Safety and Performance, Why Not Both? Bi-Objective Optimized Model Compression against Heterogeneous Attacks Toward AI Software Deployment

Jie Zhu, Leye Wang, Xiao Han, Anmin Liu, and Tao Xie, *Fellow, IEEE*

Abstract—The size of deep learning models in artificial intelligence (AI) software is increasing rapidly, hindering the large-scale deployment on resource-restricted devices (*e.g.*, smartphones). To mitigate this issue, AI software compression plays a crucial role, which aims to compress model size while keeping high performance. However, the intrinsic defects in a big model may be inherited by the compressed one. Such defects may be easily leveraged by adversaries, since a compressed model is usually deployed in a large number of devices without adequate protection. In this article, we aim to address the safe model compression problem from the perspective of safety-performance co-optimization. Specifically, inspired by the test-driven development (TDD) paradigm in software engineering, we propose a test-driven sparse training framework called *SafeCompress*. By simulating the attack mechanism as safety testing, SafeCompress can automatically compress a big model to a small one following the dynamic sparse training paradigm. Then, considering two kinds of representative and heterogeneous attack mechanisms, *i.e.*, black-box membership inference attack and white-box membership inference attack, we develop two concrete instances called BMIA-SafeCompress and WMIA-SafeCompress. Further, we implement another instance called MMIA-SafeCompress by extending SafeCompress to defend against the occasion when adversaries conduct black-box and white-box membership inference attacks simultaneously. We conduct extensive experiments on five datasets for both computer vision and natural language processing tasks. The results show the effectiveness and generalizability of our framework. We also discuss how to adapt SafeCompress to other attacks besides membership inference attack, demonstrating the flexibility of SafeCompress.

Index Terms—AI software safe compression, test-driven development, heterogeneous membership inference attack.

◆

1 INTRODUCTION

IN the last decade, artificial intelligence (AI) software, especially software based on deep neural networks (DNN), has attracted much attention and made a significant influence [1]. Currently, AI software, with DNN as representatives, is recognized as an emerging type of software artifact (sometimes known as "software 2.0" [2]). Notably, the size of DNN-based AI software has increased rapidly in recent years (mostly because of a trained deep neural network model). For instance, a state-of-the-art model of computer vision contains more than 15 billion parameters [3]. A recent natural language model, GPT-3, is even bigger, surpassing 175 billion parameters; this situation requires nearly 1TB of space to store only the model [4]. Such a big model hinders realistic applications such as autonomous driving when the software is required to be deployed in resource-restricted devices such as wearable devices or edge nodes. To this end, a new branch is derived from the traditional area of software

compression [5], [6], called *AI software compression* (especially DNN *model compression*[1]), and has attracted a lot of research interest.

Model compression aims to compress a big DNN model to a smaller one given specific requirements, *e.g.*, parameter numbers, model sparsity, and compression rate. Rashly compressing a model may lead to severe degeneration in the AI software's task performance such as classification accuracy. To balance memory storage and task performance, many compression approaches have been proposed and deployed [7], [8]. For example, Han *et al.* [8] prune AlexNet [1] and reduce its size by 9 times while losing only 0.01% accuracy in image classification. Jiao *et al.* [7] reduce the size of BERT [9] by about 2 times via knowledge distillation while losing 0.10% average score in the GLUE [10] official benchmark.

While a compressed model aims to mimic the original model's behavior, its defects may also be inherited. As a representative case, big deep models are verified to be able to memorize training data [11], thus leading to private data leakage when facing threats such as membership inference attack [12]; such a vulnerability would probably remain in the compressed model. More seriously, model compression is often used for AI software deployment on a large number of

- Jie Zhu, Leye Wang, Anmin Liu, and Tao Xie are with Key Lab of High Confidence Software Technologies (Peking University), Ministry of Education, China, and School of Computer Science, Peking University, China.
  E-mail: zhujie@stu.pku.edu.cn, leyewang@pku.edu.cn, anminliu@stu.pku.edu.cn, taoxie@pku.edu.cn
- Xiao Han is with Shanghai University of Finance and Economics, China.
  E-mail: xiaohan@shufe.edu.cn
  Corresponding author: Leye Wang and Xiao Han.

---

1. In the rest of the article, without incurring ambiguity, we use 'model compression' to indicate 'DNN model compression in AI software' for brevity.

edge devices (smartphones, wearable devices, *etc*.) [13], [14]; compared to big models (often stored in a well-maintained server), an adversary[2] thereby has much more opportunities to gain access to compressed models and attack them. For example, an adversary may act as a normal user to download a compressed model on her/his own device. A typical scenario is using a virtual shopping assistant in a smartphone. If software developers neglect the privacy risk brought by the AI model in the virtual shopping assistant and directly deploy the assistant, when an adversary obtains a victim's smartphone, the adversary could attack the AI model in a certain way and analyze the output to acquire the victim's purchase preference and financial status. *In a word, a compressed model inherits the vulnerabilities of the original big model, while facing even higher risks of being attacked*[3]. Hence, studying how to conduct **safe model compression** is urgently required.

An intuitive solution to safe model compression is directly combining two streams of techniques, forming a *two-step solution*: (i) model compression and (ii) model protection. For instance, we can first obtain a small model using existing compression techniques, and then apply protection techniques (*e.g.*, differential privacy [15] and knowledge distillation [16]) to improve the model safety against certain attacks. *However, the two-step solution may suffer from poor model performance*[4] *and low safety because protection techniques could fail to consider the potential variability introduced by model compression.* For example, Yuan and Zhang [17] find that pruning makes the divergence of prediction confidence and sensitivity increase and vary widely among different classes under prediction. This phenomenon may not be sensed by defenders due to lack of interaction but can be manipulated by adversaries.

In this article, we aim to address the problem of safe model compression from the perspective of *performance-safety co-optimization*. Specifically, inspired by *test-driven development* [18] and *dynamic sparse training* [19], we propose a framework of *test-driven sparse training* for *safe model compression*, called ***SafeCompress***. By simulating the attack mechanism to defend, SafeCompress can automatically compress a big model into a small one (with required sparsity) to optimize both model performance and safety. SafeCompress generally follows an iterative optimization mechanism. For initialization, SafeCompress randomly prunes a big model to a sparse one, which serves as the input of the first iteration. In each iteration, SafeCompress applies various compression strategies to the input model to derive more sparse models. Then SafeCompress launches a performance-safety co-optimization mechanism that applies both task-performance evaluation and simulated-attack-based safety

testing on the derived sparse models to select the best compression strategy. Then, the sparse model with the selected (best) strategy becomes the input of the next iteration. The iterative process will terminate after a predefined maximum number of iterations or a new model has little improvement in performance evaluation and safety testing.

Utilizing the SafeCompress framework, our work addresses the challenge of safeguarding compressed AI models in software against heterogeneous attacks that frequently occur in reality. Specifically, in this work, we first consider a single attack and then extend to multiple heterogeneous attacks. In the case of a single attack, we specifically consider two representative privacy attacks, namely, black-box [12] and white-box membership inference attacks (MIAs) [20], and implement two concrete safe model compression mechanisms against them, denoted as ***BMIA-SafeCompress*** and ***WMIA-SafeCompress***, respectively, which effectively mitigate the risks associated with the respective MIA. Note that we choose MIA as the attack example because MIA is representative of evaluating AI model safety, especially from the aspect of privacy leakage [21], [22]. For multiple heterogeneous attacks, our experiments (Section 6) show that our SafeCompress framework is highly flexible and generalizable, naturally supporting extensions to defend against multiple heterogeneous attacks. Specifically, we incorporate additional attack simulation techniques to SafeCompress and implement a concrete instance called ***MMIA-SafeCompress*** to defend against *heterogeneous* black-box and white-box MIAs[5] concurrently. Furthermore, we integrate adversarial training into SafeCompress by including attack models in the training process to enhance a compressed model's defense capability.

This article makes the following main contribution:

• To the best of our knowledge, this work is the first effort toward the problem of **safe model compression**, which is critical for today's large-scale AI software deployment on edge devices such as smartphones.

• To address the problem of safe model compression, we propose a general framework called SafeCompress, which can be configured to protect against a pre-specified attack mechanism. SafeCompress adopts a test-driven process to iteratively update model compression strategies to co-optimize model performance and safety.

• We take into account a single attack and multiple heterogeneous attacks, and choose representative MIAs as attack mechanisms [24]. For a single attack, we develop two concrete instances of SafeCompress, *i.e.*, BMIA-SafeCompress and WMIA-SafeCompress. For multiple heterogeneous attacks, we extend SafeCompress and implement MMIA-SafeCompress. Additionally, we integrate adversarial training into SafeCompress to further enhance a compressed model's defense capability. We also discuss how to adapt SafeCompress to other attacks.

• Using BMIA-SafeCompress, WMIA-SafeCompress, and MMIA-SafeCompress as showcases of SafeCompress, we

---

2. Without incurring ambiguity, we use 'adversary' to indicate a person who conducts the attack and use 'attacker' to indicate the attack model an adversary utilizes.

3. On one hand, accessing a compressed model is easier for adversaries. On the other hand, in our experiments, we observe that some compressed models, *e.g.*, ones compressed by knowledge distillation (KD-AdvReg, Sparsity=0.2) in Figure 1, suffer higher attack accuracy (75.54%) than uncompressed ones (67.33%), indicating that compression could potentially result in a more vulnerable model.

4. We use 'performance' as a general term to describe any specific task-dependent metric. For instance, if a model is developed for image classification, the model performance can be measured by the classification accuracy metric.

5. Heterogeneous black-box and white-box MIAs are those that construct attack models with different structures [22], [23] (as illustrated in Figure 2) so as to operate under varying degrees of knowledge about a target model. Specifically, the black-box setting posits that only model outputs are available to the adversary while in the white-box setting, besides model outputs, extra knowledge about the target model (*e.g.*, hidden layer parameters [20] and gradient descents in training epochs [22]) is also obtainable.

conduct extensive experiments on five datasets of two domains (three tasks of computer vision and two tasks of natural language processing). The experimental results show that our framework[6] significantly outperforms baseline solutions that integrate state-of-the-art compression and MIA defense approaches. The code of the three SafeCompress instances is available as open source at https://github.com/JiePKU/SafeCompress.

This article is an extension of a previous conference edition [25]. Compared to the conference edition, this article includes six main improvements: (1) configuring Safecompress to defend against white-box membership inference attack, (2) extending SafeCompress to handle multiple heterogeneous attacks, (3) incorporating adversarial training into SafeCompress, (4) performing robustness analyses on the safety and performance trade-off metric, (5) adding a new baseline called KL-Div, and (6) involving more recent work discussed in the related work.

2 BACKGROUND

In this section, we briefly introduce contextual knowledge related to our work

2.1 Test-Driven Development

Test-driven development (TDD) [18] is a programming paradigm where test code plays a vital role during the whole software development process. With TDD, before writing code for software functionality, programmers could write the corresponding test suite in advance; then, the test suite can justify whether the software functionality is implemented properly or not. In general, TDD leads to an *iterative coding-testing* process to improve the correctness and robustness of the software; TDD has become a widely adopted practice in software development (*e.g.*, SciPy [26]). Inspired by TDD, we develop a test-driven framework of safe model compression called SafeCompress. Similar to the iterative coding-testing process in TDD, SafeCompress adopts an *iterative compressing-testing* process. In particular, by specifying the attacks to fight against, SafeCompress can automatically update the compression strategies step by step to optimize model performance and safety simultaneously.

2.2 Dynamic Sparse Training

Dynamic sparse training (DST) is a sparse-to-sparse training paradigm to learn a sparse (small) DNN model based on a dense (big) one [19], [27]. Specifically, DST starts training with a sparse model structure initialized from a dense model. As training progresses, it modifies the architecture iteratively by pruning some neural network connections and growing new connections based on certain strategies. This manner enables neural networks to explore self-structure until finding the most suitable one for training data. SafeCompress's iterative updating strategy for optimizing a compressed model structure is inspired by the DST paradigm.

---

6. Without incurring ambiguity, we use 'framework' to indicate 'general framework' and 'approach' to indicate 'constructed instance based on SafeCompress or other specific approaches'.

2.3 Membership Inference Attack

Prior research has extensively shown that DNN models often exhibit different behaviors on training data records (*i.e.*, members) versus test data records (*i.e.*, non-members); the main reason is that models can memorize training data during the repeated training process lasting for a large number of epochs [11], [22]. For instance, a DNN model can generally give a higher prediction confidence score to a training data record than a test one, as the model may remember the training data's labels. Based on such observations, membership inference attack (MIA) [12] is proposed to build attack model to infer whether one data record belongs to training data or not. When it comes to sensitive data, personal privacy is exposed to a great risk. For example, if MIA learns that a target user's electronic health record (EHR) data is used to train a model related to a specific disease (*e.g.*, to predict the length of stay in ICU [28]), then the adversary knows that the target user has the disease. In this work, we use MIA to empirically demonstrate the feasibility and effectiveness of SafeCompress, because MIA has become one of the representative attacks to evaluate the safety of DNN models both theoretically and empirically [21], [22]. Moreover, compressed models in edge devices are easily accessible, *e.g.*, by downloading, allowing adversaries to snoop and collect more model information. Hence, when compressed models are deployed in edge devices, MIA could inflict great harm to personal privacy.

2.4 Adversarial Training

Adversarial training [24], [29] is a competitive learning approach that involves a target model and adversarial models with the aim of equipping the target model with the ability to withstand adversarial attacks. To achieve this ability, this training approach typically incorporates weighted objective terms of adversarial models into the loss function of the target model. Throughout the training process, the goal is to optimize the objective terms of the adversarial models while simultaneously minimizing the overall loss including both the target and adversarial objectives. More generally, this approach is essentially a min-max game [30] that has been widely used in computer vision [29], [31] and natural language processing [32], [33]. In this work, this useful skill is combined with our SafeCompress framework to further help improve a compressed model's defense ability.

