

TAG: Gradient Attack on Transformer-based Language Models

Jieren Deng^{1§}, Yijue Wang^{1§}, Ji Li², Chenghong Wang³, Chao Shang⁴, Hang Liu⁵
Sanguthevar Rajasekaran¹, Caiwen Ding¹

¹University of Connecticut, ²Microsoft, ³Duke University, ⁴JD AI Research,

⁵Stevens Institute of Technology

{jieren.deng, yijue.wang, sanguthevar.rajasekaran, caiwen.ding}@uconn.edu

{changzhouliji, chaoshangcs}@gmail.com, chenghong.wang552@duke.edu, hliu77@stevens.edu

Abstract

Although distributed learning has increasingly gained attention in terms of effectively utilizing local devices for data privacy enhancement, recent studies show that publicly shared gradients in the training process can reveal the private training data (gradient leakage) to a third party. However, so far there hasn't been any systematic study of the gradient leakage mechanism of the Transformer based language models. In this paper, as the *first* attempt, we formulate the gradient attack problem on the Transformer-based language models and propose a gradient attack algorithm, TAG, to recover the local training data. Experimental results on Transformer, TinyBERT₄, TinyBERT₆, BERT_{BASE}, and BERT_{LARGE} using GLUE benchmark show that compared with DLG (Zhu et al., 2019), TAG works well on more weight distributions in recovering private training data and achieves $1.5\times$ Recover Rate and $2.5\times$ ROUGE-2 over prior methods without the need of ground truth label. TAG can obtain up to 88.9% tokens and up to 0.93 cosine similarity in token embeddings from private training data by attacking gradients on CoLA dataset. In addition, TAG is stronger than previous approaches on larger models, smaller dictionary size, and smaller input length.

1 Introduction

Distributed training has gained popularity in reducing training time on large-scale deep learning models and datasets (Dean et al., 2012; Li et al., 2014a,b; Baruch et al., 2019; Gurevin et al., 2021; Wu et al., 2020; Lin et al., 2020). In such systems, multiple devices or participators collaborate in training one task and synchronize via exchanging gradients, allowing participants at different location for model training with their own data. It is widely believed that sharing gradients

between participants will not leak the private training data. On the other hand, large scale contextual representation models, such as ELMo (Peters et al., 2018), BERT (Devlin et al., 2019), XLNet (Yang et al., 2019), T5 (Raffel et al., 2019a), and GPT-3 (Brown et al., 2020) have significantly promoted natural language processing (NLP) in the last decade. Thus, the ability of parallel training has helped propel using distributed learning on a large scale NLP, for efficient training.

Recent studies show that private training data can be recovered through the deep learning model by gradients (Zhu et al., 2019; Chen et al., 2020; Geiping et al., 2020; Wang et al., 2021). For instance, recent work "Deep Leakage from Gradient" (DLG) (Zhu et al., 2019) showed the shared gradients could leak private training data in image classification. Another recent work IG (Geiping et al., 2020) shows that it is possible to recover images by gradients in a trained network.

Despite the success, there are several limitations on current works: (i) Lack of generalizability on different weight distribution. They only succeeded in the early training phase at certain weight distribution. (ii) Lack of formal problem formulation and gradient attack evaluation in the field of NLP. Existing work only show the reconstruction difference at the sentence level without quantitative analysis. (iii) There is little investigation on the impact of different attention heads, model architectures on Transformer gradient attack.

In this paper, we propose a novel algorithm, to recover private training data of Transformer-based language model from the shared gradients. As shown in Figure 1, our TAG adversary obtains the transformer model gradients ∇W from a distributed learning participator and updates initialized dummy data (X' , Y') by comparing the difference between the participator's gradients ∇W and adversary's gradients $\nabla W'$. Eventually, the adversary recovers the dummy data (X' , Y') and acquires the informa-

[§]Equal contribution

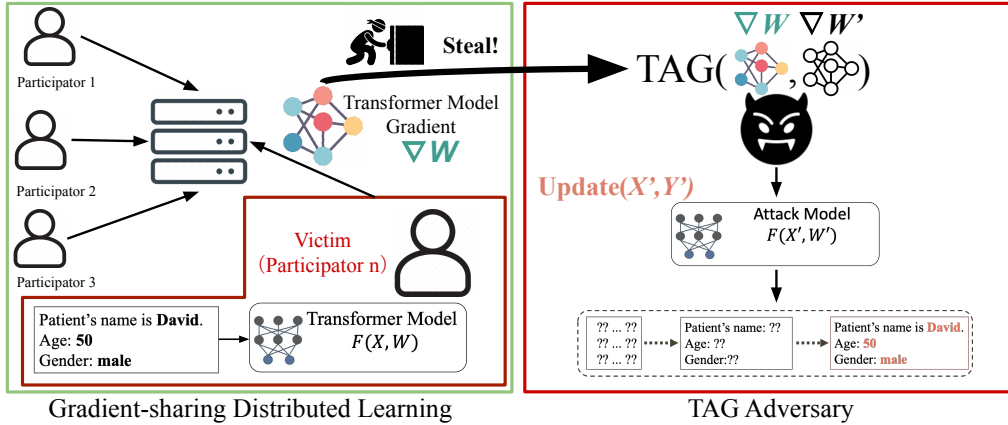


Figure 1: Gradient transformer attack process.

tion from the participant’s private training data X such as name, age, and gender. Our contributions are summarized as follows:

