

# DOMBA: Double Model Balancing for Access-Controlled Language Models via Minimum-Bounded Aggregation

Tom Segal, Asaf Shabtai, Yuval Elovici

Ben-Gurion University of the Negev  
tomsega@post.bgu.ac.il, shabtaia@bgu.ac.il, elovici@bgu.ac.il

## Abstract

The utility of large language models (LLMs) depends heavily on the quality and quantity of their training data. Many organizations possess large data corpora that could be leveraged to train or fine-tune LLMs tailored to their specific needs. However, these datasets often come with access restrictions that are based on user privileges and enforced by access control mechanisms. Training LLMs on such datasets could result in exposure of sensitive information to unauthorized users. A straightforward approach for preventing such exposure is to train a separate model for each access level. This, however, may result in low utility models due to the limited amount of training data per model compared to the amount in the entire organizational corpus. Another approach is to train a single LLM on all the data while limiting the exposure of unauthorized information. However, current exposure-limiting methods for LLMs are ineffective for access-controlled data, where sensitive information appears frequently across many training examples. We propose DOMBA—double model balancing—a simple approach for training and deploying LLMs that provides high utility and access-control functionality with security guarantees. DOMBA aggregates the probability distributions of two models, each trained on documents with (potentially many) different access levels, using a “min-bounded” average function (a function that is bounded by the smaller value, e.g., harmonic mean). A detailed mathematical analysis and extensive evaluation show that DOMBA safeguards restricted information while offering utility comparable to non-secure models.

**Code and datasets** — <https://github.com/ppol/DOMBA>

## Introduction

Organizations can benefit greatly from training dedicated LLMs, such as coding assistants, email writers or question-answering models, on their data (Tiwari et al. 2023). While the benefits can be substantial, such data often contains restricted information, and an access-control mechanism ensuring users can only access information according to their access rights is usually in place. However, LLMs inherently lack such access-control mechanisms, which can lead to the exposure of sensitive information to unauthorized users (Carlini et al. 2021; Kandpal et al. 2024; Pan et al.

2020). A basic approach for introducing access control to LLMs is to train a separate LLM for each access level (Tiwari et al. 2023). However, as our experiments show, this approach can substantially reduce model utility, since the amount of data for each access level is limited. For example, training a model on emails from only one department in an organization may be insufficient for constructing effective organizational emails. To overcome this limitation, sufficient data (including restricted data) must be included in the model’s training set. This means that any secure and high utility method should limit the exposure of the training data to users of the model (according to their access rights).

In this paper, we propose DOMBA, a method for training and deploying LLMs that incorporates an access-control mechanism while maintaining high utility. An overview of DOMBA is presented in Figure 1. To protect sensitive information, DOMBA “balances” two submodels (trained on two different data partitions, each including different access levels) during inference, using a min-bounded average function; intuitively, each submodel “knows” different restricted information. During text generation, the min-bounded function makes it unlikely for information known to just one submodel to be generated. Assuming that restricted information is not shared between the two partitions, this ensures that no restricted information is likely to be generated. It is crucial to separate access levels into partitions in a way that groups access levels with shared sensitive information into the same partition. If we represent access levels as nodes in a graph and connect nodes with edges whenever they share sensitive information, such separation is achievable if the graph is disconnected (and, for practical purposes, does not contain an overly large connected component).

To analyze the privacy protection provided by DOMBA, we formalize the notion of “exposure of secrets” by developing a framework called “token exposure.” Token exposure is based on the change in probability of a token (relative to other tokens) between two language models. The greater the change in a token’s probability (relative to other tokens) when using one model over the other, the greater the exposure of that token. Using our approach, the exposure of DOMBA over both submodels is bounded by the best possible value (i.e., replacing DOMBA by any other model would not improve the bound). This means that DOMBA limits the exposure of “secrets” that appear in just one partition (since

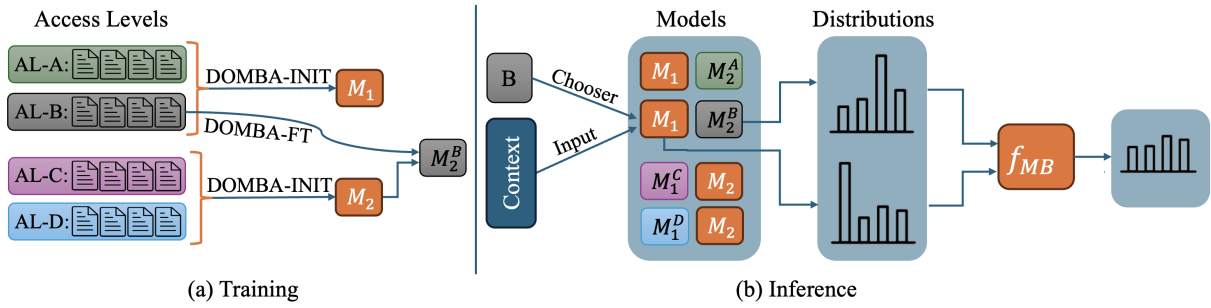


Figure 1: Two main phases of our proposed method. (a) Training: Documents are grouped and divided into two partitions according to their access levels (each access level is depicted in a different color). Training includes two steps: DOMBA-INIT: A submodel is trained on each partition, resulting in  $M_1$  and  $M_2$ . DOMBA-FT: To achieve a model dedicated to access level B (AL-B),  $M_2$  is fine-tuned on AL-B documents, resulting in  $M_2^B$  (note that AL-B documents are included in the training sets of both  $M_1$  and  $M_2^B$ ). (b) Inference: Given a context and access level, the corresponding submodels are selected, and their probability distributions are aggregated using a “min-bounded” aggregation function  $f_{MB}$ .

one of the submodels does not “know” them).