3 PROBLEM FORMULATION

Given a big model $\mathcal{F}(;\theta)$ parameterized by θ, we aim to find a sparse model $\mathcal{F}(;\hat{\theta})$ (most elements in $\hat{\theta}$ are zero) under certain memory restriction Ω and the sparse one can defend against a pre-specified attack mechanism f_A. We use G_{f_A} to denote the attack gain of f_A. We restrict the compression ratio, also named as model sparsity, below Ω (*i.e.*, the percentage of non-zero parameters in the sparse model over the original model). We aim to *minimize* both the

task performance loss $\mathcal{L}$ and the attack gain G_{f_A} over the sparse model $\mathcal{F}(; \hat{\theta})$:

$$\min_{\hat{\theta}} \sum_{x,y} \mathcal{L}(\mathcal{F}(x; \hat{\theta}), y) \qquad (1)$$

$$\min_{\hat{\theta}} G_{f_A}(\mathcal{F}(; \hat{\theta})) \qquad (2)$$

$$\text{s.t. } \frac{\|\hat{\theta}\|_0}{\|\theta\|} \leq \Omega, \qquad (3)$$

where x is a sample and y is the corresponding label, and $\mathcal{L}$ represents a task-dependent loss function. $\|\cdot\|_0$ counts the number of non-zero elements, and $\|\cdot\|$ calculates the number of all the elements. This formulation is a general one without specifying the attack mechanism f_A and is a bi-objective optimization problem regarding both model performance (Eq. 1) and safety (Eq. 2) when compressing big models.

Threat Model. We take MIA, being the representative privacy attack, as a threat model to evaluate a victim model (*i.e.*, a compressed model) about its leaked privacy. MIA usually trains a binary classification model f_A. In this work, we use a neural network classifier as f_A.

MIA Gain. The gain G_{f_A} for MIA is then formulated as follows. Given training samples (D_{tr}) and non-training samples ($D_{\neg tr}$), the expected gain of the threat model f_A is

$$G_{f_A}(\mathcal{F}(; \hat{\theta})) = \sum_{x,y} (\mathbb{1}(x \in D_{tr}) \log(f_A(\mathcal{F}(x; \hat{\theta}), y)) + \\ \mathbb{1}(x \in D_{\neg tr}) \log(1 - (f_A(\mathcal{F}(x; \hat{\theta}), y)))), \qquad (4)$$

where $\mathbb{1}(x \in Z)$ is 1 if sample x belongs to Z; otherwise 0.

4 APPROACH

In this section, we first describe our design principles. Then we present our SafeCompress framework. Finally, we show how to use our framework to defend against a single attack and multiple heterogeneous attacks, respectively.

4.1 Key Design Principles

We illustrate two key principles driving our design, *i.e.*, *attack configurability* and *task adaptability* before elaborating on design details.

Attack Configurability. In reality, adversaries may conduct various types of attacks [34]. Hence, a practical solution should be able to do safe compression against an arbitrary pre-specified attack mechanism. In other words, the proposed solution should be easily configured to fight against a given attack mechanism.

Task Adaptability. AI (especially deep learning) techniques have been applied to various task domains including computer vision (CV), natural language processing (NLP), *etc*. To this end, a useful solution is desired to be able to adapt to heterogeneous AI tasks (*e.g.*, CV or NLP) very easily.

Our design of SafeCompress just follows these two principles, thus ensuring practicality in various scenarios of AI software deployment. Next, we describe the design details.

4.2 SafeCompress: A General Framework for Safe Model Compression

Figure 1 presents an overview of SafeCompress, which contains three stages. We describe each below.

Stage 1. Sparsity-Aware Model Initialization. During this stage, we follow the paradigm of dynamic sparse training to prioritize meeting the goal of sparsity (*i.e.*, compression ratio). Specifically, based on a given big model (an arbitrary deep model for various tasks such as CV or NLP), we initialize a sparser version of the model to align with specific memory requirements. After initialization, the sparse model is sent to Stage 2.

Stage 2. Candidate Sparse Model and Simulated-attacker Generation. During this stage, the sparse model is first trained until reaching the stopping criteria. Then, the trained model is fed to two branches. The first branch is called *dynamic sparse update* where different combinations of pruning and growth strategies are performed on the input model, producing new candidate model variants with diverse sparse architectures. To ensure that the derived model always meets the given sparsity requirement, we keep the number of removed connections and reactivated ones the same in each update. The second branch is called *attack mechanism simulation*. In this branch, we simulate an external attacker that aims to attack candidate sparse models. The simulated attacker and candidate sparse models are sent to Stage 3.

Stage 3. Safety Test-driven Model Selection. A safety testing is performed on these input sparse models using the simulated attacker. The candidate model that performs the best in this testing process will be selected and sent back to Stage 2, starting a new iteration. The whole process will terminate after running for a predefined number of iterations.

Algorithm 1 SafeCompress Framework Procedure

Input: A big model $\mathcal{M}_L$; a sparsity requirement Ω; a set of update strategies $\mathcal{U}$ whose size is N; training stopping criteria for a sparse model $\mathcal{T}$; total epochs for termination Eps;

Output: A well-trained sparse model $\mathcal{M}_S$;

1: Initialize $\mathcal{M}_L$ as a sparse one to meet the sparsity requirement Ω; we denote this sparse model as $\mathcal{M}_S$
2: Train $\mathcal{M}_S$ until condition $\mathcal{T}$ is satisfied **then do**:
3: **for** each $\mathcal{U}_i$ in $\mathcal{U}$ **do**:
4:      Update $\mathcal{M}_S$ via $\mathcal{U}_i$ denoted as $\mathcal{M}_S^i$
5: **end for**
6: Obtain a candidate sparse model set $\mathcal{C} = \{\mathcal{M}_S^1 ... \mathcal{M}_S^N\}$
7: Simulate an external attacker as $\mathcal{A}$ (with the help of $\mathcal{M}_S$)
8: $\mathcal{M}_S^{best} \leftarrow$ Pick the best from {safety testing($\mathcal{C}, \mathcal{A}$)}
9: $\mathcal{M}_S \leftarrow \mathcal{M}_S^{best}$
10: **if not** achieve total epochs Eps:
11:      **go to** 2
12: **return** $\mathcal{M}_S$

The pseudo-code of SafeCompress is in Algorithm 1. Note that we do not restrict the type of attacks in SafeCompress. It thus has the potential to prevent various attacks toward AI software and models.

4.3 Defense against Single Attack

We consider scenarios where defenses are mounted against black-box MIA and white-box MIA, respectively.
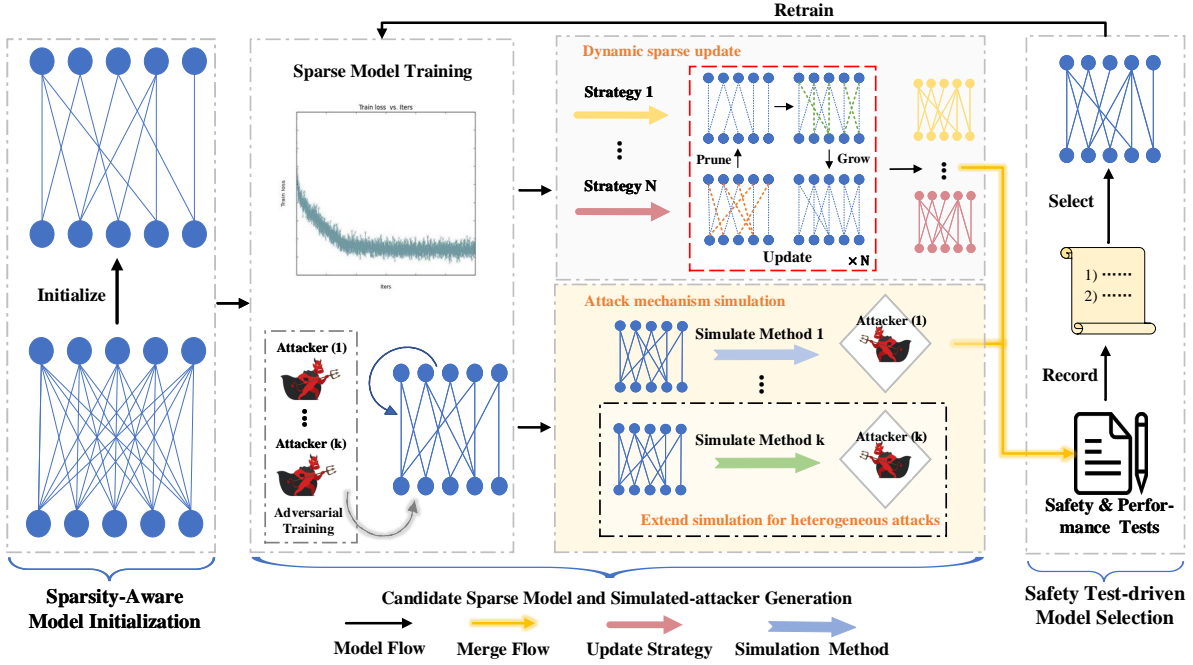
Fig. 1. An overview of our SafeCompress framework.

4.3.1 BMIA-SafeCompress: Defending Black-box Membership Inference Attack based on SafeCompress

We implement a concrete mechanism of safe model compression against black-box MIA[7], called **BMIA-SafeCompress** based on SafeCompress.

Stage 1. Sparsity-Aware Model Initialization. Given a big model, we adopt the Erdös–Rényi [35] initialization approach to reach a predefined sparsity requirement, *i.e.*, removing a number of model connections (assigning zeros to the connection weights). Specifically, for the k-th layer with n_k neurons, we collect them in a vector and denote it as $V^k = [v_1^k, v_2^k, v_3^k, ...v_{n_k}^k]$. Usually, V^k in the k-th layer is connected with the previous layer V^{k-1} via a weight matrix $W^k \in \mathbb{R}^{n_k \times n_{k-1}}$. In the sparse setting, the matrix degenerates to a Erdös–Rényi random graph, where the probability of connection between neuron V^k and V^{k-1} is decided by

$$P(W_{ij}) = \frac{\epsilon \times (n_k + n_{k-1})}{n_k \times n_{k-1}}, \quad (5)$$

where ϵ is a coefficient that is adjusted to meet a target sparsity. In general, this distribution is inclined to allocate higher sparsity (more zeros) to layers with more parameters as the probability tends to decline while the quantity of parameters increases[8].

Stage 2. Candidate Sparse Model and Simulated-attacker Generation. During this stage, we train the input sparse model for a predefined number of iterations (following the setting in previous work [36], the number of iterations is set to 4,000). We can use a vanilla training strategy or adopt advanced adversarial training (we introduce it later in the

end of this subsection). Once finished, the well-trained sparse model, denoted as $\mathcal{M}_S$, is fed into the two branches.

Branch 1. Dynamic sparse update. In the first branch, *i.e.*, the branch of dynamic sparse update, we apply two state-of-the-art pruning strategies and two growth strategies to operate on $\mathcal{M}_S$, leading to four $(2 * 2)$ different sparse typologies. The two pruning strategies are *magnitude-based* pruning [37] and *threshold-based* pruning [8]. Magnitude-based pruning removes connections with the smallest weight magnitudes (weight values); threshold-based pruning removes connections whose weight magnitudes are below a given threshold. The two pruning strategies are effective as small weight magnitudes often contribute little to the final output. The two growing strategies are *gradient-based* growth [38] and *random-based* growth [19]. The gradient-based growth reactivates a connection (weight) that has a large gradient $|\frac{\partial \mathcal{L}}{\partial w}|$, which indicates that the weight is extremely eager to be updated. The random-based growth randomly reactivates connections. This manner may help prevent the sparse model from getting stuck in a local optimum. Afterward, we fine-tune these derived sparse models and generate four candidate sparse models $\mathcal{M}_S^1, ..., \mathcal{M}_S^4$[9].

Branch 2. Attack mechanism simulation. We try to simulate an external attacker in preparation for safety testing in the second branch. Specifically, we follow previous MIA work [24] and simulate a black-box MIA attacker with a fully connected neural network, as depicted in Figure 2 (a). The simulated attacker contains three parts: *probability stream*, *label stream*, and *fusion stream*. The probability stream processes the output probability from the target sparse model $\mathcal{M}_S(x)$. The label stream deals with the label y of the sample x. Then, the fusion stream fuses the features extracted from the two preceding streams and outputs a probability to indicate whether a sample is used in training or not. To improve the

---

7. The black-box setting assumes that only model outputs are available to an adversary.

8. For the weights of remaining connections, we can keep those of the big model, or simply do random initialization. In our experiments, we find that keeping big model weights performs not better than random initialization. Hence, we adopt random initialization in our implementation.

9. As more advanced pruning or growth strategies may be proposed in the future, SafeCompress is easy to incorporate them by adding to/replacing existing strategies.
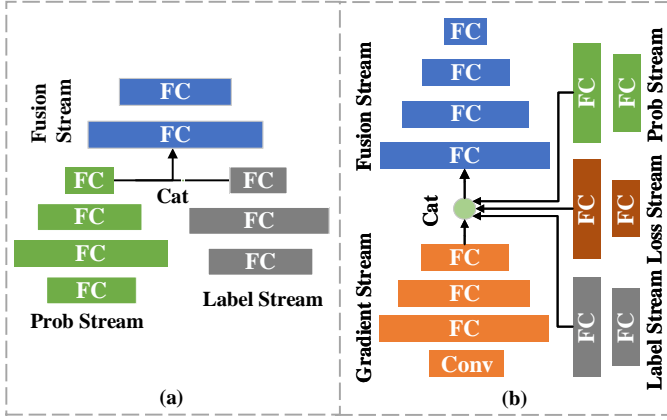
Fig. 2. The architecture of the simulated black-box MIA network model (a) and white-box one (b). FC is a fully connection layer. Conv is a convolution layer.

efficiency of attacker simulation, we do not independently train an attacker for each candidate sparse model $\mathcal{M}_s^i$ in Branch 1. Instead, we first train a simulated attacker $\mathcal{A}$ based on $\mathcal{M}_S$; then, for each $\mathcal{M}_s^i$, we fine-tune $\mathcal{A}$ for multiple epochs to ensure its attack effectiveness toward $\mathcal{M}_s^i$. This strategy of training acceleration is sensible as $\mathcal{M}_s^i$ often deviates from $\mathcal{M}_s$ marginally.

Finally, four sparse model variants (Branch 1) and corresponding well-trained simulated attackers (Branch 2) are sent to the next stage for safety testing.