- As the *first* attempt in the field of NLP, we propose a general gradient attack algorithm, TAG, to recover the private training data on Transformer-based language models. Compared to the existing methods, TAG works on more realistic weight distributions, including both pre-trained models and normally initialized models.
- We develop a quantitative evaluation method on the NLP gradient attack problem while the existing work shows the recovered texts. We use a set of metrics (Recover Rate, ROUGE-1(%), ROUGE-2(%), ROUGE-L(%), and runtime) to evaluate the effectiveness of the proposed attack algorithm. With these metrics, TAG achieves $1.5 \times$ Recover Rate and $2.5 \times$ ROUGE-2 over prior methods. TAG can also obtain up to 88.9% tokens and up to 0.93 cosine similarity in token embeddings from private training data.
- We conduct a comprehensive analysis of different weight distribution, dataset, vocabulary dictionary size, and model size on Transformer, TinyBERT₄, TinyBERT₆, BERT_{BASE}, and BERT_{LARGE}, and we observe that TAG has a stronger adversary on large models than on small ones. In addition, models with a smaller dictionary size and smaller input sequence length are riskier in leaking the private training data.

2 Related Work

2.1 Privacy leakage problem

Privacy leakage is studied in the training phase and prediction phase. Privacy attack from gradient

and model inversion (MI) attack (Fredrikson et al., 2015) aim at the training phase by constructing the features of the private training data by using the correlation between the private training data and the model output. The authors in (Fredrikson et al., 2015) showed that it is possible to infer individual genomic data via access to a linear model for personalized medicine. Recent works extend MI attack to recover features of private training data of Deep Neural Networks (DNNs). Privacy attack from gradients is different from previous MI attack. It recovers the private training data exploiting their gradients in a machine learning model. The process of privacy leakage from gradients is shown at Figure 1.

2.2 Distributed learning

Distributed learning is a popular framework for large-scale model training (Das et al., 2016; Dean et al., 2012; Li et al., 2014a,b; Baruch et al., 2019) that leverage the computation power of many devices by aggregating the local models trained on the devices. Instead of training a model with all the data at a server, each device trains a local model with a different chunk of the dataset and shares the final gradients. A popular distributed learning algorithm is *Synchronous Stochastic Gradient Descent* (sync-SGD) (Li et al., 2014a,b) which contains a single server and n local devices. Each device trains a local model and shares the gradient with the server. The server then aggregates the gradients of the different devices and starts another round by sharing the aggregated result with the devices.

2.3 Prior arts on gradients-based attack

Although a distributed learning system protects privacy by not sharing private training data, re-

search works have shown that it is possible to infer the information of private training data from the shared gradients in both language tasks and computer vision tasks. (Melis et al., 2019) enables the identification of words used in the training tokens by analyzing the gradients of the embedding layer. (Goodfellow et al., 2014) proposes an attack algorithm to synthesize images mimicking the real training images by Generative Adversary Network (GAN) models. Besides the works that recover certain properties of the private training data, DLG (Zhu et al., 2019) is a more recent work that shows that it is possible to recover private training data with pixel-wise accuracy for images and token-wise matching for texts by gradient matching. DLG (Zhu et al., 2019) achieves the recovery of images from different datasets on LeNet-5. However, DLG (Zhu et al., 2019) has limitations on evaluating the performance thoroughly on different weight distribution settings, various networks, and different training stages (pre-trained versus initialized). To the best of our knowledge, there is no existing work that comprehensively investigates gradient-based attacks for transformer-based language models with benchmark dataset and standard metric.

3 Approach

In this section, we first formulate the gradient attack in NLP, and the proposed algorithm is introduced afterward.

3.1 Transformer-based NLP models

Transformer (Vaswani et al., 2017) is the fundamental architecture for many popular pre-trained language models, e.g., BERT (Devlin et al., 2019). Scaled dot-product self-attention is the underlying key mechanism inside Transformer, which is calculated as

$$sdpsAttention(q, k, v) = v \cdot softmax(\frac{q \cdot k^T}{\sqrt{d_k}}) \quad (1)$$

where q , k , and v represent the query, key, and value, respectively, and $1/\sqrt{d_k}$ is a scaling factor. Multi-head attention is applied to first calculate attention using Eq. 1 in the subspace of embeddings and then concatenate to form the final output.

A typical flow is to first pre-train the Transformer with objectives like masked language modeling on huge amounts of unlabeled data to get a pre-trained model like BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019), and then finetune the

pre-trained model on specific downstream tasks using the labeled data.

In spite of the high accuracy achieved by the Transformer based models, the large size and high latency of such models make them less appealing to resource constrained edge devices. Accordingly, various knowledge distillation and model compression techniques have been proposed to effectively cut down the model size and inference latency with minimum impact on the accuracy.

Without any loss of generality, we consider the Transformer (Vaswani et al., 2017), BERT (Devlin et al., 2019), and TinyBERT (Jiao et al., 2020) as the representatives of encoder-decoder Transformers, decoder only pre-trained large Transformers, and compressed pre-trained Transformers. Our approach can be extended to other similar models, such as RoBERTa (Liu et al., 2019), UniLMv2 (Bao et al., 2020), and DistilBERT (Sanh et al., 2019).

3.2 NLP gradient attack formulation

We assume that an adversary cannot access the private training data (\mathbf{X}, \mathbf{Y}) in local training directly, but the adversary can gain the gradients that the local devices shared, and the current global model $\mathcal{F}(\mathbf{X}, \mathbf{W})$ in any timestamps during training, where \mathbf{X} is input tokens and \mathbf{Y} is the output labels and \mathbf{W} is the model weights.