Evaluating sensitive information exposure is challenging due to the fact that defining what constitutes sensitive information is a complex and context-dependent task, varying based on organizational policy and other factors (Brown et al. 2022). Measuring average-case sensitive information exposure is inadequate, since an adversary could devise a prompt that causes the model to substantially deviate from average behavior (Wolf et al. 2024). To address these challenges, we introduce three new security metrics for evaluating sensitive information exposure. The first metric assesses worst-case and “extreme-case” exposure across data records. In the second metric, we evaluate the probabilities assigned to certain tokens that should not be exposed by a secure model. The third new metric involves a “secret inference attack” (SIA) that is based on membership inference attack (MIA) techniques. In addition to evaluating DOMBA on the three new metrics, we also evaluate it using the canary technique (Carlini et al. 2019), in which specific phrases (canaries) are injected into the training data, and the model’s inclination to generate them compared to similar phrases not included in the training data is measured.

The contributions of this work are as follows:

1. Providing the first practical and comprehensive solution for access control in LLMs, a solution that provides high utility and security by safely utilizing restricted data.
2. Developing the mathematical framework of token exposure to analyse language models’ exposure of sensitive information and establishing bounds for DOMBA.
3. Creating three novel empirical evaluation metrics for assessing sensitive information exposure and employing these metrics in our evaluation of DOMBA.

## Related Work

Very few studies have addressed the use of LLMs in the access-control scenario: Tiwari et al. (2023) proposed using mixture of experts (MoE) in conjunction with training a separate model for each access level in order to support

users with multiple access rights. However, this approach relies solely on non-restricted documents, which substantially reduces its utility. Wutschitz et al. (2023) proposed using retrieval-augmented generation (RAG) with access rights, which prevents the retrieval of unauthorized documents. As illustrated by Tiwari et al. (2023), using RAG by itself, without training the model on the data, may be insufficient to achieve a substantial adjustment of LLM behaviors, such as altering writing styles or tones and incorporating new domain-specific knowledge and terminology. We consider RAG to be complementary to DOMBA, with the choice between them, or a combination of both, being highly dependent on the specific task and dataset. Determining the optimal approach for each case is beyond the scope of this work.

Several studies have explored the use of differential private training algorithms for deep learning (Yu et al. 2022; Abadi et al. 2016). These algorithms are designed to provide privacy guarantees when each “secret” appears in only one or a few data records. As demonstrated by Ginart et al. (2022), scaling these algorithms to safeguard a partition of sensitive documents on the same topic (which is critical in the access control scenario) results in impractical utility. Differential privacy (DP) federated learning (Geyer, Klein, and Nabi 2018) is a specific application of DP in deep learning in which contributions from different clients are aggregated during training, and noise is added to maintain DP. While federated learning might seem promising for the access-control scenario (with each access level treated as a client), Wei et al. (2020) demonstrated that performance drops substantially when using a small amount of clients (access levels). In contrast, DOMBA inherently does not depend on the number of access levels.

Papernot et al. (2017) proposed PATE, a differential private framework for training machine learning models for classification tasks. PATE is not suitable when there is a large amount of classes, and it requires an unlabeled non-private dataset. Given that all next-token prediction datasets are naturally labeled, and the number of tokens (which could be seen as classification classes) is very large for modern LLMs, PATE is not applicable in our case.

Ginart et al. (2022) introduced SUBMIX, an inference-time partition-level differential private model ensemble, however their ensemble requires many models to provide meaningful privacy guarantees, resulting in both costly text generation and a degradation in utility.

Beyond DP-based approaches, various techniques have been proposed to empirically mitigate the exposure of sensitive information. Such approaches include sanitization (Lison et al. 2021), prompt engineering (Chen et al. 2023), reinforcement learning from human feedback (RLHF) (Ouyang et al. 2022) and using a privacy regularized loss function (Miresghallah et al. 2021). However, these techniques lack theoretical guarantees, and some have been demonstrated to be vulnerable to attacks such as data extraction and “jail-breaking” (Liu et al. 2024; Chen et al. 2023; Nasr et al. 2023; Wolf et al. 2024).

## Methodology

In this section, we define the concept of exposing a secret and describe DOMBA’s training and aggregation processes. Using formal mathematical language, we establish a bound to DOMBA’s exposure of secrets. We also show that no other aggregation method could ever achieve a better bound.

### Training

DOMBA-INIT: Let  $d_1, \dots, d_k$  be the datasets corresponding to access levels  $1, \dots, k$ . We randomly assign each access level to one of two data partitions and train a submodel on each partition separately, denoted as  $M_1$  and  $M_2$ .

DOMBA-FT: For each access level  $AL$ , let  $M_1$  and  $M_2$  be the resulting submodels of DOMBA-INIT. If  $AL$  was assigned to  $M_1$  during DOMBA-INIT, we fine-tune  $M_2$  on  $d_{AL}$ . Otherwise, we fine-tune  $M_1$  on  $d_{AL}$ . We then save the states of  $M_1$  and  $M_2$ , which will be used during inference for users with access level  $AL$ . If a user has multiple access rights, an MoE can be used as demonstrated by Tiwari et al. (2023) (this scenario is not explored in this study). We note that PEFT (parameter-efficient fine-tuning) methods such as LORA (Hu et al. 2021) can be used to efficiently train and store the different states of  $M_1$  and  $M_2$ .

### Preliminaries

Let  $\Sigma$  be a set of tokens, and let  $n = |\Sigma|$ . We use  $t$  to refer to a token and  $c$  to refer to a context (i.e., a sequence of tokens preceding  $t$ ). We use  $M, M_1, M_2$  to refer to next-token prediction language models and denote the probability assigned by  $M$  to token  $t$  given context  $c$  as  $p_M(t|c)$ . We use  $\sum$  (sum) without specification to indicate summation over all tokens (i.e.,  $\sum_{t \in \Sigma}$ ). We note that the theorems in the subsections that follow commonly refer to arbitrary language models  $M_1$  and  $M_2$ , but it may be helpful to think of  $M_1$  and  $M_2$  as the outputs of DOMBA-INIT or DOMBA-FT.

