

PROVABLE UNLEARNING IN TOPIC MODELING AND DOWNSTREAM TASKS

Stanley Wei*, Sadhika Malladi*, Sanjeev Arora
Princeton University
{stanley.wei, smalladi, arora}@princeton.edu

Amartya Sanyal
University of Copenhagen
amsa@di.ku.dk

ABSTRACT

Machine unlearning algorithms are increasingly important as legal concerns arise around the provenance of training data, but verifying the success of unlearning is often difficult. Provable guarantees for unlearning are often limited to supervised learning settings. In this paper, we provide the first theoretical guarantees for unlearning in the pre-training and fine-tuning paradigm by studying topic models, simple bag-of-words language models that can be adapted to solve downstream tasks like retrieval and classification. First, we design a provably effective unlearning algorithm for topic models that incurs a computational overhead independent of the size of the original dataset. Our analysis additionally quantifies the deletion capacity of the model – *i.e.*, the number of examples that can be unlearned without incurring a significant cost in model performance. Finally, we formally extend our analyses to account for adaptation to a given downstream task. In particular, we design an efficient algorithm to perform unlearning after fine-tuning the topic model via a linear head. Notably, we show that it is easier to unlearn pre-training data from models that have been fine-tuned to a particular task, and one can unlearn this data without modifying the base model.

1 INTRODUCTION

Modern-day machine learning has shifted from single-stage supervised learning on manually constructed datasets to a paradigm in which models are pre-trained and subsequently fine-tuned (Bommasani et al., 2022). In this setting, a model initially learns a good representation of the data using a self-supervised objective on a large unstructured corpus. The resulting pre-trained model is later adapted to solve specific tasks for which it is difficult or costly to curate a large dataset. This blueprint has yielded strong performance in text (*e.g.*, Devlin et al., 2019; Brown et al., 2020), vision (*e.g.*, Oquab et al., 2024; He et al., 2022), and multimodal (*e.g.*, Radford et al., 2021; Zhai et al., 2023) settings. It is well-known that the scale of the pre-training data is strongly correlated with the final performance of the model (Hoffmann et al., 2022), leading to the construction of larger datasets via broad internet scrapes (Gao et al., 2020; Schuhmann et al., 2022; Soldaini et al., 2024; Penedo et al., 2023). Such datasets have been found to often inadvertently include private, sensitive, and unsafe data (Birhane et al., 2021; Longpre et al., 2024; He et al., 2024).

Unsafe data can generally degrade model performance and introduce biases, making the model less useful for various applications (McKenna et al., 2023; Birhane & Prabhu, 2021; Choenni et al., 2021; Naous et al., 2024). Using private and sensitive data, even unknowingly, poses legal risks (Bommasani et al., 2022; Henderson et al., 2023). In particular, recent works have shown that models can memorize and thus permit the extraction of training data (Somepalli et al., 2023; Carlini et al., 2021; 2023). Moreover, one may be requested to remove data in accordance with GDPR’s *right to be forgotten* (European Parliament & Council of the European Union), or as part of a copyright-related lawsuit (*Tremblay v. OpenAI, Inc.*, 2023; *DOE 1 v. GitHub, Inc.*, N.D. Cal. 2022).

Therefore, there is great empirical interest in developing machine unlearning algorithms that can surgically remove portions of the training data from an already learned model without harming performance. The gold standard for machine unlearning is for the model to behave as though it had

*Equal contribution.

never been trained on that datapoint (Cao & Yang, 2015). As it is often undesirable to completely retrain models, especially as they grow larger, many works have proposed computationally cheaper heuristics for solving this problem (e.g., Jang et al., 2023; Foster et al., 2024; Kurmanji et al., 2023; Zhang et al., 2024b; Eldan & Russinovich, 2023; Gandikota et al., 2023). In the absence of theoretical guarantees, it is common to use empirics to measure the success of these algorithms. However, recent works have shown that such evaluations often overestimate the success of these unlearning methods (Hayes et al., 2024; Shi et al., 2024; Maini et al., 2024) and thus it has proven difficult to confidently ascertain whether the proposed methods meet the necessary compliance standards. In this context, it is highly desirable to design efficient unlearning algorithms with well-motivated guarantees that are salient to the pre-training and finetuning paradigm (Thudi et al., 2022; Lee et al., 2024).

While there are some instances of such algorithms for linear models (Guo et al., 2020; Izzo et al., 2021; Mahadevan & Mathioudakis, 2023), general convex models (Ullah et al., 2021; Sekhari et al., 2021; Neel et al., 2021), Bayesian models (Nguyen et al., 2020), and GANs (Liu et al., 2024), there are no works on the paradigm of pre-training and fine-tuning algorithms. One of the most classical such algorithms is topic modeling (Hofmann et al., 1999; Blei et al., 2003; Blei & Lafferty, 2006; Li & McCallum, 2006), which can also be thought of as the simplest language model. *In this paper, we present the first provably effective and efficient unlearning algorithms for topic models.*

Topic models are generally pre-trained to extract latent structure (i.e., a small set of underlying topics) from a large corpus of documents. This feature extractor is then used for a variety of downstream applications, including retrieval, classification, and recommendation (Boyd-Graber et al., 2017). Despite their simplicity, topic models can be used to effectively solve many real-world natural language problems — see a survey in Churchill & Singh (2022).

1.1 OVERVIEW OF RESULTS

We focus on the setting in Arora et al. (2012b), because it admits an efficient learning algorithm with provable guarantees (Arora et al., 2012a). The corpus is assumed to contain r underlying topics, where each topic defines a distribution over words. Let \mathcal{D} be a distribution over topic distributions. Then, each document d is generated by sampling a topic distribution $W_d \sim \mathcal{D}$ over topics, and then sampling words according to W_d .

The dataset of m documents is a matrix $M \in \mathbb{R}^{n \times m}$, where M permits a non-negative matrix factorization $M = A^* X$. Here, $A^* \in \mathbb{R}^{n \times r}$ is the distribution of words in each of the r unknown underlying topics, and $X \in \mathbb{R}^{r \times m}$ is the sampled distribution of topics in each document. In particular, A^* , X have columns on the probability simplex. We seek to learn the embedding function A^* and the topic-topic covariance $R^* = \mathbb{E}_{\mathcal{D}}[X X^\top]$.

To derive provable guarantees on the success of unlearning, we adapt the notion of (ϵ, δ) -unlearning introduced in Sekhari et al. (2021) to the topic modeling setting. The unlearned model is required to behave indistinguishably from a model that was retrained on the modified dataset. We define a notion of *utility-preserving unlearning* that combines this condition with an analysis on the *deletion capacity* – i.e., the number of datapoints that can be unlearned without performance degradation (Definition 4). We now state our main result on utility-preserving unlearning in topic models.

Main Result 1 (Informal version of Theorem 2). Suppose we trained a topic model A^S, X^S on a training set S containing m documents. Algorithm 1 can perform utility-preserving unlearning of

$$m_U = \tilde{O}\left(\frac{m}{r^2 \sqrt{nr}}\right)$$

documents from the pre-trained topic model, where $\tilde{O}(\cdot)$ hides constants depending on the learning and unlearning algorithm.

To adapt a topic model to a downstream topic classification task, we learn a head $w \in \mathbb{R}^r$ on top of A to minimize a strongly convex loss function (Definition 2). When A and w are both released, one would necessarily have to first unlearn from A , which makes unlearning just as hard as it was in pre-training (Theorem 3). This setting is rather unrealistic, because there is no obvious case in which one would want to use w without A or vice versa. We thus advocate for viewing fine-tuned model $B = Aw$ as a whole i.e. it is not allowed to access outputs of A solely, and we show that it is easier to perform utility-preserving unlearning of pre-training data in this case.

Main Result 2 (Informal version of Theorem 4). After adapting the model to a downstream task (Definitions 1 and 2), Algorithm 2 can perform utility-preserving unlearning of $\tilde{\Omega}\left(\frac{mq}{r\sqrt{nr}}\right)$ documents, where $q \in [1/r, 1]$ is a task-dependent quantity, without modifying the base model \mathbf{A} . Simpler downstream tasks have a larger q , increasing the separation from the pre-training result.

We demonstrate that our unlearning algorithms run substantially faster than retraining the model (Table 1). Overall, our results imply the following takeaways in the context of topic models — (1) It is possible to effectively and efficiently unlearn datapoints from a pre-trained model without retraining it (Algorithm 1 and Theorem 2), (2) One can effectively unlearn more pre-training data from a model that has been adapted to a downstream task without harming the utility of the base and fine-tuned models (Theorem 4), and (3) One can unlearn pre-training data from a fine-tuned model without modifying the base model (Algorithm 2 and Theorem 4).

2 TOPIC MODELS

As we previously discussed, topic models can be considered as one of the simplest language models that one can pre-train in a self-supervised fashion and later fine-tune for other language-related tasks. This pipeline mirrors the modern-day paradigm of pre-training large language models to build a general understanding of natural language and later fine-tuning them to solve a variety of tasks ranging from classification to code generation.

2.1 PROBLEM DESCRIPTION

Topic modeling is a classical, bag-of-words method to discover structure in a corpus of documents (Hofmann et al., 1999). One assumes that each document contains a convex combination of topics, each of which can be described in terms of a distribution over the vocabulary. Different assumptions on the structure of this distribution and the topics have yielded a variety of topic modeling methodologies (Blei & Lafferty, 2006; Li & McCallum, 2006) – perhaps most famous among these is the latent Dirichlet allocation (LDA, Blei et al. (2003)). Many early works established the statistical learnability of topic models under such assumptions, but the learning algorithms generally were not efficient in real-world settings (Arora et al., 2012b; Recht et al., 2012).

Our paper focuses on the setting in Arora et al. (2012b), for which Arora et al. (2012a) provided an empirically efficient learning algorithm. The dataset consists of a set of m documents d_1, \dots, d_m , where each document contains L words from a vocabulary \mathcal{V} with $|\mathcal{V}| = n$.¹ The corpus contains r different underlying topics, each of which defines a distribution over words. Each word in document d is generated by: (1) sampling a distribution over topics $W_d \sim \mathcal{D}$, and then (2) sampling L words independently according to W_d .

We represent the corpus as a matrix $\mathbf{M} \in \mathbb{R}^{n \times m}$, where \mathbf{M} permits a non-negative matrix factorization $\mathbf{M} = \mathbf{A}^* \mathbf{X}$. Here, $\mathbf{A}^* \in \mathbb{R}^{n \times r}$ is the distribution of words in each of the r topics, $\mathbf{X} \in \mathbb{R}^{r \times m}$ is the distribution of topics in each document, and hence \mathbf{M} is the distribution of words in each document. While there are several algorithms for learning the feature extractor \mathbf{A}^* , it is well-known that it is hard to recover \mathbf{X} exactly (Arora et al., 2012b). Instead, it is desirable to learn how the topics co-occur together, denoted as $\mathbf{R}^* = \mathbb{E}_{\mathcal{D}}[\mathbf{X} \mathbf{X}^\top]$. This quantity is termed the *topic-topic covariance*. Further discussion of this has been included in Appendix A.

The topic modeling setting generally determines \mathcal{D} (e.g., in LDA, \mathcal{D} is a Dirichlet distribution). In order to recover \mathbf{A}^* and \mathbf{R}^* efficiently and accurately from an observed corpus $\mathbf{M} \sim \mathcal{D}$, we need to make the following assumption on the underlying data distribution.

Assumption 1 (p -separability, Arora et al. (2012b)). *The topic matrix \mathbf{A}^* is p -separable for $p > 0$ if for every topic $k \in [r]$, there exists a word $i \in [n]$ such that $\mathbf{A}_{i,k}^* \geq p$ and $\mathbf{A}_{i,k'}^* = 0$ for all $k' \neq k$. Such words are called anchor words.*

Without this separability assumption, maximum likelihood estimation of a topic model is NP-hard (Arora et al., 2012b). Assumption 1 requires that \mathbf{A}^* contains a diagonal matrix, up to row

¹Without loss of generality, we assume $L = 2$. The same asymptotic results hold for constant sized L , and it is straightforward to modify our analysis to account for this.

permutations; intuitively, the appearance of an anchor word in a document perfectly indicates the document has nonzero probability of the corresponding topic. As we will detail in Section 4, this observation inspires a two-phase learning algorithm, whereby one first approximates the anchor words for each topic and then leverages them to identify patterns among the topics.

2.2 DOWNSTREAM ADAPTATION

Topic models are frequently trained on a general corpus, and the embeddings can be later used to classify documents. The classification problem usually involves only a subset of topics. For example, after training a topic model on a large corpus of news articles with diverse topics (e.g., sports, politics, technology, finance, etc.), one relevant downstream task is to classify the subject of a given news article as sports or politics. We formalize the topic classification task below.

Definition 1 (Topic Classification Task). A topic classification task $\mathcal{T} = (\mathbb{T}_{\text{clf}}, \mathbf{w}^*)$ is defined by a subset of topics $\mathbb{T}_{\text{clf}} \subset [r]$ on which the task is defined and a ground-truth labelling vector $\mathbf{w}^* \in \mathbb{R}^r$ with bounded norm. Importantly, \mathbf{w}^* only has non-zero coordinates in the positions corresponding to \mathbb{T}_{clf} .

The classification task is defined on the latent features of a given document, so it is necessary to first identify the salient topics as they occur in the text. Fitting a topic model to the corpus yields such a feature extractor \mathbf{A} that embeds a document into the r -dimensional topic space. In order to adapt a topic model to a particular classification task, we perform head tuning on the feature extractor \mathbf{A} .

Definition 2 (τ -Head Tuning). For a given labelled document classification dataset $\mathbb{D}_{\text{clf}} = \{(d_i, y_i)\}$ representing a topic classification task \mathcal{T} , embed each document d_i as a vector $\mathbf{x}_i \in \mathbb{R}^r$ containing the word counts in the document. To perform head tuning on a pre-trained topic model \mathbf{A} , we learn $\mathbf{w}_\tau \in \mathbb{R}^r$ that is a τ -optimal point of

$$\ell_{\mathcal{T}}(\mathbf{w}; \mathbf{A}) = \frac{1}{|\mathbb{D}_{\text{clf}}|} \sum_{(\mathbf{x}, y) \in \mathbb{D}_{\text{clf}}} f(\mathbf{x}^\top \mathbf{A} \mathbf{w}, y)$$

where $\ell_{\mathcal{T}}$ is strongly convex in \mathbf{w} . In other words, \mathbf{w}_τ satisfies $\ell_{\mathcal{T}}(\mathbf{w}_\tau; \mathbf{A}) - \min_{\mathbf{w}} \ell_{\mathcal{T}}(\mathbf{w}; \mathbf{A}) \leq \tau$.

One example of f is the logistic loss with ℓ_2 regularization. For ease of exposition, we primarily consider binary classification tasks, but we point out that the definition can extend to multi-class tasks solved via the one-vs-all scheme (Rifkin & Klautau, 2004).

We note that head tuning, also referred to as linear probing, is a simpler adaptation technique than fine-tuning \mathbf{A} alongside \mathbf{w} . Nonetheless, recent works on popular language models have demonstrated that head tuning can substantially improve the ability of general pre-trained language models to solve complex classification tasks (Malladi et al., 2023a;b). Head tuning thus serves as a convenient yet effective adaptation method that avoids updating the pre-trained model, which is often desirable. For example, if a single pre-trained model needs to be separately adapted to solve many different tasks, then it is desirable to minimize the number of parameters that are fine-tuned to minimize the memory needed to store all of the adapted models.²

3 UNLEARNING

As mentioned previously, there is increased interest in machine unlearning due to the growing scale of modern datasets and the difficulty of manually inspecting each datapoint. Theoretically, the gold standard for unlearning is that the model should behave identically to one that was trained without the datapoint in its corpus (Cao & Yang, 2015). We first define what it means for two models $\theta_1, \theta_2 \in \Theta$ to behave *almost* identically, where Θ denotes the parameter space of a hypothesis class. Due to randomness in learning, θ_1 and θ_2 are random variables.

Definition 3 ((ϵ, δ) -indistinguishable models, Dwork et al. (2014)). Two models denoted by random variables $\theta_1, \theta_2 \in \Theta$ are (ϵ, δ) -indistinguishable if for all possible subsets of models $T \subseteq \Theta$,

$$\Pr(\theta_1 \in T) \leq e^\epsilon \Pr(\theta_2 \in T) + \delta \quad \text{and} \quad \Pr(\theta_2 \in T) \leq e^\epsilon \Pr(\theta_1 \in T) + \delta$$

²This motivation has driven widespread development and adoption of parameter-efficient fine-tuning methods for large language models. Liu et al. (2021) contains a survey of such techniques.

We adapt the definitions from Sekhari et al. (2021) to the topic modeling setting. A learning algorithm \mathcal{A} takes in a set of m documents, denoted as S , and returns a topic model $\theta = (\mathbf{A}, \mathbf{R})$. Analogously, an unlearning algorithm \mathcal{U} takes in the learned topic model θ , a set of documents to unlearn $S_f \subseteq S$, some statistics on the training set $T(S)$, and outputs a model. The set of datapoints to unlearn S_f is often referred to as the *forget set*. With this in mind, we now define a notion of utility-preserving unlearning, whereby the unlearning algorithm needs to not only effectively simulate retraining the model from scratch but also maintain the model’s performance.