The objective of the attack is to recover the valuable and private training data using the shared gradients. For computer vision models, the objective is to retrieve the original pixels in the training images. As mentioned in Section 2, most prior works fall into this category. In this work, we focus on modern Transformer-based models for NLP applications, and our goal is to recover the original tokens in the training set.

Attacking NLP applications is more challenging than computer vision applications, and the reasons are threefold. First, the range of possible values at each pixel is usually smaller than the range of possible tokens at each position, and it is generally more difficult to find the exact item from a larger candidate space. Second, the valuable information carried in an image can be retrieved from a region of pixels, whereas for NLP data, the sensitive information could be carried by several specific tokens, e.g., name, time, and location, and it is required to achieve an exact match on the tokens at certain positions to get the key information from the orig-

inal text. Third, humans can tolerate more errors at pixel values in an image, whereas a small error in the retrieved token id leads to irrelevant token strings in most cases.

Without any loss of generality, the attack can happen at any training stage of the shared global model, and we consider the two most common weight initialization methods, including random initialization for non-pre-trained models and specific learned values for pre-trained models. More formally, the problem is formulated as:

$$\begin{aligned} & \text{Constructing } (\mathbf{X}^*, \mathbf{Y}^*) \\ \text{s.t. } & \frac{\partial \mathcal{L}(\mathbf{W}, \mathbf{X}^*; \mathbf{Y}^*)}{\partial \mathbf{W}} = \frac{\partial \mathcal{L}(\mathbf{W}, \mathbf{X}; \mathbf{Y})}{\partial \mathbf{W}} \end{aligned} \quad (2)$$

where $(\mathbf{X}^*, \mathbf{Y}^*)$ are the recovered data, i.e., tokens and labels for language tasks.

3.3 Proposed algorithm

3.3.1 Recovered token initialization

To recover the language data, we first randomly initialize a dummy data $(\mathbf{X}', \mathbf{Y}')$, where \mathbf{X}' is called dummy input and \mathbf{Y}' is called dummy label. Then we get the corresponding dummy gradient as:

$$\nabla \mathbf{W}' = \frac{\partial \ell(\mathcal{F}(\mathbf{W}, \mathbf{X}'); \mathbf{Y}')}{\partial \mathbf{W}} \quad (3)$$

The next step is to optimize dummy gradient, $\nabla \mathbf{W}'$, to ground truth gradient $\nabla \mathbf{W}$, as closer as possible. In this case, we need to define a differentiable distance function $\mathcal{D}(\mathbf{W}, \mathbf{W}')$, so that we can obtain the best \mathbf{X}' and \mathbf{Y}' as $(\mathbf{X}^*, \mathbf{Y}^*)$ follows:

$$(\mathbf{X}^*, \mathbf{Y}^*) = \arg \min_{(\mathbf{X}', \mathbf{Y}')} \mathcal{D}(\nabla \mathbf{W}', \nabla \mathbf{W}) \quad (4)$$

3.3.2 Distance function for gradient matching

Our experimental observation shows that with different weight initialization, for the same private training data, \mathbf{X} , the NLP model may have distinctly different gradients, $\nabla \mathbf{W}$. For example, with initialized weights from a normal distribution, the gradient of the model may be larger in magnitude than with initialized weights from a uniform distribution (two distributions have the similar intervals). Besides, the $\nabla \mathbf{W}$ gathers near-zero values more heavily with weights from normal distribution than with weights from uniform distribution. We have a similar observation that the $\nabla \mathbf{W}$ gathers considerable near zero values with weights from

a pre-trained model. We consider a matrix with substantial near-zero values as a sparse matrix.

If we use the Euclidean distance (which is used in DLG (Zhu et al., 2019)) to measure the difference between $\nabla \mathbf{W}'$ and $\nabla \mathbf{W}$, the recovery of the ground truth data is driven by large gradients at the early stages. As a result, this might cause a problem when using Euclidean distance under a normal weight initialization since most of the gradients gather around zero while a small proportion of gradients have large values.

To overcome this problem, instead of using the Euclidean distance for $\nabla \mathbf{W}'$ and $\nabla \mathbf{W}$ as the distance function, we consider a combined distance of L2 norm (Euclidean distance) and L1 norm (Manhattan distance), and a coefficient parameter, α , as our distance function to measure the difference between $\nabla \mathbf{W}'$ and $\nabla \mathbf{W}$:

$$\begin{aligned} & \mathcal{D}(\nabla \mathbf{W}', \nabla \mathbf{W}) \\ & = \|\nabla \mathbf{W}' - \nabla \mathbf{W}\|_2 + \alpha(\nabla \mathbf{W}) \|\nabla \mathbf{W}' - \nabla \mathbf{W}\| \end{aligned} \quad (5)$$

where $\alpha(\nabla \mathbf{W})$ is a factor specified for each layer's $\nabla \mathbf{W}$ and its value decreases along with the order of the layer. By doing this, we put larger weights on the gradient differences on the front layers as they are closer to the input private training data. The value of $\alpha(\nabla \mathbf{W})$ is crucial and needs to be suitable for different weight settings.

3.4 The framework of the algorithm

Our complete proposed algorithm is shown in Algorithm 1, and the highlights of our algorithm are as follows. We initialize dummy input (*dummy token embeddings*) and dummy label, $(\mathbf{X}', \mathbf{Y}')$, as dummy data in line 2. During the iterations, started from line 3, we first obtain the dummy gradient, $\nabla \mathbf{W}'$, of the current dummy input. Then we use the distance function in Eq. 5 to measure the difference, $\mathcal{D}(\nabla \mathbf{W}, \nabla \mathbf{W}'_i)$, between dummy gradient $\nabla \mathbf{W}'$ and ground truth gradient $\nabla \mathbf{W}$. At the end of each iteration, we update the $(\mathbf{X}', \mathbf{Y}')$ by the calculated difference, $\mathcal{D}(\nabla \mathbf{W}, \nabla \mathbf{W}'_i)$ in line 7 and line 8. When a pre-set maximum number of iterations is reached, or in 200 iterations, or the number of recovered tokens in ground truth does not change, the algorithm will eventually output the optimized $(\mathbf{X}', \mathbf{Y}')$ as $(\mathbf{X}^*, \mathbf{Y}^*)$ after the iterative recovery process.