### Token Exposure

We begin by defining exposing a secret in a formal sense. As highlighted by Brown et al. (2022), secrecy is relative – something is deemed secret if it is known by some but

unknown to others. Therefore, our concept of secrecy involves comparing probabilities assigned to a token by two models. One possible approach is to use the ratio of the probabilities assigned by the models to assess secrecy. However, this method has drawbacks. Consider the following probability distributions over the tokens  $a, b, c, d$ :  $p_1 = (0.7, 0.1, 0.1, 0.1)$  and  $p_2 = (0.97, 0.01, 0.01, 0.01)$ . The probabilities’ ratios ( $p_1/p_2$ ) are  $(0.72, 10, 10, 10)$ . This implies that tokens  $b, c$ , and  $d$  are “secret” in  $p_1$  compared to  $p_2$ . However, it seems more appropriate to consider  $a$  as secret in  $p_2$  compared to  $p_1$ , because  $p_2$  assigns  $a$  a probability that is 97 higher than all other tokens, whereas  $p_1$  assigns it a probability that is only seven times higher. To address this, we compare the probability ratio of a token  $t$  (between two models) to a “typical probability ratio” (TPR).

**Definition 1** (Geometric mean). *Let  $f : \Sigma \rightarrow \mathbb{R}^+$ . The geometric mean of  $f$  is  $GM(f(t)) := \exp(\frac{1}{n} \sum \log(f(t)))$ .*

**Definition 2** (TPR). *Let  $c$  be a context, and let  $M_1, M_2$  be language models. We define the “TPR at  $c$ ” of  $M_1, M_2$  as  $tpr_c(M_1, M_2) = GM(\frac{p_{M_1}(t|c)}{p_{M_2}(t|c)})$ .*

**Definition 3** (Token exposure). *Let  $c$  be a context,  $t$  be a token, and  $M_1, M_2$  be language models. We call  $t$  “ $\alpha$ -exposed by  $M_1$  over  $M_2$  at  $c$ ” if  $\frac{p_{M_1}(t|c)}{p_{M_2}(t|c) \cdot tpr_c(M_1, M_2)} = \alpha$ . We also say that  $t$  is “ $\leq \alpha$ -exposed” if  $t$  is  $\beta$ -exposed for some  $\beta \leq \alpha$ .*

In other words, instead of directly dividing  $p_{M_1}(t|c)$  by  $p_{M_2}(t|c)$ , we adjust  $p_{M_2}(t|c)$  by multiplying it by the TPR. In the example discussed, the TPR is 5.18, which results in the following exposures of tokens  $a, b, c, d$ :  $M_1$  over  $M_2$ :  $(0.14, 1.93, 1.93, 1.93)$  and  $M_2$  over  $M_1$ :  $(7.19, 0.52, 0.52, 0.52)$ . These values better reflect our intuition that  $a$  is secret and not  $b, c$ , and  $d$ .

### Token Exposure Properties

In this subsection, we explore certain properties of token exposure that are essential for later discussions.

**Definition 4** (Typical and relative probability). *Let  $c$  be a context, and let  $M$  be a language model. We define the “typical probability at  $c$ ” of  $M$  as  $tp_c(M) = GM(p_M(t|c))$ . Let  $t$  be a token, we further define the “relative probability of  $t$  at  $c$  by  $M$ ” as  $rp_M(t|c) := \frac{p_M(t|c)}{tp_c(M)}$ .*

**Lemma 1.**  $tpr_c(M_1, M_2) = \frac{tp_c(M_1)}{tp_c(M_2)}$ .

*Proof.*  $tpr_c(M_1, M_2) = \exp(\frac{1}{n} \sum \log(\frac{p_{M_1}(t|c)}{p_{M_2}(t|c)})) = \frac{\exp(\frac{1}{n} \sum \log(p_{M_1}(t|c)))}{\exp(\frac{1}{n} \sum \log(p_{M_2}(t|c)))} = \frac{tp_c(M_1)}{tp_c(M_2)}$ .  $\square$

By this lemma,  $\alpha$ -exposed is equivalent to  $\frac{rp_{M_1}(t|c)}{rp_{M_2}(t|c)} = \alpha$ .

**Lemma 2** (Token exposure multiplicity). *If  $t$  is  $\alpha$ -exposed by  $M_1$  over  $M_2$  at  $c$  and  $\beta$ -exposed by  $M_2$  over  $M_3$  at  $c$ , then  $t$  is  $\alpha\beta$ -exposed by  $M_1$  over  $M_3$  at  $c$ .*

*Proof.*  $\frac{rp_{M_1}(t|c)}{rp_{M_3}(t|c)} = \frac{rp_{M_1}(t|c)}{rp_{M_2}(t|c)} \cdot \frac{rp_{M_2}(t|c)}{rp_{M_3}(t|c)} = \alpha\beta$ .  $\square$

## Aggregation

In this subsection, we provide a formal definition of the notion of a min-bounded function and describe how DOMBA aggregates two submodels.

**Definition 5** (Proper-avg function). *Let  $f : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ . We call  $f$  a proper-avg function if  $\forall x, y : \min(x, y) \leq f(x, y) \leq \max(x, y)$ .*

**Definition 6** (Min-bounded function). *Let  $f$  be a proper-avg function. we call  $f$  min-bounded if  $\forall x, y, f(x, y) \leq \lambda_f \min(x, y)$  for some constant  $\lambda_f$ .*

In practice we use the generalized mean (Sykora 2009) with  $\alpha < 0$  for min-bounded functions, that is,  $f(x, y) = (\frac{1}{2}(x^\alpha + y^\alpha))^{\frac{1}{\alpha}}$ ,  $\lambda_f = 2^{-\frac{1}{\alpha}}$ . Two special cases are:

1.  $\alpha \rightarrow -\infty$  (Minimum):  $f(x, y) = \min(x, y)$ ,  $\lambda_f = 1$ .
2.  $\alpha = -1$  (Harmonic mean):  $f(x, y) = \frac{2xy}{x+y}$ ,  $\lambda_f = 2$ .

Note that the arithmetic mean ( $\frac{x+y}{2}$ ) is not min-bounded.