Definition 4 (Utility-preserving (ϵ, δ) -Unlearning with Deletion Capacity). Let $m \in \mathbb{N}$ be a constant that depends on the topic modeling distribution \mathcal{D} satisfying Assumption 1. For any training dataset $S \stackrel{\text{i.i.d.}}{\sim} \mathcal{D}$ of size at least m , and $\epsilon, \delta > 0$, we say that a pair of learning and unlearning algorithms $(\mathcal{A}, \mathcal{U})$ performs *utility-preserving unlearning with deletion capacity* $T_{\epsilon, \delta}^{\mathcal{A}, \mathcal{U}}(m)$ if

1. With probability at least 0.9 over draws from \mathcal{D} , for any forget set $S_f \subseteq S$ of size at most $T_{\epsilon, \delta}^{\mathcal{A}, \mathcal{U}}(m)$, model trained on $S \setminus S_f$ is indistinguishable from that resulting from unlearning with \mathcal{U} . Here, $\stackrel{\epsilon, \delta}{\approx}$ denotes (ϵ, δ) -indistinguishability.

$$\mathcal{U}(S_f, \mathcal{A}(S), T(S)) \stackrel{\epsilon, \delta}{\approx} \mathcal{U}(\emptyset, \mathcal{A}(S \setminus S_f), T(S \setminus S_f))$$

2. Even for an adversarially chosen S_f , the unlearned model does not suffer a large performance degradation. Formally,

$$\mathbb{E}_{\mathcal{A}, \mathcal{U}} \left[\max_{|S_f| \leq T_{\epsilon, \delta}^{\mathcal{A}, \mathcal{U}}(m)} h(\mathcal{U}(S_f, \mathcal{A}(S), T(S))) - h^* \right] \leq 0.01$$

where $h : \Theta \rightarrow \mathbb{R}$ is the loss of the topic model, and $h^* = \min_{w \in \mathcal{W}} h(w)$ is the irreducible loss.

The above definition can be applied to both the pre-training and the downstream adaptation stages of training a topic model. Of particular note is that (1) does not guarantee (2), since the former only concerns indistinguishability between the unlearned and retrained models, while the latter is a statement about utility preservation. Moreover, unless $T(S)$ contains the entire dataset, we note that the unlearning algorithm \mathcal{U} cannot be as simple as retraining the model. In this paper, we will design an unlearning algorithm for topic models that satisfies this definition of provable unlearning, and the number of statistics $T(S)$ will not depend on the initial dataset size m .

To show (ϵ, δ) -indistinguishability, we utilize the Gaussian mechanism, a classic tool from differential privacy. Given a particular function, the Gaussian mechanism essentially prescribes how much noise one must add to the output in order for the input to be indistinguishable from a similar one. The guarantee of the Gaussian mechanism is described in the following lemma.

Lemma 1 (Gaussian Mechanism, Dwork et al. (2014)). *Let f be an arbitrary d -dimensional function, and define its ℓ_2 -sensitivity to be $\Delta_2 f := \max_{\text{adjacent } x, y} \|f(x) - f(y)\|_2$. Then, for $c^2 > 2 \log \frac{1-2\delta}{\delta}$, the Gaussian mechanism with parameter $\sigma \geq c\Delta_2 f / \epsilon$ is (ϵ, δ) -differentially private.*

In our case, we define adjacent inputs (i.e., training datasets) as the case where y is a superset of x .

4 LEARNING AND UNLEARNING TOPIC MODELS

In this section, we present the learning and unlearning algorithms and guarantees for topic models.

Notation. We use \mathbf{A}^* to refer to the ground-truth topic model, \mathbf{A}^S to refer to a topic model trained on S , and \mathbf{A}^F to denote a topic model retrained with the forget set removed $S \setminus S_f$. We also use $\hat{\mathbf{A}}$ to denote the unlearned topic model before applying the Gaussian mechanism and $\tilde{\mathbf{A}}$ to denote the model after the mechanism is applied. Analogous notations are used for \mathbf{R} .

4.1 LEARNING ALGORITHM AND GUARANTEES

Per Arora et al. (2012a), the learning algorithm $\mathcal{A}_{\text{base}}$ takes in a corpus of documents $S = \{d_1, \dots, d_m\}$ and consists of the following three phases to learn a topic model $\theta = (\mathbf{A}^S, \mathbf{R}^S)$.

1. **Measure the word co-occurrences.** Compute the word co-occurrence matrix $\mathbf{Q} \in \mathbb{R}^{n \times n}$, where Q_{ij} is the number of times word i appears in the same document as word j . We also compute $\bar{\mathbf{Q}}$, which normalizes the rows of \mathbf{Q} to sum to 1. A detailed discussion of the construction of \mathbf{Q} and its relationship to the factorization $\mathbf{M} = \mathbf{A}^* \mathbf{X}$ is included in Appendix A.
2. **Identify the anchor words P .** Recall that in order to be able to learn topic models efficiently, there must exist a set of anchor words P with $|P| = r$, and each anchor word must appear exclusively in a single topic (Assumption 1). This subroutine uses $\bar{\mathbf{Q}}$ to approximately identify the r anchor words P .
3. **Learn the feature extractor \mathbf{A}^S and the topic-topic covariance \mathbf{R}^S .** The algorithm uses the anchor words P and the word co-occurrences $\bar{\mathbf{Q}}$ to learn \mathbf{A}^S and \mathbf{R}^S . Each word is expressed as a convex combination of anchor words, and thus, topics. With appropriate normalization and by cross-referencing information with the co-occurrence matrix, one can recover \mathbf{A}^* , \mathbf{R}^* in the infinite data limit.

We sketch how this algorithm recovers the ground truth \mathbf{A}^* , \mathbf{R}^* when one has infinitely many documents in Appendix A. Arora et al. (2012a) gives the following finite-document guarantee.

Theorem 1 (Learning Guarantee). *Running \mathcal{A}_{base} on a dataset S of size m , where m is at least*

$$\max \left\{ \mathcal{O} \left(\frac{ar^3 \log n}{L(\gamma p)^6 \epsilon_0} \right), \mathcal{O} \left(\frac{a^3 r^3 \log n}{L \epsilon_0^3 (\gamma p)^4} \right), \mathcal{O} \left(\frac{r^2 \log r}{L \epsilon_0^2} \right) \right\}$$

recovers \mathbf{A}^S and \mathbf{R}^S with entrywise additive error up to ϵ_0 from the ground truth \mathbf{A}^ , \mathbf{R}^* , respectively. Here, a is the topic imbalance parameter, and γ is the condition number of the ground truth \mathbf{R}^* . Formally, we have $a = \max_{i,j \in [r]} \Pr_{\mathcal{D}}[z = i] / \Pr_{\mathcal{D}}[z = j]$.*

Approximating the anchor words. We defer a precise description of the anchor word identification algorithm to Appendix A and instead focus here on the intuitions driving its design and the guarantees we will use throughout the paper. First, we note the relationship between $\bar{\mathbf{Q}}$ and the set of anchor words. If we had infinitely many documents, then the convex hull of the rows in $\bar{\mathbf{Q}}$ will be a simplex with vertices corresponding to the anchor words, because each anchor word corresponds to a topic, and each topic prescribes a distribution over words. However, in the finite document setting, each row of $\bar{\mathbf{Q}}$ only approximates their expected value, and so one must approximate the vertices of a convex hull when given access to a perturbation of the points that define it.

We start by requiring that each topic is distinctly different from any mixture on the other topics. Formally, this requires that the simplex is robust, in that each vertex (i.e., anchor word) is sufficiently far from any combination of the other topics. Most topic modeling settings define lower bounds on the robustness of the simplex. By a result in Arora et al. (2012b), the simplex defined by the r anchor word rows of the population $\bar{\mathbf{Q}}$ is γp -robust. We can now define exactly the sense in which a $\bar{\mathbf{Q}}$ computed on a finite dataset approximates the population co-occurrence matrix.

Definition 5. Let $\{a_i\}_{i=1}^n$ be a set of points whose convex hull P is a simplex with vertices $\{v_i\}_{i=1}^r$. We say a set of r points is ϵ -close to the vertex set $\{v_i\}_{i=1}^r$ if each of the r points is ϵ -close in ℓ_2 distance to a different vertex in P . Moreover, we say that a simplex P is β -robust if for every vertex v of P , the ℓ_2 distance between v and the convex hull of the rest of the vertices is at least β .

In the context of this definition, P corresponds to the ground truth convex hull, and the finite sample $\bar{\mathbf{Q}}$ can be seen as a perturbation to it. In particular, Arora et al. (2012a) used this to establish a guarantee on the accuracy of anchor word recovery.

Lemma 2 (Approximation Guarantee on Anchor Words). *Suppose each row of $\bar{\mathbf{Q}}$ is at most δ distance away from the ground truth γp -robust simplex \mathbf{Q}^* in ℓ_2 norm. If $20r\delta/(\gamma p)^2 < \gamma p$, then the set of anchor words found by the algorithm is $\mathcal{O}(\delta/\gamma p)$ -close to the ground truth anchor words.*

We now describe how to use the recovered approximate anchor words to learn the topic model.

Learning the topic model from anchor words. We are given the set of anchor words P , the word co-occurrence matrix $\mathbf{Q} \in \mathbb{R}^{n \times n}$, and the normalized co-occurrence matrix $\bar{\mathbf{Q}}$. Our goal is to use these quantities to learn $\mathbf{A} \in \mathbb{R}^{n \times r}$ and $\mathbf{R} \in \mathbb{R}^{r \times r}$. We will do so by first expressing each word $i \in [n]$ as a convex combination of the anchor words (and thus, the topics). In particular, for each word i , we learn the coefficients $\mathbf{C}_i \in \Delta_r$ as

$$\mathbf{C}_i = \arg \min_{v \in \Delta_r} \|\bar{\mathbf{Q}}_i - v^\top \bar{\mathbf{Q}}_P\|^2$$

Algorithm 1 Unlearning algorithm ($\mathcal{U}_{\text{base}}$)

Input: Forget set $S_f \subseteq S$, statistics $T(S)$ which include $\{C_i^S\}_{i=1}^n$, \mathbf{Q}^S, P , normalization constants \mathbf{p}^S

Output: Unlearned model $\tilde{\mathbf{A}}, \tilde{\mathbf{R}}$

Compute the updated co-occurrence matrix \mathbf{Q}^F by subtracting documents in S_f

Store the updated normalization constants $\mathbf{p}^F = \mathbf{Q}^F \mathbf{1}$

for i in $1, \dots, n$ **do**

 Newton step update on C_i 's:

$$\bar{C}_i^F \leftarrow C_i^S - H_{C_i^S}^{-1} \nabla \mathcal{L}(C_i^S, S \setminus S_f)$$

$$\bar{C}_i^F \leftarrow \text{proj}_{\Delta_r}(\bar{C}_i^F)$$

 where $\mathcal{L}(\mathbf{v}, S \setminus S_f) := \|\bar{\mathbf{Q}}_i^F - \mathbf{v}^\top \bar{\mathbf{Q}}_P^F\|^2$ and $H_{C_i^S} = \nabla^2 \mathcal{L}(C_i^S, S \setminus S_f)$

end for

$\bar{\mathbf{A}}' = \text{diag}(\mathbf{p}^F) \bar{\mathbf{C}}$

$\bar{\mathbf{A}} =$ column normalized $\bar{\mathbf{A}}'$

$\bar{\mathbf{R}} = \bar{\mathbf{A}}^\dagger \mathbf{Q}^F \bar{\mathbf{A}}^{\dagger\top}$ where $\bar{\mathbf{A}}^\dagger$ is the pseudoinverse of $\bar{\mathbf{A}}$

Sample $\nu_{\mathbf{A}}, \nu_{\mathbf{R}}$ from normal distribution defined by Gaussian mechanism guarantee

$\hat{\mathbf{A}} =$ Project each column of $\bar{\mathbf{A}} + \nu_{\mathbf{A}}$ to Δ_n .

$\hat{\mathbf{R}} =$ Project $\bar{\mathbf{R}} + \nu_{\mathbf{R}}$ onto the set of PSD matrices.

return The unlearned topic model $\hat{\mathbf{A}}, \hat{\mathbf{R}}$

where $\bar{\mathbf{Q}}_P$ is the P rows of $\bar{\mathbf{Q}}$ corresponding to the anchor words. Arora et al. (2012a) showed the following approximation guarantee for C_i compared to the ground-truth coefficients.

Lemma 3. When $20r\delta/(\gamma p)^2 < \gamma p$, for every word i , C_i has entrywise error $O(\delta/(\gamma p)^2)$ from C_i^* .

We then normalize this C_i by the total number of co-occurrences that word i is involved in. Note that the C_i can be assembled into a matrix $\mathbf{C} \in \mathbb{R}^{n \times r}$. We set \mathbf{A} to be \mathbf{C} after normalizing the columns sum to 1, since the columns represent the topic-conditioned distribution over the vocabulary. We finally compute $\mathbf{R} = \mathbf{A}^\dagger \mathbf{Q} \mathbf{A}^{\dagger\top}$, where \mathbf{A}^\dagger denotes the pseudoinverse of \mathbf{A} .

4.2 UNLEARNING ALGORITHM AND GUARANTEES

Learning Phase	Retrain Time	Unlearning Update	Unlearning Time
Co-occurrence matrix computation	$\mathcal{O}(m)$	Updating frequencies	$\mathcal{O}(m_U)$
Identify anchor words	$\mathcal{O}(n^2 + nr/\epsilon_0^2)$	Use learned anchor words	$\mathcal{O}(1)$
Recover topics from anchors	$\mathcal{O}(n^2 r + nr^2/\epsilon_0^2)$	Projected Newton step	$\mathcal{O}(nr^2)$
Head tuning \mathbf{w} (Definition 2)	ERM	Newton step	$\mathcal{O}(r^3)$

Table 1: Our unlearning algorithms generally have a runtime shorter than the retraining procedure. ERM denotes empirical risk minimization, and we note the training time relies on the error tolerance.

We describe our unlearning algorithm $\mathcal{U}_{\text{base}}$ to forget a set S_f from a trained model (Algorithm 1), which crucially updates C_i with a Newton step. We then compute $\hat{\mathbf{A}}$ from the modified C_i and apply the Gaussian mechanism to ensure indistinguishability. We describe our formal guarantee on the unlearning algorithm below, sketching out our utility preserving guarantees with respect to \mathbf{A}^* . The arguments for \mathbf{R}^* follow analogously; we defer the discussion to the appendix.

Theorem 2 (Utility-Preserving Unlearning on the Base Model). *Let $\mathcal{A}_{\text{base}}$ be the learning algorithm described in the prior sections and $\mathcal{U}_{\text{base}}$ be the unlearning algorithm in Algorithm 1. Then, $(\mathcal{A}_{\text{base}}, \mathcal{U}_{\text{base}})$ performs utility-preserving unlearning with deletion capacity*

$$T_{\epsilon, \delta}^{\mathcal{A}_{\text{base}}, \mathcal{U}_{\text{base}}}(m) \geq c \cdot \frac{m}{r^2 \sqrt{rn}}$$

where m is the number of training documents, r is the number of topics, and c is a constant dependent on ϵ, δ , and \mathcal{D} . The loss function h used in the utility-preserving definition is the maximum entrywise error from the ground truth topic model \mathbf{A}^* .

Proof sketch. The full proof can be found in Appendix B.2. We delete $m_U \leq \frac{0.001m\epsilon_0(\gamma p)^3}{a^2 r^2}$ points. This upper bound ensures that the anchor words are likely unchanged per Lemma 2. Recall that utility-preserving unlearning requires: (1) that the unlearned model is indistinguishable from the retrained model, and (2) that the unlearned model is not too far from the ground-truth model.

Indistinguishability. The Gaussian mechanism introduced in Lemma 1 allows us to make two models with a given ℓ_2 -sensitivity (ϵ, δ) -indistinguishable from each other. We bound the ℓ_2 -sensitivity of the feature extractor \mathbf{A} by noting that $\bar{\mathbf{A}}$ is a rescaled version of $\bar{\mathbf{C}}$.

Lemma 4. For $\epsilon, \delta > 0$, the following holds for the $\bar{\mathbf{C}}$ and the topic matrix $\bar{\mathbf{A}}$:

$$\|\bar{\mathbf{C}} - \mathbf{C}^F\|_\infty \leq c \cdot \frac{ar m_U}{m\epsilon_0 \gamma p} \quad \|\bar{\mathbf{A}} - \mathbf{A}^F\|_\infty \leq (ar) \cdot \|\bar{\mathbf{C}} - \mathbf{C}^F\|_\infty$$

Applying the Gaussian mechanism with noise $\sigma = \frac{\Delta}{\epsilon} \sqrt{2 \log(1.25/\delta)}$, where $\Delta = c\sqrt{nr} \cdot \frac{(ar)^2 m_U}{m\epsilon_0 \gamma p}$ and followed by projecting the columns of $\bar{\mathbf{A}} + \nu_{\mathbf{A}}$ back to Δ_n yields the desired result.

Utility Preservation. We first apply Lemma 2 to show that, with high probability, the anchor words do not change when unlearning m_U documents. Then, we use Lemma 8 to bound the distance between the unlearned $\bar{\mathbf{C}}_i$ and the ground truth \mathbf{C}_i^* . Accounting for the noise added via the Gaussian mechanism completes the proof.