4 Experimental setup

All of our experiments are conducted on servers with Intel(R) Xeon(R) Gold 5218 (64 virtual CPUs

Algorithm 1 TAG

```
1: Input:  $\nabla\mathbf{W}$ : ground truth gradient;  $\mathcal{F}(\mathbf{X}, \mathbf{W}')$ : NLP model;  $\eta$ : learning rate;  $\mathbf{W}'$ : parameter weights
2: Initial:  $\mathbf{X}' \sim \mathcal{N}(0, 1)$ ,  $\mathbf{Y}' \sim \mathcal{N}(0, 1)$ 
3: for the  $i$ -th iteration do
4:    $\nabla\mathbf{W}'_i \leftarrow \partial\ell(\mathcal{F}(\mathbf{X}', \mathbf{W}')/\partial\mathbf{W}')$  //get dummy gradient by TAG
5:    $\mathcal{D}(\nabla\mathbf{W}, \nabla\mathbf{W}'_i) \leftarrow \|\nabla\mathbf{W}'_i - \nabla\mathbf{W}\|_2 + \alpha(\nabla\mathbf{W})\|\nabla\mathbf{W}'_i - \nabla\mathbf{W}\|$ 
6:   update  $(\mathbf{X}', \mathbf{Y}')$ :
7:    $\mathbf{X}' \leftarrow \mathbf{X}' - \eta \frac{\partial\mathcal{D}(\nabla\mathbf{W}, \nabla\mathbf{W}'_i)}{\partial\nabla\mathbf{X}'}$ ,
8:    $\mathbf{Y}' \leftarrow \mathbf{Y}' - \eta \frac{\partial\mathcal{D}(\nabla\mathbf{W}, \nabla\mathbf{W}'_i)}{\partial\nabla\mathbf{Y}'}$ 
9: end for
10: Output: Recovered Data  $\mathbf{X}^*, \mathbf{Y}^*$ 
```

Models	Layers	Hidden Units	Attention Heads	Filter Size
Transformer	2	100	4	200
TinyBERT ₄	4	312	6	1,200
TinyBERT ₆	6	768	12	3,072
BERT _{BASE}	12	768	12	3,072
BERT _{LARGE}	24	1,024	16	4,096

Table 1: Model structures of Transformer, TinyBERT₄, TinyBERT₆, BERT_{BASE}, BERT_{LARGE}.

with 504 GB memory) and 8 NVIDIA Quadro RTX 6000 GPUs (24GB memory) by PyTorch 1.5.1, Python 3.6, and CUDA 10.2.

4.1 Datasets

We evaluate our algorithm on the following tasks from the General Language Understanding Evaluation (GLUE) (Wang et al., 2019) benchmark.

CoLA. The Corpus of Linguistic Acceptability (Warstadt et al., 2019) consists of English acceptability judgments drawn from book and journal articles on linguistic theory. Each example is a sequence of words annotated with whether it is a grammatical English sentence.

SST-2. The Stanford Sentiment Treebank (Socher et al., 2013) consists of sentences from movie reviews and human annotations of their sentiment. The task is to predict the sentiment of a given sentence. We use the two-way (positive/negative, 1/0) class split and use only sentence-level labels.

RTE. The Recognizing Textual Entailment (RTE) (Dagan et al., 2005) datasets come from a series of annual textual entailment challenges. This dataset is constructed based on news and Wikipedia text with a combination of RTE1-3, and RTE5.

We select these three datasets because they contain sentences of different lengths. Typically, sentences are about 5 to 15 words for CoLA, 10 to 30 words for SST-2, and 50 to 100 words for RTE. In fact, our algorithm is data agnostic, which can

work on any text inputs from any benchmark, or even any sentence from any source.

4.2 Model settings

We conduct experiments using three popular transformer-based networks, including the basic transformer model (Vaswani et al., 2017), TinyBERT (Jiao et al., 2020) and BERT (Devlin et al., 2019). The basic transformer contains two transformer encoders and one transformer decoder. The number of heads in the self-attention layers is four, and the dimension of the feed-forward network model is 200. The activation function is Gaussian Error Linear Units (GELU) (Hendrycks and Gimpel, 2016). We also applied our algorithm to two different sizes TinyBERT and two different sizes BERT. The TinyBERT₄ is with four layers, 312 hidden units, feed-forward filter size of 1200 and 6 attention heads. The TinyBERT₆ is with 6 layers, 768 hidden units, feed-forward filter size of 3072 and 12 attention heads. In addition, we use the configurations from (Devlin et al., 2019) for BERT. The BERT_{BASE} has 12 layers, 768 hidden units, 3072 feed-forward filter size, and 12 attention heads. The BERT_{LARGE} has 24 layers, 1024 hidden units, 4096 feed-forward filter size and, 16 attention heads. Table 1 summarizes the model structures explored in this work.

4.3 Experiment parameters settings

For each task and dataset of interest, we use the same set of hyperparameters: BertAdam optimizer (Devlin et al., 2019) with learning rate 0.05. For every single sentence recovering, we set the max iteration as 1,000 for our algorithm.