Stage 3. Safety Test-driven Model Selection. Safety test-driven model selection aims to choose a candidate sparse model with the best trade-off in task performance and safety protection. Specifically, we employ the simulated attacker to conduct safety testing of black-box MIA on four candidate sparse models and then record the attack accuracy. Subsequently, we can simply select the one whose attack accuracy is the lowest (strongest defense ability). However, this manner ignores the task performance, thereby missing a comprehensive consideration. Hence, we also conduct task-performance evaluation on four sparse models (*e.g.*, image classification accuracy for CV tasks). Then, by considering the attack accuracy (safety) and task performance together, we can select the best candidate model. In our implementation, we use a newly defined **TM-score** (*i.e.*, **T**ask performance divided by **M**IA attack accuracy, with details illustrated later) for model selection. The candidate sparse model with the highest TM-score is selected. The selected model is then sent back to Stage 2 and a new iteration starts. The whole process will terminate after repeating a predefined number of iterations.

Metric. For performance, we evaluate model accuracy on classification tasks, represented by *Task Acc*. The higher the Task Acc is, the better the model performs. For safety, we adopt the accuracy of membership inference attacks, denoted as *MIA Acc*, to reflect defense ability. The lower the MIA Acc is, the stronger the model's defense ability is. Besides, we take both performance and safety into consideration, and design a metric called *TM-score* that aims to directly evaluate the performance-safety trade-off:

$$TM\text{-}score = \frac{(Task\ Acc)^{\lambda}}{MIA\ Acc}, \qquad (6)$$

where λ is a coefficient to control this trade-off and we set $\lambda = 1$ in our approach for simplicity if not specified. If a model keeps high Task Acc and low MIA Acc, the TM-score will be high. In brief, this metric is in line with our goal to seek a model with both good performance and strong safety.

Enhancement with Adversarial Training. A target model and adversarial models (both usually involved in adversarial training) compete mutually during training. In this way, the target model can improve its capability against adversarial models. Following this idea, we regard the external attacker as the adversarial model and leverage adversarial training based on BMIA-SafeCompress to further improve the target model's defense ability. Specifically, we additionally construct an adversarial model, *i.e.*, a black-box MIA model f_A^B following the structure of Figure 2 (a). The attack gains are denoted as $G_{f_A^B}$. Thus, we formalize joint membership privacy attack (gain) and performance objectives in the following min-max optimization problem:

$$\min_{\mathcal{F}} \left\{ \mathcal{L}(\mathcal{F}) + \beta * \max_{f_A^B} G_{f_A^B}(\mathcal{F}) \right\} \qquad (7)$$

$\mathcal{L}(\mathcal{F})$ is our performance objective of BMIA-SafeCompress, where $\mathcal{F}$ is our target sparse model. β is a weighted coefficient. We set it to 0.1 by default.

Finally, we may need to point out the difference between the adversarial model and the simulated attacker in Stage 2 of BMIA-SafeCompress to remove possibly caused confusion. The simulated attacker in Stage 2 of BMIA-SafeCompress is used for safety testing. They are trained offline and thereby cannot impact the gradient backpropagation when training the target model. In contrast, the adversarial model is trained together with the target model. This process influences the gradient backpropagation and regularizes the target model.

4.3.2 WMIA-SafeCompress: Defending White-box Membership Inference Attack based on SafeCompress

Different from black-box MIA, white-box MIA assumes that an adversary has extra knowledge about a target model beyond outputs (*e.g.*, hidden layer parameters [20] and gradient descents in training epochs [22]). Hence, white-box MIA could be regarded as a stronger attack manner. In this part, we aim to configure SafeCompress to white-box MIA. We name this instance of SafeCompress framework as **WMIA-SafeCompress**.

Configuration. To implement WMIA-SafeCompress, concretely, we include the extra information as inputs to simulate a white-box attacker mechanism $\mathcal{A}$ proposed in previous work [22], [39] while maintaining other processes of BMIA-SafeCompress. In other words, the main adjustment is to replace the simulated black-box attacker in BMIA-SafeCompress with a white-box one to participate in the safety testing. This alteration offers a notable advantage in terms of ease and convenience. As illustrated in Figure 2 (b), the simulated attacker contains five parts: *probability stream, loss stream, gradient stream, label stream, and fusion stream*. The first four streams are designed to process four different inputs, respectively. They are a target sample's predicted probability, loss value (*e.g.*, classification loss), gradients of the parameters of the target model's previous layer, and one-hot encoding of the target sample's true label. The fifth

part, as the name suggests, fuses the extracted features from the first four streams and outputs a probability to indicate whether the sample is used in training or not.

Metric. We simply follow the same metric, *i.e.*, TM-score, in BMIA-SafeCompress.

We also discuss how to configure SafeCompress for other attacks including attribute inference attack and model inversion attack (attack configurability) and task adaptability using semantic segmentation [40] as illustrated in **Appendix A**.

4.4 Defense against Multiple Heterogeneous Attacks

In a more practical sense, besides a single attack, AI software could face multiple heterogeneous attacks, which can exacerbate privacy leakage. For example, a potential adversary may employ black-box and/or white-box MIAs. In such scenarios, it becomes difficult for a defender to predict the attack that the adversary may adopt, being black-box, white-box, or both. This situation is more intractable than a single attack stated earlier. To alleviate this issue, we are supposed to take into account all possible attacks, namely, both black-box and white-box MIAs in this context, and introduce our solution below.

4.4.1 MMIA-SafeCompress: Extend SafeCompress against Multiple Heterogeneous Membership Inference Attacks

To defend against multiple heterogeneous attacks, we extend our SafeCompress framework. We list two reasons to support its feasibility. (1) The first branch of Stage 2 in SafeCompress, *i.e.*, *dynamic sparse update*, is independent as it does not rely on the second branch, *i.e.*, *attack mechanism simulation*[10]. **This characteristic enables us to add additional attack mechanisms to simulate while causing little influence on the *dynamic sparse update* branch**. (2) Our SafeCompress framework strictly follows the *attack configurability* principle. **In other words, the branch of *attack mechanism simulation* can be configured with various attacks.** Hence, our SafeCompress framework exhibits high flexibility to seamlessly accommodate multiple heterogeneous attacks through the incorporation of additional simulated attack mechanisms in Stage 2, as well as corresponding safety testing in Stage 3. Further, based on the extension, we can implement various instances against different combinations of attacks toward AI software. Below, we implement a concrete instance called ***MMIA-SafeCompress*** to defend against black-box and white-box MIAs simultaneously.

Configuration. To implement MMIA-SafeCompress, the key is to add simulated black-box and white-box attack mechanisms. In fact, we can build MMIA-SafeCompress based on BMIA-SafeCompress or WMIA-SafeCompress as they already have one simulated attack mechanism. Here, to ease illustration, we choose to build MMIA-SafeCompress based on BMIA-SafeCompress. An extra simulated white-box attacker illustrated in Figure 2 (b) is first added in Stage 2, in parallel with the simulated black-box attacker. Then, corresponding white-box MIA testing is added in Stage 3. Consequently, we shall perform two kinds of safety testing,

i.e., black-box and white-box MIA testing, to select the best candidate model that will be sent back to Stage 2, starting a new iteration.

Metric. We leverage a straightforward manner by using the weighted sum to combine two metrics[11]:

$$TM\text{-}score_M = \alpha * TM\text{-}score_B + (1-\alpha) * TM\text{-}score_W \quad (8)$$

$TM\text{-}score_B$ is the score of black-box safety testing while $TM\text{-}score_W$ is the score of white-box safety testing. α is a coefficient to balance the two metrics. We set α to 0.5 as we deem that these two metrics are equally important.

5 EXPERIMENTAL SETUP

In this section, we present an overview of five used datasets, five used models, and eight baselines compared in our work.

5.1 Datasets and Models

We conduct experiments on five datasets that are representative with high citation count according to Google Scholar and widely used for membership inference in previous work [12], [41], [42]. The partition setting of training/test data follows the previous work that proposed the respective datasets. For each dataset, we select as the original big model one model that is famous and highly cited in the domain of computer vision or natural language processing.

CIFAR10 and **CIFAR100** [43] are two benchmark datasets for image classification. Both of them have $50,000$ training images and $10,000$ test images. CIFAR10 has 10 categories while CIFAR100 has 100 categories. The size of every image is 32×32 pixels. We adopt the AlexNet model [1] for CIFAR10 and the VGG16 model [44] for CIFAR100.

Tiny ImageNet [45] is an image dataset that contains 200 categories. Each category includes 500 training images and 50 test images. The size of each image is 64×64 pixels. We adopt the ResNet18 model [46] for Tiny ImageNet.

AG News [47] is a topic classification dataset that has 4 categories. For each category, AG News contains 30,000 training and 1,900 test texts. We adopt the RoBERTa model [48] for AG News.

Yelp-5 [49] is a review dataset for sentiment classification (5 categories). Yelp-5 includes $130,000$ training and $10,000$ test texts per category. We adopt the BERT model [50] for Yelp-5.

AlexNet [1], VGG16 [44], and ResNet18 [46] are convolution neural networks (CNN). BERT [50] and RoBERTa [48] are transformer-based neural networks.

Dataset Splits for Attacker Simulation. We split experimental datasets as shown in Table 1 to evaluate a model's defense ability. Following previous work [16], [51], we assume that a simulated attacker knows 50% of a (target) model's training data and 50% of test data (non-training data). These known data are used for training the attack model. The other data are adopted for evaluation. We put more details of training in **Appendix B**.

---

10. The second branch may need candidate models derived from the first branch to help implementation, *e.g.*, our BMIA-SafeCompress and WMIA-SafeCompress. Safety interaction is performed in Stage 3.

11. More advanced strategies, such as bootstrapping according to different attacks' severity and commonality [34], may deserve future research.

TABLE 1
Number of samples in dataset splits.

Datasets	Attack Training		Attack Evaluation	
	D_{train}^{known}	D_{test}^{known}	$D_{train}^{unknown}$	$D_{test}^{unknown}$
CIFAR10	25,000	5,000	25,000	5,000
CIFAR100	25,000	5,000	25,000	5,000
Tiny ImageNet	50,000	5,000	50,000	5,000
AG News	60,000	3,800	60,000	3,800
Yelp-5	325,000	25,000	325,000	25,000

TABLE 2
Two-step baseline approaches.

Baselines	Model Compression		MIA Defense		
	Pruning	KD	DP	AdvReg	DMP
Pr-DP	✓		✓		
Pr-AdvReg	✓			✓	
Pr-DMP	✓				✓
KD-DP		✓	✓		
KD-AdvReg		✓		✓	
KD-DMP		✓			✓

5.2 Baselines

Given that no existing studies focus on safe model compression, we formulate multiple *two-step* baseline approaches by combining state-of-the-art model compression and MIA defense techniques.

In the first step, we choose one **compression** technique $\mathcal{C}$. We consider two widely used and effective choices to compress big models, *i.e.*, *pruning* [37] and *knowledge distillation (KD)* [52]. For pruning, we adopt the 'pretrain→prune→fine-tune' paradigm. We first pretrain a big model from scratch. Afterward, we leverage magnitude-based pruning and compress this full model to a certain sparsity. Finally, we fine-tune the pruned model to recover model performance. For KD, we design a small dense model as a student model given the sparsity requirement. Then, we use a well trained big model as a teacher model to help the student model's training.

In the second step, we select one MIA **defense** technique $\mathcal{D}$. We test three defense techniques including *differential privacy (DP)* [15], *adversary regularization (AdvReg)* [24], and *distillation for membership privacy (DMP)* [16]. DP adds noise to gradients while training a target model. AdvRes adds membership inference gain G_A of an attacker to the loss function of a target model as a regularization term, forming a min-max game. DMP first trains an unprotected (vanilla) teacher model. Then, DMP utilizes extra unlabeled data and adopts knowledge distillation to force a student model *i.e.*, a target model, to simulate the output of the teacher.

Finally, we group two-step techniques by enhancing the training or fine-tuning process in $\mathcal{C}$ with the defense technique in $\mathcal{D}$. For example, the Pr-DP baseline adopts the 'pretrain→prune→DP-based fine-tune' process; KD-AdvReg trains a student (a compressed model) with the help of a big teacher model in a min-max game manner. All the six two-step baseline approaches are listed in Table 2.

Baseline 7 (MIA-Pr). MIA-Pr is originally proposed by previous work [51] to defend MIA, and its main technique is pruning. Hence, although the previous work [51] does not specify the usage for model compression, MIA-Pr can naturally reduce a model's size. We use MIA-Pr as a baseline
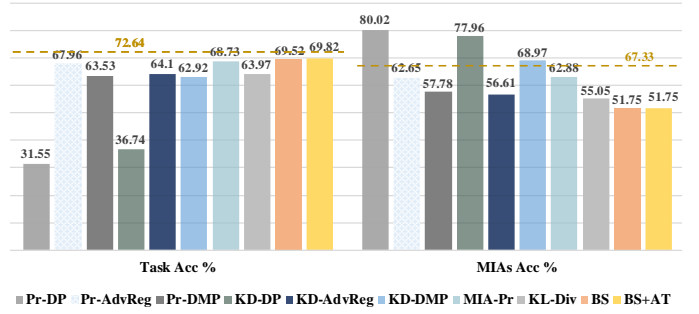

Fig. 3. Results on CIFAR100 with sparsity 0.05 (BS: BMIA-SafeCompress, AT: Adversarial Training, dash line: uncompressed VGG16).

to prune a big model until satisfying a given sparsity requirement. Note that when pruning the big model, MIA-Pr optimizes the defense effect of MIA, while the two-step baseline Pr-X (the first step is pruning) optimizes the task performance (*e.g.*, classification accuracy).

Baseline 8 (KL-Div [17]). KL-Div is initially proposed by previous work [17] to defend against a special attacker that leverages potential defects in pruning. The main idea is to use KL-divergence to align the posterior predictions of different input samples, aiming at improving a compressed model's defense ability. However, this approach does not consider optimizing model performance. In this work, we also regard this approach as our baseline and follow the setting stated in the previous work [17].

6 EXPERIMENTAL RESULTS

In this section, we present the results of our BMIA-SafeCompress, WMIA-SafeCompress, and MMIA-SafeCompress, respectively, to show the effectiveness of our SafeCompress framework.

6.1 Experimental Results on BMIA-SafeCompress

To show the effectiveness of BMIA-SafeCompress, we conduct experiments on CV and NLP datasets, respectively.

