Differentially Private Pre-Trained Model Fusion using Decentralized Federated Graph Matching

Qian Chen, Yiqiang Chen, *Senior Member, IEEE,* Xinlong Jiang, *Member, IEEE,*
Teng Zhang, Weiwei Dai, Wuliang Huang, Zhen Yan, and Bo Ye

Abstract—Model fusion is becoming a crucial component in the context of model-as-a-service scenarios, enabling the delivery of high-quality model services to local users. However, this approach introduces privacy risks and imposes certain limitations on its applications. Ensuring secure model exchange and knowledge fusion among users becomes a significant challenge in this setting. To tackle this issue, we propose PrivFusion, a novel architecture that preserves privacy while facilitating model fusion under the constraints of local differential privacy. PrivFusion leverages a graph-based structure, enabling the fusion of models from multiple parties without necessitating retraining. By employing randomized mechanisms, PrivFusion ensures privacy guarantees throughout the fusion process. To enhance model privacy, our approach incorporates a hybrid local differentially private mechanism and decentralized federated graph matching, effectively protecting both activation values and weights. Additionally, we introduce a perturbation filter adapter to alleviate the impact of randomized noise, thereby preserving the utility of the fused model. Through extensive experiments conducted on diverse image datasets and real-world healthcare applications, we provide empirical evidence showcasing the effectiveness of PrivFusion in maintaining model performance while preserving privacy. Our contributions offer valuable insights and practical solutions for secure and collaborative data analysis within the domain of privacy-preserving model fusion.

Index Terms—pre-trained model fusion, local differential privacy, federated graph matching, hybrid perturbation mechanism.

Fig. 1. Users select a pre-trained model from the "Model Service Provider" and subsequently fine-tune and update it using local data. The updated model is then contributed back to the pool, enabling other users to leverage it within their own processes.

I. INTRODUCTION

DATA-driven artificial intelligent systems have gradually evolved into a new paradigm that is model-driven [1]. So far, as a provider of Model-as-a-Service (MaaS), there are nearly 81k models on HuggingFace [2] and about 800 models on ModelScope [3] that can be downloaded and executed in some common ways(e.g., fine-tuning [4], domain adaptation [5], knowledge distillation [6], model fusion [7], etc.) for

Qian Chen, Yiqiang Chen, Xinlong Jiang, Teng Zhang, Wuliang Huang, and Zhen Yan are with the Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China, and also with the University of Chinese Academy of Sciences, Beijing, China. (e-mail: {chenqian20b, yqchen, jiangxinlong, zhangteng9s, huangwuliang19b, yanzhen21s}@ict.ac.cn).
Weiwei Dai is with Changsha Aier Eye Hospital, Changsha, Hunan, China (e-mail: daiweiwei@aierchina.com).
Bo Ye is with Nanchang Aier Eye Hospital, Nanchang, Jiangxi, China (e-mail: yebo@aierchina.com).
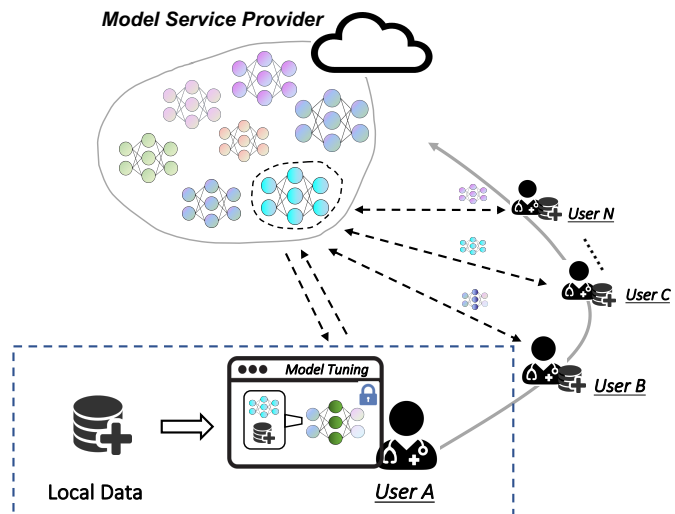*Corresponding authors: Yiqiang Chen and Xinlong Jiang.*