Lemma 5. For $\epsilon, \delta > 0$, denote the unlearned model after the Gaussian mechanism described above as $\tilde{\mathbf{A}}$. Then, $\tilde{\mathbf{A}}$ satisfies:

$$\mathbb{E} \left[\|\tilde{\mathbf{A}} - \mathbf{A}^*\|_\infty \right] \leq c \cdot \frac{(ar)^2 m_U}{m\epsilon_0 \gamma p} \cdot \left(\sqrt{nr} \cdot \sqrt{\log(nr)} \cdot \frac{\sqrt{\log(1/\delta)}}{\epsilon} + 1 \right)$$

Each of the two terms in the above equation yield a constraint on m_U . In particular, $m_U \leq \min \left\{ \tilde{\mathcal{O}} \left(\frac{m}{r^2 \sqrt{nr}} \right), \mathcal{O} \left(\frac{m}{r^2} \right) \right\}$, so setting $m_U \leq \tilde{\mathcal{O}} \left(\frac{m}{r^2 \sqrt{nr}} \right)$ completes the proof.

5 UNLEARNING WITH RESPECT TO A DOWNSTREAM TASK

We are interested in unlearning a set of pre-training documents $S_f \subseteq S$. A topic classification task is usually defined on a subset of the topics in the dataset — for example, if the pre-training corpus contained diverse news articles, one plausible downstream task is to classify the content of a given document as containing politics or sports. Definition 1 formalizes this: a topic classification task $\mathcal{T} = (\mathbb{T}_{\text{clf}}, \mathbf{w}^*)$ is defined on a subset of the topics \mathbb{T}_{clf} and a r -length ground-truth labelling vector $\mathbf{w}^* \in \mathcal{W}_{\text{head}}$, where \mathbf{w}^* only has non-zero values in positions corresponding to \mathbb{T}_{clf} . We describe two possible settings under which we can show utility-preserving unlearning. For the sake of exposition, we will assume for now $\tau = 0$ in downstream head tuning; the extension to inexact head tuning ($\tau > 0$), which is a more realistic regime, will be deferred to the appendix.

5.1 NAIVE SETTING

In the first setting, the learning algorithm $\mathcal{A}_{\text{head, naive}}$ returns the pre-trained feature extractor \mathbf{A} and the head \mathbf{w} separately. So, we must ensure that the forget set $S_f \subseteq S$ cannot be recovered from either \mathbf{A} or \mathbf{w} . As such, we must necessarily perform unlearning on \mathbf{A} as described in Algorithm 1, which means that unlearning the fine-tuned model is exactly as difficult as unlearning the base model.

Theorem 3 (Unlearning when releasing \mathbf{A} and \mathbf{w}). For a downstream task \mathcal{T} with loss function $\ell_{\mathcal{T}}$, consider the unlearning algorithm $\mathcal{U}_{\text{head, naive}}$ that first runs Algorithm 1 to compute $\tilde{\mathbf{A}} = \mathcal{U}_{\text{base}}(S_f, \mathcal{A}_{\text{base}}(S), T(S))$, where $(\mathcal{A}_{\text{base}}, \mathcal{U}_{\text{base}})$ performs utility-preserving unlearning (Theorem 2). Then, it fits a head $\mathbf{w} = \arg \min_{\mathbf{w} \in \mathcal{W}_{\text{head}}} \ell_{\mathcal{T}}(\mathbf{w}; \tilde{\mathbf{A}})$ and returns $\tilde{\mathbf{A}}$ and \mathbf{w} . We assert that $(\mathcal{A}_{\text{head, naive}}, \mathcal{U}_{\text{head, naive}})$ performs utility-preserving unlearning (Definition 4).

Algorithm 2 Unlearning algorithm for task \mathcal{T} (\mathcal{U}_{head})

Input: Document deletion requests $S_f \subseteq S$, statistics $T(S)$ which include $\mathbf{A}^S, \{\mathbf{C}_i^S\}_{i=1}^n, \mathbf{Q}^S, P, \text{diag}(\mathbf{p}^S), \mathbf{w}^S = \arg \min_{\mathbf{w} \in \mathcal{W}_{head}} \ell_{\mathcal{T}}(\mathbf{w}; \mathbf{A}^S)$
 $\bar{\mathbf{A}}, \bar{\mathbf{R}} = \text{Run Algorithm 1 } (\mathcal{U}_{base})$ up to the Gaussian mechanism
 $\bar{\mathbf{w}} = \mathbf{w}^S - \mathbf{H}_{\mathbf{w}^S}^{-1} \nabla_{\mathbf{w}} \ell_{\mathcal{T}}(\mathbf{w}^S; \bar{\mathbf{A}})$ where $\mathbf{H}_{\mathbf{w}^S} = \nabla_{\mathbf{w}}^2 \ell_{\mathcal{T}}(\mathbf{w}^S; \bar{\mathbf{A}})$
return $(\mathbf{A}^S)^\dagger \bar{\mathbf{A}} \bar{\mathbf{w}} + \boldsymbol{\xi}$, in accordance with the Gaussian mechanism

Given the guarantee on $\bar{\mathbf{A}}$ from Theorem 2, we show that this result extends to \mathbf{w} by the well-known fact: for $\epsilon, \delta > 0$, post-processing indistinguishable quantities (Definition 3) preserves (ϵ, δ) -indistinguishability (Dwork et al., 2014). The full proof of utility preservation can be found in Appendix C, which essentially boils down to a Lipschitz condition. However, there are some downsides to this algorithm. First, it requires retraining the head \mathbf{w} for each unlearning request, but we want to perform unlearning without access to \mathbb{D}_{clf} . Second, repeatedly noising the base model via the Gaussian mechanism will erode its utility. We address these issues in the realistic setting.

5.2 REALISTIC SETTING

There is little reason to release \mathbf{A} and \mathbf{w} separately after fine-tuning the model, because it is unclear why one would want to use \mathbf{A} without \mathbf{w} or vice versa. One can obtain \mathbf{A} directly after pre-training instead of relying on a fine-tuned model, and there is little use for \mathbf{w} alone, because it is highly sensitive to the specific topics extracted by \mathbf{A} and their ordering. As such, we argue for releasing the fine-tuned model as a single matrix³ $\mathbf{B} = \mathbf{A}\mathbf{w}$, where $\mathbf{B} \in \mathbb{R}^{n \times 1}$.

Theorem 4 (Utility-Preserving Unlearning on the Downstream Task). *Suppose that the downstream task \mathcal{T} only depends on a subset of topics $\mathbb{T}_{\text{clf}} \subseteq [r]$; that is, $\mathbf{w}^* = \arg \min_{\mathbf{v} \in \mathcal{W}_{base}} \ell_{\mathcal{T}}(\mathbf{v}; \mathbf{A}^*)$ has non-zero entries only in the index set \mathbb{T}_{clf} . Denote $q := \min_{k \in \mathbb{T}_{\text{clf}}} \Pr_{\mathcal{D}}[z = k]$, and let \mathcal{A}_{head} be the head tuning algorithm (Definition 2) and \mathcal{U}_{head} be Algorithm 2. Then, $(\mathcal{A}_{head}, \mathcal{U}_{head})$ performs utility-preserving unlearning with deletion capacity*

$$T_{\epsilon, \delta}^{\mathcal{A}_{head}, \mathcal{U}_{head}}(m) \geq c' \cdot \frac{mq}{r\sqrt{nr}}$$

where c' is a constant dependent on $\epsilon, \delta, \mathcal{D}$, and \mathcal{T} .

The full proof is in Appendix C, including the worst case of $\mathbb{T}_{\text{clf}} = [r]$. When the task relies heavily on every single topic (i.e., $q = 1/ar$), the above guarantee is equivalent to the one in the pre-training phase. However, in most realistic settings, the downstream task will only depend on a subset of the latent topics in the corpus. In this case, $q > 1/ar$, and we can unlearn more points without degrading the utility of the model. Intuitively this makes sense too; the more reliance \mathcal{T} has on a rare topic, the less adversarial deletion it can tolerate.

Proof sketch. We again assume that we are deleting $m_U \leq \frac{0.001m\epsilon_0(\gamma p)^3}{a^2 r^2}$ points. For any modification made to \mathbf{A} , there is an equivalent modification that can be made to \mathbf{w} instead such that $\mathbf{B} = \mathbf{A}\mathbf{w}$ is preserved, so we do not need to update \mathbf{A} . We look for $\mathbf{v} \in \mathcal{W}_{head}$ such that $\mathbf{A}^S \mathbf{v} = \mathbf{A}^F \mathbf{w}^F$, where \mathbf{w}^F is the head learned on \mathbf{A}^F . It can be shown that $\bar{\mathbf{A}}^S$ has a unique pseudoinverse since it is full rank; naturally, we set $\mathbf{v} = \bar{\mathbf{A}}^{S\dagger} \mathbf{A}^F \mathbf{w}^F$, thereby ensuring privacy even if one recovers a part of \mathbf{A} from $\mathbf{B} = \mathbf{A}\mathbf{w}$. We furthermore define $\bar{\mathbf{v}}$ that is fit to the unlearned model before the Gaussian mechanism, $\bar{\mathbf{v}} = \bar{\mathbf{A}}^{S\dagger} \bar{\mathbf{A}} \bar{\mathbf{w}}$. We now need to show \mathbf{v} and $\bar{\mathbf{v}}$ satisfy both the indistinguishability and utility preservation conditions in Definition 4.

Indistinguishability. Let $\bar{\mathbf{w}}^* = \arg \min_{\mathbf{v} \in \mathcal{W}_{head}} \ell_{\mathcal{T}}(\mathbf{v}; \bar{\mathbf{A}})$ denote the result of head tuning $\bar{\mathbf{A}}$, and let $\bar{\mathbf{w}}$ be the result of taking a Newton step on \mathbf{w} (see Algorithm 2). Then by triangle inequality,

$$\|\bar{\mathbf{A}}\bar{\mathbf{w}} - \mathbf{A}^F \mathbf{w}^F\|_2 \leq \|\bar{\mathbf{A}}\bar{\mathbf{w}} - \bar{\mathbf{A}}\bar{\mathbf{w}}^*\|_2 + \|\bar{\mathbf{A}}\bar{\mathbf{w}}^* - \mathbf{A}^F \bar{\mathbf{w}}^*\|_2 + \|\mathbf{A}^F \bar{\mathbf{w}}^* - \mathbf{A}^F \mathbf{w}^F\|_2$$

³One can generalize this to the case where the downstream task is a C -way classification, in which case $\mathbf{B} \in \mathbb{R}^{n \times C}$.

Informally, the first term is controlled by the error in the Newton step approximation, and the third term is bounded by the error to the retrained \mathbf{w}^F . The remaining term can be rewritten as $\|(\bar{\mathbf{A}} - \mathbf{A}^F)(\bar{\mathbf{w}}^* - \mathbf{w}^*) + (\bar{\mathbf{A}} - \mathbf{A}^F)\mathbf{w}^*\|$, where the first term can be bounded using the same technique used to prove Lemmas 4 and 5. The second term can be bounded by noting that \mathbf{w}^* is sparse, which yields the below lemma that plays a crucial role in establishing the improved deletion capacity.

Lemma 6 (Modification of Lemma 4 for downstream task). *For $\epsilon, \delta > 0$,*

$$\|\bar{\mathbf{A}} - \mathbf{A}^F\|_\infty \leq \frac{1}{q} \cdot \|\bar{\mathbf{C}} - \mathbf{C}^F\|_\infty = c \cdot \frac{1}{q} \cdot \frac{arm_U}{m\epsilon_0\gamma p}$$

As in the pre-training case, we can now set the noise scale in the Gaussian mechanism and complete the proof. In the worst case, when the downstream task depends on *every* topic, then $q = 1/ar$, and we recover Lemma 4; however, this is unlikely to happen in practice.

Utility Preservation. We compare the value of \mathbf{v} after the Gaussian mechanism $\tilde{\mathbf{v}} = \bar{\mathbf{v}} + \nu_{\bar{\mathbf{v}}}$ to what it would be for the ground-truth model $\mathbf{v}^* = (\mathbf{A}^S)^\dagger \mathbf{A}^* \mathbf{w}^*$. We again rely on the sparsity of \mathbf{w}^* and bound $\mathbb{E}[\|\tilde{\mathbf{v}} - \mathbf{v}^*\|_\infty]$ in Lemma 31.

6 RELATED WORKS

Provable unlearning. One ideally wants the unlearned model to behave identically to one that was retrained from scratch with the forget set removed from the training data (Cao & Yang, 2015; Bourtole et al., 2021; Gupta et al., 2021). This is difficult to achieve in many settings, so there are several notions of approximate unlearning (Ginart et al., 2019; Guo et al., 2020; Neel et al., 2021) reminiscent of differential privacy (Dwork et al., 2014). Most relevant to our work is the notion of (ϵ, δ) -unlearning introduced in Sekhari et al. (2021), which we adapt to construct Definition 4. Our work focuses on deriving unlearning guarantees in the pre-training and fine-tuning pipeline. Golatkar et al. (2020) is closest to our work. They show considerably weaker guarantees on unlearning information with respect to probes fit to the weights. In contrast, our work is focused on realistic topic classification tasks and demonstrates strong guarantees (Definition 4). Recent works have extended notions of certified unlearning to nonconvex settings. Zhang et al. (2024a); Mu & Klabjan (2024); Chien et al. (2024) provide unlearning algorithms without deletion capacity guarantees. Qiao et al. (2024) also proposes an unlearning method for non-convex settings but analyzes its deletion capacity in a convex setting. Our work extends beyond the convex setting to provide provable unlearning methods and corresponding deletion capacity analysis for non-convex models.

Theoretical analysis of pre-training and fine-tuning. Our downstream task definition (Section 2.2) is inspired by works on transfer learning in language models (Saunshi et al., 2021; Wei et al., 2021; Wu et al., 2023; Kumar et al., 2022), contrastive learning (Lee et al., 2021; HaoChen & Ma, 2023), and meta-learning (Chua et al., 2021; Collins et al., 2022; Yüksel et al., 2024).

7 CONCLUSION

This work uses topic models to develop the first provable guarantees on unlearning in the modern-day pre-training and fine-tuning paradigm. We propose two unlearning algorithms that can effectively and efficiently unlearn from both the pre-trained model (Algorithm 1 and Theorem 2) and the fine-tuned model (Algorithm 2 and Theorem 4). Notably, we find that it is easier, in terms of the deletion capacity (Definition 4), to unlearn pre-training data from the fine-tuned model, and we can do so without modifying the pre-trained base model. Our findings suggest that task-specific unlearning is easier than full model unlearning, providing a promising path forward to design efficient algorithms for large-scale models.

The most notable limitation of our work is that our usage of topic models, which permit a tractable analysis but cannot capture interesting features of modern-day language models (e.g., their autoregressive nature). Moreover, with the growing popularity of foundation models, there is scholarly discussion around meaningful definitions of unlearning and how they can be measured (Thudi et al., 2022; Lee et al., 2024). Our work focuses on traditional notions of unlearning centered on differential privacy (see Definition 4), but we hope to extend these definitions to capture additional features of generative models that are salient to their real-world uses.

REFERENCES

- Sanjeev Arora, Rong Ge, Yoni Halpern, David Mimno, Ankur Moitra, David Sontag, Yichen Wu, and Michael Zhu. A practical algorithm for topic modeling with provable guarantees, 2012a. URL <https://arxiv.org/abs/1212.4777>.
- Sanjeev Arora, Rong Ge, and Ankur Moitra. Learning topic models - going beyond svd, 2012b.
- Abeba Birhane and Vinay Uday Prabhu. Large image datasets: A pyrrhic win for computer vision? In *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1536–1546. IEEE, 2021.
- Abeba Birhane, Vinay Uday Prabhu, and Emmanuel Kahembwe. Multimodal datasets: misogyny, pornography, and malignant stereotypes. *arXiv preprint arXiv:2110.01963*, 2021.
- David M. Blei and John D. Lafferty. Dynamic topic models. In *Proceedings of the 23rd International Conference on Machine Learning, ICML '06*, pp. 113–120, New York, NY, USA, 2006. Association for Computing Machinery. ISBN 1595933832. doi: 10.1145/1143844.1143859. URL <https://doi.org/10.1145/1143844.1143859>.
- David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.
- 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, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohith Kudipudi, 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 Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avaniika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, Julian Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Rob Reich, Hongyu Ren, Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishnan 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 Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. On the opportunities and risks of foundation models, 2022. URL <https://arxiv.org/abs/2108.07258>.
- Lucas Bourtole, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning. In *2021 IEEE Symposium on Security and Privacy (SP)*, pp. 141–159. IEEE, 2021.
- Jordan Boyd-Graber, Yuening Hu, David Mimno, et al. Applications of topic models. *Foundations and Trends® in Information Retrieval*, 11(2-3):143–296, 2017.
- 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 (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 1877–1901. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfbcb4967418bfb8ac142f64a-Paper.pdf.