6.1.1 Results on CV Datasets

CIFAR100 (VGG16). We first report Task Acc and MIA Acc on CIFAR100 with sparsity 0.05 using VGG16. As illustrated in Figure 3, the results indicate that BMIA-SafeCompress outperforms eight baselines in Task Acc while alleviating MIA risks remarkably. Specifically, BMIA-SafeCompress produces 51.75% for MIA Acc (just a bit higher than random guess), decreasing by 28.27%, 10.9%, 6.03%, 26.21%, 4.86%, 17.22%, 11.13%, and 3.30% compared with eight baselines, respectively. Moreover, equipping BMIA-SafeCompress with Adversarial Training leads to an additional improvement of 0.30% in Task Acc while not increasing privacy risk. Interestingly, when compared to the uncompressed VGG16, SafeCompress reduces MIA risks to a large extent while sacrificing Task Acc only a little. This result highlights the practicality of SafeCompress in generating a small and safe model with competitive task performance.

Further, to validate the effectiveness on different sparsity requirements, we report the results in Figure 4 with sparsity 0.1 and 0.2. Consistent with the results of sparsity 0.05,

Fig. 4. Results on CIFAR100 with sparsity 0.1 and 0.2 (BS: BMIA-SafeCompress, AT: Adversarial Training, dash line: uncompressed VGG16).
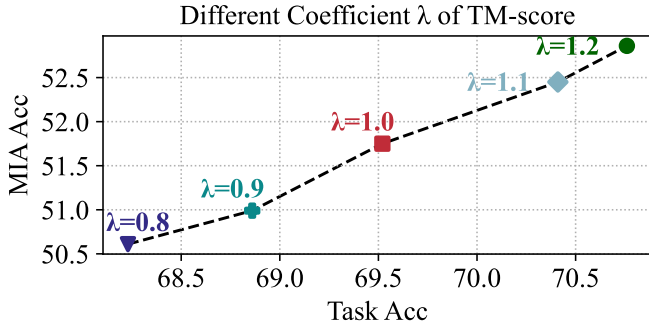

Fig. 5. Different coefficient λ on BMIA-SafeCompress with sparsity=0.05.

TABLE 3
Defend against black-box MIA. TM-score$_B$ on CIFAR100. The best results are marked in **bold**.

Approach	Sparsity		
	0.05	0.1	0.2
VGG16 (uncompressed)	1.08	1.08	1.08
Pr-DP	0.39	0.41	0.47
Pr-AdvReg	1.08	0.99	1.06
Pr-DMP	1.10	1.22	1.21
KD-DP	0.47	0.47	0.48
KD-AdvReg	1.13	1.13	0.90
KD-DMP	0.91	0.93	0.84
MIA-Pr	1.09	1.06	1.05
KL-Div	1.16	1.12	1.17
BMIA-SafeCompress	1.34	1.35	1.22
+ Adversarial Training	**1.35**	**1.38**	**1.27**

BMIA-SafeCompress still beats all the baselines by achieving higher Task Acc and lower MIA Acc. We also observe that when sparsity increases (*i.e.*, the number of non-zero parameters increases), Task Acc of BMIA-SafeCompress increases from 69.52% (sparsity 0.05) to 71.96% (sparsity 0.2); meanwhile, MIA Acc also goes up from 51.75% (sparsity 0.05) to 58.85% (sparsity 0.2), indicating that the compressed model becomes more vulnerable with increasing sparsity. Such a phenomenon may be incurred by the extra information (*i.e.*, more parameters) kept in the compressed model when sparsity rises — some of the extra information may be generalizable so that task performance is enhanced; other information might be specific to training data, thus leading to higher MIA risks.

Finally, to gain insight into the performance-safety trade-off in one shot, we report TM-score (Task Acc divided by MIA Acc, Eq. (6)) in Table 3. BMIA-SafeCompress (+ Adversarial Training) outperforms all the baselines significantly and consistently under three sparsity settings by making a better trade-off between performance and safety.

In practice, users may have different requirements for safety and performance. Fortunately, it can be easily achieved in our approach by adjusting the coefficient λ of TM-score. For example, a user may think that safety is more important than performance, and therefore she can decrease λ to a smaller value (*e.g.*, 0.9 or 0.8); otherwise, she can increase λ to a bigger value (*e.g.*, 1.1 or 1.2). We illustrate how the coefficient λ controls the trade-off in Figure 5. It is seen that a larger λ (performance is important) results in higher performance while a smaller λ results in stronger defense capability.

We also explore six other aspects of our framework based on BMIA-SafeCompress in **Appendix C.1.1**. The first is time consumption where the time consumption of BMIA-

SafeCompress is comparable to that of other baselines. We further find that, in BMIA-SafeCompress, attack mechanism simulation consumes more time than that of sparse model training. The second is the rationality of selecting the candidate sparse model that performs the best in the safety testing. We empirically demonstrate that compared to enumerating all candidate sparse models, this strategy in SafeCompress is reasonable in light of the performance and significantly reduces time consumption. The third is the flexibility of SafeCompress to incorporate other training tricks where we show that SafeCompress is compatible with most training tricks including dropout, data augmentation, adversary regularization, *etc*. The fourth is to perform SafeCompress on a bigger model (*e.g.*, ResNet50 [46]) and a larger dataset (*e.g.*, ImageNet [53]). The result indicates that our framework is scalable. The fifth is the comparison between BMIA-SafeCompress and baselines on TM-scores with different λ where SafeCompress is shown to be robust against metric alteration. The last is a statistical analysis with multiple runs, showing that our approach is statistically superior to others. Besides, we conduct experiments on CIFAR10 and Tiny ImageNet in **Appendix C.1.2** where the superior results over other baselines further show the effectiveness of BMIA-SafeCompress.

6.1.2   Results on NLP Datasets

AG News (RoBERTa). To validate the effectiveness and generalization of BMIA-SafeCompress, we conduct experiments on AG News. As indicated in Table 4, BMIA-SafeCompress achieves 87.50% for Task Acc, slightly inferior to the highest 88.10% (KD-AdvReg). However, our approach decreases MIA

TABLE 4
Defend against black-box MIA. Task Acc (performance), MIA Acc (safety), and TM-score$_B$ on AG News.

Approach	Sparsity=0.5		
	Task Acc	*MIA Acc*	*TM-score*
RoBERTa (Uncompressed)	89.20%	57.59%	1.55
Pr-DP	87.28%	56.07%	1.56
Pr-AdvReg	87.24%	57.58%	1.52
Pr-DMP	86.38%	56.31%	1.53
KD-DP	82.64%	56.13%	1.47
KD-AdvReg	**88.10%**	56.56%	1.56
KD-DMP	87.24%	56.64%	1.54
MIA-Pr	87.91%	57.04%	1.54
KL-Div	83.26%	**54.09%**	1.54
BMIA-SafeCompress	87.50%	55.28%	1.58
+ Adversarial Training	87.63%	54.77%	**1.60**

Acc to 55.28%, much lower than KD-AdvReg (56.56%). BMIA-SafeCompress outperforms all the baselines in MIA Acc, resulting in a high TM-score (1.58). Moreover, incorporating BMIA-SafeCompress with Adversarial Training leads to a further enhancement in all the metrics. The results on Yelp dataset in **Appendix C.1.3** further show the superiority of BMIA-SafeCompress over other baselines and demonstrate its effectiveness.

6.2 Experimental Results on WMIA-SafeCompress

CIFAR100 (VGG16). The results are reported in Table 5. When model sparsity is set to 0.05, our approach, namely WMIA-Safecompress, produces 67.51% for Task Acc, achieving a competitive classification accuracy among all the baselines. In addition, our approach also decreases MIA Acc by 15.36% compared to the uncompressed VGG16. Although it is a little bit inferior (2.21% lower) in MIA Acc than Pr-DMP, our approach maintains higher performance (3.98% higher) than Pr-DMP. We also calculate the TM-score for each approach to show its trade-off degree. It is observed that our approach obtains the highest TM-score (1.20). Similar results can be observed for sparsity 0.1 where WMIA-Safecompress also achieves the highest TM-score (1.32). In total, this result is reasonable as our framework, SafeCompress, targets a bi-objective (safety and performance) optimization process. Simultaneously, this result indicates that our framework can be flexibly configured for other, even stronger attacks. Leveraging Adversarial Training leads to a further descent in privacy risk, especially for sparsity 0.05.

Besides CIFAR100, the experimental results on CIFAR10 and Tiny ImageNet (CV) in **Appendix C.2.1** and AG News (NLP) in **Appendix C.2.2** show the effectiveness of WMIA-SafeCompress.

6.3 Experimental Results on MMIA-SafeCompress

CIFAR100 (VGG16). As shown in Table 6, our approach, namely MMIA-SafeCompress, achieves the most competitive classification performance in both sparsity settings. Specifically, it produces 68.13% (sparsity=0.05) for Task Acc, just 0.60% lower than MIA-Pr (68.73%). When sparsity is 0.1, it obtains 70.48%, outperforming all the baselines in Task Acc. At the same time, it holds the lowest MIA Acc$_B$ among all the approaches under comparison. For MIA Acc$_W$, although our approach is a little bit inferior to Pr-DMP,

our approach obtains relatively low MIA Acc$_W$ in both sparsity settings compared to the defense effects of the baselines. Finally, our approach produces 1.23 (sparsity=0.05) and 1.25 (sparsity=0.1) for TM-score, outperforming all the baselines. The results indicate that our approach is able to handle heterogeneous membership inference attacks and make the performance-safety trade-off, demonstrating the feasibility of extending our SafeCompress framework against heterogeneous attacks.

When our approach is equipped with Adversarial Training, MIA Acc$_B$ further decreases in both sparsity settings. We also observe that Adversarial Training has an inconsistent influence on Task Acc and MIA Acc$_W$ for different sparsity settings. Specifically, Adversarial Training improves Task Acc when sparsity is 0.05 and brings a slight descent when sparsity is 0.1. With Adversarial Training, MIA Acc$_W$ increases when sparsity is 0.05 and decreases when sparsity is 0.1. This phenomenon is interesting. Since it is not our main interest, we leave it for future exploration. Adversarial Training further brings improvements to TM-Score for both sparsity settings, showing its effectiveness.

Another observation is that our approach seems to pay more attention to defending against black-box MIA than white-box MIA. We speculate that white-box MIA is harder to defend against than black-box one as white-box MIA involves more information (*e.g.*, hidden layer parameters [20] and gradient descents in training epochs [22]). So when the two heterogeneous attackers attack our model simultaneously, our approach is inclined to defend against the black-box attacker more. It could be an advantage to some extent as the black-box attack is easier to perform compared to the white-box one, which needs to collect extra information. Similar situations can be seen in some baselines such as KD-AdvReg and MIA-Pr. We consider that this phenomenon may be general and it is interesting to find the potential reason behind it in the future.

Besides CIFAR100, in **Appendix C.3.1** and **Appendix C.3.2**, we also show the performance of our approach on CIFAR10, Tiny ImageNet, and AG News, respectively. These results further indicate the effectiveness and generalization of SafeCompress in defending against multiple heterogeneous attacks.

7 RELATED WORK

In this section, we discuss related work from three aspects including membership inference attack, membership inference defense, and model compression.

7.1 Membership Inference Attack

Membership inference attack (MIA) [12], [20], which aims to infer whether a data record is used to train a model or not, has the potential to raise severe privacy risks to individuals. A prevalent attack fashion is to train a neural network via multiple shadow training [12], converting the task of recognizing member and non-member of training datasets to a binary classification problem. Unlike binary classifier-based MIA, another attack form, metric-based MIA that just computes metrics (*e.g.*, prediction correctness [21] or confidence [54], and entropy [55]) is simpler and less

TABLE 5
Defend against white-box MIA. Task Acc (performance), MIA Acc (safety) and TM-score$_W$ results on CIFAR100. The best results are marked in **bold**.

Approach	Sparsity=0.05			Sparsity=0.1		
	Task Acc	MIA Acc	TM-score	Task Acc	MIA Acc	TM-score
VGG16 (uncompressed)	72.64%	71.58%	1.01	72.64%	71.58%	1.01
Pr-DP	31.55%	84.18%	0.37	34.99%	82.73%	0.42
Pr-AdvReg	67.96%	67.05%	1.01	68.85%	70.65%	0.97
Pr-DMP	63.53%	**54.01%**	1.18	65.81%	**52.28%**	1.26
KD-DP	36.74%	77.40%	0.47	37.25%	78.17%	0.48
KD-AdvReg	64.10%	61.27%	1.05	65.66%	63.87%	1.03
KD-DMP	62.92%	63.30%	0.99	64.55%	61.67%	1.05
MIA-Pr	**68.73%**	66.77%	1.03	69.89%	68.59%	1.02
KL-Div	63.97%	56.54%	1.13	63.89%	57.73%	1.11
WMIA-SafeCompress	67.51%	56.22%	1.20	**70.53%**	53.44%	**1.32**
+ Adversarial Training	**68.73%**	55.63%	**1.24**	70.51%	53.25%	**1.32**

TABLE 6
Defend against black-box and white-box MIAs. Task Acc (performance), MIA Acc (safety) and TM-score$_M$ results on CIFAR100. The best results are marked in **bold**.

Approach	Sparsity=0.05				Sparsity=0.1			
	Task Acc	MIA Acc$_B$	MIA Acc$_W$	TM-score	Task Acc	MIA Acc$_B$	MIA Acc$_W$	TM-score
VGG16 (uncompressed)	72.64%	67.33%	71.58%	1.05	72.64%	67.33%	71.58%	1.05
Pr-DP	31.55%	80.02%	84.18%	0.38	34.99%	84.40%	82.73%	0.42
Pr-AdvReg	67.96%	62.65%	67.05%	1.05	68.85%	69.46%	70.65%	0.98
Pr-DMP	63.53%	57.78%	**54.01%**	1.14	65.81%	53.84%	**52.28%**	1.24
KD-DP	36.74%	77.96%	77.40%	0.47	37.25%	79.71%	78.17%	0.47
KD-AdvReg	64.10%	56.61%	61.27%	1.09	65.66%	58.11%	63.87%	1.08
KD-DMP	62.92%	68.97%	63.30%	0.95	64.55%	69.72%	61.67%	0.99
MIA-Pr	68.73%	62.88%	66.77%	1.06	69.89%	65.74%	68.59%	1.04
KL-Div	63.97%	55.05%	56.54%	1.15	63.89%	57.14%	57.73%	1.11
MMIA-SafeCompress	68.13%	52.32%	59.01%	1.23	**70.48%**	51.79%	61.60%	1.25
+ Adversarial Training	**69.17%**	**51.23%**	61.17%	**1.24**	70.31%	**51.66%**	58.66%	**1.28**

computational. However, the performance of metric-based attacks is inferior to that of classifier-based attacks. Liu *et al.* [56] exploit membership information from the whole training process of a target model for improving attack performance. Recently, membership inference attack has been widely extended to more fields. For example, MIA is investigated on generative models [57] and mainly focuses on GANs [29]. Song and Raghunathan [58] introduce the first MIA on word embedding and sentence embedding models that map raw objects (*e.g.*, words, sentences, and graphs) to real-valued vectors. Duddu *et al.* [59] introduce the first MIA on graph embedding models. Tseng *et al.* [60] present the first MIA on self-supervised speech models under black-box access. M4I [61] is the first to investigate membership inference on multi-modal models. Liu *et al.* [42] propose the first membership inference approach called EncoderMI for contrastive-learning based models.