various downstream tasks, such as finance, transportation, smart buildings, especially collaborative healthcare. Collaborative healthcare applications [8] leverage artificial intelligence and medical expert knowledge to enable model deployment and adaptation across multiple healthcare organizations. Local fine-tuning of pre-trained models enables local control over medical data and allows model providers to offer assistance, as illustrated in Fig. 1. However, collaborative healthcare applications within the context of MaaS present inherent privacy risks and application limitations. Firstly, users are constrained to rely on a presumed trusted third-party entity as mandated by the MaaS framework for the fusion and sharing of model knowledge, which poses challenges in terms of practical guarantees. Additionally, uploading locally adapted models into a less trustworthy model pool can compromise the confidentiality of sensitive local medical information. Consequently, if direct data exchange is unfeasible, the pressing question arises: **Is there a decentralized approach that can ensure secure model sharing and knowledge fusion among users?**

Pre-trained model fusion emerges as a promising and feasible approach to address the aforementioned challenges. The pre-trained model fusion aims to achieve local model updates by averaging parameters with minimal training data.

The redundancy of neural network parameterization makes it challenging to find a one-to-one correspondence between the weights of different networks, hindering the effectiveness of vanilla averaging, which performs poorly on trained networks with nonlinear differences in weights. To achieve better performance of the fused model, it is necessary to obtain a mapping relationship between the two neural networks by sharing their structure and parameters. This information is shared to calculate the similarity or affinity of weights to get the permutation matrix and align the neurons for layer-by-layer fusion with other models. Recently, there have been several works that have delved into the field of model fusion. Singh et al. [7] suggest optimal transport to merge models from multiple sources, demonstrating its effectiveness in image classification tasks. Liu et al. [9] propose a deep neural network fusion approach based on graph matching to improve model ensemble and federated learning. It achieves state-of-the-art results on several datasets. Ainsworth et al. [10] present a method for merging models that are identical up to permutation symmetries, which is shown to improve performance in certain tasks. Akash et al. [11] use Wasserstein Barycenter to fuse multiple models and explore the linear mode connectivity of neural networks in the context of model fusion. Jin et al. [12] present a dataless knowledge fusion approach for language models, which combines the weights of multiple pre-trained models to improve their performance on downstream tasks. In scenarios involving cross-institutional collaboration, these methods inevitably require sharing information about the model, including its neural network structure, weights, and hyper-parameters, to varying degrees. Compared to traditional ensemble-based methods [13]–[15], pre-trained model fusion can avoid the need to save all trained models, thereby reducing the space resources required. Compared to distillation-based methods [6], [16], by incorporating a teacher network as a guide, the process of pre-trained model fusion can be accomplished without the need for retraining using training samples or pseudo samples. Compared to federated learning methods [17], [18], pre-trained model fusion can effectively address the dilemma of unavailable model training samples. Considering the current technological advantages, the integrated and updated model has the potential to be shared among multiple healthcare organizations, fostering a positive cycle of model development and improvement within the MaaS community.