**Definition 7** (DOMBA aggregation). *Let  $M_1, M_2$  be language models, and let  $f$  be a min-bounded function. We define  $DAGG_f(M_1, M_2)$  (denoted as  $M$ ) as a model that assigns probabilities as follows:  $p_M(t|c) = \frac{M(t|c)}{\sum_{t' \in \Sigma} M(t'|c)}$ , where  $M(t|c) = f(rp_{M_1}(t|c), rp_{M_2}(t|c))$ .*

We note that DOMBA uses  $f$  to average the **relative probabilities**. In contrast, averaging the probabilities would lead to inferior bounds in the subsequent subsection.

## Bounding the Exposure of DOMBA

In this subsection, we establish the bounds on DOMBA's exposure over both submodels (Theorem 2) as well as over any other model (Corollary 1). We begin by introducing several definitions and lemmas that will be used for proving the main theorem later on.

**Definition 8.** *Let  $M_1, M_2$  be language models, and let  $f$  be a min-bounded function. Let  $M = DAGG_f(M_1, M_2)$ . We define  $\bar{f}_c(M_1, M_2) = GM(M(t|c)^{-1})$ .*

While it might be unclear how to interpret  $\bar{f}_c(M_1, M_2)$ , it relates to a notion of "mean exposure" between  $M_1$  and  $M_2$ :

**Definition 9** (Mean absolute exposure). *Let  $c$  be a context and  $M_1, M_2$  be language models. We define the "mean absolute exposure between  $M_1$  and  $M_2$  at  $c$ " as  $MAE_c(M_1, M_2) = GM(\max(\frac{rp_{M_1}(t|c)}{rp_{M_2}(t|c)}, \frac{rp_{M_2}(t|c)}{rp_{M_1}(t|c)}))$ .*

**Lemma 3.**  $\bar{f}_c(M_1, M_2) \leq \sqrt{MAE_c(M_1, M_2)}$ .

*Proof.* Let  $x := \sum \log(\min(rp_{M_1}(t|c), rp_{M_2}(t|c)))$ ,  $y := \sum \log(\max(rp_{M_1}(t|c), rp_{M_2}(t|c)))$ . We observe that  $x + y = \sum \log(rp_{M_1}(t|c)) + \sum \log(rp_{M_2}(t|c)) = 0 + 0 = 0$ . by definition,  $y - x = n \cdot \log(MAE_c(M_1, M_2))$ , which implies,  $x = -\frac{n}{2} \log(MAE_c(M_1, M_2))$ , we conclude that  $\bar{f}_c(M_1, M_2) \leq \exp(-\frac{x}{n}) = \sqrt{MAE_c(M_1, M_2)}$ .  $\square$

**Lemma 4.**  $rp_M(t|c) = M(t|c) \cdot \bar{f}_c(M_1, M_2)$ .

*Proof.*  $rp_M(t|c) = \frac{p_M(t|c)}{tp_c(M)} = \frac{p_M(t|c)}{\exp(\frac{1}{n} \sum \log(p_M(t|c)))} = \frac{M(t|c)}{\exp(\frac{1}{n} \sum \log(M(t|c)))} = M(t|c) \cdot \bar{f}_c(M_1, M_2)$   $\square$

In the following theorem, we provide a lower bound to the minimum token exposure achievable over two models.

**Theorem 1.** *Let  $c$  be a context and  $M, M_1, M_2$  be language models. There exists a token  $t$  that is  $\geq \sqrt{MAE_c(M_1, M_2)}$ -exposed by  $M$  over either  $M_1$  or  $M_2$ .*

*Proof.* By the proof of lemma 3:  $\sqrt{MAE_c(M_1, M_2)} = \exp(-\frac{x}{n}) = GM(\frac{rp_M(t|c)}{\min(rp_{M_1}(t|c), rp_{M_2}(t|c))})$ . The right end side is an average over tokens. Therefore there exists a token for which the term inside is greater than or equal to the left end side, which finishes the proof.  $\square$

In the following theorem, we demonstrate that using DOMBA provides a bound on the exposure that is a constant multiple of the best possible bound (Theorem 1). This constant is solely dependent on  $f$  and can even reach a value of 1 ( $f = \text{Minimum}$ ).

**Theorem 2.** *Let  $f$  be a min-bounded function,  $t$  a token and  $M = DAGG_f(M_1, M_2)$ .  $t$  is  $\leq \gamma$ -exposed over both  $M_1$  and  $M_2$  for  $\gamma = \lambda_f \bar{f}_c(M_1, M_2) \leq \lambda_f \sqrt{MAE_c(M_1, M_2)}$ .*

*Proof.*  $\frac{rp_M(t|c)}{rp_{M_1}(t|c)} = \frac{M(t|c) \cdot \bar{f}_c(M_1, M_2)}{rp_{M_1}(t|c)} \leq \frac{\lambda_f \min(rp_{M_1}(t|c), rp_{M_2}(t|c)) \cdot \bar{f}_c(M_1, M_2)}{rp_{M_1}(t|c)} \leq \lambda_f \bar{f}_c(M_1, M_2)$ .  $\square$

We note that by assuming  $M_1$  and  $M_2$  assign similar relative probabilities to most tokens, we can anticipate the mean absolute exposure to be low. Essentially, we achieve average case behavior for all tokens.

In the following corollary, we informally think of  $M_b$  as our base model (although the corollary holds in general).

**Corollary 1.** *Let  $c$  be a context, and let  $M_1, M_2, M_b$  be language models. Let  $t$  be a token that is  $\alpha$ -exposed by  $M_1$  over  $M_b$  at  $c$  and  $\beta$ -exposed by  $M_2$  over  $M_b$  at  $c$ . Let  $M := DAGG_f(M_1, M_2)$ . Then  $t$  is  $\leq \gamma \min(\alpha, \beta)$ -exposed by  $M$  over  $M_b$  at  $c$  for  $\gamma$  as in theorem 2.*

*Proof.* Follows directly from lemma 2 and theorem 2.  $\square$

Stating the corollary in other words, if we fix a context  $c$ , for any token  $t$ , the exposure of the aggregated model  $M$  over any model  $M_b$  is bounded by the minimum of the exposures of the submodels  $M_1$  and  $M_2$  over  $M_b$ , multiplied by a small value. This implies that if the exposure of a token  $t$  by either submodel over  $M_b$  is small (i.e.,  $t$  is not substantially exposed by at least one submodel), then the exposure by the aggregated model over  $M_b$  cannot be too large (i.e.,  $t$  will not be substantially exposed by the aggregated model). Given that in DOMBA, each submodel is trained on separate access levels, and assuming that access levels with shared secrets are assigned to the same partition, it is expected that each secret will not be substantially exposed by at least one of the submodels, and thus, DOMBA will provide a defense against the exposure of these secrets.