7.2 Membership Inference Defense

Considering that membership inference attack has the potential to raise severe privacy risks to individuals, studying how to defend against it becomes important. For example, differential privacy [15], [62] (known as DP) interrupts an attack model by adding noise to the learning object or output of a target model. But the cost between utilization and defense is unacceptable [63]. To improve the model utility, Jia *et al.* [64] introduce Memguard, which adds carefully computed noise to the output of a target model, aiming to defend an attack model while keeping performance. However, Memguard is vulnerable to threshold-based attack [55].

Adversarial regularization [24], known as AdvReg, combines a target model's training process with that of an attack model, formulating a min-max game. Yang *et al.* [65] propose Purifier to defend inference attacks via transforming the vectors of confidence score predicted by a target classifier. More recently, model compression technologies (*e.g.*, knowledge distillation [52] and pruning [37]) have been employed to protect member privacy. Based on knowledge distillation, Shejwalkar and Houmansadr [16] propose Distillation For Membership Privacy (DMP) defense, where a teacher is trained to label an unlabeled reference dataset, and those with low prediction entropy are selected to train a target model. Further, Zheng *et al.* [66] propose complementary knowledge distillation (CKD) and pseudo-complementary knowledge distillation (PCKD). Tang *et al.* [67] propose to leverage distillation to distill a training dataset via a novel ensemble architecture. However, as knowledge distillation is an indirect learning strategy [68], some critical information may be lost while mimicking the teacher. In more recent work [51], pruning is adopted to mitigate MIA while reducing model size simultaneously. But that work mainly focuses on preventing membership inference attack without explicitly considering the memory restriction.

Our work aims to address the problem of safe model compression, being different from the preceding previous work, which mainly focuses on defending against MIA. More specifically, our work aims to decrease the risk of privacy attacks (*e.g.*, MIA) and keep excellent performance when compressing a deep neural model (*e.g.*, DNN). Essentially,

the target problem in our work is a bi-objective optimization problem.

7.3 Model Compression

Due to limited memory and computation resources, model compression [37] plays a crucial role, especially when transformer-based big models [50], [69] become the mainstream. To alleviate the issue, various approaches have been proposed. For example, pruning [37], as a direct and effective approach, removes unimportant weights or structures according to certain criteria (*e.g.*, weight magnitude). Knowledge distillation [52], known as KD, transforms knowledge from a big model (named as a teacher) to a small model (named as a student) during training. Quantification [8] converts a long storage width in memory to a shorter one, *e.g.*, converting float (64 bit) to 8-bit integer [70], even to binary (1 bit) [71]. In addition, dynamic sparse training [19] (denoted as DST) is proposed as a new compression approach, and achieves surprising performance, attracting much attention from researchers. Follow-up work further introduces weight redistribution [72], [73], gradient-based weight growth [38], [72], and extra weights update in the backward pass [74], [75] to improve the sparse training performance. More recently, the DST [76] paradigm has shown great potential to deploy neural network models into edge devices at a large scale. Lin *et al*. [77] design an algorithm-system co-design framework to make on-device training possible with only 256KB of memory. While there exists numerous work in model compression to balance size and performance, privacy safety is not well considered. Differing from the preceding work, our work considers both safety and performance during model compression.

8 LIMITATIONS AND FUTURE WORK

As a study toward safe model compression, we have also identified a number of possibilities of future work that may attract more effort to this direction.

Extending SafeCompress to more heterogeneous attacks. In practice, AI software could face various types of attacks and even their combinations. The SafeCompress framework follows the *attack configurability* principle to be able to be configured against different attacks and has the potential to be extended to defend against heterogeneous attacks. In this work, we consider black-box MIA and white-box MIA, two heterogeneous attacks, to conduct experiments across computer vision and natural language processing. It is interesting and worth expecting to include more heterogeneous attacks in more domains as the SafeCompress framework is flexible and easily extensible.

Benchmarking performance-safety trade-offs between models. In this work, for any dataset, we use only one state-of-the-art model as the original big model. Prior research [22] has pointed out that, when two models do the same task with similar performance, the model with more parameters may face higher MIA risks. Then, if we compress the two models with the same sparsity restrictions, could the statement still stand or be overturned? Addressing this question could significantly help the model selection process in AI software deployment.

9 CONCLUSION

In this article, we have presented a performance-safety co-optimization framework, called SafeCompress, to address the problem of **safe model compression**, as it is critical for current large-scale AI software deployment. SafeCompress is a test-driven sparse training framework, which can be easily configured to fight against pre-specified attack mechanisms. Specifically, by simulating the attack mechanisms, SafeCompress performs both safety testing and performance evaluation and iteratively updates the compressed sparse model. Based on SafeCompress, we have implemented three concrete instances called the BMIA-SafeCompress, WMIA-SafeCompress, and MMIA-SafeCompress approaches against heterogeneous membership inference attacks. Extensive experiments have been conducted on five datasets including three computer vision tasks and two natural language processing tasks. The results demonstrate the effectiveness and generalization of our framework. We have also elaborated on how to adopt SafeCompress to other attacks, incorporate other training tricks into BMIA-SafeCompress, perform SafeCompress on a bigger model and a larger dataset, and conduct robustness analysis on the metric of TM-score, showing the flexibility, scalability, and robustness of SafeCompress. As a study toward safe model compression, we expect that our work can attract more effort to this promising direction in the new era when AI software is increasingly prevalent.

ACKNOWLEDGMENT

APPENDIX A

A.1 Configuring SafeCompress to Other Attacks

To adapt BMIA-SafeCompress (or WMIA-SafeCompress) to other attacks, the main modification is using another attack mechanism instead of black-box (or white-box) MIA in Branch 2 of Stage 2. Two examples are listed as follows.

Attribute Inference Attack. Attribute inference attack (AIA) aims to infer private information of a user (sample), such as age and location [34]. Suppose that a big model (*e.g.*, auto-encoder [78]) is trained to obtain a representation from a user's movie ratings to serve downstream tasks such as recommendation; this representation may leak the user's private information [79]. Then, we can simulate an attacker $\mathcal{A}$ as a specific AIA approach (*e.g.*, logistic regression [79]), or a basket of multiple AIA approaches [34]. Note that the safety and performance metrics also need to be modified. For safety, the metric depends on the private attribute — *e.g.*, to protect age inference, the metric could be set to mean absolute error. For performance, we can list several downstream tasks and evaluate the learned representations' effects on these tasks.

Model Inversion Attack. ModInv [80], [81] aims to reconstruct data samples from a target ML model. In other words, they allow an adversary to directly learn information about a training dataset. For example, a ModInv adversary could perform such an attack in a facial recognition system, aiming to learn the facial data of a victim whose data is used to train the model in the system. Then, we can simulate an attacker $\mathcal{A}$ as a specific ModInv approach (*e.g.*, to synthesize the training dataset relying on GANs [29], [81]). Also, the safety and performance metrics need to be modified.

A.2 Task Adaptability

SafeCompress can work not only for classification tasks. While classification models are widely used as targets in membership inference attack [42], [82], [83], [84], we adopt classification task in our work following them. Nevertheless, it is worth emphasizing that SafeCompress can be utilized for other tasks as long as those tasks meet two specific criteria: 1) Big models are built by deep neural networks, *e.g.*, DNN, CNN, RNN, LSTM, Transformer, *etc.*, to enable dynamic sparse training that is operated on the neural network connection (model weight). 2) The attack mechanism can be simulated by a defender as SafeCompress needs to use the simulated attacker to perform a safety testing. We conduct a preliminary experiment by taking a medical segmentation task for example. We use a representative UNet [85] model to perform segmentation on Breast Ultrasound Images Dataset [86]. We set sparsity to 0.1 with a black-box setting and adopt segmentation membership inference [87] in our experiment. BMIA-SafeCompress obtains 64.67% Task Acc (mIoU) and 52.83% MIA Acc, resulting in 1.22 of TM-score which is superior to UNet (uncompressed) under the same 100 training epochs whose TM-score is 1.20 (68.24% Task Acc and 56.66% MIA Acc). The result indicates that our framework also can be extended to other tasks.

APPENDIX B

B.1 Implementation Details

We perform all the experiments using Pytorch 1.8. on Ubuntu 20.04. For CV tasks, we use NVIDIA 1080Ti with the CPU of Intel Xeon Gold 5118 (4 cores, 2.3GHz) and 16GB memory. We train big models with batch size 128 for 200 epochs. For NLP tasks, we use NVIDIA 3090 with the CPU of Intel Xeon Gold 5218 (6 cores, 2.3GHz) and 40 GB memory. The batch size is 256 and the training lasts for 10 epochs.

In SafeCompress, for dynamic sparse training, we simply follow previous work [36] and use its default hyperparameters, *e.g.*, SGD optimizer, learning rate 0.1, and multi-step decay strategy for the learning rate. We only prolong the training epochs from 250 to 300 to help a sparse model enumerate more sparse topology structures. In the fine-tuning process, we set the batch size to 128 with Adam optimizer whose weight decay is set to 0.05 and betas is set to $(0.9, 0.999)$ by default. We leverage a small learning rate $5e-4$ and train until it is converged (usually $5 \sim 10$ epochs). To simulate the attacker, following previous work [20], [24], [39], we adopt ReLu as the activation function in our attack neural network models. All the network weights are initialized with normal distribution, and all biases are set to 0 by default. The batch size is 128. We use the Adam optimizer with a learning rate of 0.001 and we train the attack models 100 epochs. During the training process, we ensure that every training batch contains the same number of member and non-member data samples, aiming to prevent the attack model from being biased toward either side.

APPENDIX C

C.1 Experimental Results on BMIA-SafeCompress

C.1.1 Potential Tapping

We also explore other six aspects of our framework based on BMIA-SafeCompress. **Time Consumption.** Table 7 reports the time consumption for each approach to get a compressed model (sparsity = 0.05) on CIFAR100. In general, almost all the approaches (except Pr-DP) can finish compression in 1–5 hours, and the time consumption of BMIA-SafeCompress is comparable to that of others. Considering that BMIA-SafeCompress can achieve the best performance-safety balance (from our previous experiment results), the time consumption of BMIA-SafeCompress is generally acceptable in practice. Note that we only need to run BMIA-SafeCompress once to get a compressed model, and the model can then be deployed in hundreds of thousands of devices repeatedly.

Further, we provide a detailed analysis of training overhead below. In our framework, we only perform Stage 1 once, which requires little time (around $2s$ for VGG16) and thereby can be ignored. Stage 3 is simply about testing and also does not need much time. Hence, almost the whole training time is occupied by Stage 2 including sparse model training and attack mechanism simulation. Here, the time consumption of dynamic sparse update could be incorporated into sparse model training as the two steps are close: dynamic sparse update is triggered when the pre-defined iteration of sparse model training is over. We take VGG16 on CIFAR100 in Table 7 for example to illustrate the constitution of time consumption for each part. In our experiment, we find attack mechanism simulation consumes 2.8 hours (70%). This consumption is more than 1.2 hours (30%) of sparse model training. This result could be reasonable as attack mechanism simulation is reemployed for each update. It could be helpful to design new strategies of attack mechanism simulation to

TABLE 7
Time consumption on CIFAR100. VGG16 is the uncompressed model. Others are compressed with sparsity 0.05.

Approach	VGG16 (uncompressed)	Pr-DP	Pr-AdvReg	Pr-DMP	KD-DP	KD-AdvReg	KD-DMP	MIA-Pr	KL-Div	**BMIA-SafeCompress**
Time (h)	1.05	16.90	2.92	4.75	4.89	2.22	1.42	2.37	2.27	3.98

TABLE 8
Using other training tricks in *BMIA-SafeCompress*.

Trick	Vanilla	Dropout	Aug	AdvReg	BigPara	KD
Task Acc	69.52%	69.43%	62.03%	68.86%	68.52%	**70.02%**
MIA Acc	**51.75%**	53.94%	65.46%	52.96%	52.05%	53.10%
TM-score	**1.35**	1.29	0.95	1.30	1.32	1.32

TABLE 9
Different settings of λ for TM-score$_B$ on CIFAR100.

Approach	Coefficient λ				
	0.8	0.9	1	1.1	1.2
VGG16 (uncompressed)	0.46	0.70	1.08	1.66	2.54
Pr-DP	0.20	0.28	0.39	0.56	0.79
Pr-AdvReg	0.47	0.71	1.08	1.65	2.52
Pr-DMP	0.48	0.73	1.10	1.67	2.52
KD-DP	0.23	0.33	0.47	0.68	0.97
KD-AdvReg	0.49	0.75	1.13	1.72	2.60
KD-DMP	0.40	0.60	0.91	1.38	2.09
MIA-Pr	0.47	0.72	1.09	1.67	2.55
KL-Div	0.51	0.77	1.16	1.76	2.67
BMIA-SafeCompress	**0.58**	**0.88**	**1.34**	**2.05**	**3.14**

alleviate this issue. Since it is not our main focus, we will leave it as future work.