4.4 Experiment evaluation

Recover Rate. This metric is defined as the maximum percentage of tokens in ground truth recovered by TAG.




Figure 2: The average Loss vs Iteration curve of models TinyBERT₄, TinyBERT₆, BERT_{BASE}, BERT_{LARGE} on data CoLA, SST-2 and RTE. The loss decreases at the first 200 iterations and becomes stable after 200 iterations.

ROUGE. Recall-Oriented Understudy for Gisting Evaluation (Lin, 2004), is a set of metrics used for evaluating automatic summarization and machine translation in natural language processing. We use ROUGE-1, ROUGE-2, and ROUGE-L to evaluate the similarity between the sentence generated from gradient attacks and the original sentences. More specifically speaking, ROUGE-1 and ROUGE-2 refer to the overlap of unigrams and bigrams between the recovered text and reference, respectively, and ROUGE-L measures the longest matching subsequence of tokens.

Runtime. This metric is the average of elapsed system time to complete the attack.

5 Result Analysis and Visualization

In this section, we conduct carefully designed experiments to evaluate the proposed TAG on various datasets mentioned in Section 4.1 using the metrics defined in Section 4.4. We have four highlighted results for our evaluation.

Our algorithm shows stable and distinct convergence for NLP models. Here, we measure the distance between the dummy gradient, $\nabla \mathbf{W}'$, and the ground truth gradient, $\nabla \mathbf{W}$, via the aforementioned distance function. We define this distance as the loss of the algorithm. We normalize the loss of all selected data at log scale between 0.8 to 1.2 as shown in Fig. 2. The loss is continuously decreasing for different models and we can observe a stable and distinct convergence from the loss curve, especially for the first 200 iterations.

The TAG attacking process can be visualized in token embeddings level (Fig. 3) and in sentence level (Fig. 4). In tokens embeddings level (Fig. 3), we first reduce the dimension of token embeddings for both dummy input and ground truth by Princi-

pal Component Analysis (PCA). We use the cosine similarity (Li et al., 2020) to evaluate similarity of the dimension reduced token embeddings. In Fig. 3 (a), the cosine similarity of the token embeddings between dummy data and ground truth is 0.42 at the 5-th iteration which means we can observe a 0.93 cosine similarity of those two token embeddings after 200 iterations. As the number of iterations increases, the increasing cosine similarity indicates that TAG iteratively recovers the data at token embeddings level.

In sentence level (Fig. 4), we convert the dummy input (*dummy token embeddings*), \mathbf{X}' , to dummy tokens by the embedding matrix and then a tokenizer can help us to map the tokens with words. In the Fig. 4(a), the dummy data seems random compared to the ground truth (Recover Rate 0%). After 20 iterations, in the Fig. 4(b), the dummy data contains two tokens (Recover Rate 22.22%) from the ground truth, "rocks" and "the". After 50 iterations, the algorithm has recovered 7 of 9 tokens (Recover Rate 77.78%) in the ground truth, and one more token has been recovered when it reached 200 iterations (Recover Rate 88.89%).

Larger model leaks more information. Table 2 summarizes the averaged metrics of TinyBERT₄, TinyBERT₆, BERT_{BASE} and BERT_{LARGE} on the mixture of datasets mentioned in Section 4.1, i.e., RTE, SST-2, and CoLA, with the same vocabulary dictionary. According to Table 1, the size of model structure is sequentially increasing from TinyBERT₄, TinyBERT₆, BERT_{BASE} to BERT_{LARGE}. From Table 2, we observe that larger models leak more information than the smaller ones. For Recover Rate, the BERT_{LARGE} leaks 30% more comparing to the TinyBERT₄, 20% more comparing to the TinyBERT₆ and 10% more

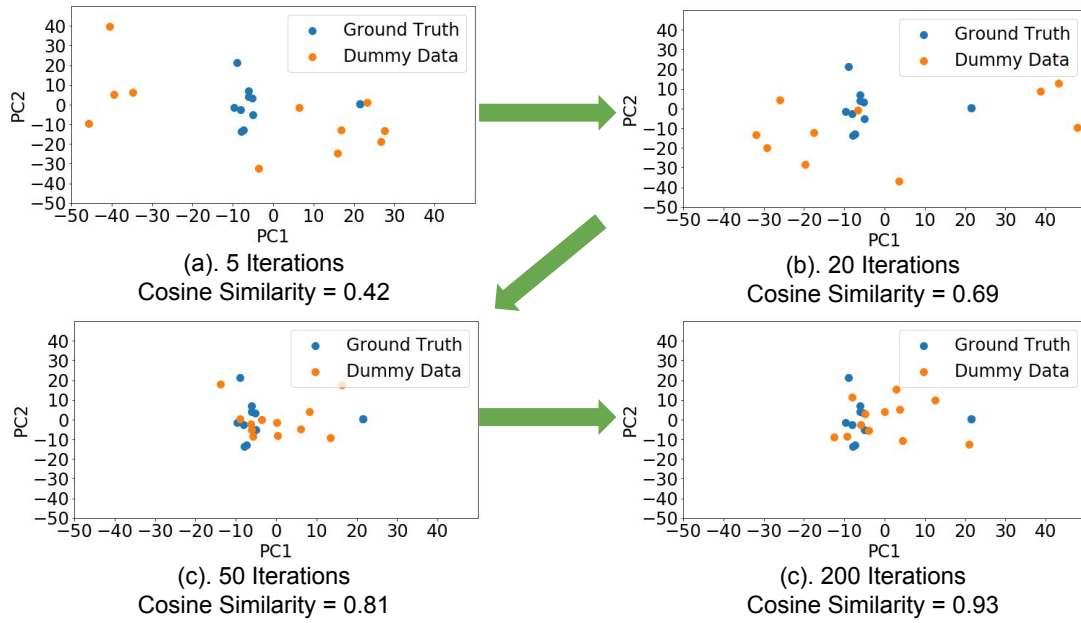


Figure 3: PCA 2D plot for dimension reduced token embeddings of TinyBERT₄ on CoLA. The cosine similarity of dimension reduced token embeddings between dummy data and ground truth increase with training iterations.