However, privacy poses a substantial challenge in the domain of real-world applications and decentralized model sharing, particularly when it comes to the exchange of models among their respective owners. Even during the process of model fusion, sharing the structure and weights of neural networks can still lead to a privacy breach of local data. This can occur through various types of attacks, including model extraction attacks [19], [20], model inversion attacks [21], [22], or membership inference attacks [23], [24]. The sensitivity of healthcare data and legal constraints on data sharing highlights the necessity for healthcare models to incorporate privacy protection measures. To address the privacy challenges arising from knowledge sharing among different parties, various privacy-preserving techniques have been proposed, including Differential Privacy (DP) [25], [26], Homomorphic Encryption (HE) [27], [28], Secure Multi-Party Computation (SMPC) [29], [30], and hybrid methods [31], [32]. DP provides a strong privacy guarantee by adding noise to the output of a computation, making it difficult for attackers to extract sensitive information from the result. HE allows computation on encrypted data without decrypting it, which can prevent unauthorized access to sensitive data. SMPC enables multiple parties to jointly compute a function while keeping their inputs private, but it can be computationally expensive and communication intensive. Hybrid methods combine the benefits of multiple privacy-preserving techniques and can achieve better performance than individual methods alone. For privacy-preserving model fusion, DP and its variants are suitable choices as it provides a rigorous privacy guarantee with a relatively low computational cost compared to other methods. DP is well-suited for scenarios where healthcare data is sensitive and model privacy protection is critical. In addition, DP has been widely adopted and extensively studied in the privacy community, making it a mature and trusted technique for privacy preservation.

In order to address the aforementioned issues and challenges, we propose a novel decentralized architecture, PrivFusion, for model fusion that prioritizes privacy preservation. The PrivFusion addresses the challenge of balancing privacy and utility in sharing trained models. To achieve this, we adopt a graph-based representation of the trained neural network, where nodes represent neurons/channels and edges represent the weights connecting them. By leveraging a decentralized federated approach and employing randomized graph matching mechanisms that adhere to the principles of local differential privacy, models from different parties can be effectively fused into a unified model with a consistent structure. This approach establishes trust among multiple parties involved in the sharing process and promotes secure collaboration. The main contributions of this paper are summarized as follows:

- We propose PrivFusion, a novel and privacy-preserving architecture for decentralized model fusion across multiple parties. To the best of our knowledge, we are the first to address the privacy concerns associated with the fusion of pre-trained models.
- We propose a decentralized federated graph matching approach that aligns neurons of neural networks layer by layer to ensure effective model fusion. Moreover, we achieve a trade-off between privacy and utility by leveraging a hybrid local differentially private perturbation mechanism.
- We integrate model fusion strategies based on activation values and weights, analyzing their contributions to the fused model through individual perturbed performance assessment. To improve the utility of private model fusion, we employ a perturbation filter adapter to alleviate the impact of randomized noise on the fused model's performance.
- Experiments show the effectiveness of PrivFusion in both privacy goals and model performance. Additionally, we have also validated the PrivFusion in some real-world

collaborative health tasks.

The rest of this paper is organized as follows. Section 2 provides a brief review of related work. Section 3 introduces preliminaries and problem formulation, while Section 4 presents the detail of the PrivFusion. Section 5 provides a detailed explanation and analysis of the experimental results. Finally, Section 6 concludes the paper.

II. RELATED WORK

A. Model Fusion

Model fusion or ensemble methods have garnered significant attention due to their capacity to enhance the performance of machine learning models. Traditional ensemble learning is a well-established approach that combines multiple models to achieve superior performance compared to individual models. This technique entails training a collection of base models on the same dataset using diverse algorithms or settings and aggregating their predictions to generate a final prediction. Prominent ensemble methods include bagging [13], stacking [14], and boosting [15], each employing distinct techniques for combining model predictions. Bagging involves training multiple models on different subsets of the data and combining their predictions through averaging or voting. Boosting, on the other hand, trains models sequentially, with each model emphasizing examples that previous models struggled with. Stacking, a versatile approach, combines predictions from multiple models using another model, known as a meta-learner, which learns to weigh the predictions based on their accuracy. These ensemble methods offer various strategies for aggregating model predictions and can be applied to a wide range of tasks in different domains.