- Yinzhi Cao and Junfeng Yang. Towards making systems forget with machine unlearning. In *2015 IEEE symposium on security and privacy*, pp. 463–480. IEEE, 2015.
- Nicholas Carlini, Florian Tramer, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Ulfar Erlingsson, et al. Extracting training data from large language models. In *30th USENIX Security Symposium (USENIX Security 21)*, pp. 2633–2650, 2021.
- Nicolas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwal, Florian Tramer, Borja Balle, Daphne Ippolito, and Eric Wallace. Extracting training data from diffusion models. In *32nd USENIX Security Symposium (USENIX Security 23)*, pp. 5253–5270, 2023.
- Eli Chien, Haoyu Wang, Ziang Chen, and Pan Li. Langevin unlearning: A new perspective of noisy gradient descent for machine unlearning, 2024. URL <https://arxiv.org/abs/2401.10371>.
- Rochelle Choenni, Ekaterina Shutova, and Robert van Rooij. Stepmothers are mean and academics are pretentious: What do pretrained language models learn about you? In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 1477–1491, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.111. URL <https://aclanthology.org/2021.emnlp-main.111>.
- Kurtland Chua, Qi Lei, and Jason D Lee. How fine-tuning allows for effective meta-learning. *Advances in Neural Information Processing Systems*, 34:8871–8884, 2021.
- Rob Churchill and Lisa Singh. The evolution of topic modeling. *ACM Comput. Surv.*, 2022.
- Liam Collins, Aryan Mokhtari, Sewoong Oh, and Sanjay Shakkottai. Maml and anil provably learn representations. In *International Conference on Machine Learning*, pp. 4238–4310. PMLR, 2022.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In Jill Burstein, Christy Doran, and Thamar Solorio (eds.), *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)*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>.
- Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- Ronen Eldan and Mark Russinovich. Who’s harry potter? approximate unlearning in llms, 2023. URL <https://arxiv.org/abs/2310.02238>.
- DOE 1 v. GitHub, Inc.* 4:22-cv-06823, N.D. Cal. 2022.
- Tremblay v. OpenAI, Inc.*, 23-cv-03416-AMO, (N.D. Cal.), 2023.
- European Parliament and Council of the European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council. URL <https://data.europa.eu/eli/reg/2016/679/oj>.
- Jack Foster, Stefan Schoepf, and Alexandra Brintrup. Fast machine unlearning without retraining through selective synaptic dampening. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 12043–12051, 2024.
- Rohit Gandikota, Joanna Materzynska, Jaden Fiotto-Kaufman, and David Bau. Erasing concepts from diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2426–2436, 2023.

- 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. URL <https://arxiv.org/abs/2101.00027>.
- Antonio Ginart, Melody Guan, Gregory Valiant, and James Y Zou. Making ai forget you: Data deletion in machine learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/cb79f8fa58b91d3af6c9c991f63962d3-Paper.pdf.
- Aditya Golatkar, Alessandro Achille, and Stefano Soatto. Eternal sunshine of the spotless net: Selective forgetting in deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9304–9312, 2020.
- Chuan Guo, Tom Goldstein, Awni Hannun, and Laurens Van Der Maaten. Certified data removal from machine learning models. In *International Conference on Machine Learning*, pp. 3832–3842. PMLR, 2020.
- Varun Gupta, Christopher Jung, Seth Neel, Aaron Roth, Saeed Sharifi-Malvajerdi, and Chris Waites. Adaptive machine unlearning. *Advances in Neural Information Processing Systems*, 34:16319–16330, 2021.
- Jeff Z. HaoChen and Tengyu Ma. A theoretical study of inductive biases in contrastive learning. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=AuEgNlEAmEd>.
- Jamie Hayes, Iliia Shumailov, Eleni Triantafyllou, Amr Khalifa, and Nicolas Papernot. Inexact unlearning needs more careful evaluations to avoid a false sense of privacy, 2024. URL <https://arxiv.org/abs/2403.01218>.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16000–16009, 2022.
- Luxi He, Yangsibo Huang, Weijia Shi, Tinghao Xie, Haotian Liu, Yue Wang, Luke Zettlemoyer, Chiyuan Zhang, Danqi Chen, and Peter Henderson. Fantastic copyrighted beasts and how (not) to generate them. *arXiv preprint arXiv:2406.14526*, 2024.
- Peter Henderson, Xuechen Li, Dan Jurafsky, Tatsunori Hashimoto, Mark A Lemley, and Percy Liang. Foundation models and fair use. *Journal of Machine Learning Research*, 24(400):1–79, 2023.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.
- Thomas Hofmann et al. Probabilistic latent semantic analysis. In *UAI*, volume 99, pp. 289–296, 1999.
- Zachary Izzo, Mary Anne Smart, Kamalika Chaudhuri, and James Zou. Approximate data deletion from machine learning models, 2021.
- Joel Jang, Dongkeun Yoon, Sohee Yang, Sungmin Cha, Moontae Lee, Lajanugen Logeswaran, and Minjoon Seo. Knowledge unlearning for mitigating privacy risks in language models, 2023. URL <https://openreview.net/forum?id=zAxuIJLb38>.
- Ananya Kumar, Aditi Raghunathan, Robbie Matthew Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pretrained features and underperform out-of-distribution. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=UYneFzXSJWh>.

- Meghdad Kurmanji, Peter Triantafillou, Jamie Hayes, and Eleni Triantafillou. Towards unbounded machine unlearning. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=OveBaTtUAT>.
- Jason D Lee, Qi Lei, Nikunj Saunshi, and Jiacheng Zhuo. Predicting what you already know helps: Provable self-supervised learning. *Advances in Neural Information Processing Systems*, 34:309–323, 2021.
- Katherine Lee, A. Cooper, Christopher Choquette-Choo, Ken Liu, Matthew Jagielski, Niloofar Mireshghallah, Lama Ahmed, James Grimmelmann, David Bau, Christopher De Sa, Fernando Delgado, Vitaly Shmatikov, Katja Filippova, Seth Neel, Miranda Bogen, Amy Cyphert, Mark Lemley, and Nicolas Papernot. Extended abstract: Machine unlearning doesn’t do what you think, 04 2024.
- Wei Li and Andrew McCallum. Pachinko allocation: Dag-structured mixture models of topic correlations. In *Proceedings of the 23rd international conference on Machine learning*, pp. 577–584, 2006.
- Jiaqi Liu, Jian Lou, Zhan Qin, and Kui Ren. Certified minimax unlearning with generalization rates and deletion capacity. *Advances in Neural Information Processing Systems*, 36, 2024.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing, 2021. URL <https://arxiv.org/abs/2107.13586>.
- Shayne Longpre, Robert Mahari, Anthony Chen, Naana Obeng-Marnu, Damien Sileo, William Brannon, Niklas Muennighoff, Nathan Khazam, Jad Kabbara, Kartik Perisetla, et al. A large-scale audit of dataset licensing and attribution in ai. *Nature Machine Intelligence*, 6(8):975–987, 2024.
- Ananth Mahadevan and Michael Mathioudakis. Cost-effective retraining of machine learning models. *arXiv preprint arXiv:2310.04216*, 2023.
- Pratyush Maini, Zhili Feng, Avi Schwarzschild, Zachary C. Lipton, and J. Zico Kolter. Tofu: A task of fictitious unlearning for llms, 2024. URL <https://arxiv.org/abs/2401.06121>.
- Sadhika Malladi, Tianyu Gao, Eshaan Nichani, Alex Damian, Jason D. Lee, Danqi Chen, and Sanjeev Arora. Fine-tuning language models with just forward passes. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023a. URL <https://openreview.net/forum?id=Vota6rFhBQ>.
- Sadhika Malladi, Alexander Wettig, Dingli Yu, Danqi Chen, and Sanjeev Arora. A kernel-based view of language model fine-tuning. In *International Conference on Machine Learning*, pp. 23610–23641. PMLR, 2023b.
- Nick McKenna, Tianyi Li, Liang Cheng, Mohammad Javad Hosseini, Mark Johnson, and Mark Steedman. Sources of hallucination by large language models on inference tasks. *arXiv preprint arXiv:2305.14552*, 2023.
- Siqiao Mu and Diego Klabjan. Rewind-to-delete: Certified machine unlearning for nonconvex functions, 2024. URL <https://arxiv.org/abs/2409.09778>.
- Tarek Naous, Michael Ryan, Alan Ritter, and Wei Xu. Having beer after prayer? measuring cultural bias in large language models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 16366–16393, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.862. URL <https://aclanthology.org/2024.acl-long.862>.
- Seth Neel, Aaron Roth, and Saeed Sharifi-Malvajerdi. Descent-to-delete: Gradient-based methods for machine unlearning. In *Algorithmic Learning Theory*, pp. 931–962. PMLR, 2021.
- Quoc Phong Nguyen, Bryan Kian Hsiang Low, and Patrick Jaillet. Variational bayesian unlearning. *Advances in Neural Information Processing Systems*, 33:16025–16036, 2020.

- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khilodov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL <https://openreview.net/forum?id=a68SUt6zFt>.
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtessam Almazrouei, and Julien Launay. The refinedweb dataset for falcon llm: outperforming curated corpora with web data, and web data only. *arXiv preprint arXiv:2306.01116*, 2023.
- Xinbao Qiao, Meng Zhang, Ming Tang, and Ermin Wei. Efficient and generalizable certified unlearning: A hessian-free recollection approach, 2024. URL <https://arxiv.org/abs/2404.01712>.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Ben Recht, Christopher Re, Joel Tropp, and Victor Bittorf. Factoring nonnegative matrices with linear programs. *Advances in neural information processing systems*, 25, 2012.
- Ryan Rifkin and Aldebaro Klautau. In defense of one-vs-all classification. *The Journal of Machine Learning Research*, 5:101–141, 2004.
- Nikunj Saunshi, Sadhika Malladi, and Sanjeev Arora. A mathematical exploration of why language models help solve downstream tasks. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=vVjIW3sEcls>.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022.
- Ayush Sekhari, Jayadev Acharya, Gautam Kamath, and Ananda Theertha Suresh. Remember what you want to forget: Algorithms for machine unlearning, 2021. URL <https://arxiv.org/abs/2103.03279>.
- Weijia Shi, Jaechan Lee, Yangsibo Huang, Sadhika Malladi, Jieyu Zhao, Ari Holtzman, Daogao Liu, Luke Zettlemoyer, Noah A. Smith, and Chiyuan Zhang. Muse: Machine unlearning six-way evaluation for language models, 2024. URL <https://arxiv.org/abs/2407.06460>.
- Luca Soldaini, Rodney Kinney, Akshita Bhagia, Dustin Schwenk, David Atkinson, Russell Authur, Ben Bogin, Khyathi Chandu, Jennifer Dumas, Yanai Elazar, Valentin Hofmann, Ananya Jha, Sachin Kumar, Li Lucy, Xinxin Lyu, Nathan Lambert, Ian Magnusson, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew Peters, Abhilasha Ravichander, Kyle Richardson, Zejiang Shen, Emma Strubell, Nishant Subramani, Oyvind Tafjord, Evan Walsh, Luke Zettlemoyer, Noah Smith, Hannaneh Hajishirzi, Iz Beltagy, Dirk Groeneveld, Jesse Dodge, and Kyle Lo. Dolma: an open corpus of three trillion tokens for language model pretraining research. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15725–15788, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.840. URL <https://aclanthology.org/2024.acl-long.840>.
- Gowthami Somepalli, Vasu Singla, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Diffusion art or digital forgery? investigating data replication in diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6048–6058, 2023.

- Anvith Thudi, Hengrui Jia, Iliia Shumailov, and Nicolas Papernot. On the necessity of auditable algorithmic definitions for machine unlearning. In *31st USENIX Security Symposium (USENIX Security 22)*, pp. 4007–4022, Boston, MA, August 2022. USENIX Association. ISBN 978-1-939133-31-1. URL <https://www.usenix.org/conference/usenixsecurity22/presentation/thudi>.
- Enayat Ullah, Tung Mai, Anup Rao, Ryan A Rossi, and Raman Arora. Machine unlearning via algorithmic stability. In *Conference on Learning Theory*, pp. 4126–4142. PMLR, 2021.
- Colin Wei, Sang Michael Xie, and Tengyu Ma. Why do pretrained language models help in downstream tasks? an analysis of head and prompt tuning. *Advances in Neural Information Processing Systems*, 34:16158–16170, 2021.
- Chenwei Wu, Holden Lee, and Rong Ge. Connecting pre-trained language model and downstream task via properties of representation. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=YLOJ4aKAKa>.
- Oğuz Kaan Yüksel, Etienne Boursier, and Nicolas Flammarion. First-order ANIL provably learns representations despite overparametrisation. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=if2vRbS8Ew>.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 11975–11986, 2023.
- Binchi Zhang, Yushun Dong, Tianhao Wang, and Jundong Li. Towards certified unlearning for deep neural networks, 2024a. URL <https://arxiv.org/abs/2408.00920>.
- Ruiqi Zhang, Licong Lin, Yu Bai, and Song Mei. Negative preference optimization: From catastrophic collapse to effective unlearning. In *First Conference on Language Modeling*, 2024b. URL <https://openreview.net/forum?id=MXLBXjQkmb>.

A PRECISE DESCRIPTION OF $\mathcal{A}_{\text{BASE}}$

A.1 COMPLETE DESCRIPTION

Algorithm 3 High level learning algorithm (\mathcal{A})

Input: document corpus $S = \{d_i\}_{i=1}^m$, anchor word tolerance ϵ_0
Output: matrices \mathbf{A}, \mathbf{R}
 \mathbf{Q} = word co-occurrences
 $\bar{\mathbf{Q}}$ = row-normalized \mathbf{Q}
 $P = \text{RecoverAnchors}(\{\bar{\mathbf{Q}}_1, \dots, \bar{\mathbf{Q}}_n\})$
 $\mathbf{A}, \mathbf{R} = \text{RecoverTopics}(\mathbf{Q}, S)$
return \mathbf{A}, \mathbf{R}

Algorithm 4 RecoverAnchors, same as Arora et al. (2012a)

Input: Row-normalized co-occurrence matrix $\bar{\mathbf{Q}}$ and ϵ_0 tolerance parameter
Output: r points of this perturbed simplex close to the vertices of the actual simplex
 Project the rows to a randomly chosen $4 \log n / \epsilon_0^2$ dimensional subspace
 $S \leftarrow \{\bar{\mathbf{Q}}_i\}$ where $\bar{\mathbf{Q}}_i$ is the furthest point from the origin
for i in $1, \dots, r-1$ **do**
 Let $\bar{\mathbf{Q}}_j$ be the row of $\bar{\mathbf{Q}}$ with largest distance to $\text{span}(S)$
 $S \leftarrow S \cup \{\bar{\mathbf{Q}}_j\}$
end for $S = \{\bar{\mathbf{Q}}_{s_1}, \dots, \bar{\mathbf{Q}}_{s_r}\}$
for i in $1, \dots, r$ **do**
 Let $\bar{\mathbf{Q}}_j$ be the point that has largest distance to $\text{span}(S \setminus \{\bar{\mathbf{Q}}_{s_i}\})$
 Remove $\bar{\mathbf{Q}}_{s_i}$ from S and insert $\bar{\mathbf{Q}}_j$ into S
end for
return S

Algorithm 5 Recover Topics, from Arora et al. (2012a)

Input: Co-occurrence matrix \mathbf{Q} , anchor words $P = \{s_1, \dots, s_k\}$, tolerance parameter ϵ_0
Output: Matrices \mathbf{A}, \mathbf{R}
 $\bar{\mathbf{Q}}$ = row normalized \mathbf{Q}
 Store the normalization constants $\mathbf{p} = \mathbf{Q}\mathbf{1}$
for i in $1, \dots, n$ **do**
 Solve $C_i = \arg \min_{v \in \Delta_r} \|\bar{\mathbf{Q}}_i - v^\top \bar{\mathbf{Q}}_P\|^2$
 up to ϵ_0 accuracy
end for
 $\mathbf{A}' = \text{diag}(\mathbf{p})\mathbf{C}$
 \mathbf{A} = column-sum-one normalized \mathbf{A}'
 $\mathbf{R} = \mathbf{A}^\dagger \mathbf{Q} \mathbf{A}^{\dagger\top}$ where \mathbf{A}^\dagger is the pseudoinverse of \mathbf{A}
return \mathbf{A}, \mathbf{R}

More formally, the co-occurrence matrix is constructed as follows. For each document, let $H_d \in \mathbb{R}^n$ be the frequency vector of each word in the document; the sum of its entries should be L . Then, for a document d , consider the matrix

$$\mathbf{G}_d := \tilde{\mathbf{H}}_d \tilde{\mathbf{H}}_d^\top - \hat{\mathbf{H}}_d$$

where

$$\tilde{\mathbf{H}}_d := \frac{\mathbf{H}_d}{\sqrt{L(L-1)}}$$

$$\hat{\mathbf{H}}_d := \frac{\text{diag}(\mathbf{H}_d)}{L(L-1)}$$

In particular, the denominator term $L(L - 1)$ is precisely the number of co-occurrences in each document, by simple combinatorics, and it can be seen that the sum of the entries of \mathbf{G}_d is always 1. Our co-occurrence matrix \mathbf{Q} is defined to be

$$\mathbf{Q} := \frac{1}{m} \sum_{d=1}^m \mathbf{G}_d$$

so that \mathbf{Q} also has entries that sum to 1. By linearity of expectation, we have

$$\mathbb{E}[\mathbf{Q}] = \mathbb{E}[\mathbf{G}_d] = \mathbf{A}^* \mathbb{E}[\mathbf{X}_d \mathbf{X}_d^\top] \mathbf{A}^{*\top}$$

which implies that as the number of documents increases, \mathbf{Q} concentrates around $\mathbf{A} \mathbb{E}[\mathbf{X} \mathbf{X}^\top] \mathbf{A}^\top = \mathbb{E}[\mathbf{M} \mathbf{M}^\top]$. Therefore, we should expect $\mathbf{A}^\dagger \mathbf{Q} \mathbf{A}^{\dagger\top}$ to concentrate around $\mathbb{E}[\mathbf{X} \mathbf{X}^\top] = \mathbf{R}^*$.