Models	Recover Rate(%)	ROUGE-1(%)	ROUGE-2(%)	ROUGE-L(%)	Runtime (Seconds)
TinyBERT₄	29.45	27.07	3.12	22.41	503.24
TinyBERT₆	38.37	34.95	6.54	30.87	526.01
BERT_{BASE}	40.84	41.95	7.77	38.08	1278.62
BERT_{LARGE}	49.62	48.67	15.03	53.09	1672.52

Table 2: The average values of Recover Rate, ROUGE-1, ROUGE-2, ROUGE-L and Runtime. The results are obtained from TinyBERT₄, TinyBERT₆, BERT_{BASE} and BERT_{LARGE} on CoLA, SST-2, RTE datasets.

	Example 1	Example 2
TAG	We monitoring the the global pandemic and will and update the conference plans of of the the conference dates dates.	The area chairs reviewers reviewers will and area of conference broad expertise expertise cover machines or cases
DLG (Zhu et al., 2019)	We we students monitoring monitoring the pandemic and of pandemic plans plans as needed closer to the conference dates .	The we chairs chairs written work will will people expertise expertise longer cases cases .
Ground Truth	We are monitoring the ongoing global pandemic and will update the conference plans as needed closer to the conference dates.	The area chairs and reviewers in each area will have broad expertise to cover these cases.

Table 3: Recover comparison of DLG (Zhu et al., 2019) and TAG on sample texts with basic transformer language model. The sentences are selected randomly from online source. Compared to DLG (Zhu et al., 2019), TAG recovers up to 2× words.

Ground Truth: [CLS] the sailors rode the breeze clear of the rocks .
Dummy: ufo ☹ ##ub 999 12 hostages strictly ##ouse cool writing nonstop
(a). 5 iterations (Dummy data contains 0 tokens in Ground Truth) ↓
Dummy: rocks . hydroelectric ari jamie cornerstone greenfield herrera rocks . cares the
(b). 20 iterations (Dummy data contains 2 out of 9 tokens in Ground Truth) ↓
Dummy: rocks [CLS] . . . the rode breeze the . clear the
(c). 50 iterations (Dummy data contains 7 out of 9 tokens in Ground Truth) ↓
Dummy: rocks [CLS] sailors . . the rode breeze the . clear the
(d). 200 iterations (Dummy data contains 8 out of 9 tokens in Ground Truth)

Figure 4: Recover progress of TAG on a sentence example of CoLA. In token level, TAG eventually recovers 8 of 9 tokens (88.89% Recover Rate) from ground truth which comes from a sentence of CoLA dataset.

comparing to the $BERT_{BASE}$. A similar result can be found in ROUGE-1. As for ROUGE-2, the information leaked from $BERT_{LARGE}$ is $5\times$, $2.5\times$, and $2\times$ compared to $TinyBERT_4$, $TinyBERT_6$, and $BERT_{BASE}$, respectively. For ROUGE-L, the largest model $BERT_{LARGE}$ leaks the most information, which is $2.5\times$, $1.8\times$, and $1.5\times$ larger than $TinyBERT_4$, $TinyBERT_6$, and $BERT_{BASE}$.

Researchers indicate that *to obtain a better result in NLP, we should use a larger model on a larger dataset* in their paper (Raffel et al., 2019b). Based on our experiments, smaller NLP models tend to be more resilient against gradient-based attacks. Information and data security could be another dimension adding to the current tradeoffs among accuracy, latency, and memory footprint.

Larger model requires more time for recovering. We evaluate the runtime performance of our proposed algorithm among different models under 1,000 iterations. A larger model generates more gradients, and in order to recover the data, we need to build the same structure model as the adversarial model to apply our algorithm. Hence, in Table 2, we can see that runtime increases as the model gets larger. $BERT_{LARGE}$ costs $3\times$ runtime as compared to the $TinyBERT_4$, and $BERT_{BASE}$ takes $2.5\times$ more runtime as compared with $TinyBERT_4$.

Our algorithm achieves $2.7\times$ in ROUGE-2 to prior art. We also compare our algorithm with the prior art DLG (Zhu et al., 2019). In Table 3, we apply our algorithm and DLG (Zhu et al., 2019) on Transformer (Vaswani et al., 2017) and attack a sentence from online source. Compared to the DLG (Zhu et al., 2019), our proposed algorithm recovers more than $2\times$ words and compares to

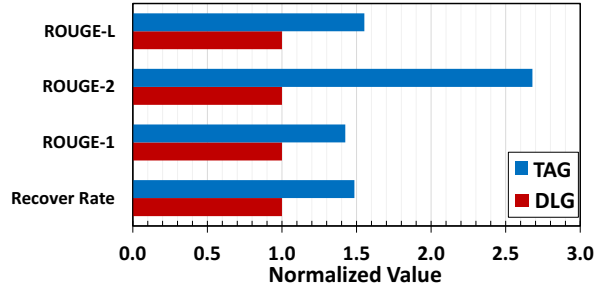


Figure 5: Normalized values of ROUGE-1, ROUGE-2, ROUGE-L, and Recover Rate of TAG and DLG (Zhu et al., 2019) (values normalized to DLG metrics). Especially for ROUGE-2, TAG is more than $2.5\times$ to DLG (Zhu et al., 2019).