In recent years, several methods have been proposed to address scenarios where training data is unavailable due to privacy concerns or other reasons. One notable approach is OT-Fusion [7], which tackles model fusion as a linear assignment problem focusing on weight alignment. By utilizing Wasserstein barycenters, OTFusion improves upon simple averaging methods. However, it has a limitation in neglecting second-order weight similarity, leading to potential degradation in performance. To address this limitation, Liu et al. [9] propose GAMF, a model fusion method that formulates the problem as graph matching. GAMF aligns channels in each layer to maximize weight similarity, taking into account second-order similarities. This efficient graduated assignment-based method overcomes challenges associated with the quadratic assignment problem. GAMF has been extended to handle multi-model fusion and has demonstrated superior performance compared to state-of-the-art baselines in tasks such as compact model ensemble and federated learning across various datasets. In parallel, Ainsworth et al. [10] introduce three algorithms that permute hidden units of one model to align with a reference model, effectively merging them in weight space. This process generates functionally equivalent weights within an approximately convex basin. Subsequent works, such as [11] and [12], have built upon these foundations. Akash et al. propose a fusion method that employs Wasserstein distance to combine multiple models, while Jin et al. focus on merging weights of language models to enhance their capabilities. These advancements in model fusion techniques have contributed to enhancing the performance of pre-trained models, addressing challenges related to unavailable training data.

In summary, the selection of a model fusion or ensemble approach should consider the specific requirements of the problem, including privacy considerations, data availability, and computational resources. It is essential to prioritize the safeguarding of sensitive information related to private models, such as neural network architectures, weights and proprietary data, throughout the model fusion process.

B. Local Differential Privacy

Differential Privacy [25], [26] is a robust framework for quantifying the privacy guarantees offered by data release mechanisms. It provides a rigorous and provable definition of privacy by ensuring that an individual's data remains indistinguishable, regardless of their participation in a dataset. This framework has gained significant attention as a rigorous and well-founded privacy definition, leading to extensive research in recent years. Local Differential Privacy (LDP) [26], [33] is a specific instantiation of differential privacy that has emerged as a popular method for preserving privacy while collecting sensitive data and computing statistical queries. LDP enables the calculation of aggregate functions, such as mean, count, and histogram, while preserving the privacy of individual data points. Notably, major technology companies including Google, Apple, and Microsoft have embraced LDP for conducting large-scale private data analytics [34]. The adoption of differential privacy and LDP has paved the way for conducting privacy-preserving data analysis, ensuring the protection of individuals' sensitive information while enabling valuable insights to be derived from the data. The rigorous nature of these privacy-preserving techniques offers a strong foundation for addressing privacy concerns in data-driven applications and has motivated further advancements in the field.