## Evaluation

### Datasets

Since access-controlled datasets are not publicly available, we required datasets that mimic the access-control scenario. These datasets need to be divided into different topics (which serve as access levels), with many data records per topic. Additionally, to use two of our security evaluation metrics, the data records should contain phrases that we refer to as “sensitive-mimicking phrases” – phrases unique to the topic that could be considered sensitive / secret.

**Movie Reviews** The first dataset we utilized is the IMDB Spoiler reviews dataset (Misra 2019). We randomly selected 50 reviews of different movies released after 2015 and considered the movie of each review selected as an access level. Then, we collected all of the reviews for each of the 50 movies. We note that some reviews contain details about the movie’s plot, cast members, or characters, which mimic sensitive information. We utilized the Movies Metadata dataset (Banik 2018) to retrieve cast members’ names and used them as sensitive-mimicking phrases. The number of reviews totaled 22,742, with 10% of each movie’s reviews set aside for evaluation. The number of reviews per movie ranged between 160 and 751.

**Recipes** The second dataset used is the Food.com Recipes and Interactions dataset (Li 2019). We utilized class labels of the Food-101 dataset (Bossard, Guillaumin, and Van Gool 2014) to partition the recipes into multiple sets. Each set includes recipes with titles containing a specific class label (e.g., pizza). We selected the 10 most frequent classes as the access levels. We note that the recipes include specific details about the process of creating each dish, which can mimic, for example, sensitive detailed descriptions of product manufacturing processes. We use the ingredients of each recipe as sensitive-mimicking phrases. However, since some ingredients are common among many classes, we only consider ingredients that appear in recipes of a certain class with a frequency at least 10 times greater than the frequency in all of the recipes. The number of recipes totaled 10,829, with 10% of each class put aside for evaluation. The number of recipes per class ranged between 408 and 2283.

### Training

For training we used LORA (Hu et al. 2021). LORA is a fine-tuning technique that uses a small number of trainable parameters. Training is relatively fast with LORA, and the resulting model requires minimal storage space. These qualities were crucial for our experiments, as we conducted numerous trials with limited computational resources. However, it is important to note that the theoretical analysis is not dependent on the training method, and we anticipate that the experiments can be replicated using other training techniques as well. The base model used was OpenAI-GPT (Radford et al. 2018) which has 117 million parameters and a vocabulary size of 40,478. This model’s original training data is a books dataset from 2015 (Zhu et al. 2015). This limits the prior knowledge the model possesses regarding movies and recipes. Since recent LLMs are trained

on more recent and diverse datasets, evaluating them on sensitive information from the movie and recipe datasets would be challenging, as the models are probably familiar with some of the information. We note that although recent LLMs are larger and perform better than OpenAI-GPT, many are still based on the same underlying principles. Our theoretical analysis and proposed approach generalizes to any language model based on next-token prediction and does not rely on the specifics of any particular architecture. Regarding training parameters, we conducted experiments with varying numbers of training epochs (1, 2, and 4). The hyperparameters for LORA were set to default values and were not explored:  $r=64$ ,  $lora\_alpha=32$ ,  $lora\_dropout=0.05$ ,  $optimizer=paged\_adamw\_32bit$ ,  $learning\ rate=5e-4$ , and  $warmup\_ratio=0.03$ . All experiments were conducted on an NVIDIA A100-SXM4-40GB GPU.

### Compared Models

**Non-secure models (NSec)** In these models, which serve as baselines, no attempt is made to secure sensitive information. **FT-ALL**: OpenAI-GPT fine-tuned on the entire training dataset. **AGG-A**: Similar to DOMBA-INIT, but using arithmetic mean, a non-min-bounded function.

**Secure models (Sec)** While these models are trained on all the data with an effort made to secure sensitive information, they do not include an access-control mechanism. **SUB-MIX**: A DP aggregated model constructed using the method of Ginart et al. (2022), with three submodels (two parts + the base model). For a meaningful comparison, we tuned the privacy parameter  $\beta$  to 0.3, which resulted in utility comparable to DOMBA on the movies dataset. **D-I-H**: DOMBA-INIT (without DOMBA-FT), using harmonic mean for aggregation. **D-I-M**: DOMBA-INIT (without DOMBA-FT), using minimum for aggregation.

**Access-controlled models (AC)** These models are designed to secure sensitive information while providing an access-control mechanism. **Per-AL**: A separate model for each access level, achieved by fine-tuning OpenAI-GPT only on data records of that access level. **DOMBA**: Our full method, using the minimum function for aggregation.

### Metrics

In this section, we describe the metrics used to evaluate the models’ utility and security. For utility we use perplexity, which measures the model’s ability to predict the next token in a text. For security we use four different metrics: exposure, secret perplexity, a secret inference attack AUC-ROC, and the canary technique score (Carlini et al. 2019). We note that for access-controlled models, we evaluate the security of each variant (corresponding to an access level) using data with a different access level than the one that the variant was trained for.

