, m\}$  do
24:    $\Delta'(j) = \pi(\Delta(j), S_j)$ 
25: end for
26: return  $\Delta'$ 
27:
28: UserUpdateFedAvg( $k, \theta^0, \text{ClipFn}$ ):
29: Parameters:  $B, E, \eta$ 
30:  $\theta \leftarrow \theta^0$ 
31: for each local epoch  $i$  from 1 to  $E$  do
32:    $\mathcal{B} \leftarrow$  ( $k$ 's data split into size  $B$  batches)
33:   for batch  $b \in \mathcal{B}$  do
34:      $\theta \leftarrow \theta - \eta \nabla l(\theta; b)$ 
35:      $\theta \leftarrow \theta^0 + \text{ClipFn}(\theta - \theta^0)$ 
36:   end for
37: end for
38: return update  $\Delta_k = \theta - \theta^0$  // Already clipping
39:
40: UserUpdateFedSGD( $k, \theta^0, \text{ClipFn}$ ):
41: Parameters:  $B, \eta$ 
42: select a batch  $b$  of size  $B$  from  $k$ 's examples
43: return update  $\Delta_k = \text{ClipFn}(-\eta \nabla l(\theta; b))$ 

```

---