Best-performing Selection Strategy in Safety Test. In SafeCompress, we only keep the best-performing model in the safety testing due to time consumption. Given that we keep two models each time, the number of models sent back to "Stage 2" will be double, *e.g.*, 2, 4, 8, 16, and the time consumption for training will be exponentially growing. The cost is unbearable as the whole training process involves about 30 times of updates, resulting in 2^{30} models left. Hence, we keep only the best-performing model. This choice will potentially result in the local optima. Nevertheless, considering the saved time consumption and that our approach performs the best among all solutions in various experiments, it is bearable for possibly being local optima. We also conduct an experiment by keeping no more than 4 candidate models after the safety testing using VGG16 in CIFAR100 with sparsity set to 0.05. The whole time consumption is 15.36 hours, more than 4 times of our default setting (3.98 h). For the performance, this strategy produces 69.65% for Task Acc and 51.72% for MIA Acc. Compared to our default setting (69.52% Task Acc and 51.75% MIA Acc), the improvement is marginal, empirically demonstrating that the best-performing selection strategy in SafeCompress is more reasonable in light of the performance and significantly saved time consumption.

Flexibility with Other Training Tricks. It is worth noting that SafeCompress is also flexible to incorporate other training tricks not mentioned in previous experiments. To illustrate this flexibility, we try to incorporate five other widely-used training tricks in BMIA-SafeCompress, including *dropout* [88], *data augmentation* (*Aug*, including random cropping, resizing, and flipping), *adversary regularization* (*AdvReg*), *sparse model initialization with the big model's parameter values (BigPara)*, and *knowledge distillation (KD)* where the sparse model's prediction score approximates the big model's [52]. Table 8 shows the results of CIFAR100.

Compared to our BMIA-SafeCompress implementation in main experiments, we find that most training tricks achieve similar TM-scores. The result indicates that SafeCompress is compatible with most training tricks. Then, in practice, given a specific dataset/task, we could enumerate different combinations of training tricks and find the best one in the SafeCompress framework.

Another interesting observation is that *data augmentation significantly increases MIA Acc* (*i.e.*, reduces safety) for the final compressed model. The possible reason is that data augmentation allows training samples to be memorized more easily, as the variants of original samples (*e.g.*, flipping and resizing) are also used for training. Notably, while not focusing on compressed sparse models, Yu *et al.* also point out that data augmentation may lead to a significantly higher MIA risk [89]. Inspired by the work [89], SafeCompress may also be extended to a useful framework to benchmark different training tricks' impacts on compressed models' safety. This task could be an interesting future direction.

Scale to Big Model and Large Dataset. In SafeCompress, we use CIFAR10, CIFAR100, and Tiny ImageNet in our experiment as they are widely used in membership inference attacks in previous work [42], [82], [83], [84]. To demonstrate the scalability, we scale model size by using ResNet50 [46] and train on a large dataset ImageNet [53]. We set sparsity to 0.1 following previous work [36] with a black-box setting. In our experiment, BMIA-SafeCompress obtains 73.80% Task Acc and 56.67% MIA Acc, resulting in 1.3 of TM-score which is superior to ResNet50 (uncompressed) under the same 100 training epochs whose TM-score is 1.16 (76.13% Task Acc and 65.23% MIA Acc). The result indicates that our approach is also scalable.

Robustness on Metric Variants over Baselines. We can achieve different requirements for safety and performance in our approach by adjusting the coefficient λ (*e.g.*, 0.8, 0.9, 1.0, 1.1, and 1.2) of the TM-score. To show whether our approach is robust to this situation over the other eight baselines, we illustrate the results of CIFAR100 in Table 9. It is observed that BMIA-SafeCompress performs the best over other baselines among all the settings. This result indicates that BMIA-SafeCompress is capable of meeting different degrees of safety and performance balance over baselines. This result also implies that our SafeCompress may naturally have the property of robustness against metric alteration to some extent.

TABLE 10
Results on CIFAR100 with sparsity 0.05. We repeat three times and report the mean result and standard error.

Approach	Task Acc	MIA Acc	TM-score
VGG16 (uncompressed)	72.70 ± 0.06	67.87 ± 0.52	1.07 ± 0.0081
Pr-DP	31.64 ± 0.11	80.14 ± 0.11	0.39 ± 0.0014
Pr-AdvReg	67.93 ± 0.05	62.75 ± 0.16	1.08 ± 0.0023
Pr-DMP	63.44 ± 0.15	57.78 ± 0.24	1.10 ± 0.0024
KD-DP	36.81 ± 0.24	78.08 ± 0.12	0.47 ± 0.0025
KD-AdvReg	64.25 ± 0.14	56.73 ± 0.14	1.13 ± 0.0008
KD-DMP	62.98 ± 0.17	69.08 ± 0.19	0.91 ± 0.0003
MIA-Pr	68.75 ± 0.19	62.84 ± 0.24	1.09 ± 0.0014
KL-Div	63.98 ± 0.10	55.10 ± 0.32	1.16 ± 0.0050
BMIA-SafeCompress	69.63 ± 0.14	52.04 ± 0.21	1.34 ± 0.0048

TABLE 11
Defend against black-box MIA. Task Acc (performance), MIA Acc (safety) and TM-score$_B$ results on CIFAR10. The best results are marked in **bold**.

Approach	Sparsity=0.05			Sparsity=0.1		
	Task Acc	MIA Acc	TM-score	Task Acc	MIA Acc	TM-score
AlexNet (uncompressed)	87.41%	58.34%	1.50	87.41%	58.34%	1.50
Pr-DP	54.57%	68.35%	0.80	62.87%	66.49%	0.95
Pr-AdvReg	82.72%	52.74%	1.57	86.19%	54.83%	1.57
Pr-DMP	81.66%	55.63%	1.47	83.73%	54.05%	1.55
KD-DP	72.09%	59.94%	1.20	73.27%	60.42%	1.21
KD-AdvReg	80.47%	54.38%	1.48	**86.84%**	55.59%	1.56
KD-DMP	82.23%	61.63%	1.33	86.26%	63.45%	1.36
MIA-Pr	82.26%	53.15%	1.55	86.11%	56.12%	1.53
KL-Div	63.89%	**52.00%**	1.23	72.19%	53.21%	1.36
BMIA-SafeCompress	83.94%	52.97%	**1.59**	85.31%	**52.37%**	**1.63**
+ Adversarial Training	**84.20%**	53.12%	**1.59**	85.38%	52.52%	**1.63**

TABLE 12
Defend against black-box MIA. Task Acc (performance), MIA Acc (safety) and TM-score$_B$ results on Tiny ImageNet. The best results are marked in **bold**.

Approach	Sparsity=0.05			Sparsity=0.1		
	Task Acc	MIA Acc	TM-score	Task Acc	MIA Acc	TM-score
ResNet18 (uncompressed)	65.48%	69.73%	0.94	65.48%	69.73%	0.94
Pr-DP	19.26%	71.07%	0.27	24.56%	74.17%	0.33
Pr-AdvReg	60.10%	57.62%	1.04	61.16%	63.93%	0.96
Pr-DMP	55.56%	60.03%	0.92	59.61%	64.56%	0.92
KD-DP	17.10%	52.32%	0.33	17.45%	52.44%	0.33
KD-AdvReg	52.34%	53.27%	0.98	54.48%	53.60%	1.02
KD-DMP	53.71%	57.18%	0.94	57.16%	55.65%	1.03
MIA-Pr	58.36%	57.92%	1.01	60.91%	61.03%	1.00
KL-Div	48.22%	55.26%	0.87	56.82%	57.19%	0.99
BMIA-SafeCompress	**63.81%**	52.46%	1.22	**65.15%**	52.79%	**1.24**
+ Adversarial Training	63.12%	**51.38%**	**1.23**	64.24%	**52.01%**	**1.24**

Statistical Analysis with Multiple Runs. Due to the stochastic nature of optimization algorithms and DNN training procedures, we conduct multiple runs in a case to present a statistical analysis. We use BMIA-Safecompress using VGG16 on CIFAR100 with sparsity set to 0.05. We run each approach three times with random seeds and report the results in Tab 10. One can see that our approach is statistically superior to others. To further support our conclusion, we perform a statistical test using the student's t test (considering the sample numbers) and compare the P value with that of the best-performing baseline on each metric. For Task Acc, we compare BMIA-SafeCompress with MIA-Pr, and the P value is 0.007, smaller than the default alpha level (0.01). For MIA Acc, we compare BMIA-SafeCompress with KL-Div,

and the P value is 0.0004 ($0.0004 < 0.01$). As for TM-Score, we compare BMIA-SafeCompress with KL-Div. The P value ($1.2e-5$) is also smaller than 0.01. These results show our conclusion is significant.

C.1.2 Results on CV Datasets

CIFAR10 (AlexNet). The results on CIFAR10 are presented in Table 11. It can be seen that BMIA-SafeCompress obtains almost the best performance in Task Acc and maintains a pretty strong defensive ability when sparsity is 0.05. Thanks to the excellent effectiveness in both aspects, BMIA-SafeCompress produces the highest TM-score, showing its outstanding ability to make the performance-safety trade-off. When the sparsity is set to 0.1, our approach produces 85.31% for Task

TABLE 13
Defend against black-box MIA. Task Acc, MIA Acc, and TM-score$_B$ on Yelp-5.

Approach	Sparsity=0.5		
	Task Acc	MIA Acc	TM-score
BERT (Uncompressed)	62.21%	71.15%	0.87
Pr-DP	60.77%	64.79%	0.94
Pr-AdvReg	61.17%	73.34%	0.83
Pr-DMP	61.32%	66.41%	0.92
KD-DP	58.21%	65.36%	0.89
KD-AdvReg	61.38%	67.40%	0.91
KD-DMP	59.69%	70.25%	0.85
MIA-Pr	61.80%	73.85%	0.84
KL-Div	60.03%	**64.24%**	0.93
BMIA-SafeCompress	61.66%	64.37%	0.96
+ Adversarial Training	**62.63%**	64.90%	**0.97**

ACC, slightly inferior to the best result (86.84%) produced by KD-AdvReg. However, BMIA-SafeCompress decreases MIA Acc to 52.37%, much lower than KD-AdvReg (55.59%), leading to the highest TM-score again. Finally, leveraging Adversarial Training leads to further improvement in Task Acc at the cost of a slight increase in MIA Acc.

Tiny ImageNet (ResNet18). We present the results of Tiny ImageNet in Table 12. Excitingly, BMIA-SafeCompress produces the best accuracy for Task Acc in both sparsity 0.05 and 0.1, outperforming all the baselines by a large margin. In addition, BMIA-SafeCompress obtains competitive MIA Acc in both sparsity settings compared to the best defense effects of baselines. Specifically, BMIA-SafeCompress achieves 52.46% (sparsity 0.05) and 52.79% (sparsity 0.1) in MIA Acc; The best MIA Acc results among baselines are 52.32% (sparsity 0.05) and 52.44% (sparsity 0.1). The gaps are tiny (0.14% and 0.35%). Finally, BMIA-SafeCompress (+ Adversarial Training) achieves the best TM-score in both sparsity settings.

C.1.3 Results on NLP Dataset

Yelp-5 (BERT). The results are reported in Table 13. It can be seen that BMIA-SafeCompress produces a competitive Task Acc. At the same time, it decreases MIA Acc by 6.78% compared to the uncompressed model, outperforming all the baselines. BMIA-SafeCompress produces 0.96 for the TM-score, indicating its ability to balance task performance and safety. Besides, when attaching Adversarial Training to BMIA-SafeCompress, we further improve the Task Acc, leading to the highest TM-score.

C.2 Experimental Results on WMIA-SafeCompress

C.2.1 Results on CV Datasets

CIFAR10 (AlexNet). We provide results as shown in Table 14 to further show the effectiveness. It can be seen that WMIA-SafeCompress obtains a very competitive performance in Task Acc and maintains a pretty strong defensive ability when sparsity is 0.05. As a result, WMIA-SafeCompress produces the highest TM-score, showing its outstanding ability to make the performance-safety trade-off. When the sparsity is set to 0.1, our approach produces 83.08% for Task ACC, inferior to the best result (86.84%) produced

by KD-AdvReg. However, WMIA-SafeCompress decreases MIA Acc to 52.38%, much lower than KD-AdvReg (56.95%), leading to the highest TM-score again (the same as KD-DMP). Finally, incorporating Adversarial Training with WMIA-SafeCompress leads to a further improvement in Task Acc at the cost of a relatively slight increase of MIA Acc, especially for sparsity 0.05.

Tiny ImageNet (ResNet18). The results on Tiny ImageNet here as shown in Table 15. Excitingly, WMIA-SafeCompress produces the best accuracy for Task Acc in both sparsity 0.05 and 0.1, outperforming all the baselines by a large margin. Additionally, WMIA-SafeCompress obtains relatively low MIA Acc in both sparsity settings compared to the defense effects of baselines. Specifically, the best MIA Acc results among baselines are 51.09% (sparsity 0.05) and 51.17% (sparsity 0.1) that are both produced by KD-DP costing at intolerable performance drop; The worst MIA Acc results among baselines are 63.77% (sparsity 0.05) and 70.37% (sparsity 0.1) that are both produced by Pr-AdvReg; In contrast, WMIA-SafeCompress performs slightly well, achieving 53.09% (sparsity 0.05) and 54.23% (sparsity 0.1) in MIA Acc and obtaining the best TM-score in both sparsity settings. Finally, equipping with Adversarial Training greatly decreases privacy risk, especially for sparsity 0.1 where the TM-score rises to 1.18, outperforming other baselines by a large margin.

C.2.2 Results on NLP Dataset

AG News (RoBERTa). To further validate the effectiveness and generalization of WMIA-SafeCompress, we conduct experiments on AG News. As indicated in Table 16, WMIA-SafeCompress obtains a competitive performance in Task Acc achieving 87.31%, slightly inferior to the highest 88.10% produced by KD-AdvReg. However, our approach decreases MIA Acc to 55.75%, lower than KD-AdvReg (56.33%). When comparing with KD-DP, we observe that WMIA-SafeCompress is a little bit inferior to KD-DP in defense, but outperforms KL-Div by a large margin in Task Acc. As a result, WMIA-SafeCompress achieves the highest TM-score, showing its great trade-off capability. Finally, leveraging Adversarial Training leads to a further enhancement in all the metrics.