**Utility Evaluation** We evaluate utility in terms of perplexity on two evaluation sets as follows: 1. **HOPPL**: perplexity on held out data with access levels that were not used for training. This metric provides a “fair” way of comparing secure and non-secure models, as the non-secure models are

Type	Model	Utility metrics				Security metrics							
		HOPPL ↓		AUPPL ↓		EXP ↓		SPPL ↑		SIA ↓		CAN ↓	
		R	M	R	M	R	M	R	M	R	M	R	M
NSec	FT-ALL	19.9	48.39	15.61	41.55	-	-	15.31	61.81	0.81	0.83	14.47	28.03
	AGG-A	22.34	49.37	17.83	43.55	207.9	1699	19.78	78.93	0.82	0.82	14.82	23.88
Sec	SUBMIX	29.58	50.84	25.4	48.54	5.1	17.36	49.95	550	0.76	0.8	4.73	4.01
	D-I-H	<b>23.57</b>	<b>50.83</b>	<u>20.09</u>	<u>48.37</u>	2.38	2.84	50.87	895.1	0.64	0.66	3.49	<u>2.43</u>
	D-I-M	<u>24.54</u>	51.89	21.14	49.79	<b>1.77</b>	<b>2.31</b>	<u>61.99</u>	<b>1161</b>	<u>0.6</u>	<u>0.64</u>	<u>3.21</u>	<b>2.21</b>
AC	PER-AL	45.15	63.39	27.87	54.18	-	-	-	-	-	-	-	-
	DOMBA	25.19	52.22	<b>16.85</b>	<b>42.48</b>	<u>1.78</u>	<u>2.37</u>	<b>74</b>	<u>1127</u>	<b>0.54</b>	<b>0.62</b>	<b>2.88</b>	2.44

Table 1: Results with two epochs of training. For each metric and model, the results for both the recipe dataset (R) and movie reviews dataset (M) are presented. The best values for secure and access-controlled models are in bold, and the second best values are underlined. (Note that PER-AL is trivially secure and therefore, no results are presented for its security; the EXP metric is only meaningful for methods which aggregate two submodels and therefore, it is not presented for FT-ALL.)

not expected to gain by “knowing” restricted information. 2. **AUPPL**: perplexity on held out data of the access levels used for training (for access-controlled models - the corresponding variant is used for each access level). The main purpose of this metric is to compare the utility of secure and access-controlled models. We expect the access-controlled models to gain utility by “knowing” authorized restricted information. For both metrics above, we calculate the perplexity as:  $perp_M(D_e) = \exp(\frac{1}{|D_e|} \sum_{r \in D_e} \sum_i -\log(p_M(r_i|r_{<i})))$ , where  $|D_e|$  is the amount of tokens in  $D_e$ ,  $r$  is a data record,  $r_i$  is the  $i$ 'th token in the record,  $r_{<i}$  are the tokens preceding it, and  $p_M$  is the probability assigned by the model.

**Exposure (EXP)** In our theoretical analysis (Theorem 2), we established that the exposure of  $M = DAGG_f(M_1, M_2)$  over both  $M_1$  and  $M_2$  is bounded for any token by  $\lambda_f \bar{f}_c(M_1, M_2) \leq \lambda_f \sqrt{MAE_c(M_1, M_2)}$ . To validate this, we measure “extreme case” exposure of  $M$  over  $M_1$  and  $M_2$ . We report the maximum and 99th percentile exposure ( $= \frac{rp_M(t|c)}{\min(rp_{M_1}(t|c), rp_{M_2}(t|c))}$ ) for all tokens observed in the data, given the previous tokens as context.

**Secret Perplexity (SPPL)** One way of measuring the model’s ability to handle sensitive information is by evaluating perplexity specifically on sensitive-mimicking phrases. Given a model  $M$ , we measure the perplexity of each instance of a sensitive-mimicking phrase in the evaluation dataset. Specifically, let  $x := x_1, \dots, x_k$  be the token representations of a sensitive-mimicking phrase and  $c$  be the tokens preceding this phrase, we measure  $perp_M(x|c) = \exp(\frac{1}{k} \sum_i -\log(p_M(x_i|c, x_1, \dots, x_{i-1})))$ . We report the average of the mean perplexity of each access level. This metric aims to provide a basic, rough evaluation of a model’s ability to handle sensitive information.

**Secret Inference Attack (SIA)** This attack is based on a membership inference attack with a reference model (Miresghallah et al. 2022; Kumar Murakonda, Shokri, and Theodorakopoulos 2021). The original attack works as follows: Given a reference model  $M_b$ , a target model  $M$ , and a potential training data record  $r$  of  $M$ , measure the log ratio of the probabilities of  $r$  according to  $M$  and  $M_b$ , that

is  $\log(\frac{p_M(r)}{p_{M_b}(r)})$ . If this value is above a certain threshold, consider  $r$  as belonging to the training data of  $M$ . In our scenario, instead of inferring the membership of any data record, the adversary tries to infer secrets. Therefore, we only consider probabilities assigned to sensitive-mimicking phrases: cast members’ names for the movie reviews dataset and secret ingredients for the recipe dataset. The attack dataset consists of tuples  $(c, t, label)$ , where  $c$  is a context,  $t$  is a phrase, and  $label$  is *true* if  $t$  is sensitive and *false* otherwise. To obtain data points labeled *false*, we replace each sensitive-mimicking phrase  $t$  by  $t'$ , which is another phrase of the same type (cast member name or ingredient) that is not a sensitive-mimicking phrase. For every data point  $(c, t, true)$ , we have a data point  $(c, t', false)$ . We report the AUC-ROC of the attack.

**The Canary Technique (CAN)** We adapt the attack proposed by Carlini et al. (2019) to the access-control scenario. For each access level, we insert 30 repetitions of a phrase (canary) consisting of seven randomly chosen words into the training set for that access level (the number of repetitions and phrase length were selected arbitrarily). This canary mimics sensitive information for the access level. We report the median attack score across access levels. An attack score of  $s$  means that only  $(\frac{1}{2})^s$  of phrases of the same length have a higher probability of being generated by the model. A score near one suggests that the model did not memorize the canary.

## Results

The results with two epochs of training are shown in Table 1. As expected, FT-ALL achieved the best utility across both metrics. Among secure and access-controlled models, D-I-H had the highest HOPPL utility, while DOMBA excelled on the AUPPL metric, demonstrating the value of the DOMBA-FT step. Comparing access-controlled models, DOMBA substantially outperformed PER-AL across both datasets and both utility metrics.