LDP operates by introducing random noise to the data prior to its disclosure, rendering it challenging for an attacker to deduce the original data. The application of LDP has been extensively explored in the context of graph-structured data representations, particularly when pre-trained models can be transformed into such representations. Protecting sensitive information in graphs has garnered increasing attention in recent years, leading to the exploration of differential privacy (DP) and LDP in this domain. One common strategy involves utilizing graph-based mechanisms to inject noise into various aspects of the data, including node features, graph structure, and edge weights [35]–[37]. This approach aims to preserve the privacy of graph data while allowing for meaningful analysis. Another approach leverages Graph Neural Networks (GNNs) to learn representations of the data, enabling downstream tasks without compromising sensitive information [38]–[40]. In this approach, the randomization of the graph occurs without explicitly focusing on the model optimization process. These advancements in applying DP and LDP to
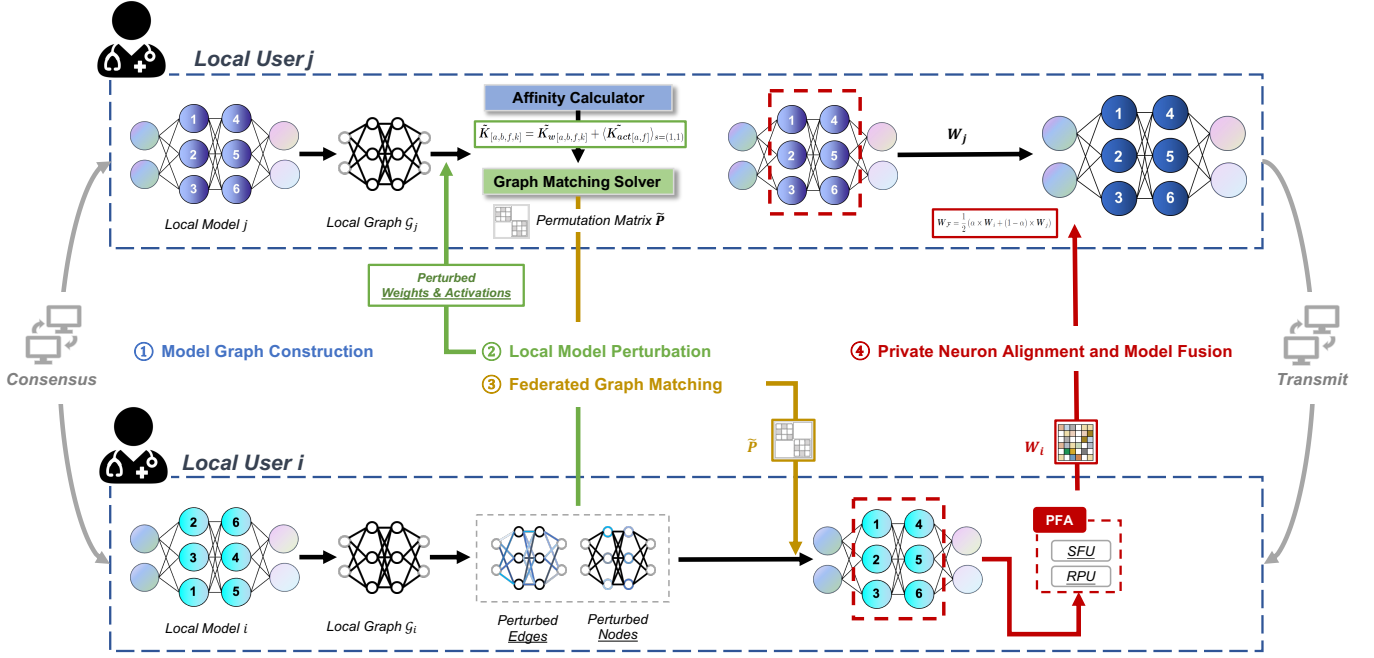
Fig. 2. The framework and workflow of the PrivFusion.

graph data have opened up new avenues for privacy-preserving analysis, enabling the protection of sensitive information while extracting valuable insights from graph-structured datasets. The integration of LDP and GNNs showcases the potential for achieving a balance between privacy preservation and effective data utilization in graph-related applications.

III. PRELIMINARIES

A. Notations and Problem Formulation

We formulate the pre-trained model fusion as a graph matching and publishing problem without losing generality. Many neurons in a trained neural network model can interact with one another and recognize a certain pattern. Based on this association structure, the model may be viewed as data with a graph structure, naturally, sharing models are turned into transmitted graph data.

The trained neural network can be described by a graph in which nodes represent the neurons, edges represent the relations between them, and features associated with the edges represent the weights of the neural network. We firstly construct m trained neural networks as graphs $\mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_m$. Each graph $\mathcal{G}_i$ contains n_i nodes, and its neighboring layers are fully connected. For a given graph $\mathcal{G}_i = (\mathbf{V}_i, \mathbf{X}_i, \mathbf{A}_i)$, $\mathbf{V}_i$ is the node/neuron set, and $\mathbf{X}_i \in \mathbb{R}^{N \times d}$ denotes d-dimensional node features of the i^{th} network. The node features can be considered for activations of neurons by running inference over a set of samples from different classes. $\mathbf{A}_i \in \mathbb{R}^{N \times N}$ represents the adjacency matrix encodes the weighted connectivity of the graph $\mathcal{G}_i$. The adjacency matrix is relatively sparse since the cross-layer links are meaningless except for the connections of neurons in the neighboring layer. Obviously, $\mathcal{G}_i$ is a directed graph whose direction is pointed at the model from input to output.