C.3 Experimental Results on MMIA-SafeCompress

C.3.1 Results on CV Datasets

CIFAR10 (AlexNet). We present results on CIFAR10 in Table 17. It is observed that MMIA-SafeCompress achieves competitive classification performance among all baselines in both sparsity. It also decreases MIA Acc$_B$ and MIA Acc$_W$ compared to the uncompressed model, greatly improving the defense capability of the compressed model. Further, we can see that MMIA-SafeCompress produces the highest TM-score, demonstrating its outstanding capability to make the performance-safety trade-off. Besides, equipping MMIA-SafeCompress with Adversarial Training leads to great improvement in Task Acc and generally strengthens model defense ability. Consequently, these enhancements contribute to further boosting in TM-score.

Tiny ImageNet (ResNet18). The results on Tiny ImageNet are presented in Table 18. It is easily seen that MMIA-SafeCompress produces 62.23% (sparsity=0.05) and 63.53%

TABLE 14
Defend against white-box MIA. Task Acc (performance), MIA Acc (safety) and TM-score$_W$ results on CIFAR10. The best results are marked in **bold**.

Approach	Sparsity=0.05			Sparsity=0.1		
	Task Acc	MIA Acc	TM-score	Task Acc	MIA Acc	TM-score
AlexNet (uncompressed)	87.41%	59.13%	1.48	87.41%	59.13%	1.48
Pr-DP	54.57%	52.83%	1.03	62.87%	54.23%	1.16
Pr-AdvReg	82.72%	54.84%	1.51	86.19%	57.82%	1.49
Pr-DMP	81.66%	53.69%	1.52	83.73%	54.84%	1.53
KD-DP	72.09%	**50.95%**	1.41	73.27%	**50.88%**	1.44
KD-AdvReg	80.47%	54.91%	1.47	**86.84%**	56.95%	1.52
KD-DMP	82.23%	54.27%	1.52	86.26%	54.09%	1.59
MIA-Pr	82.26%	55.48%	1.48	86.11%	58.14%	1.48
KL-Div	63.89%	52.32%	1.22	72.19%	53.32%	1.35
WMIA-SafeCompress	81.33%	53.03%	1.53	83.08%	52.38%	**1.59**
+ Adversarial Training	**83.38%**	53.32%	**1.56**	84.70%	53.34%	**1.59**

TABLE 15
Defend against white-box MIA. Task Acc (performance), MIA Acc (safety) and TM-score$_W$ results on Tiny ImageNet. The best results are marked in **bold**.

Approach	Sparsity=0.05			Sparsity=0.1		
	Task Acc	MIA Acc	TM-score	Task Acc	MIA Acc	TM-score
ResNet18 (uncompressed)	65.48%	80.15%	0.82	65.48%	80.15%	0.82
Pr-DP	19.26%	60.88%	0.32	24.56%	66.98%	0.37
Pr-AdvReg	60.10%	63.77%	0.94	61.16%	70.37%	0.87
Pr-DMP	55.56%	57.72%	0.96	59.61%	58.72%	1.02
KD-DP	17.10%	**51.09%**	0.33	17.45%	**51.17%**	0.34
KD-AdvReg	52.34%	58.08%	0.90	54.48%	54.81%	0.99
KD-DMP	53.71%	55.65%	0.97	57.16%	54.45%	1.05
MIA-Pr	58.36%	58.66%	0.99	60.91%	69.47%	0.88
KL-Div	48.22%	54.78%	0.88	56.82%	54.96%	1.03
WMIA-SafeCompress	**60.86%**	53.09%	**1.15**	**61.57%**	54.23%	1.14
+ Adversarial Training	59.50%	51.85%	**1.15**	60.81%	51.72%	**1.18**

TABLE 16
Defend against white-box MIA. Task Acc, MIA Acc, and TM-score$_W$ on AG News.

Approach	Sparsity=0.5		
	Task Acc	MIA Acc	TM-score
RoBERTa (Uncompressed)	89.20%	57.43%	1.55
Pr-DP	87.28%	55.96%	1.56
Pr-AdvReg	87.24%	56.12%	1.55
Pr-DMP	86.38%	55.28%	1.56
KD-DP	82.64%	**53.02%**	1.56
KD-AdvReg	**88.10%**	56.33%	1.56
KD-DMP	87.24%	55.95%	1.56
MIA-Pr	87.91%	57.25%	1.54
KL-Div	83.26%	54.18%	1.54
WMIA-SafeCompress	87.31%	55.75%	1.57
+ Adversarial Training	87.63%	55.06%	**1.59**

(sparsity=0.1) for Task Acc, outperforming all baselines by a large margin. In addition, MMIA-SafeCompress maintains considerable defense capability. For example, it holds the lowest MIA Acc$_B$ and the second-lowest MIA Acc$_W$ when the sparsity is set to 0.05. Thus, MMIA-SafeCompress achieves the highest TM-score in both sparsity and even outperforms the other eight approaches by a large margin. When combined with Adversarial Training, MMIA-SafeCompress further enhances defense ability via costing at

a slight performance drop and brings little bad influence to the final TM-score.

C.3.2   Results on NLP Datasets

AG News (RoBERTa). We further conduct experiments on AG News to validate the effectiveness and generalization of MMIA-SafeCompress. The results are reported in Table 19. It can be seen that MMIA-SafeCompress achieves a competitive performance in Task Acc among all the approaches. Moreover, MMIA-SafeCompress maintains considerable defense capability. For example, it holds the second-lowest MIA Acc$_B$ and the third-lowest MIA Acc$_W$. Consequently, these advantages lead MMIA-SafeCompress to the highest TM-score. Further, combining MMIA-SafeCompress with Adversarial Training improves the task performance and defense ability against the black-box attack while causing a slightly increased risk for white-box attack. Finally, such alterations bring little bad influence on TM-score.

APPENDIX D

D.1   Discussion about trade-off between performance, safety, compressed models, and time costs

First, the size to which a big model should be compressed is determined by the scenario or the requirements. For example, to deploy AI software on smartphones, when one hopes the

TABLE 17
Defend against black-box and white-box MIAs. Task Acc (performance), MIA Acc (safety) and TM-score$_M$ results on CIFAR10. The best results are marked in **bold**.

Approach	Sparsity=0.05				Sparsity=0.1			
	Task Acc	MIA Acc$_B$	MIA Acc$_W$	TM-score	Task Acc	MIA Acc$_B$	MIA Acc$_W$	TM-score
AlexNet (uncompressed)	87.41%	58.34%	59.13%	1.49	87.41%	58.34%	59.13%	1.49
Pr-DP	54.57%	68.35%	52.83%	0.92	62.87%	66.49%	54.23%	1.05
Pr-AdvReg	82.72%	52.74%	54.84%	1.53	86.19%	54.83%	57.82%	1.53
Pr-DMP	81.66%	55.63%	53.69%	1.49	83.73%	54.05%	54.84%	1.54
KD-DP	72.09%	59.94%	**50.95%**	1.31	73.27%	60.42%	**50.88%**	1.33
KD-AdvReg	80.47%	54.38%	54.91%	1.52	**86.94%**	55.59%	56.95%	1.54
KD-DMP	82.23%	61.63%	54.27%	1.42	86.26%	63.45%	54.09%	1.48
MIA-Pr	82.26%	53.15%	55.48%	1.52	86.11%	56.12%	58.14%	1.51
KL-Div	63.89%	52.00%	52.32%	1.22	72.19%	53.21%	53.32%	1.36
MMIA-SafeCompress	81.64%	51.79%	54.98%	1.53	83.98%	**52.81%**	55.86%	1.55
+ Adversarial Training	**83.29%**	**50.28%**	54.70%	**1.59**	85.12%	54.21%	54.94%	**1.56**

TABLE 18
Defend against black-box and white-box MIAs. Task Acc (performance), MIA Acc (safety) and TM-score$_M$ results on Tiny ImageNet. The best results are marked in **bold**.

Approach	Sparsity=0.05				Sparsity=0.1			
	Task Acc	MIA Acc$_B$	MIA Acc$_W$	TM-score	Task Acc	MIA Acc$_B$	MIA Acc$_W$	TM-score
ResNet18 (uncompressed)	65.48%	69.73%	80.15%	0.88	65.48%	69.73%	80.15%	0.88
Pr-DP	19.26%	71.07%	60.88%	0.29	24.56%	74.17%	66.98%	0.35
Pr-AdvReg	60.10%	57.62%	63.77%	0.99	61.16%	63.93%	70.37%	0.91
Pr-DMP	55.56%	60.30%	57.72%	0.94	59.61%	64.56%	58.72%	0.97
KD-DP	17.10%	52.32%	**51.09%**	0.33	17.45%	**52.44%**	**51.17%**	0.34
KD-AdvReg	52.34%	53.27%	58.08%	0.94	54.48%	53.60%	54.81%	1.01
KD-DMP	53.71%	57.18%	55.65%	0.95	57.16%	55.65%	54.45%	1.04
MIA-Pr	58.36%	57.92%	58.66%	1.00	60.91%	61.03%	69.47%	0.94
KL-Div	48.22%	55.26%	54.78%	0.88	56.82%	57.19%	54.96%	1.01
MMIA-SafeCompress	**62.23%**	50.59%	54.13%	**1.19**	**63.53%**	53.09%	56.34%	**1.16**
+ Adversarial Training	61.76%	**50.28%**	53.55%	**1.19**	62.51%	52.77%	55.94%	1.15

TABLE 19
Defend against black-box and white-box MIAs. Task Acc, MIA Acc$_B$, MIA Acc$_W$, and TM-score$_M$ on AG News.

Approach	Sparsity=0.5			
	Task Acc	MIA Acc$_B$	MIA Acc$_W$	TM-score
RoBERTa (Uncompressed)	89.20%	57.59%	57.43%	1.55
Pr-DP	87.28%	56.07%	55.96%	1.56
Pr-AdvReg	87.24%	57.58%	56.12%	1.53
Pr-DMP	86.38%	56.31%	55.28%	1.55
KD-DP	82.64%	56.13%	**53.02%**	1.52
KD-AdvReg	**88.10%**	56.56%	56.33%	1.56
KD-DMP	87.24%	56.64%	55.95%	1.55
MIA-Pr	87.91%	57.04%	57.25%	1.54
Kl-Div	83.26%	**54.09%**	54.18%	1.54
MMIA-SafeCompress	87.04%	55.92%	55.04%	**1.57**
+ Adversarial Training	87.31%	55.45%	55.65%	**1.57**

important. We also conduct experiments with different λ presented in Figure 5 in our article. As for time costs, it is related to the training times of SafeCompress. Theoretically, we can stop at any time. In general, the more the training time is (if we do not consider overfitting), the better the performance is. Hence, we could predefine a number of training iterations to balance the time consumption and performance.

REFERENCES

[1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 1106–1114, 2012.

[2] Z. Zhang, Y. Li, J. Wang, B. Liu, D. Li, X. Chen, Y. Guo, and Y. Liu, "ReMoS: Reducing Defect Inheritance in Transfer Learning via Relevant Model Slicing," in *Proceedings of the 44th International Conference on Software Engineering (ICSE)*, 2022, pp. 1856–1868.

[3] C. Riquelme, J. Puigcerver, B. Mustafa, M. Neumann, R. Jenatton, A. Susano Pinto, D. Keysers, and N. Houlsby, "Scaling Vision with Sparse Mixture of Experts," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 8583–8595, 2021.

[4] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language Models are Few-Shot Learners," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 1877–1901, 2020.

[5] W. W. Pugh, "Compressing Java Class Files," in *Proceedings of the 1999 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, 1999, pp. 247–258.

deployed size is 10% of an original uncompressed one, the sparsity is 0.1. Secondly, to balance the trade-off between performance and safety, we can simply adjust the coefficient λ of the TM-score in Eq (6). For example, a user may think safety is more important than performance, and therefore she can decrease λ to a smaller value; Otherwise, she can increase λ to a bigger value. In most experiments, we set $\lambda = 1$ by default as we consider performance and safety to be equally

[6] M. Drinić, D. Kirovski, and H. Vo, "PPMexe: Program Compression," *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 2007.

[7] X. Jiao, Y. Yin, L. Shang, X. Jiang, X. Chen, L. Li, F. Wang, and Q. Liu, "TinyBERT: Distilling BERT for Natural Language Understanding," *arXiv preprint arXiv:1909.10351*, 2019.

[8] S. Han, J. Pool, J. Tran, and W. Dally, "Learning both Weights and Connections for Efficient Neural Networks," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 1135–1143, 2015.

[9] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," *arXiv preprint arXiv:1810.04805*, 2018.

[10] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman, "GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding ," *arXiv preprint arXiv:1804.07461*, 2018.

[11] N. Carlini, C. Liu, Ú. Erlingsson, J. Kos, and D. Song, "The Secret Sharer: Evaluating and Testing Unintended Memorization in Neural Networks," in *28th USENIX Security Symposium (USENIX Security)*, 2019, pp. 267–284.

[12] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership Inference Attacks Against Machine Learning Models," *2017 IEEE Symposium on Security and Privacy (SP)*, pp. 3–18, 2017.

[13] L. Deng, G. Li, S. Han, L. Shi, and Y. Xie, "Model Compression and Hardware Acceleration for Neural Networks: A Comprehensive Survey," *Proceedings of the IEEE*, pp. 485–532, 2020.

[14] Y. Liu, S. Gao, A. Huang, J. Zhu, L. Xu, and A. Nathan, "Ensemble Learning-Based Technique for Force Classifications in Piezoelectric Touch Panels," *IEEE Sensors Journal*, pp. 9540–9549, 2020.

[15] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep Learning with Differential Privacy," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2016, pp. 308–318.

[16] V. Shejwalkar and A. Houmansadr, "Membership Privacy for Machine Learning Models Through Knowledge Transfer," *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pp. 9549–9557, 2021.

[17] X. Yuan and L. Zhang, "Membership Inference Attacks and Defenses in Neural Network Pruning," in *31st USENIX Security Symposium (USENIX Security)*, 2022, pp. 4561–4578.

[18] K. L. Beck, "Test-driven Development - By example," in *The Addison-Wesley Signature Series*, 2003.

[19] D. C. Mocanu, E. Mocanu, P. Stone, P. H. Nguyen, M. Gibescu, and A. Liotta, "Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science," *Nature Communications*, pp. 1–12, 2018.

[20] A. Sablayrolles, M. Douze, C. Schmid, Y. Ollivier, and H. Jégou, "White-box vs Black-box: Bayes Optimal Strategies for Membership Inference," in *Proceedings of International Conference on Machine Learning (ICML)*, 2019, pp. 5558–5567.