Regarding security, non-secure models performed substantially worse, compared to secure models, on all metrics. Among the secure and access-controlled models, SUBMIX obtained the worst values for all metrics and datasets, D-I-




Figure 2: Exposure (log scale, lower is better) of different models with 1, 2, or 4 training epochs. The colored bars represent the 99th percentile exposure, while the dashed bars represent the maximal observed exposure.

M and DOMBA obtained the best values, and D-I-H was slightly worse. Although secure models provide substantially better security compared to non-secure models, they are not perfect. For instance, a perfectly secure model would score 0.5 on SIA and one on CAN. This does not imply that secure models are not actually secure. For example, the values obtained by all secure and access-controlled models for the canary technique metric are considered impractical for extracting useful information (Carlini et al. 2019).

Regarding the impact of the choice of min-bounded function, D-I-M achieved better security and worse utility compared to D-I-H, as expected. We also tested replacing the min-bounded function used in DOMBA from minimum to harmonic mean and the effect was similar.

Figure 2 shows the worst-case and 99th percentile exposure of models employing different aggregation methods on the recipe and movie reviews datasets for 1, 2, and 4 training epochs. The maximum exposure of DOMBA, D-I-H, and D-I-M is 4.69. In comparison, SUBMIX reaches a maximum exposure of  $8.5e4$  and AGG-A reaches a maximum exposure of  $1.3e10$ . We observe that DOMBA’s 99th percentile exposure is similar to its maximum exposure, supporting the theoretical bound established by our analysis (Theorem 2). Regarding the effect of the number of epochs, increasing it generally leads to higher exposure. However, the increase in exposure for DOMBA, D-I-H, and D-I-M is moderate compared to AGG-A for both datasets, while for SUBMIX, the change in exposure is inconsistent between the two datasets.

Figure 3 illustrates the trade-off between utility and security for different methods across both datasets. For most models, as the number of training epochs increases, security tends to worsen while utility improves. However, non-secure models experience a much greater decline in security. DOMBA achieves the best trade-off, providing superior security while maintaining utility levels similar to those of the non-secure models.

## Discussion

To analyze DOMBA, we developed the token exposure framework, offering an alternative to DP for theoretically analyzing the security of language models. The key idea behind token exposure is to evaluate the difference between a safe model and a potentially unsafe one. In contrast, DP evaluates the differences between models trained on nearly




Figure 3: Utility-security trade-off (for both metrics lower is better, with the security metric on a log scale). Each line represents a model trained for different numbers of epochs (1, 2, and 4). The point representing 1 epoch is always the leftmost and highest point on the line.

identical datasets (in both approaches, a smaller difference indicates stronger security). Additionally, token exposure utilizes relative probabilities over absolute ones. We find the token exposure approach more intuitive and practical for analyzing DOMBA. Future work could further explore the token exposure framework and investigate its relationships and potential synergies with DP.

DOMBA relies on a strict separation of access levels into two distinct partitions that do not share sensitive information. Achieving this separation requires the graph of shared secrets between access levels to be disconnected. We hypothesize that this condition is usually met, as organizations often handle documents across diverse topics. Future research could validate this hypothesis and investigate approaches for scenarios where the graph remains connected.

While the training time of DOMBA scales linearly with the dataset size, similar to training a single LLM, and storage requirements can be kept low by using LORA, inference incurs additional resource overhead due to the deployment of two LLMs instead of one. This overhead may render the approach impractical for certain applications. One potential solution is to employ DOMBA as a teacher model to train a student model via knowledge distillation (Xu et al. 2024; Papernot et al. 2017), where the student model serves as a deployed model mimicking DOMBA.

## Conclusion

In this paper we proposed DOMBA, a novel approach for training and deploying access-controlled LLMs with high utility. We formalized the concept of exposed secrets by developing the token exposure framework and bounded DOMBA’s exposure. We evaluated DOMBA’s performance on two access-controlled datasets, mimicking real world organizations’ needs. Our evaluation showed that DOMBA achieves a better security-utility trade-off than existing methods, across both datasets, two utility metrics and four security metrics. Finally, we believe that the principles of min-bounded aggregation and relative probabilities, which serve as DOMBA’s core, have substantial potential to serve as foundational elements in a wide range of future machine learning research, extending beyond the scope of security.

## Acknowledgments

The authors would like to thank Dr. Avi Segal for their valuable insights and thoughtful feedback, which contributed to the refinement of this manuscript.