A.2 SKETCH: POPULATION ANALYSIS

To understand this algorithm, consider the setting where we have infinitely many documents. Specifically, consider two words w_1, w_2 in a document and their respective topics z_1, z_2 . Then, this population co-occurrence matrix \mathbf{Q} will have elements $Q_{i,j} = \Pr[w_1 = i, w_2 = j]$, and the row-normalized co-occurrence matrix $\bar{\mathbf{Q}}$ will have entries $\bar{Q}_{i,j} = \Pr[w_2 = j | w_1 = i]$. Moreover, we have that $\mathbf{A}_{i,k} = \Pr[w_1 = i | z_1 = k] = \Pr[w_2 = i | z_2 = k]$.

Consider the set of anchor words $P = \{s_1, \dots, s_r\} \subseteq [n]$, where s_k is the anchor word for topic k . Then, observe that for an anchor word row s_k of $\bar{\mathbf{Q}}$, it holds that

$$\begin{aligned} \bar{Q}_{s_k,j} &= \Pr[w_2 = j | w_1 = s_k] = \sum_{k'} \Pr[z_1 = k' | w_1 = s_k] \Pr[w_2 = j | w_1 = s_k, z_1 = k'] \\ &= \Pr[w_2 = j | w_1 = s_k, z_1 = k] \\ &= \Pr[w_2 = j | z_1 = k] \end{aligned}$$

where the second line follows from only $\Pr[z_1 = k | w_1 = s_k] = 1$ in the summation, and the last line follows from w_2, w_1 are conditionally independent given z_1 . Furthermore, for non-anchor word rows i of $\bar{\mathbf{Q}}$, it holds that

$$\bar{Q}_{i,j} = \sum_k \Pr[z_1 = k | w_1 = i] \Pr[w_2 = j | z_1 = k]$$

where again we use that w_2, w_1 are conditionally independent z_1 . For a word i , let $\mathbf{C}_i \in \mathbb{R}^r$ be the vector such that $\mathbf{C}_{i,k} := \Pr[z_1 = k | w_1 = i]$. Then, it holds that $\bar{\mathbf{Q}}_i = \mathbf{c}_i^\top \bar{\mathbf{Q}}_S$, where $\bar{\mathbf{Q}}_S$ is the submatrix of $\bar{\mathbf{Q}}$ constrained to the anchor word rows. In other words, for every word i , $\bar{\mathbf{Q}}_i$ is a convex combination of rows of $\bar{\mathbf{Q}}_S$.

In the algorithm, one can see that $\mathbf{A}'_{i,k} = \mathbf{C}_{i,k} \mathbf{p}_i$. Normalizing this along each column, we obtain

$$\mathbf{A}_{i,k} = \frac{\mathbf{C}_{i,k} \mathbf{p}_i}{\sum_{i'} \mathbf{C}_{i',k} \mathbf{p}_{i'}} = \frac{\Pr[z_1 = k | w_1 = i] \Pr[w_1 = i]}{\sum_{i'} \Pr[z_1 = k | w_1 = i'] \Pr[w_1 = i']} = \Pr[w_1 = i | z_1 = k]$$

Hence, in the infinite document limit, this algorithm recovers the ground truth $\mathbf{A}^*, \mathbf{R}^*$.

B FROM PROPERTIES OF THE LEARNING ALGORITHM TO THE PROOF OF THEOREM 2

We first give the formal statement of Theorem 2.

Theorem 5 (Formal statement of Theorem 2). *Let \mathcal{A}_{base} be the learning algorithm described in the prior sections and \mathcal{U}_{base} be the unlearning algorithm in Algorithm 1. Then, $(\mathcal{A}_{base}, \mathcal{U}_{base})$ performs utility-preserving unlearning with deletion capacity*

$$T_{\epsilon, \delta}^{\mathcal{A}_{base}, \mathcal{U}_{base}}(m) \geq c \cdot \min \left\{ \frac{m\epsilon}{r^2 \sqrt{rn \log 1/\delta}}, \frac{0.001m}{r^2} \right\}$$

where m is the number of training documents, r is the number of topics, and c is a constant dependent on \mathcal{D} . The loss function h used in the utility-preserving definition is the maximum entrywise error from the ground truth topic model \mathbf{A}^* .

B.1 PRELIMINARIES

When the norm is not specified, we assume that it is the Euclidean norm $\|\cdot\|_2$. We now start off with a technical assumption on the precision of the learning algorithm.

Assumption 2. $\epsilon_0 \leq O(1/\sqrt{nr})$.

Assumption 3. *Every word appears with probability $\epsilon_0/4ar$ without loss of generality; see discussion in Arora et al. (2012b). Essentially, less probable words can be combined in a sense to form a single category of "rare" words.*

We recall the definitions from Arora et al. (2012a).

Definition 6 (β -robust simplex). A simplex P is β -robust if for every vertex v of P , the ℓ_2 distance between v and the convex hull of the rest of the vertices is at least β .

Definition 7. Let $\{a_i\}_{i=1}^n$ be a set of points whose convex hull is a simplex with vertices $\{v_i\}_{i=1}^r$. We say a set of r points is ϵ -close to the vertex set $\{v_i\}_{i=1}^r$ if each of the r points is ϵ -close in ℓ_2 distance to a different vertex in this vertex set.

The following result will be used throughout our proof.

Proposition 1 (Arora et al. (2012b)). \bar{Q}_P^* in population is γp -robust.

We now list the high probability events we condition on throughout our proof. These follow from previous results in Arora et al. (2012a); they concern the properties of the output of the learning algorithm.

Proposition 2. *With high probability, in our regime of m , the following hold:*

- *The correct anchor words are selected.*
- *Each word appears at least $O(\frac{m\epsilon_0}{4ar})$ times.*
- *The error in the empirical matrix \hat{Q} is entrywise at most $\tilde{O}(1/\sqrt{m})$ from the population Q^* .*

We also utilize the following two key lemmas from Arora et al. (2012a) that we touched upon in the main paper.

Lemma 7 (Approximation Guarantee on Anchor Words). *Suppose each row of \bar{Q} is at most δ distance away from the ground truth γp -robust simplex Q^* in ℓ_2 norm. If $20r\delta/(\gamma p)^2 < \gamma p$, then the set of anchor words found by the algorithm is $O(\delta/\gamma p)$ -close to the ground truth anchor words.*

Lemma 8. *When $20r\delta/(\gamma p)^2 < \gamma p$, it holds for every word i that C_i has entrywise error $O(\delta/(\gamma p)^2)$ from C_i^* .*

B.2 PROOF OF THEOREM 2

The following are lemmas bounding the relation between $\bar{Q}_i^S, \bar{Q}_i^F, \bar{Q}_i^*$.

Lemma 9. *After training, the error of each row of \bar{Q}^S is at most $\delta_2 := O(\sqrt{\frac{4ar}{m\epsilon_0}})$. That is, $\|\bar{Q}_i^S - \bar{Q}_i^*\| \leq \delta_2$ for all words i .*

Importantly, note that

$$20r\delta_2/(\gamma p)^2 < \gamma p$$

This implies that the anchor words of \bar{Q}_i^S are $O(\delta_2/(\gamma p))$ close to the anchor words of \bar{Q}_i^ .*

Consequently, it holds that

$$\|C^S - C^*\|_\infty \leq O(\delta_2/(\gamma p)^2)$$

Proof. The first part follows directly from the fact that if the number of documents $m = \tilde{\Omega}(1/\epsilon_Q^2)$, then $\|\bar{Q}_i^S - \bar{Q}_i^*\| \leq \delta_2$ for each row i . To show that

$$20r\delta_2/(\gamma p)^2 < \gamma p$$

we note that by the sample complexity guarantee,

$$m\epsilon_0 \geq \tilde{O}\left(\frac{ar^3}{(\gamma p)^6}\right)$$

which implies that

$$\delta_2 \leq \tilde{O}\left(\frac{(\gamma p)^3}{r}\right)$$

as desired. \square

Lemma 10. *When we delete $m_U \leq \frac{0.001m\epsilon_0(\gamma p)^3}{a^2r^2}$, it holds that*

$$\|\bar{Q}_i^F - \bar{Q}_i^S\| \leq \frac{m_U}{m\epsilon_0/4ar} = \frac{4arm_U}{m\epsilon_0}$$

In particular, this is smaller than

$$\frac{0.001m\epsilon_0(\gamma p)^3}{a^2r^2} \cdot \frac{1}{m\epsilon_0/4ar} = \frac{0.004(\gamma p)^3}{ar}$$

Proof. For a word i , consider the change in \bar{Q}_i after deletion requests. Let F be the initial sum of the the i th row of Q . Each coordinate $j \in [n]$ will change as follows:

$$\delta_j = \frac{f_j - t_j}{F - m_U} - \frac{f_j}{F} = \frac{m_U f_j - F t_j}{F(F - m_U)}$$

where f_j is the initial number of cooccurrences of words i, j and t_j is the number of documents removed that have this cooccurrence. Moreover, F is the number of initial occurrences of word i , and T is the number of deletions of the word i . From the previous lemma, it holds that $F \geq m\epsilon_0/4ar$, and that $m_U \geq \sum_{j=1}^n t_j$. Hence, it follows that the squared Euclidean norm of the change is:

$$\sum_{j=1}^n \delta_j^2 = \frac{1}{F^2(F - T)^2} \sum_{j=1}^n (m_U f_j - F t_j)^2 \leq \frac{2F^2 m_U^2}{F^2(F - m_U)^2} \leq 2\left(\frac{m_U}{F - m_U}\right)^2$$

Hence, for the regime where $m_U \leq \frac{0.001m\epsilon_0(\gamma p)^3}{a^2r^2}$, we have

$$\|\bar{Q}_i^S - \bar{Q}_i^F\| \leq \sqrt{2} \frac{m_U}{F - m_U} \lesssim \frac{m_U}{F} \lesssim \frac{4arm_U}{m\epsilon_0}$$

Of particular notice is that when m_U is taken as large as possible, this is at most

$$\frac{0.001m\epsilon_0(\gamma p)^3/a^2r^2}{m\epsilon_0/4ar} = 0.004(\gamma p)^3/ar$$

\square

We now combine the above two with triangle inequality.

Lemma 11. *Hence, it holds that*

$$\|\bar{Q}_i^F - \bar{Q}_i^*\| \leq \frac{4arm_U}{m\epsilon_0} + \delta_2 = \frac{4arm_U}{m\epsilon_0} + O\left(\sqrt{\frac{4ar}{m\epsilon_0}}\right) =: \delta'_2$$

Importantly, note that

$$20r\delta'_2/(\gamma p)^2 < \gamma p$$

This implies that the anchor words of \bar{Q}_i^F are $O(\delta'_2/(\gamma p))$ close to the anchor words of \bar{Q}_i^ .*

Consequently, it holds that

$$\|C^F - C^*\|_\infty \leq O(\delta'_2/(\gamma p)^2)$$

Proof. The first part follows from triangle inequality, and the second part follows from lemma B.1 from Arora et al. (2012a). \square

We now bound what happens to $\|C^F - C^S\|_\infty$. First, we have that the perturbed simplex \bar{Q}_P^S is $\gamma p/2$ -robust.

Lemma 12. *The perturbed simplex \bar{Q}_P^S is $\gamma p/2$ -robust.*

Proof. This is because of Lemma A.1 in Arora et al. (2012a). Since $10\sqrt{r}\delta_2 < \gamma p$, the result of that lemma applies. \square

Hence, we will apply Lemma B.1 from Arora et al. (2012a) on C^S to say something about $\|C^F - C^S\|_\infty$.

Lemma 13. *Recall that when we delete $m_U \leq \frac{0.001m\epsilon_0(\gamma p)^3}{a^2r^2}$, it holds that*

$$\|\bar{Q}_i^F - \bar{Q}_i^S\| \leq \frac{m_U}{m\epsilon_0/4ar} = \frac{4arm_U}{m\epsilon_0}$$

Importantly, note that

$$20r \left(\frac{4arm_U}{m\epsilon_0} \right) / (\gamma p/2)^2 < \gamma p/2$$

This implies that the anchor words of \bar{Q}_i^F are $\frac{4arm_U/m\epsilon_0}{\gamma p/2}$ close to the anchor words of \bar{Q}_i^S . By lemma B.1 from Arora et al. (2012a), it holds that

$$\|C^F - C^S\|_\infty \leq O\left(\frac{4arm_U}{m\epsilon_0} / (\gamma p/2)^2\right)$$

Observe that this is smaller than $O((\gamma p)/ar)$.

We now deal with the Hessian step that we had took to prevent retraining the C_i 's. In particular, we will denote \bar{C} to be our estimated new C .

First, a lemma to say that our Hessian step is full rank and has a lower bound on its minimum singular value.

Lemma 14. *When we delete $m_U \leq \frac{0.001m\epsilon_0(\gamma p)^3}{a^2r^2}$ samples, it holds that the minimum eigenvalue of $\bar{Q}_P^F \bar{Q}_P^{F\top}$ is at least $\gamma p/2$.*

Proof. Follows from Lemma A.3 in Arora et al. (2012a). \square

Lemma 15. *When we delete $m_U \leq \frac{0.001m\epsilon_0(\gamma p)^3}{a^2r^2}$ samples, it holds for all i ,*

$$\|C_i^F - \bar{C}_i^F\| \leq \frac{4}{\gamma p} \left(\delta_2 + \frac{4arm_U}{m\epsilon_0} \right)$$

Proof. For the case of $d(\cdot, \cdot)$ being the squared loss, we will denote the following:

$$\begin{aligned} C_{i,\text{uncon}} &:= \arg \min_C \|\bar{Q}_P^{F\top} C - \bar{Q}_i^{F\top}\|^2 = (\bar{Q}_P^F \bar{Q}_P^{F\top})^{-1} \bar{Q}_P^F \bar{Q}_i^{F\top} \\ \bar{C}_i^F &:= \text{proj}_{\Delta_r}(C_{i,\text{uncon}}) \\ C_i^F &:= \arg \min_{C \in \Delta_r} \|\bar{Q}_P^{F\top} C - \bar{Q}_i^{F\top}\|^2 \end{aligned}$$

In particular, the Newton step plus projection outputs $C_{i,\text{proj}}$. First, observe that by one of the anchor word lemmas,

$$\min_C \|\bar{Q}_P^{F\top} C - \bar{Q}_i^{F\top}\| = \|\bar{Q}_P^{F\top} C_{i,\text{uncon}} - \bar{Q}_i^{F\top}\| \leq \|\bar{Q}_P^{F\top} C_i^F - \bar{Q}_i^{F\top}\| \leq \delta_2 + \frac{4arm_U}{m\epsilon_0}$$

The last inequality follows from the fact that $\bar{\mathbf{Q}}_P^F$ is a perturbed version of $\bar{\mathbf{Q}}_P^S$, and $\bar{\mathbf{Q}}_P^S$ is a perturbed version of $\bar{\mathbf{Q}}_P^*$. Hence, we will bound

$$\begin{aligned} \|\bar{\mathbf{C}}_i^F - \mathbf{C}_i^F\| &= \|\text{proj}_{\Delta_r}(\mathbf{C}_{i,\text{uncon}}) - \text{proj}_{\Delta_r}(\mathbf{C}_i^F)\| \\ &\leq \|\mathbf{C}_{i,\text{uncon}} - \mathbf{C}_i^F\| \\ &\leq \frac{1}{\sigma_{\min}} \|\bar{\mathbf{Q}}_P^{F\top} (\mathbf{C}_{i,\text{uncon}} - \mathbf{C}_i^F)\| \\ &\leq \frac{1}{\sigma_{\min}} (\|\bar{\mathbf{Q}}_i^{F\top} - \bar{\mathbf{Q}}_P^{F\top} \mathbf{C}_i^F\| + \|\bar{\mathbf{Q}}_P^{F\top} \mathbf{C}_{i,\text{uncon}} - \bar{\mathbf{Q}}_i^{F\top}\|) \\ &\leq \frac{2}{\sigma_{\min}} \left(\delta_2 + \frac{4arm_U}{m\epsilon_0} \right) \end{aligned}$$

where σ_{\min} is the smallest singular value of $\bar{\mathbf{Q}}_i^{F\top}$, which is guaranteed to be full rank per the previous lemma. Due to a result in Arora et al. (2012a), this $\sigma_{\min} \geq (\gamma p)/2$. This gives us that the whole quantity is at most

$$\frac{4}{\gamma p} \left(\delta_2 + \frac{4arm_U}{m\epsilon_0} \right)$$

□

Corollary 1. *We have that*

$$\|\mathbf{C}^F - \bar{\mathbf{C}}^F\|_{\infty} \leq \frac{4}{\gamma p} \left(\delta_2 + \frac{4arm_U}{m\epsilon_0} \right)$$

since the ℓ_{∞} norm is upper bounded by the ℓ_2 norm.