[21] S. Yeom, I. Giacomelli, M. Fredrikson, and S. Jha, "Privacy Risk in Machine Learning: Analyzing the Connection to Overfitting," in *Proceedings of the 2018 IEEE 31st Computer Security Foundations Symposium (CSF)*, 2018, pp. 268–282.

[22] M. Nasr, R. Shokri, and A. Houmansadr, "Comprehensive Privacy Analysis of Deep Learning: Passive and Active White-box Inference Attacks against Centralized and Federated Learning," in *Proceedings of the 2019 IEEE Symposium on Security and Privacy (SP)*, 2019, pp. 739–753.

[23] H. Hu, Z. Salcic, L. Sun, G. Dobbie, P. S. Yu, and X. Zhang, "Membership Inference Attacks on Machine Learning: A Survey," *ACM Computing Surveys (CSUR)*, pp. 235:1–235:37, 2022.

[24] M. Nasr, R. Shokri, and A. Houmansadr, "Machine Learning with Membership Privacy using Adversarial Regularization," *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, pp. 634–646, 2018.

[25] J. Zhu, L. Wang, and X. Han, "Safety and Performance, Why not Both? Bi-Objective Optimized Model Compression toward AI Software Deployment," in *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2022, pp. 1–13.

[26] P. V. et al., "Scipy 1.0: fundamental algorithms for scientific computing in python," *Nature Methods*, pp. 261 – 272, 2020.

[27] D. C. Mocanu *et al.*, *Network computations in artificial intelligence*. Technische Universiteit Eindhoven, 2017.

[28] L. Ma, X. Ma, J. Gao, X. Jiao, Z. Yu, C. Zhang, W. Ruan, Y. Wang, W. Tang, and J. Wang, "Distilling Knowledge from Publicly Available Online EMR Data to Emerging Epidemic for Prognosis," in *Proceedings of the Web Conference (WWW)*, 2021, pp. 3558–3568.

[29] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," *Communications of the ACM*, pp. 139–144, 2020.

[30] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards Deep Learning Models Resistant to Adversarial Attacks," *arXiv preprint arXiv:1706.06083*, 2017.

[31] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 2223–2232.

[32] T. Miyato, A. M. Dai, and I. Goodfellow, "Adversarial Training Methods for Semi-Supervised Text Classification," *arXiv preprint arXiv:1605.07725*, 2016.

[33] A. Shafahi, M. Najibi, M. A. Ghiasi, Z. Xu, J. Dickerson, C. Studer, L. S. Davis, G. Taylor, and T. Goldstein, "Adversarial Training for Free!" *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

[34] X. Han, H. Huang, and L. Wang, "F-PAD: Private Attribute Disclosure Risk Estimation in Online Social Networks," *IEEE Transactions on Dependable and Secure Computing (TDSC)*, pp. 1054–1069, 2019.

[35] D. C. Mocanu, E. Mocanu, P. Stone, P. H. Nguyen, M. Gibescu, and A. Liotta, "Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science," *Nature Communications*, pp. 1–12, 2018.

[36] S. Liu, L. Yin, D. C. Mocanu, and M. Pechenizkiy, "Do We Actually Need Dense Over-Parameterization? In-Time Over-Parameterization in Sparse Training," in *Proceedings of the 38th International Conference on Machine Learning (ICML)*, 2021, pp. 6989–7000.

[37] S. Han, H. Mao, and W. J. Dally, "Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding," *arXiv preprint arXiv:1510.00149*, 2015.

[38] U. Evci, T. Gale, J. Menick, P. S. Castro, and E. Elsen, "Rigging the Lottery: Making All Tickets Winners," in *Proceedings of International Conference on Machine Learning (ICML)*, 2020, pp. 2943–2952.

[39] Y. Liu, R. Wen, X. He, A. Salem, Z. Zhang, M. Backes, E. De Cristofaro, M. Fritz, and Y. Zhang, "ML-Doctor: Holistic Risk Assessment of Inference Attacks Against Machine Learning Models," *arXiv preprint arXiv:2102.02551*, 2021.

[40] J. Zhu, H. Huang, B. Li, and L. Wang, "E-CRF: Embedded Conditional Random Field for Boundary-caused Class Weights Confusion in Semantic Segmentation," in *Proceedings of the Eleventh International Conference on Learning Representations (ICLR)*, 2022.

[41] L. Melis, C. Song, E. De Cristofaro, and V. Shmatikov, "Exploiting Unintended Feature Leakage in Collaborative Learning," in *Proceedings of the 2019 IEEE Symposium on Security and Privacy (SP)*, 2019, pp. 691–706.

[42] H. Liu, J. Jia, W. Qu, and N. Z. Gong, "EncoderMI: Membership Inference against Pre-trained Encoders in Contrastive Learning," in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2021, pp. 2081–2095.

[43] A. Krizhevsky, G. Hinton *et al.*, "Learning Multiple Layers of Features from Tiny Images," 2009.

[44] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, 2015.

[45] Y. Le and X. Yang, "Tiny ImageNet Visual Recognition Challenge," *CS 231N*, p. 3, 2015.

[46] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.

[47] G. M. Del Corso, A. Gulli, and F. Romani, "Ranking A Stream of News," in *Proceedings of the 14th international conference on World Wide Web (WWW)*, 2005, pp. 97–106.

[48] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "RoBERTa: A Robustly Optimized BERT Pretraining Approach," *arXiv preprint arXiv:1907.11692*, 2019.

[49] X. Zhang, J. Zhao, and Y. LeCun, "Character-level Convolutional Networks for Text Classification," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 649–657, 2015.

[50] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," *arXiv preprint arXiv:1810.04805*, 2018.

[51] Y. Wang, C. Wang, Z. Wang, S. Zhou, H. Liu, J. Bi, C. Ding, and S. Rajasekaran, "Against Membership Inference Attack: Pruning is All You Need," in *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI)*, 2021, pp. 3141–3147.

[52] G. Hinton, O. Vinyals, J. Dean *et al.*, "Distilling the Knowledge in a Neural Network," *arXiv preprint arXiv:1503.02531*, 2015.

[53] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 248–255.

[54] A. Salem, Y. Zhang, M. Humbert, P. Berrang, M. Fritz, and M. Backes, "ML-Leaks: Model and Data Independent Membership Inference Attacks and Defenses on Machine Learning Models," in *Proceedings of Network and Distributed Systems Security (NDSS) Symposium*, 2019.

[55] L. Song and P. Mittal, "Systematic Evaluation of Privacy Risks of Machine Learning Models," in *30th USENIX Security Symposium (USENIX Security)*, 2021, pp. 2615–2632.

[56] Y. Liu, Z. Zhao, M. Backes, and Y. Zhang, "Membership Inference Attacks by Exploiting Loss Trajectory," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2022, pp. 2085–2098.

[57] J. Hayes, L. Melis, G. Danezis, and E. De Cristofaro, "LOGAN: Membership Inference Attacks Against Generative Models," *Proceedings on Privacy Enhancing Technologies Symposium (PoPETs)*, pp. 133–152, 2019.

[58] C. Song and A. Raghunathan, "Information Leakage in Embedding Models," in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2020, pp. 377–390.

[59] V. Duddu, A. Boutet, and V. Shejwalkar, "Quantifying Privacy Leakage in Graph Embedding," in *Proceedings of the 17th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services (MobiQuitous)*, 2020, pp. 76–85.

[60] W.-C. Tseng, W.-T. Kao, and H.-y. Lee, "Membership Inference Attacks Against Self-supervised Speech Models," *arXiv preprint arXiv:2111.05113*, 2021.

[61] P. Hu, Z. Wang, R. Sun, H. Wang, and M. Xue, "M^4I: Multi-modal Models Membership Inference," *Advances in Neural Information Processing Systems*, vol. 35, pp. 1867–1882, 2022.

[62] M. A. Rahman, T. Rahman, R. Laganière, N. Mohammed, and Y. Wang, "Membership Inference Attack against Differentially Private Deep Learning Model," *Transactions on Data Privacy (TDP)*, pp. 61–79, 2018.

[63] B. Jayaraman and D. Evans, "Evaluating Differentially Private Machine Learning in Practice," in *28th USENIX Security Symposium (USENIX Security)*, 2019, pp. 1895–1912.

[64] J. Jia, A. Salem, M. Backes, Y. Zhang, and N. Z. Gong, "MemGuard: Defending against Black-Box Membership Inference Attacks via Adversarial Examples," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2019, pp. 259–274.

[65] Z. Yang, L. Wang, D. Yang, J. Wan, Z. Zhao, E.-C. Chang, F. Zhang, and K. Ren, "Purifier: Defending Data Inference Attacks via Transforming Confidence Scores," in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2023, pp. 10871–10879.

[66] J. Zheng, Y. Cao, and H. Wang, "Resisting membership inference attacks through knowledge distillation," *Neurocomputing*, pp. 114–126, 2021.

[67] X. Tang, S. Mahloujifar, L. Song, V. Shejwalkar, M. Nasr, A. Houmansadr, and P. Mittal, "Mitigating Membership Inference Attacks by Self-Distillation Through a Novel Ensemble Architecture," in *31st USENIX Security Symposium (USENIX Security)*, 2022, pp. 1433–1450.

[68] E. J. Crowley, G. Gray, and A. J. Storkey, "Moonshine: Distilling with Cheap Convolutions," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018, pp. 2893–2903.

[69] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," *Proceedings of International Conference on Learning Representations (ICLR)*, 2021.

[70] B. Jacob, S. Kligys, B. Chen, M. Zhu, M. Tang, A. Howard, H. Adam, and D. Kalenichenko, "Quantization and Training of Neural Networks for Efficient Integer-Arithmetic-Only Inference," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 2704–2713.

[71] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio, "Binarized neural networks," *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.

[72] T. Dettmers and L. Zettlemoyer, "Sparse Networks from Scratch: Faster Training without Losing Performance," *arXiv preprint arXiv:1907.04840*, 2019.

[73] H. Mostafa and X. Wang, "Parameter Efficient Training of Deep Convolutional Neural Networks by Dynamic Sparse Reparameterization," in *Proceedings of International Conference on Machine Learning (ICML)*, 2019, pp. 4646–4655.

[74] S. Jayakumar, R. Pascanu, J. Rae, S. Osindero, and E. Elsen, "Top-KAST: Top-K Always Sparse Training," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 20744–20754, 2020.

[75] M. A. Raihan and T. Aamodt, "Sparse Weight Activation Training," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 15625–15638, 2020.

[76] S. Liu, D. C. Mocanu, A. R. R. Matavalam, Y. Pei, and M. Pechenizkiy, "Sparse evolutionary deep learning with over one million artificial neurons on commodity hardware," *Neural Computing and Applications (NCAA)*, pp. 2589–2604, 2021.

[77] J. Lin, L. Zhu, W.-M. Chen, W.-C. Wang, C. Gan, and S. Han, "On-device Training under 256KB Memory," *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 22941–22954, 2022.

[78] Y. Bengio, A. Courville, and P. Vincent, "Representation Learning: A Review and New Perspectives," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, pp. 1798–1828, 2013.

[79] M. Kosinski, D. Stillwell, and T. Graepel, "Private traits and attributes are predictable from digital records of human behavior," *Proceedings of the National Academy of Sciences (PNAS)*, pp. 5802–5805, 2013.

[80] M. Fredrikson, S. Jha, and T. Ristenpart, "Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2015, pp. 1322–1333.

[81] Y. Zhang, R. Jia, H. Pei, W. Wang, B. Li, and D. Song, "The Secret Revealer: Generative Model-Inversion Attacks Against Deep Neural Networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 253–261.

[82] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership Inference Attacks Against Machine Learning Models," in *Proceedings of the 2017 IEEE Symposium on Security and Privacy (SP)*, 2017, pp. 3–18.

[83] A. Salem, Y. Zhang, M. Humbert, P. Berrang, M. Fritz, and M. Backes, "Ml-leaks: Model and data independent membership inference attacks and defenses on machine learning models," *arXiv preprint arXiv:1806.01246*, 2018.

[84] Y. Kaya and T. Dumitras, "When Does Data Augmentation Help With Membership Inference Attacks?" in *Proceedings of International Conference on Machine Learning (ICML)*, 2021, pp. 5345–5355.

[85] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2015, pp. 234–241.

[86] W. Al-Dhabyani, M. Gomaa, H. Khaled, and A. Fahmy, "Dataset of breast ultrasound images," *Data in Brief*, p. 104863, 2020.

[87] Y. He, S. Rahimian, B. Schiele, and M. Fritz, "Segmentations-Leak: Membership Inference Attacks and Defenses in Semantic Image Segmentation," in *Proceedings of European Conference on Computer Vision (ECCV)*, 2020, pp. 519–535.

[88] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research (JMLR)*, pp. 1929–1958, 2014.

[89] D. Yu, H. Zhang, W. Chen, J. Yin, and T.-Y. Liu, "How Does Data Augmentation Affect Privacy in Machine Learning?" in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2021, pp. 10746–10753.

Jie Zhu is currently pursuing the Ph.D. degree in the School of Computer Science, Peking University, China.

Leye Wang received the Ph.D. degree in computer science from TELECOM SudParis and University Paris 6, France, in 2016. He is currently an assistant professor with the Key Lab of High Confidence Software Technologies, Peking University, MOE, and the School of Computer Science, Peking University, China. He was a postdoctoral researcher with the Hong Kong University of Science and Technology. His research interests include ubiquitous computing, mobile crowdsensing, and urban computing.

Xiao Han received the Ph.D. degree in computer science from Pierre and Marie Curie University and the Institut Mines-TELECOM/TELECOM SudParis in 2015. She is currently a Full Professor with the Shanghai University of Finance and Economics, China. Her research interests include social network analysis, fintech, and privacy protection.

Anmin Liu is currently pursuing the Ph.D. degree in the School of Computer Science, Peking University, China.

Tao Xie received the BS degree in computer science from Fudan University, Shanghai, China, the MS degree in computer science from Peking University, Beijing, China, and the PhD degree in computer science from the University of Washington at Seattle, WA. He is a Peking University Chair Professor, where he specializes in software engineering, system software, software security, and trustworthy AI. He is a Foreign Member of Academia Europaea, and a Fellow of ACM, IEEE, and AAAS.