Training the neural network with stochastic gradient descent and different training sets shuffles the neurons. Furthermore, the model fused by simply averaging is ineffective if the weights of neurons are in unaligned positions. Since graph matching can solve the alignment problem of neuron positions, for simplicity, we formulate the matching and fusing with two fully connected networks with two hidden layers without bias. The goal of matching is to find a suitable permutation of the nodes(neurons or channels) between any two graphs $\mathcal{G}_i, \mathcal{G}_j$, $i, j \in 1, 2, \ldots, m$, which is equivalent to permuting $\mathcal{G}_i$ to fit $\mathcal{G}_j$ by graph matching. For example, fusing two fully-connected networks through graph matching can be represented as follows:

$$\max_{\boldsymbol{P}} \sum_{a=0}^{d_\Sigma-1} \sum_{b=0}^{d_\Sigma-1} \sum_{e=0}^{d_\Sigma-1} \sum_{f=0}^{d_\Sigma-1} \boldsymbol{P}_{[a,b]} \boldsymbol{K}_{[a,b,f,k]} \boldsymbol{P}_{[f,k]} \quad (1)$$

where $\boldsymbol{P}$ is the permutations of two hidden layers, and the neurons of input/output layers are assumed to have been aligned and need not be permuted. These constraints guarantee the existence of one-to-one mapping relationships between neurons that are within the same layer. The similarity of two graphs can be denoted as $\boldsymbol{K} \in \mathbb{R}^{d_\Sigma \times d_\Sigma \times d_\Sigma \times d_\Sigma}$, which is a 4-dimensional affinity tensor. Element $\boldsymbol{K}_{[a,b,f,k]}$ measures the affinity between the edges (a, f) and (b, k) and denotes the similarity between the weight matrices of two graphs in the model fusion problem. The measurement of similarity can be done by common means such as Gaussian kernel, cosine similarity, correlation distance, etc.

B. Privacy Threat and Protection

In a decentralized architecture, where there is no involvement of a trusted third-party server, model owners (local users) are considered semi-trusted, meaning they are honest

but curious. The adversary attempts to extract model details and speculate on private training samples. Consequently, it is unsafe to directly share the actual weight values for calculating affinity. Our ultimate goal is to obtain a set of mechanisms that allows model owners to aggregate multiple models' knowledge from other parties while preserving the privacy of the models. Additionally, model owners should be able to benefit from other participants over the collected noisy graph and improve their own models with the best possible generalization capability after optimizing. To ensure the privacy of the entire process of model fusion, we use a privacy-preserving scheme based on differential privacy, which is defined below.

As a rigorous and provable definition of privacy, differential privacy [25] has received a lot of attention and has been extensively researched in the last few years. Differential Privacy provides privacy protection for data with adversary background knowledge maximization, and it is a notion of privacy that addresses the privacy problem of statistical databases.

Definition 1 (ϵ-Differential Privacy): *A random mechanism $\mathcal{M}$ satisfies ϵ-differential privacy if for any pair of neighboring datasets X and $X' \in X$ that differ by a single data instance and for any set of outcomes $S \subseteq R$,*

$$\boldsymbol{Pr}[\mathcal{M}(X) \in \mathcal{S}] \leq \exp(\epsilon)\boldsymbol{Pr}\left[\mathcal{M}\left(X'\right) \in \mathcal{S}\right] \qquad (2)$$

where ϵ is referred to as the privacy budget, and it quantifies the privacy risk of mechanism $\mathcal{M}$. The parameter ϵ is also used to tune utility versus privacy: a smaller (resp. larger) ϵ leads to stronger (weaker) privacy guarantees, but lower (higher) utility. $\boldsymbol{Pr}[\cdot]$ represents the randomness of the $\mathcal{M}$ on the datasets X and X'. This definition ensures that the presence or absence of an individual will not significantly affect the output. Differential privacy has some lemmas. In the differentially private computing [25], we can take advantage of the following lemmas:

Lemma 1 Post-processing. *Any calculation of output under differential privacy does not increase privacy loss.*

Lemma 2 Serialized combination theorem. *Serialized combination of differential privacy mechanisms still satisfies differential privacy protection.*

One major problem with the notion of differential privacy is that users still have to trust a central authority to keep their privacy. This is not feasible in the setting of the cross-client model fusion problem. In order to be able to give users stronger privacy guarantees the concept of local differential privacy(LDP) was introduced. We use the definition of LDP given by [33].