Lemma 16. *The following are true.*

- $\|\mathbf{C}^F - \bar{\mathbf{C}}^F\|_{\infty} \leq \frac{4}{\gamma p} \left(\delta_2 + \frac{4arm_U}{m\epsilon_0} \right)$
- $\|\bar{\mathbf{C}}^F - \mathbf{C}^*\|_{\infty} \leq \|\bar{\mathbf{C}}^F - \mathbf{C}^F\|_{\infty} + \|\mathbf{C}^F - \mathbf{C}^*\|_{\infty} \leq \frac{4}{\gamma p} \left(\delta_2 + \frac{4arm_U}{m\epsilon_0} \right) + O(\delta'_2/(\gamma p)^2)$

From this, we can bound the errors on the topic matrix.

Lemma 17. *The following are true.*

- $\|\mathbf{A}^F - \bar{\mathbf{A}}\|_{\infty} \leq O(ar\|\mathbf{C}^F - \bar{\mathbf{C}}^F\|_{\infty})$
- $\|\bar{\mathbf{A}} - \mathbf{A}^*\|_{\infty} \leq O(ar\|\bar{\mathbf{C}}^F - \mathbf{C}^*\|_{\infty})$
- $\|\mathbf{A}^S - \mathbf{A}^F\|_{\infty} \leq O(ar\|\mathbf{C}^F - \mathbf{C}^S\|_{\infty})$

Proof. Note that entries $\mathbf{A}_{i,k}$ are

$$\mathbf{A}_{i,k} = \frac{\mathbf{C}_{i,k} \Pr[w=i]}{\Pr[z=k]}$$

Therefore, the perturbation in \mathbf{A} will be the perturbation in \mathbf{C} multiplied by ar , since the denominator is lower bounded by $1/ar$ due to the topic imbalance constant. □

Now, we give a new lemma.

Proposition 3. *When $m_U \geq \Omega(\sqrt{\frac{m\epsilon_0}{4ar}})$, we have that*

$$\delta'_2 = \delta_2 + \frac{4arm_U}{m\epsilon_0} = \sqrt{\frac{4ar}{m\epsilon_0}} + \frac{4arm_U}{m\epsilon_0} \leq O\left(\frac{arm_U}{m\epsilon_0}\right)$$

Now, we analyze what happens given that $\Omega(\sqrt{\frac{m\epsilon_0}{4ar}}) \leq m_U \leq \frac{0.001m\epsilon_0(\gamma p)^3}{a^2r^2}$.

Lemma 18. For $\epsilon, \delta > 0$, the deletion capacity satisfies

$$T_{\epsilon, \delta}^{\mathcal{A}, \mathcal{U}}(m) \geq \tilde{\Omega}\left(\frac{m}{r^2 \sqrt{nr}}\right)$$

Proof. Recall that

$$\|\bar{\mathbf{A}} - \mathbf{A}^*\|_\infty \leq O(ar\delta'_2(1/\gamma p + 1/(\gamma p)^2)) \leq O\left(\frac{(ar)^2 m_U}{m\epsilon_0 \gamma p}\right)$$

Moreover, we also have that

$$\begin{aligned} \|\bar{\mathbf{A}} - \mathbf{A}^F\|_\infty &\leq O(ar\|\mathbf{C}^F - \bar{\mathbf{C}}^F\|_\infty) \\ &\leq O\left(\frac{4ar\delta'_2}{\gamma p}\right) \\ &\leq O\left(\frac{(ar)^2 m_U}{m\epsilon_0 \gamma p}\right) \end{aligned}$$

Note that \mathbf{A} has ℓ_2 sensitivity $O\left(\sqrt{nr} \frac{(ar)^2 m_U}{m\epsilon_0 \gamma p}\right)$. We now apply the Gaussian mechanism to the matrix \mathbf{A} entrywise with noise

$$\sigma = \frac{O\left(\sqrt{nr} \frac{(ar)^2 m_U}{m\epsilon_0 \gamma p}\right)}{\epsilon} \sqrt{2 \log(1.25/\delta)}$$

From this, we obtain that

$$\begin{aligned} \mathbb{E}\left[\|\tilde{\mathbf{A}} - \mathbf{A}^*\|_\infty\right] &\leq \mathbb{E}\left[\max_{i,k} |\nu_{i,k}|\right] + \mathbb{E}\left[\|\bar{\mathbf{A}} - \mathbf{A}^*\|_\infty\right] \\ &\leq O\left(\sqrt{nr} \cdot \frac{(ar)^2 m_U}{m\epsilon_0 \gamma p} \cdot \sqrt{\log(nr)} \cdot \frac{\sqrt{\log(1/\delta)}}{\epsilon}\right) + O\left(\frac{(ar)^2 m_U}{m\epsilon_0 \gamma p}\right) \end{aligned}$$

Finally, this says that when

$$m_U \leq \tilde{\Omega}\left(\frac{m}{r^2 \sqrt{nr}}\right)$$

we have that the utility is preserved up to constant amount, say 0.01. \square

This proves Theorem 2. It is straightforward to continue the perturbation analysis for the topic-topic covariance matrix \mathbf{R}^* and prove similar deletion capacity rates.

C DOWNSTREAM TASK PROOFS

Recall the algorithm for learning the downstream task head.

Algorithm 6 Learning algorithm for task $\mathcal{T}(\mathcal{A}_{head})$

Input: document corpus $S = \{d_i\}_{i=1}^m$, anchor word tolerance ϵ_0
 $\mathbf{A}, \mathbf{R} = \mathcal{A}_{base}(S)$
return $\arg \min_{\mathbf{w} \in \mathcal{W}_{head}} \ell_{\mathcal{T}}(\mathbf{w}; \mathbf{A})$

Assumption 4. For any \mathbf{A} , $\ell_{\mathcal{T}}$ is λ -strongly convex with respect to w .

Since our topic matrix \mathbf{A} , can only take on a bounded support (i.e. the set of matrices where each row is on the probability simplex), it is natural to say that the set of values $w^*(\mathbf{A})$ takes on over all topic matrices \mathbf{A} is bounded in a certain sense. As such, we also assume the following:

Assumption 5. For any base model \mathbf{A} , the vector \mathbf{v} such that $\mathbf{v} = \arg \min_{\mathbf{w}} \ell_{\mathcal{T}}(\mathbf{w}; \mathbf{A})$ satisfies $\|\mathbf{v}\|_2 \leq B$.

Assumption 6. For any \mathbf{A} , $\ell_{\mathcal{T}}$ is L -Lipschitz with respect to \mathbf{w} and the ℓ_2 norm, and is L_2 -Hessian Lipschitz with respect to \mathbf{w} and the ℓ_2 norm. In other words,

$$\begin{aligned} \|\ell_{\mathcal{T}}(\mathbf{A}, \mathbf{w}_1) - \ell_{\mathcal{T}}(\mathbf{A}, \mathbf{w}_2)\|_2 &\leq L\|\mathbf{w}_1 - \mathbf{w}_2\|_2 \\ \|\nabla_{\mathbf{w}}^2 \ell_{\mathcal{T}}(\mathbf{A}, \mathbf{w}_1) - \nabla_{\mathbf{w}}^2 \ell_{\mathcal{T}}(\mathbf{A}, \mathbf{w}_2)\|_2 &\leq L_2\|\mathbf{w}_1 - \mathbf{w}_2\|_2 \end{aligned}$$

Assumption 7. For any w , $\nabla_w \ell_{\mathcal{T}}$ is L_{∞} -Lipschitz with respect to \mathbf{A} and the ℓ_{∞} norm; that is,

$$\|\nabla_w \ell_{\mathcal{T}}(\mathbf{A}, \mathbf{w}) - \nabla_w \ell_{\mathcal{T}}(\tilde{\mathbf{A}}, \mathbf{w})\|_2 \leq L_{\infty} \|\mathbf{A} - \tilde{\mathbf{A}}\|_{\infty}$$

We give a helper lemma that (ϵ, δ) -indistinguishability is immune to post processing.

Lemma 19 (Post-processing immunity). *Consider two random variables $\theta_1, \theta_2 \in \Theta$ that are (ϵ, δ) -indistinguishable. Then, for any arbitrary mapping $f : \Theta \rightarrow \Theta'$, it holds that $f(\theta_1), f(\theta_2) \in \Theta'$ are (ϵ, δ) -indistinguishable.*

Proof. Consider an arbitrary set $T' \subseteq \Theta'$; let $T = \{r \in \Theta : f(r) \in T'\}$. Then, it holds that

$$\begin{aligned} \Pr[f(\theta_1) \in T'] &= \Pr[\theta_1 \in T] \\ &\leq e^{\epsilon} \Pr[\theta_2 \in T] + \delta \\ &= e^{\epsilon} \Pr[f(\theta_2) \in T'] + \delta \end{aligned}$$

as desired. \square

We now give a certifiable unlearning guarantee for the most naive retraining algorithm for the downstream task, which we mentioned in the main text as Theorem 3.

Theorem 6 (Unlearning when releasing \mathbf{A} and \mathbf{w}). *For a downstream task \mathcal{T} with loss function $\ell_{\mathcal{T}}$, consider the unlearning algorithm $\mathcal{U}_{\text{head, naive}}$ that first runs Algorithm 1 to compute $\tilde{\mathbf{A}} = \mathcal{U}_{\text{base}}(S_f, \mathcal{A}_{\text{base}}(S), T(S))$, where $(\mathcal{A}_{\text{base}}, \mathcal{U}_{\text{base}})$ perform utility-preserving unlearning (Theorem 2). Then, it fits a head $\mathbf{w} = \arg \min_{\mathbf{w} \in \mathcal{W}_{\text{head}}} \ell_{\mathcal{T}}(\mathbf{w}; \tilde{\mathbf{A}})$ and returns $\tilde{\mathbf{A}}$ and \mathbf{w} . We assert that $(\mathcal{A}_{\text{head, naive}}, \mathcal{U}_{\text{head, naive}})$ performs utility-preserving unlearning (Definition 4).*

Proof. Intuitively, this is a result of post processing. More precisely, consider the (ϵ, δ) -indistinguishable base models $\tilde{\mathbf{A}} := \mathcal{U}_{\text{base}}(S_f, \mathcal{A}_{\text{base}}(S), T(S))$ and $\tilde{\mathbf{A}}' := \mathcal{U}_{\text{base}}(\emptyset, \mathcal{A}_{\text{base}}(S \setminus S_f), T(S \setminus S_f))$. Then, since the head fitting is a deterministic post-processing of the original model, this proves the (ϵ, δ) -indistinguishability between the two.

To prove the utility preservation, observe that in this setting

$$\mathbb{E}[\|\tilde{\mathbf{A}} - \mathbf{A}^*\|_{\infty}] \leq 0.01$$

We thus obtain by Lemma 20

$$\begin{aligned} \mathbb{E}[\|\mathbf{w}^*(\tilde{\mathbf{A}}) - \mathbf{w}^*(\mathbf{A}^*)\|_{\infty}] &\leq \mathbb{E}[\|\mathbf{w}^*(\tilde{\mathbf{A}}) - \mathbf{w}^*(\mathbf{A}^*)\|_2] \\ &\leq \frac{L_{\infty}}{\lambda} \mathbb{E}[\|\tilde{\mathbf{A}} - \mathbf{A}^*\|_{\infty}] \end{aligned}$$

which is at most 0.01, up to constant rescaling. \square

The above result is nice, and it follows from the fact that the training algorithm of the downstream task head is just a post-processing. However, a downside is that it still requires retraining of the downstream task head. We can show something stronger: even without provable unlearning of the base model (\mathbf{A} and \mathbf{R}), we can achieve provable unlearning of the downstream task head weights when the downstream task loss is convex in the trainable weights w .

We will now consider an arbitrary task \mathcal{T} . We first give the following notation.

Definition 8. For a base model \mathbf{A} , let $\mathbf{w}^*(\mathbf{A}) := \arg \min_{\mathbf{w}} \ell_{\mathcal{T}}(\mathbf{w}; \mathbf{A})$.

First, we give the following helper lemma that will be useful later on.

Lemma 20. Consider two base models \mathbf{A}_1 and \mathbf{A}_2 . Then, it holds that

$$\|\mathbf{w}^*(\mathbf{A}_1) - \mathbf{w}^*(\mathbf{A}_2)\|_2 \leq \frac{L_\infty}{\lambda} \|\mathbf{A}_1 - \mathbf{A}_2\|_\infty$$

Proof. Observe that

$$\begin{aligned} \lambda \|\mathbf{w}^*(\mathbf{A}_1) - \mathbf{w}^*(\mathbf{A}_2)\|_2 &\leq \|\nabla_{\mathbf{w}} \ell_{\mathcal{T}}(\mathbf{w}^*(\mathbf{A}_1); \mathbf{A}_2) - \nabla_{\mathbf{w}} \ell_{\mathcal{T}}(\mathbf{w}^*(\mathbf{A}_2); \mathbf{A}_2)\|_2 \\ &= \|\nabla_{\mathbf{w}} \ell_{\mathcal{T}}(\mathbf{w}^*(\mathbf{A}_1); \mathbf{A}_2) - \nabla_{\mathbf{w}} \ell_{\mathcal{T}}(\mathbf{w}^*(\mathbf{A}_1); \mathbf{A}_1)\|_2 \\ &\leq L_\infty \|\mathbf{A}_1 - \mathbf{A}_2\|_\infty \end{aligned}$$

where the first line follows from strong convexity, the second line from the gradients being zero, and the third line from the definition of L_∞ Lipschitz constant. Dividing both sides by λ gives the desired result. \square

We now define the following notations for clarity.

- $\mathbf{w}^S := \mathbf{w}^*(\mathbf{A}^S)$
- $\mathbf{w}^F := \mathbf{w}^*(\mathbf{A}^F)$
- $\bar{\mathbf{w}}^* := \mathbf{w}^*(\bar{\mathbf{A}})$
- $\bar{\mathbf{w}} := \mathbf{w}^S - H_{\mathbf{w}^S}^{-1} \nabla_{\mathbf{w}} \ell_{\mathcal{T}}(\mathbf{w}^S; \bar{\mathbf{A}})$, which is the Newton step we take from \mathbf{w}^S to approximate $\bar{\mathbf{w}}^*$

First, we give a bound on the approximation error of the Newton step.

Lemma 21. It holds that

$$\|\bar{\mathbf{w}} - \bar{\mathbf{w}}^*\| \leq \frac{L_2 L_\infty^2}{2\lambda^3} \|\mathbf{A}^S - \bar{\mathbf{A}}\|_\infty^2$$

Proof. We aim to bound the distance of the Newton step from $\bar{\mathbf{w}}^*$:

$$\bar{\mathbf{w}} - \bar{\mathbf{w}}^* = (\mathbf{w}^S - H_{\mathbf{w}^S}^{-1} \nabla_{\mathbf{w}} \ell_{\mathcal{T}}(\bar{\mathbf{A}}, \mathbf{w}^S)) - \bar{\mathbf{w}}^*$$

where $H_{\mathbf{w}^S} = \nabla_{\mathbf{w}}^2 \ell_{\mathcal{T}}(\bar{\mathbf{A}}, \mathbf{w}^S)$. Then, it holds that

$$\begin{aligned} &\mathbf{w}^S - H_{\mathbf{w}^S}^{-1} \nabla_{\mathbf{w}} \ell_{\mathcal{T}}(\bar{\mathbf{A}}, \mathbf{w}^S) - \bar{\mathbf{w}}^* \\ &= \mathbf{w}^S - \bar{\mathbf{w}}^* - H_{\mathbf{w}^S}^{-1} (\nabla_{\mathbf{w}} \ell_{\mathcal{T}}(\bar{\mathbf{A}}, \mathbf{w}^S) - \nabla_{\mathbf{w}} \ell_{\mathcal{T}}(\bar{\mathbf{A}}, \bar{\mathbf{w}}^*)) \\ &= H_{\mathbf{w}^S}^{-1} \left(H_{\mathbf{w}^S} (\mathbf{w}^S - \bar{\mathbf{w}}^*) - \int_0^1 H_{\bar{\mathbf{w}}^* + t(\mathbf{w}^S - \bar{\mathbf{w}}^*)} (\mathbf{w}^S - \bar{\mathbf{w}}^*) dt \right) \\ &= H_{\mathbf{w}^S}^{-1} \int_0^1 (H_{\mathbf{w}^S} - H_{\bar{\mathbf{w}}^* + t(\mathbf{w}^S - \bar{\mathbf{w}}^*)}) dt \cdot (\mathbf{w}^S - \bar{\mathbf{w}}^*) \end{aligned}$$

The norm of this quantity is therefore bounded by

$$\begin{aligned} &\|H_{\mathbf{w}^S}^{-1}\|_2 \cdot \frac{L_2}{2} \|\mathbf{w}^S - \bar{\mathbf{w}}^*\| \cdot \|\mathbf{w}^S - \bar{\mathbf{w}}^*\| \\ &\leq \frac{L_2}{2\lambda} \|\mathbf{w}^S - \bar{\mathbf{w}}^*\|_2^2 \\ &\leq \frac{L_2}{2\lambda} \left(\frac{1}{\lambda} \|\nabla \ell_{\mathcal{T}}(\bar{\mathbf{A}}, \mathbf{w}^S) - \nabla \ell_{\mathcal{T}}(\bar{\mathbf{A}}, \bar{\mathbf{w}}^*)\|_2 \right)^2 \\ &\leq \frac{L_2}{2\lambda} \left(\frac{L_\infty}{\lambda} \|\bar{\mathbf{A}} - \mathbf{A}^S\|_\infty \right)^2 \end{aligned}$$

Hence, we have that

$$\|\bar{\mathbf{w}} - \bar{\mathbf{w}}^*\|_2 \leq \frac{L_2 L_\infty^2}{2\lambda^3} \|\mathbf{A}^S - \bar{\mathbf{A}}\|_\infty^2$$

\square

C.1 INSTANTIATING FOR $\mathbb{T}_{\text{CLF}} = [r]$

We first instantiate Theorem 4 for the case where $\mathbb{T}_{\text{clf}} = [r]$, or equivalently when $q = 1/ar$.