the ground truth. More importantly, we almost recover all keywords. We further apply TAG and DLG (Zhu et al., 2019) on BERT, and evaluate the results on the randomly chosen 100 sentences from CoLA and RTE dataset and calculate the averaged value for each experiment. Fig. 5 shows the results. Compared to DLG (Zhu et al., 2019), TAG demonstrates distinct advantages. For ROUGE-2, the result of TAG is about $2.7\times$ to DLG (Zhu et al., 2019). As for ROUGE-1, ROUGE-L and Recover Rate, TAG also takes a $1.5\times$ advantages to DLG (Zhu et al., 2019), which is significant.

6 Ablation Studies

In this section, we conduct ablation experiments over several parameters when we evaluate the results of our algorithm. We change the section of the following factors: the weight distributions, the pre-trained weight, the length of the sentence data, and the size of the vocabulary dictionary.

6.1 Effects of weight distributions

We evaluate the effects of weight distributions by different distributions and different standard deviations of the distributions. As shown in Table 4, we use the $TinyBERT_6$ model and choose sample data from CoLA to apply different weight distributions. For normal distribution with mean as 0, TAG can recover half words from the sentence when standard deviation is 0.01 while it can only recover one of three words from the sentence with a 0.03 standard deviation. For the uniform distribution weight initialization, the results show that TAG is able to recover more with larger initialization range.

Weight Distribution	Uniform (Initializer Range)			Normal (Mean=0) (Standard Deviation)		
	± 0.01	± 0.02	± 0.03	0.01	0.02	0.03
	Recover Rate(%)	36.21	52.17	60.25	50.12	41.57
ROUGE-1(%)	39.39	44.27	60.98	54.54	45.56	35.71
ROUGE-2(%)	14.54	15.09	23.63	30.00	1.01	0
ROUGE-L(%)	44.39	46.98	57.43	66.66	40.01	37.01

Table 4: Recover Rate, ROUGE-1, ROUGE-2, ROUGE-L values of TAG with TinyBERT₆ uniform distribution and normal distribution on sample sentence from CoLA under.

Models	Pre-trained Model		Initialized Model	
	CoLA	SST-2	CoLA	SST-2
Recover Rate(%)	48.76	43.85	34.13	33.82
ROUGE-1(%)	45.68	36.40	30.84	30.74
ROUGE-2(%)	8.01	4.26	6.41	5.45
ROUGE-L(%)	37.61	32.95	26.80	26.42

Table 5: Recover Rate, ROUGE-1, ROUGE-2, ROUGE-L values of TAG with TinyBERT₆ on weight from pre-trained model and normal initialized model.

6.2 Effects of weights from pre-trained model

We evaluate our proposed algorithm on the effects of weights from pre-trained model on two different datasets, CoLA and SST-2. In this experiment, we choose the TinyBERT₆ model and download the pre-trained version from GitHub and also initialize this model using normal distribution with mean as 0 and standard deviation as 0.02. In Table 5, for the CoLA dataset, pre-trained model demonstrates 1.5 \times better than the initialized model. Overall, the pre-trained model shows a better result than the initialized model. We consider that the pre-trained model may contain more information during the training process than the initialized model.

6.3 Performance on different datasets

To evaluate the effects of different sentence length to our proposed algorithm, we conduct experiments on datasets: RTE and CoLA. RTE is a dataset that contains longer sentences than CoLA. We choose sentences to contain more than 50 words from RTE, while sentences within ten words from CoLA as the input data for this experiment. We choose the TinyBERT₆ model with initialized normal distributed weight for this experiment. In Table 6, the results from CoLA are better than RTE, especially for ROUGE family. The ROUGE-1 and ROUGE-2 of CoLA are 3 \times better than RTE, and ROUGE-L is 2.5 \times better than RTE.

Datasets	Recover Rate(%)	R-1(%)	R-2(%)	R-L(%)
	RTE(~ 50 words)	22.70	13.40	1.09
CoLA(~ 10 words)	34.13	30.84	6.41	26.80

Table 6: Recover Rate, ROUGE-1 (R-1), ROUGE-2 (R-2), and ROUGE-L (R-L) values of TAG on comparison of different length sentences from RTE and CoLA.

Vocabulary	Small-Scale	Medium-Scale	Ratio
Total # of Tokens	21,128	30,522	0.69
Recover Rate(%)	54.61	34.13	1.60
ROUGE-1(%)	54.87	30.84	1.78
ROUGE-2(%)	11.83	6.41	1.85
ROUGE-L(%)	47.40	26.80	1.77

Table 7: Recover Rate, ROUGE-1, ROUGE-2, and ROUGE-L values of TAG on comparison of different scales of vocabulary dictionaries.

6.4 Effects of vocabulary dictionary

To evaluate the effects of vocabulary scale, we choose a small scale vocabulary from (Cui et al., 2019) and a medium scale vocabulary from BERT (Devlin et al., 2019). The total numbers of tokens in the small and medium vocabularies are 21,128 and 30,522, respectively. We use TinyBERT₆ model on CoLA and only alter the vocabulary. In Table 7, we observe that the smaller vocabulary size may result in more leakage while the larger one leaks less. For the smaller vocabulary size, the result is more than 1.6 \times improvement compared to the larger one in terms of all evaluation metrics.