Definition 2 (ϵ-Local Differential Privacy): *A random algorithm $\mathcal{A}$ satisfies ϵ-local differential privacy if for all pairs of client's values V_1 and V_2 and for all $Q \subseteq Range(\mathcal{A})$ and for $\epsilon > 0$, the follow equation holds:*

$$\boldsymbol{Pr}\left[\mathcal{A}\left(V_1\right) \in Q\right] \leq \exp(\varepsilon)\boldsymbol{Pr}\left[\mathcal{A}\left(V_2\right) \in Q\right] \qquad (3)$$

$Range(\mathcal{A})$ is the set of all possible outputs of the randomized algorithm $\mathcal{A}$.

Properties. Given the constraint of local differential privacy (LDP), there exists a lack of trust among users. To address this, a perturbation mechanism is applied by each user to their data before sharing it with the curator or other users. This mechanism ensures differential privacy, preserving the privacy of each individual's data. While the shared data may not have any meaningful interpretation on its own, it can be used for aggregating parameters and fusing models. Importantly, the privacy guarantee remains intact even if there is post-processing of the algorithm's output, as long as it adheres to the principles of LDP.

IV. PRIVFUSION: OUR PROPOSED METHOD

This section describes the main components of our proposed method, called PrivFusion, toward differentially private model fusion via decentralized federated graph matching with a decentralized architecture. The overall workflow of the proposed PrivFusion(shown in Fig 2). Once users reach a consensus on model tasks, structure, and requirements, the process mainly consists of three components: (1) *Local Model Perturbation*, (2) *Decentralized Federated Graph Matching*, and (3) *Private Neuron Alignment and Fusion*. The first component aims to perform privacy-preserving obfuscation on the graph that is converted by the pre-trained model. Each user(model owner) independently shares a variant of the model - the obfuscated graph data, including the obfuscated version of node features and the obfuscated weight matrix. The second component involves calculating layer-wise affinity on the noisy graph composed of obfuscated data shared by all users. It employs federated interactions to perform graph matching and obtain permutation matrices for neurons/channels. The third component involves applying neuron permutation based on the permutation matrix to align the neurons of each layer in the models. This is followed by averaging the corresponding weights of the models under LDP constraints. At the end of the entire workflow, the fused model is transmitted among the users.

In what follows, we first introduce PrivFusion's technical details, including the hybrid randomized mechanisms used for preserving model privacy that satisfies local differential privacy, as well as the decentralized federated graph matching approach. Then, we theoretically and empirically analyze the privacy-utility trade-off issue caused by the hybrid differentially private mechanisms.

A. Local Model Perturbation

In the setting of the PrivFusion, model owner U_i and U_j hold their local model and construct them into two graphs $\mathcal{G}_i$ and $\mathcal{G}_j$ with nodes and weighted unidirectional edges. Take the U_i as an example, graph $\mathcal{G}_i$ includes node feature $\mathbf{X}_i$, the adjacency matrix $\mathbf{A}_i$ and weight matrix $\mathbf{W}_i$. All nodes are connected layer by layer, they are only connected to nodes of neighboring layers, and there are no second-order neighbors of nodes(cross-layer connection). The mechanisms utilized to ensure the privacy of the node feature, the adjacency matrix, and the weight matrix of the model, respectively, are described in the paragraphs that follow.