Lemma 22. *Recall our retrained model for the downstream task is $A^F w^F$. Then, it holds that*

$$\|\bar{A}\bar{w} - A^F w^F\|_2 \leq O\left(\sqrt{r}\left(\frac{(ar)^2 m_U}{m\epsilon_0 \gamma p}\right)^2 + B\sqrt{nr}\frac{(ar)^2 m_U}{m\epsilon_0 \gamma p}\right)$$

Proof. We rewrite as follows.

$$\bar{A}\bar{w} - A^F w^F = (\bar{A}\bar{w} - \bar{A}\bar{w}^*) + (\bar{A}\bar{w}^* - A^F \bar{w}^*) + (A^F \bar{w}^* - A^F w^F)$$

Now, we proceed to bound the ℓ_2 norm of each of these individual terms separately. For the first term, we have that

$$\begin{aligned} \|\bar{A}\bar{w} - \bar{A}\bar{w}^*\|_2 &= \|\bar{A}(\bar{w} - \bar{w}^*)\|_2 \\ &\leq \|\bar{w} - \bar{w}^*\|_1 \\ &\leq \sqrt{r}\|\bar{w} - \bar{w}^*\|_2 \\ &\leq \sqrt{r}\frac{L_2 L_\infty^2}{2\lambda^3}\|A^S - \bar{A}\|_\infty^2 \\ &\leq \sqrt{r}\frac{L_2 L_\infty^2}{2\lambda^3}\left(\frac{(ar)^2 m_U}{m\epsilon_0 \gamma p}\right)^2 \end{aligned}$$

where second line follows from \bar{A} having column sum 1, and the fourth line follows from Lemma 20. For the third term, we have a similar analysis.

$$\begin{aligned} \|A^F \bar{w}^* - A^F w^F\|_2 &= \|A^F(\bar{w}^* - w^F)\|_2 \\ &\leq \|\bar{w}^* - w^F\|_1 \\ &\leq \sqrt{r}\|\bar{w}^* - w^F\|_2 \\ &\leq \sqrt{r}\frac{L_\infty}{\lambda}\|\bar{A} - A^F\|_\infty \\ &\leq \sqrt{r}\frac{L_\infty}{\lambda}\left(\frac{(ar)^2 m_U}{m\epsilon_0 \gamma p}\right) \end{aligned}$$

Finally, for the second term, we have that

$$\begin{aligned} \|\bar{A}\bar{w}^* - A^F \bar{w}^*\|_2 &\leq \|\bar{A} - A^F\|_2\|\bar{w}^*\|_2 \\ &\leq \|\bar{A} - A^F\|_\infty\sqrt{nr}\|\bar{w}^*\|_2 \\ &\leq O\left(\frac{(ar)^2 m_U}{m\epsilon_0 \gamma p}\sqrt{nr}B\right) \end{aligned}$$

By triangle inequality, we obtain the desired result. \square

First, we note show the following property of the learned topic model A^S .

Lemma 23. *The minimum singular value of the ground truth topic matrix A^S is at least $\Theta(p)$, since the perturbations in entries of A^* are at most $\epsilon_0 \leq O(1/\sqrt{nr})$. Hence, the singular values cannot change by more than a constant factor relative to p .*

Proof. We know that A^* is a p -separable topic model, and hence has smallest singular value at least p . For the given sample complexity of learning, A^S will have smallest singular value at least $\Theta(p)$. \square

The above result says that A^S has a unique pseudoinverse, and has largest singular value at most $O(1/p)$.

Recall that our goal for the downstream task is to approximate the \mathbf{v} such that

$$\mathbf{A}^S \mathbf{v} = \mathbf{A}^F \mathbf{w}^F$$

in order to say we have approximated the unlearned fine-tuned model. Therefore, it suffices to obtain indistinguishability of our unlearning algorithm output $\tilde{\mathbf{w}}$ with $(\mathbf{A}^S)^\dagger \mathbf{A}^F \mathbf{w}^F$. Our following claim is that we can use $(\mathbf{A}^S)^\dagger \bar{\mathbf{A}} \bar{\mathbf{w}}$ as the approximation for this.

Proposition 4. *It holds that*

$$\begin{aligned} \|(\mathbf{A}^S)^\dagger \bar{\mathbf{A}} \bar{\mathbf{w}} - (\mathbf{A}^S)^\dagger \mathbf{A}^F \mathbf{w}^F\|_2 &\leq O\left(\frac{1}{p} \|\bar{\mathbf{A}} \bar{\mathbf{w}} - \mathbf{A}^F \mathbf{w}^F\|_2\right) \\ &\leq O\left(\frac{1}{p} \cdot \left[\sqrt{r} \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right)^2 + B \sqrt{nr} \frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right]\right) \end{aligned}$$

Let $\bar{\mathbf{v}} := (\mathbf{A}^S)^\dagger \bar{\mathbf{A}} \bar{\mathbf{w}}$ and $\mathbf{v} = (\mathbf{A}^S)^\dagger \mathbf{A}^F \mathbf{w}^F$. We claim the following.

Lemma 24. *The unlearning algorithm \mathcal{U}_{head} that outputs*

$$\tilde{\mathbf{v}} := \bar{\mathbf{v}} + \nu_v$$

where ν_v is the noise defined by the Gaussian mechanism using the above sensitivity satisfies provable (ϵ, δ) unlearning. In particular, we use

$$\sigma = \frac{O\left(\frac{1}{p} \cdot \left[\sqrt{r} \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right)^2 + B \sqrt{nr} \frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right]\right)}{\epsilon} \sqrt{2 \log(1.25/\delta)}$$

where the numerator of the fraction is from the previous proposition.

Proof. This follows from Gaussian mechanism. \square

We now proceed to bound the deletion capacity. In this case, the utility is defined by the closeness of $\tilde{\mathbf{v}}$ to $(\mathbf{A}^S)^\dagger \mathbf{A}^* \mathbf{w}^*$ in ℓ_∞ norm, similar the way we defined this for the base model unlearning algorithm \mathcal{U}_{base} earlier.

First, the following lemma to bound $\mathbf{A}^F \mathbf{w}^F - \mathbf{A}^* \mathbf{w}^*$.

Lemma 25. *We have that*

$$\|\mathbf{A}^F \mathbf{w}^F - \mathbf{A}^* \mathbf{w}^*\|_2 \leq O\left(B \sqrt{nr} \frac{(ar)^2 m_U}{m \epsilon_0 \gamma p}\right)$$

Proof. We decompose as follows.

$$\mathbf{A}^F \mathbf{w}^F - \mathbf{A}^* \mathbf{w}^* = (\mathbf{A}^F \mathbf{w}^F - \mathbf{A}^F \mathbf{w}^*) + (\mathbf{A}^F \mathbf{w}^* - \mathbf{A}^* \mathbf{w}^*)$$

The first term is bounded by

$$\|\mathbf{A}^F \mathbf{w}^F - \mathbf{A}^F \mathbf{w}^*\|_2 \leq \sqrt{r} \|\mathbf{w}^F - \mathbf{w}^*\|_2 \leq O(\sqrt{r} \|\mathbf{A}^F - \mathbf{A}^*\|_\infty) \leq O\left(\sqrt{r} \frac{(ar)^2 m_U}{m \epsilon_0 \gamma p}\right)$$

The second term is bounded by

$$\|\mathbf{A}^F \mathbf{w}^* - \mathbf{A}^* \mathbf{w}^*\|_2 \leq O\left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \sqrt{nr} B\right)$$

by considering the spectral norm $\|\mathbf{A}^F - \mathbf{A}^*\|_2$. This gives the desired result. \square

As a result, the following holds.

Proposition 5. *It holds that*

$$\|(\mathbf{A}^S)^\dagger \mathbf{A}^F \mathbf{w}^F - (\mathbf{A}^S)^\dagger \mathbf{A}^* \mathbf{w}^*\|_2 \leq O\left(\frac{1}{p} \left[\sqrt{r} \frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} + B \sqrt{nr} \frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right]\right)$$

This is once again from the bounded operator norm property of $(\mathbf{A}^S)^\dagger$.

Finally, we can apply triangle inequality to get the following.

Lemma 26. *It holds that*

$$\|(\mathbf{A}^S)^\dagger \bar{\mathbf{A}}\bar{\mathbf{w}} - (\mathbf{A}^S)^\dagger \mathbf{A}^* \mathbf{w}^*\|_2 \leq \left(\frac{1}{p} \cdot \left[\sqrt{r} \left(\frac{(ar)^2 m_U}{m\epsilon_0 \gamma p} \right)^2 + B\sqrt{nr} \frac{(ar)^2 m_U}{m\epsilon_0 \gamma p} \right] \right)$$

Then, we can get the following bound on deletion capacity.

Lemma 27. *For $\epsilon, \delta > 0$, the deletion capacity satisfies*

$$T_{\epsilon, \delta}^{\mathcal{A}_{head}, \mathcal{U}_{head}}(m) \geq \tilde{\Omega} \left(\frac{m}{r^2 \sqrt{nr}} \right)$$

Proof. The calculation is as follows.

$$\begin{aligned} \mathbb{E}[\|\bar{\mathbf{v}} - (\mathbf{A}^S)^\dagger \mathbf{A}^* \mathbf{w}^*\|_\infty] &\leq \mathbb{E}[\|\nu_v\|_\infty] + \mathbb{E}[\|(\mathbf{A}^S)^\dagger \bar{\mathbf{A}}\bar{\mathbf{w}} - (\mathbf{A}^S)^\dagger \mathbf{A}^* \mathbf{w}^*\|_\infty] \\ &\leq \left(\frac{1}{p} \cdot \left[\sqrt{r} \left(\frac{(ar)^2 m_U}{m\epsilon_0 \gamma p} \right)^2 + B\sqrt{nr} \frac{(ar)^2 m_U}{m\epsilon_0 \gamma p} \right] \right) \left(\frac{\sqrt{\log r \log 1/\delta}}{\epsilon} + 1 \right) \end{aligned}$$

For this to be a small constant, we require

$$\frac{(ar)^2 m_U}{m\epsilon_0 \gamma p} \leq \tilde{O} \left(\min \left\{ \frac{1}{r^{1/4}}, \frac{1}{\sqrt{nr}} \right\} \right)$$

Therefore, we should have

$$m_U \leq \tilde{\Omega} \left(\frac{m}{r^2 \sqrt{nr}} \right)$$

□

C.2 PROOF FOR GENERAL q

The following is the formal statement of Theorem 4.

Theorem 7 (Formal version of Theorem 4). *Suppose that the downstream task \mathcal{T} only depends on a subset of topics $\mathbb{T}_{clf} \subseteq [r]$; that is, $\mathbf{w}^* = \arg \min_{\mathbf{v} \in \mathcal{W}_{base}} \ell_{\mathcal{T}}(\mathbf{v}; \mathbf{A}^*)$ has non-zero entries only in the index set \mathbb{T}_{clf} . Denote $q := \min_{k \in \mathbb{T}_{clf}} \Pr_{\mathcal{D}}[z = k]$, and let \mathcal{A}_{head} be the head tuning algorithm (Definition 2) and \mathcal{U}_{head} be Algorithm 2. Then, $(\mathcal{A}_{head}, \mathcal{U}_{head})$ performs utility-preserving unlearning with deletion capacity*

$$T_{\epsilon, \delta}^{\mathcal{A}_{head}, \mathcal{U}_{head}}(m) \geq c' \cdot \min \left\{ \frac{mq\epsilon}{r\sqrt{nr} \log 1/\delta}, \frac{0.001m}{r^2} \right\}$$

where c' is a constant dependent on \mathcal{D} , and \mathcal{T} .

Lemma 28. *Recall our retrained model for the downstream task is $\mathbf{A}^F \mathbf{w}^F$. Then, it holds that*

$$\|\bar{\mathbf{A}}\bar{\mathbf{w}} - \mathbf{A}^F \mathbf{w}^F\|_2 \leq O \left(\sqrt{r} \left(\frac{(ar)^2 m_U}{m\epsilon_0 \gamma p} \right) \right) + O \left(B\sqrt{nr} \frac{(1/q)arm_U}{m\epsilon_0 \gamma p} \right) + O \left(\left(\frac{(ar)^2 m_U}{m\epsilon_0 \gamma p} \right)^2 \sqrt{nr} \right)$$

Proof. Consider this decomposition again.

$$\bar{\mathbf{A}}\bar{\mathbf{w}} - \mathbf{A}^F \mathbf{w}^F = (\bar{\mathbf{A}}\bar{\mathbf{w}} - \bar{\mathbf{A}}\bar{\mathbf{w}}^*) + (\bar{\mathbf{A}}\bar{\mathbf{w}}^* - \mathbf{A}^F \bar{\mathbf{w}}^*) + (\mathbf{A}^F \bar{\mathbf{w}}^* - \mathbf{A}^F \mathbf{w}^F)$$

The first term is the same as old analysis; the second term is from considering q ; the third is the same as the old analysis. In particular, when $q = 1/ar$, we recover the old bound. We have that the first term is

$$\|\bar{\mathbf{A}}\bar{\mathbf{w}} - \bar{\mathbf{A}}\bar{\mathbf{w}}^*\| \leq \sqrt{r} \frac{L_2 L_\infty^2}{2\lambda^3} \left(\frac{(ar)^2 m_U}{m\epsilon_0 \gamma p} \right)^2$$

The third term is

$$\|\mathbf{A}^F \bar{\mathbf{w}}^* - \mathbf{A}^F \mathbf{w}^F\| \leq \sqrt{r} \frac{L_\infty}{\lambda} \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right)$$

The second term is

$$\begin{aligned} \|\bar{\mathbf{A}} \bar{\mathbf{w}}^* - \mathbf{A}^F \bar{\mathbf{w}}^*\| &\leq \|(\bar{\mathbf{A}} - \mathbf{A}^F) \bar{\mathbf{w}}^*\| + \|(\bar{\mathbf{A}} - \mathbf{A}^F)(\mathbf{w}^* - \bar{\mathbf{w}}^*)\| \\ &\leq O\left(B \sqrt{nr} \frac{(1/q) ar m_U}{m \epsilon_0 \gamma p}\right) + O\left(\left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p}\right)^2 \sqrt{nr}\right) \end{aligned}$$

This gives the desired result using triangle inequality. \square

Continuing, we have the following.

Proposition 6. *It holds that*

$$\begin{aligned} &\|(\mathbf{A}^S)^\dagger \bar{\mathbf{A}} \bar{\mathbf{w}} - (\mathbf{A}^S)^\dagger \mathbf{A}^F \mathbf{w}^F\|_2 \\ &\leq O\left(\frac{1}{p} \|\bar{\mathbf{A}} \bar{\mathbf{w}} - \mathbf{A}^F \mathbf{w}^F\|_2\right) \\ &\leq O\left(\frac{1}{p} \cdot \left[\sqrt{r} \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right) + B \sqrt{nr} \frac{(1/q) ar m_U}{m \epsilon_0 \gamma p} + \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right)^2 \sqrt{nr} \right]\right) \end{aligned}$$

This gives us the following.

Lemma 29. *The unlearning algorithm \mathcal{U}_{head} that outputs*

$$\tilde{\mathbf{v}} := \bar{\mathbf{v}} + \nu_v$$

where ν_v is the noise defined by the Gaussian mechanism using the above sensitivity satisfies provable (ϵ, δ) unlearning. In particular, we use

$$\sigma = \frac{O\left(\frac{1}{p} \cdot \left[\sqrt{r} \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right) + B \sqrt{nr} \frac{(1/q) ar m_U}{m \epsilon_0 \gamma p} + \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right)^2 \sqrt{nr} \right]\right)}{\epsilon} \sqrt{2 \log(1.25/\delta)}$$

where the numerator of the fraction is from the previous proposition.

Proof. This follows from Gaussian mechanism. \square

We now proceed to bound the deletion capacity. In this case, the utility is defined by the closeness of $\tilde{\mathbf{v}}$ to $(\mathbf{A}^S)^\dagger \mathbf{A}^* \mathbf{w}^*$ in ℓ_∞ norm, similar the way we defined this for the base model unlearning algorithm \mathcal{U}_{base} earlier.

First, the following lemma to bound $\mathbf{A}^F \mathbf{w}^F - \mathbf{A}^* \mathbf{w}^*$.