7 Conclusion

In this work, we propose, TAG, Transformer Attack from Gradient framework with an adversary algorithm to recover private text data from the transformer model’s gradients. We demonstrate that TAG addresses private information like name is likely to be leaked in transformer-based model. We develop a set of metrics to evaluate the effectiveness of the proposed attack algorithm quantitatively. Our experiments show that TAG works well on more different weight distributions in recovering private training data on Transformer, TinyBERT₄, TinyBERT₆, BERT_{BASE}, and BERT_{LARGE} using GLUE benchmark, and achieves 1.5 \times Recover Rate and 2.5 \times ROUGE-2 over prior methods without the need of ground truth label. TAG can obtain up to 88.9% tokens and up to 0.93 cosine similarity in token embeddings from private training data by attacking gradients on CoLA dataset. We hope

the proposed TAG will shed some light on the privacy leakage problem in Transformer-based NLP models.

8 Acknowledgements

This research is supported in part by the National Science Foundation (NSF) Grants 1743418 and 1843025. Hang Liu is in part supported by NSF CRII Award No. 2000722 and CAREER Award No. 204610.

References

- Hangbo Bao, Li Dong, Furu Wei, Wenhui Wang, Nan Yang, Xiaodong Liu, Yu Wang, Songhao Piao, Jianfeng Gao, Ming Zhou, et al. 2020. Unilmv2: Pseudo-masked language models for unified language model pre-training. *arXiv preprint arXiv:2002.12804*.
- Gilad Baruch, Moran Baruch, and Yoav Goldberg. 2019. A little is enough: Circumventing defenses for distributed learning. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Si Chen, Ruoxi Jia, and Guo-Jun Qi. 2020. [Improved techniques for model inversion attacks](#).
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Ziqing Yang, Shijin Wang, and Guoping Hu. 2019. Pre-training with whole word masking for chinese bert. *arXiv preprint arXiv:1906.08101*.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. The pascal recognising textual entailment challenge. In *Machine Learning Challenges Workshop*, pages 177–190. Springer.
- Dipankar Das, Sasikanth Avancha, Dheevatsa Mudigere, Karthikeyan Vaidynathan, Srinivas Sridharan, Dhiraj Kalamkar, Bharat Kaul, and Pradeep Dubey. 2016. Distributed deep learning using synchronous stochastic gradient descent. *arXiv preprint arXiv:1602.06709*.
- Jeffrey Dean, Greg Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Mark Mao, Marc' aurelio Ranzato, Andrew Senior, Paul Tucker, Ke Yang, Quoc Le, and Andrew Ng. 2012. Large scale distributed deep networks. In *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#).
- Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. 2015. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 1322–1333.
- Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. 2020. [Inverting gradients - how easy is it to break privacy in federated learning?](#) In *Advances in Neural Information Processing Systems*, volume 33, pages 16937–16947. Curran Associates, Inc.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680.
- Deniz Gurevin, Mikhail Bragin, Caiwen Ding, Shanglin Zhou, Lynn Pepin, Bingbing Li, and Fei Miao. 2021. [Enabling retrain-free deep neural network pruning using surrogate lagrangian relaxation](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 2497–2504. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Dan Hendrycks and Kevin Gimpel. 2016. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020. [Tinybert: Distilling bert for natural language understanding](#).
- Bohan Li, Hao Zhou, Junxian He, Mingxuan Wang, Yiming Yang, and Lei Li. 2020. On the sentence embeddings from pre-trained language models. *arXiv preprint arXiv:2011.05864*.
- Mu Li, David G. Andersen, Jun Woo Park, Alexander J. Smola, Amr Ahmed, Vanja Josifovski, James Long, Eugene J. Shekita, and Bor-Yiing Su. 2014a. Scaling distributed machine learning with the parameter server. In *Proceedings of the 11th USENIX Conference on Operating Systems Design and Implementation, OSDI'14*, page 583–598, USA. USENIX Association.
- Mu Li, David G Andersen, Alexander J Smola, and Kai Yu. 2014b. Communication efficient distributed machine learning with the parameter server. In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

- Sheng-Jie Lin, Chenghong Wang, Hongjia Li, Jieren Deng, Yanzhi Wang, and Caiwen Ding. 2020. Esmfl: Efficient and secure models for federated learning. *ArXiv*, abs/2009.01867.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Luca Melis, Congzheng Song, Emiliano De Cristofaro, and Vitaly Shmatikov. 2019. Exploiting unintended feature leakage in collaborative learning. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 691–706. IEEE.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#).
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019a. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv*, pages arXiv–1910.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019b. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [Glue: A multi-task benchmark and analysis platform for natural language understanding](#).
- Yijue Wang, Chenghong Wang, Zigeng Wang, Shanglin Zhou, Hang Liu, Jinbo Bi, Caiwen Ding, and Sanguthevar Rajasekaran. 2021. [Against membership inference attack: Pruning is all you need](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 3141–3147. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. [Neural network acceptability judgments](#).
- Xin Wu, Hao Zheng, Zuochoao Dou, Feng Chen, Jieren Deng, Xiang Chen, Shengqian Xu, Guanmin Gao, Mengmeng Li, Zhen Wang, Yuhui Xiao, Kang Xie, Shuang Wang, and Huji Xu. 2020. [A novel privacy-preserving federated genome-wide association study framework and its application in identifying potential risk variants in ankylosing spondylitis](#). *Briefings in Bioinformatics*, 22(3). Bbaa090.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5754–5764.
- Ligeng Zhu, Zhijian Liu, and Song Han. 2019. Deep leakage from gradients. In *Advances in Neural Information Processing Systems*, pages 14774–14784.