## References

- Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H. B.; Mironov, I.; Talwar, K.; and Zhang, L. 2016. Deep Learning with Differential Privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, 308–318. New York, NY, USA: Association for Computing Machinery. ISBN 9781450341394.
- Banik, R. 2018. The Movies Dataset. <https://www.kaggle.com/datasets/rounakbanik/the-movies-dataset>. Accessed: 2024-03-03.
- Bossard, L.; Guillaumin, M.; and Van Gool, L. 2014. Food101 – Mining Discriminative Components with Random Forests. In *European Conference on Computer Vision*.
- Brown, H.; Lee, K.; Mireshghallah, F.; Shokri, R.; and Tramèr, F. 2022. What Does it Mean for a Language Model to Preserve Privacy? In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, 2280–2292. New York, NY, USA: Association for Computing Machinery. ISBN 9781450393522.
- Carlini, N.; Liu, C.; Erlingsson, Ú.; Kos, J.; and Song, D. 2019. The Secret Sharer: Evaluating and Testing Unintended Memorization in Neural Networks. In *28th USENIX Security Symposium (USENIX Security 19)*, 267–284. Santa Clara, CA: USENIX Association. ISBN 978-1-939133-06-9.
- Carlini, N.; Tramèr, F.; Wallace, E.; Jagielski, M.; Herbert-Voss, A.; Lee, K.; Roberts, A.; Brown, T.; Song, D.; Erlingsson, Ú.; Oprea, A.; and Raffel, C. 2021. Extracting Training Data from Large Language Models. In *30th USENIX Security Symposium (USENIX Security 21)*, 2633–2650. USENIX Association. ISBN 978-1-939133-24-3.
- Chen, Y.; Mendes, E.; Das, S.; Xu, W.; and Ritter, A. 2023. Can Language Models be Instructed to Protect Personal Information? arXiv:2310.02224.
- Geyer, R. C.; Klein, T.; and Nabi, M. 2018. Differentially Private Federated Learning: A Client Level Perspective. arXiv:1712.07557.
- Ginart, A.; van der Maaten, L.; Zou, J.; and Guo, C. 2022. Submix: Practical Private Prediction for Large-Scale Language Models. arXiv:2201.00971.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. LoRA: Low-Rank Adaptation of Large Language Models. arXiv:2106.09685.
- Kandpal, N.; Pillutla, K.; Oprea, A.; Kairouz, P.; Choquette-Choo, C. A.; and Xu, Z. 2024. User Inference Attacks on Large Language Models. arXiv:2310.09266.
- Kumar Murakonda, S.; Shokri, R.; and Theodorakopoulos, G. 2021. Quantifying the Privacy Risks of Learning High-Dimensional Graphical Models. In Banerjee, A.; and Fukumizu, K., eds., *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, 2287–2295. PMLR.
- Li, S. 2019. Food.com Recipes and Interactions.
- Lison, P.; Pilán, I.; Sanchez, D.; Batet, M.; and Øvrelid, L. 2021. Anonymisation Models for Text Data: State of the art, Challenges and Future Directions. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 4188–4203. Online: Association for Computational Linguistics.
- Liu, Y.; Deng, G.; Xu, Z.; Li, Y.; Zheng, Y.; Zhang, Y.; Zhao, L.; Zhang, T.; Wang, K.; and Liu, Y. 2024. Jailbreaking ChatGPT via Prompt Engineering: An Empirical Study. arXiv:2305.13860.
- Mireshghallah, F.; Goyal, K.; Uniyal, A.; Berg-Kirkpatrick, T.; and Shokri, R. 2022. Quantifying Privacy Risks of Masked Language Models Using Membership Inference Attacks. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 8332–8347. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Mireshghallah, F.; Inan, H.; Hasegawa, M.; Rühle, V.; Berg-Kirkpatrick, T.; and Sim, R. 2021. Privacy Regularization: Joint Privacy-Utility Optimization in Language Models. In Toutanova, K.; Rumshisky, A.; Zettlemoyer, L.; Hakkani-Tur, D.; Beltagy, I.; Bethard, S.; Cotterell, R.; Chakraborty, T.; and Zhou, Y., eds., *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 3799–3807. Online: Association for Computational Linguistics.
- Misra, R. 2019. IMDB Spoiler Dataset. <https://www.kaggle.com/datasets/rmisra/imdb-spoiler-dataset>. Accessed: 2024-03-03.
- Nasr, M.; Carlini, N.; Hayase, J.; Jagielski, M.; Cooper, A. F.; Ippolito, D.; Choquette-Choo, C. A.; Wallace, E.; Tramèr, F.; and Lee, K. 2023. Scalable Extraction of Training Data from (Production) Language Models. arXiv:2311.17035.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P. F.; Leike, J.; and Lowe, R. 2022. Training language models to follow instructions with human feedback. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 27730–27744. Curran Associates, Inc.
- Pan, X.; Zhang, M.; Ji, S.; and Yang, M. 2020. Privacy Risks of General-Purpose Language Models. In *2020 IEEE Symposium on Security and Privacy (SP)*, 1314–1331.
- Papernot, N.; Abadi, M.; Úlfar Erlingsson; Goodfellow, I.; and Talwar, K. 2017. Semi-supervised Knowledge Transfer

for Deep Learning from Private Training Data. In *Proceedings of the International Conference on Learning Representations*.

Radford, A.; Narasimhan, K.; Salimans, T.; and Sutskever, I. 2018. Improving language understanding by generative pre-training.

Sykora, S. 2009. *Mathematical Means and Averages: Basic Properties*. *Stan's Library*, 3.

Tiwari, T.; Gururangan, S.; Guo, C.; Hua, W.; Kariyappa, S.; Gupta, U.; Xiong, W.; Maeng, K.; Lee, H.-H. S.; and Suh, G. E. 2023. Information Flow Control in Machine Learning through Modular Model Architecture. arXiv:2306.03235.

Wei, K.; Li, J.; Ding, M.; Ma, C.; Yang, H. H.; Farokhi, F.; Jin, S.; Quek, T. Q. S.; and Vincent Poor, H. 2020. Federated Learning With Differential Privacy: Algorithms and Performance Analysis. *IEEE Transactions on Information Forensics and Security*, 15: 3454–3469.

Wolf, Y.; Wies, N.; Avnery, O.; Levine, Y.; and Shashua, A. 2024. Fundamental Limitations of Alignment in Large Language Models. arXiv:2304.11082.

Wutschitz, L.; Köpf, B.; Paverd, A.; Rajmohan, S.; Salem, A.; Tople, S.; Zanella-Béguelin, S.; Xia, M.; and Rühle, V. 2023. Rethinking Privacy in Machine Learning Pipelines from an Information Flow Control Perspective. arXiv:2311.15792.

Xu, X.; Li, M.; Tao, C.; Shen, T.; Cheng, R.; Li, J.; Xu, C.; Tao, D.; and Zhou, T. 2024. A Survey on Knowledge Distillation of Large Language Models. arXiv:2402.13116.

Yu, D.; Naik, S.; Backurs, A.; Gopi, S.; Inan, H. A.; Kamath, G.; Kulkarni, J.; Lee, Y. T.; Manoel, A.; Wutschitz, L.; Yekhanin, S.; and Zhang, H. 2022. Differentially Private Fine-tuning of Language Models. arXiv:2110.06500.

Zhu, Y.; Kiros, R.; Zemel, R.; Salakhutdinov, R.; Urtasun, R.; Torralba, A.; and Fidler, S. 2015. Aligning Books and Movies: Towards Story-Like Visual Explanations by Watching Movies and Reading Books. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.