Lemma 30. *We have that*

$$\|\mathbf{A}^F \mathbf{w}^F - \mathbf{A}^* \mathbf{w}^*\|_2 \leq O\left(\sqrt{r} \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right) + B \sqrt{nr} \frac{(1/q) ar m_U}{m \epsilon_0 \gamma p} + \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right)^2 \sqrt{nr}\right)$$

Proof. We decompose as follows.

$$\mathbf{A}^F \mathbf{w}^F - \mathbf{A}^* \mathbf{w}^* = (\mathbf{A}^F \mathbf{w}^F - \mathbf{A}^F \mathbf{w}^*) + (\mathbf{A}^F \mathbf{w}^* - \mathbf{A}^* \mathbf{w}^*)$$

The first term is bounded by

$$\|\mathbf{A}^F \mathbf{w}^F - \mathbf{A}^F \mathbf{w}^*\|_2 \leq \sqrt{r} \|\mathbf{w}^F - \mathbf{w}^*\|_2 \leq O(\sqrt{r} \|\mathbf{A}^F - \mathbf{A}^*\|_\infty) \leq O\left(\sqrt{r} \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right)\right)$$

The second term is bounded by

$$\|\mathbf{A}^F \mathbf{w}^* - \mathbf{A}^* \mathbf{w}^*\|_2 \leq B \sqrt{nr} \frac{(1/q) ar m_U}{m \epsilon_0 \gamma p} + \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right)^2 \sqrt{nr}$$

Triangle inequality gives us the desired result. \square

As a result, the following holds.

Proposition 7. *It holds that*

$$\|(\mathbf{A}^S)^\dagger \mathbf{A}^F \mathbf{w}^F - (\mathbf{A}^S)^\dagger \mathbf{A}^* \mathbf{w}^*\|_2 \leq O\left(\frac{1}{p} \cdot \left[\sqrt{r} \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right) + B \sqrt{nr} \frac{(1/q) ar m_U}{m \epsilon_0 \gamma p} + \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right)^2 \sqrt{nr} \right]\right)$$

This is once again from the bounded operator norm property.

Finally, we can apply triangle inequality to get the following.

Lemma 31. *It holds that*

$$\|(\mathbf{A}^S)^\dagger \bar{\mathbf{A}} \bar{\mathbf{w}} - (\mathbf{A}^S)^\dagger \mathbf{A}^* \mathbf{w}^*\|_2 \leq O\left(\frac{1}{p} \cdot \left[\sqrt{r} \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right) + B \sqrt{nr} \frac{(1/q) ar m_U}{m \epsilon_0 \gamma p} + \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right)^2 \sqrt{nr} \right]\right)$$

Then, we can get the following bound on deletion capacity.

Lemma 32. *For $\epsilon, \delta > 0$, the deletion capacity satisfies*

$$T_{\epsilon, \delta}^{\mathbf{A}^{head}, \mathcal{U}^{head}}(m) \geq \tilde{\Omega}\left(\frac{m}{r^2 \sqrt{nr}}\right)$$

Proof. The calculation is as follows.

$$\begin{aligned} \mathbb{E}[\|\tilde{\mathbf{v}} - (\mathbf{A}^S)^\dagger \mathbf{A}^* \mathbf{w}^*\|_\infty] &\leq \mathbb{E}[\|\nu_v\|_\infty] + \mathbb{E}[\|(\mathbf{A}^S)^\dagger \bar{\mathbf{A}} \bar{\mathbf{w}} - (\mathbf{A}^S)^\dagger \mathbf{A}^* \mathbf{w}^*\|_\infty] \\ &\leq \left(\frac{1}{p} \cdot \left[\sqrt{r} \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right) + B \sqrt{nr} \frac{(1/q) ar m_U}{m \epsilon_0 \gamma p} + \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right)^2 \sqrt{nr} \right]\right) \\ &\quad \cdot \left(\frac{\sqrt{\log r \log 1/\delta}}{\epsilon} + 1 \right) \end{aligned}$$

For this to be a small constant, we require

$$\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \leq \tilde{O}\left(\min\left\{\frac{1}{r^{1/2}}, \frac{1}{(nr)^{1/4}}, \frac{arq}{\sqrt{nr}}\right\}\right)$$

When n is at least r^3 , the last of these terms will be the smallest. Therefore, we have that

$$m_U \leq \tilde{\Omega}\left(\frac{mq}{r^{1.5} n^{0.5}}\right)$$

□

C.3 DOWNSTREAM TASKS WITH INEXACT MINIMIZERS

We will consider τ -optimal minimizers for the downstream task.

Definition 9. For a base model \mathbf{A} , let $\mathbf{w}_\tau(\mathbf{A}) \in \{\mathbf{w} : \ell_\tau(\mathbf{w}; \mathbf{A}) - \ell_\tau(\mathbf{w}^*(\mathbf{A}); \mathbf{A}) \leq \tau\}$.

We will have the following technical assumption on τ to eliminate trivial cases.

Assumption 8. *We assume that $\tau = O(1/r)$.*

We have the following intermediate result.

Lemma 33. *For a base model \mathbf{A} , it holds that*

$$\|\mathbf{w}_\tau(\mathbf{A}) - \mathbf{w}^*(\mathbf{A})\|_2 \leq \sqrt{2\tau/\lambda}$$

Proof. This follows from a standard strong convexity argument. □

This gives an updated version of Lemma 20, using triangle inequality.

Lemma 34. Consider two base models \mathbf{A}_1 and \mathbf{A}_2 . Then, it holds that

$$\|\mathbf{w}_\tau(\mathbf{A}_1) - \mathbf{w}_\tau(\mathbf{A}_2)\|_2 \leq 2\sqrt{2\tau/\lambda} + \|\mathbf{w}^*(\mathbf{A}_1) - \mathbf{w}^*(\mathbf{A}_2)\| \leq 2\sqrt{2\tau/\lambda} + \frac{L_\infty}{\lambda} \|\mathbf{A}_1 - \mathbf{A}_2\|_\infty$$

We now define the following notations for clarity.

- $\mathbf{w}_\tau^S := \mathbf{w}_\tau(\mathbf{A}^S)$
- $\mathbf{w}_\tau^F := \mathbf{w}_\tau(\mathbf{A}^F)$
- $\bar{\mathbf{w}}^* := \mathbf{w}^*(\bar{\mathbf{A}})$
- $\bar{\mathbf{w}}_\tau := \mathbf{w}_\tau^S - H_{\mathbf{w}_\tau^S}^{-1} \nabla_w \ell_{\mathcal{T}}(\mathbf{w}_\tau^S; \bar{\mathbf{A}})$, which is the Newton step we take from \mathbf{w}_τ^S to approximate $\bar{\mathbf{w}}^*$

Now, consider our Newton step from $\mathbf{w}_\tau(\mathbf{A}^S)$ to approximate $\mathbf{w}_\tau(\mathbf{A}^F)$. To do so, consider

$$\bar{\mathbf{w}}_\tau := \mathbf{w}_\tau^S - H_{\mathbf{w}_\tau^S}^{-1} \nabla_w \ell_{\mathcal{T}}(\mathbf{w}_\tau^S; \bar{\mathbf{A}})$$

Then, the following holds, using a similar argument as a previous section.

$$\begin{aligned} \|\bar{\mathbf{w}}_\tau - \bar{\mathbf{w}}^*\|_2 &\leq \frac{L_2}{2\lambda} \|\mathbf{w}_\tau^S - \bar{\mathbf{w}}^*\|_2^2 \\ &\leq \frac{L_2}{\lambda} (\|\mathbf{w}_\tau^S - \mathbf{w}^S\|_2^2 + \|\mathbf{w}^S - \bar{\mathbf{w}}^*\|_2^2) \\ &\leq \frac{L_2}{\lambda} \left(\frac{2\tau}{\lambda} + \left(\frac{L_\infty}{\lambda} \|\bar{\mathbf{A}} - \mathbf{A}^S\|_\infty \right)^2 \right) \end{aligned}$$

where the second term of the last line follows from a line in the previous proof of the Newton step.

We now bound $\|\bar{\mathbf{A}}\bar{\mathbf{w}}_\tau - \mathbf{A}^F \mathbf{w}_\tau^F\|_2$. We rewrite this as follows.

$$\bar{\mathbf{A}}\bar{\mathbf{w}}_\tau - \mathbf{A}^F \mathbf{w}_\tau^F = (\bar{\mathbf{A}}\bar{\mathbf{w}}_\tau - \bar{\mathbf{A}}\bar{\mathbf{w}}^*) + (\bar{\mathbf{A}}\bar{\mathbf{w}}^* - \mathbf{A}^F \mathbf{w}^F) + (\mathbf{A}^F \mathbf{w}^F - \mathbf{A}^F \mathbf{w}_\tau^F)$$

The second term we bounded as part of a previous proof. The first term we bound using the above, and the third part we bound using the inexact minimizer lemma. By a step in a previous proof, we can say that the first term

$$\begin{aligned} \|\bar{\mathbf{A}}\bar{\mathbf{w}}_\tau - \bar{\mathbf{A}}\bar{\mathbf{w}}^*\|_2 &\leq \sqrt{r} \|\bar{\mathbf{w}}_\tau - \bar{\mathbf{w}}^*\|_2 \\ &\leq \sqrt{r} \left(\frac{2L_2\tau}{\lambda^2} + \frac{L_2L_\infty^2}{2\lambda^3} \left(\frac{(ar)^2 m_U}{m\epsilon_0\gamma p} \right)^2 \right) \end{aligned}$$

The second term, to reiterate, satisfies

$$\|\bar{\mathbf{A}}\bar{\mathbf{w}}^* - \mathbf{A}^F \mathbf{w}^F\|_2 \leq O\left(\frac{(ar)^2 m_U}{m\epsilon_0\gamma p} \sqrt{nr} B \right)$$

The third term satisfies

$$\|\mathbf{A}^F \mathbf{w}^F - \mathbf{A}^F \mathbf{w}_\tau^F\|_2 \leq \sqrt{r} \|\mathbf{w}^F - \mathbf{w}_\tau^F\|_2 \leq \sqrt{2r\tau/\lambda}$$

Therefore, the entire term is bounded by the sum of the above three terms via triangle inequality. In particular, this is

$$O\left(\sqrt{r} \left(\frac{(ar)^2 m_U}{m\epsilon_0\gamma p} \right)^2 + B\sqrt{nr} \frac{(ar)^2 m_U}{m\epsilon_0\gamma p} + \sqrt{r\tau} + \sqrt{r\tau} \right)$$

Recall that our goal for the downstream task is to approximate the \mathbf{v} such that

$$\mathbf{A}^S \mathbf{v} = \mathbf{A}^F \mathbf{w}_\tau^F$$

or in other words, $\mathbf{v} = (\mathbf{A}^S)^\dagger \mathbf{A}^F \mathbf{w}_\tau^F$. Our claim is that we can use $(\mathbf{A}^S)^\dagger \bar{\mathbf{A}}\bar{\mathbf{w}}_\tau$ as an approximation for this. The following statement formalizes this, utilizing the fact about the maximum singular value of \mathbf{A}^S being at most $\Theta(1/p)$.

Proposition 8. *It holds that*

$$\|(\mathbf{A}^S)^\dagger \bar{\mathbf{A}} \bar{\mathbf{w}}_\tau - (\mathbf{A}^S)^\dagger \mathbf{A}^F \mathbf{w}_\tau^F\|_2 \leq O\left(\frac{1}{p} \left[\sqrt{r} \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right)^2 + B \sqrt{nr} \frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} + \sqrt{r\tau} + \sqrt{r\tau} \right]\right)$$

Now, let $\bar{\mathbf{v}} := (\mathbf{A}^S)^\dagger \bar{\mathbf{A}} \bar{\mathbf{w}}_\tau$, and let $\mathbf{v} := (\mathbf{A}^S)^\dagger \mathbf{A}^F \mathbf{w}_\tau^F$. We claim the following.

Lemma 35. *The unlearning algorithm $\mathcal{U}_{\text{head}}$ that outputs*

$$\tilde{\mathbf{v}} := \bar{\mathbf{v}} + \nu_v$$

where ν_v is the noise defined by the Gaussian mechanism using the above sensitivity satisfies provable (ϵ, δ) unlearning. In particular, we use

$$\sigma = \frac{O\left(\frac{1}{p} \left[\sqrt{r} \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right)^2 + B \sqrt{nr} \frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} + \sqrt{r\tau} + \sqrt{r\tau} \right]\right)}{\epsilon} \sqrt{2 \log(1.25/\delta)}$$

where the numerator of the fraction is from the previous proposition.

Proof. This follows from Gaussian mechanism. \square

We now proceed to bound the deletion capacity. In this case, the utility is defined by the closeness of $\tilde{\mathbf{v}}$ to $(\mathbf{A}^S)^\dagger \mathbf{A}^* \mathbf{w}^*$ in ℓ_∞ norm, similar the way we defined this for the base model unlearning algorithm $\mathcal{U}_{\text{base}}$ earlier.

First, the following lemma to bound $\mathbf{A}^F \mathbf{w}_\tau^F - \mathbf{A}^* \mathbf{w}^*$.

Lemma 36. *We have that*

$$\|\mathbf{A}^F \mathbf{w}_\tau^F - \mathbf{A}^* \mathbf{w}^*\|_2 \leq O\left(B \sqrt{nr} \frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} + \sqrt{r\tau}\right)$$

Proof. We decompose as follows.

$$\mathbf{A}^F \mathbf{w}_\tau^F - \mathbf{A}^* \mathbf{w}^* = (\mathbf{A}^F \mathbf{w}_\tau^F - \mathbf{A}^F \mathbf{w}^F) + (\mathbf{A}^F \mathbf{w}^F - \mathbf{A}^F \mathbf{w}^*) + (\mathbf{A}^F \mathbf{w}^* - \mathbf{A}^* \mathbf{w}^*)$$

The second term is bounded by

$$\|\mathbf{A}^F \mathbf{w}^F - \mathbf{A}^F \mathbf{w}^*\|_2 \leq \sqrt{r} \|\mathbf{w}^F - \mathbf{w}^*\|_2 \leq O(\sqrt{r} \|\mathbf{A}^F - \mathbf{A}^*\|_\infty) \leq O\left(\sqrt{r} \frac{(ar)^2 m_U}{m \epsilon_0 \gamma p}\right)$$

The third term is bounded by

$$\|\mathbf{A}^F \mathbf{w}^* - \mathbf{A}^* \mathbf{w}^*\|_2 \leq O\left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \sqrt{nr} B\right)$$

by considering the spectral norm $\|\mathbf{A}^F - \mathbf{A}^*\|_2$.

Finally, the first term is bounded, from an above argument, by $\sqrt{2r\tau/\lambda}$. This gives the desired result. \square

As a result, the following holds.

Proposition 9. *It holds that*

$$\|(\mathbf{A}^S)^\dagger \mathbf{A}^F \mathbf{w}_\tau^F - (\mathbf{A}^S)^\dagger \mathbf{A}^* \mathbf{w}^*\|_2 \leq O\left(\frac{1}{p} \left[\sqrt{r} \frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} + B \sqrt{nr} \frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} + \sqrt{r\tau} \right]\right)$$

This is once again from the bounded operator norm property of $(\mathbf{A}^S)^\dagger$.

Finally, we can apply triangle inequality to get the following.

Lemma 37. *It holds that*

$$\|(\mathbf{A}^S)^\dagger \bar{\mathbf{A}} \bar{\mathbf{w}}_\tau - (\mathbf{A}^S)^\dagger \mathbf{A}^* \mathbf{w}^*\|_2 \leq \left(\frac{1}{p} \cdot \left[\sqrt{r} \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right)^2 + B \sqrt{nr} \frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} + \sqrt{r\tau} + \sqrt{r\tau} \right] \right)$$

Then, we can get the following bound on deletion capacity.

Lemma 38. *For $\epsilon, \delta > 0$, the deletion capacity satisfies*

$$T_{\epsilon, \delta}^{\mathcal{A}^{head}, \mathcal{U}^{head}}(m) \geq \tilde{\Omega} \left(\frac{m}{r^2 \sqrt{nr}} \right)$$

Proof. The calculation is as follows.

$$\begin{aligned} \mathbb{E}[\|\tilde{\mathbf{v}} - (\mathbf{A}^S)^\dagger \mathbf{A}^* \mathbf{w}^*\|_\infty] &\leq \mathbb{E}[\|\nu_v\|_\infty] + \mathbb{E}[\|(\mathbf{A}^S)^\dagger \bar{\mathbf{A}} \bar{\mathbf{w}}_\tau - (\mathbf{A}^S)^\dagger \mathbf{A}^* \mathbf{w}^*\|_\infty] \\ &\leq \left(\frac{1}{p} \cdot \left[\sqrt{r} \left(\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \right)^2 + B \sqrt{nr} \frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} + \sqrt{r\tau} + \sqrt{r\tau} \right] \right) \left(\frac{\sqrt{\log r \log 1/\delta}}{\epsilon} + 1 \right) \end{aligned}$$

For this to be a small constant, we require

$$\frac{(ar)^2 m_U}{m \epsilon_0 \gamma p} \leq \tilde{O} \left(\min \left\{ \frac{1}{r^{1/4}}, \frac{1}{\sqrt{nr}} \right\} \right)$$

since we already have $\tau = O(1/r)$ by assumption. Therefore, we should have

$$m_U \leq \tilde{\Omega} \left(\frac{m}{r^2 \sqrt{nr}} \right)$$

□

Finally, for general q , the same idea used in this section for inexact minimizers can be combined with the previous section.