Node Feature Perturbation. We retain the activations for all of the model's neurons after running inference over a few instances. After that, we consider the corresponding activation

values as features of the nodes and randomize these features. In this phase, we apply a variant of the Laplace mechanism [41], one of the implementation mechanisms for LDP, to secure the delivery process of the model's activations. The privacy-preserving module adds Laplacian noise to the node feature $\boldsymbol{X}_i$ which is formulated as follows:

$$\tilde{\boldsymbol{X}}_i = \boldsymbol{X}_i + Laplace \ (S(f)/\epsilon_x) \tag{4}$$

where ϵ_x is the "privacy budget" for node feature perturbation. $S(f)$ represents L_1 sensitivity which is defined as follows.

Definition 3 (L_1 Sensitivity): *The L_1 sensitivity of a function $f : D^n \to \mathbb{R}^d$ is the smallest number $S(f)$ such that for all $\boldsymbol{X}_i, \boldsymbol{X'}_i \in D^n$ which differ in a single entry,*

$$\left\| f(\boldsymbol{X}_i) - f\left(\boldsymbol{X}'_i\right) \right\|_1 \le S(f) \tag{5}$$

where $\|\cdot\|_1$ represents L_1 norm of a vector. Sensitivity is the key parameter that determines the amount of noise added, which refers to the maximum change caused by deleting any record. Moreover, *Laplace* $(\cdot)$ is a random variable that has the following probability density function:

$$\boldsymbol{Pr}[x, \lambda] = \frac{1}{2\lambda} \cdot e^{-\frac{|x|}{\lambda}}, \quad \forall x \in \mathbb{R}^d \tag{6}$$

where x presents the specific variable λ control the strength of protection (volume of noise) with a privacy budget ϵ_x bounded by $\frac{S(f)}{\lambda}$.

Theorem 1. *For an arbitrary adversary, let $f(\boldsymbol{X}_i) : D^n \to \mathbb{R}^d$ be its query function as a means of observation. If $\lambda = maxS(f)/\epsilon_x$, the the mechanism $\mathcal{M}_x$ satisfies ϵ_x-local differential privacy.*

Proof. Considering that $\boldsymbol{X}_i$ and $\boldsymbol{X'}_i$ only differ in single entry. Using the law of conditional probability and given any output $Sx = (Sx_1, Sx_2, \ldots, Sx_n)$ from the node feature perturbation mechanism $\mathcal{M}_x$,

$$\frac{\boldsymbol{Pr}\left[\mathcal{M}_x(\boldsymbol{X}_i) = Sx\right]}{\boldsymbol{Pr}\left[\mathcal{M}_x\left(\boldsymbol{X}'_i\right) = Sx\right]} = \prod_h \frac{\boldsymbol{Pr}\left[\mathcal{M}_x(\boldsymbol{X}_{i,h}) = Sx_h\right]}{\boldsymbol{Pr}\left[\mathcal{M}_x\left(\boldsymbol{X}'_{i,h}\right) = Sx_h\right]} \tag{7}$$

For each term in the product, $h = 1, 2, \ldots, n$. The conditional distributions are therefore standard Laplacians, which means we can bind each term and its product as

$$\begin{aligned}\prod_h \frac{\boldsymbol{Pr}\left[\mathcal{M}_x(\boldsymbol{X}_{i,h}) = Sx_h\right]}{\boldsymbol{Pr}\left[\mathcal{M}_x\left(\boldsymbol{X}'_{i,h}\right) = Sx_h\right]} &\le \prod_h exp(|f(\boldsymbol{X}_{i,h}) \\ &\quad - f(\boldsymbol{X}'_{i,h}|/\lambda) \\ &= exp(\left\| f(\boldsymbol{X}_i) - f\left(\boldsymbol{X}'_i\right) \right\|_1 /\lambda)\end{aligned} \tag{8}$$

Consequently, we finish the proof of the node feature perturbation mechanism using the bound $S(f) \le \lambda\epsilon_x$.

Weight Matrix Perturbation. There are no second-order neighbors of nodes in graph $\mathcal{G}_i$, and all nodes are connected to other nodes only in neighboring layers. The adjacency matrix is only used to represent the structural information of the network, however, its structure is public and very sparse. To avoid too much extra computation, we focus privacy concerns on the pre-trained network's weights themselves. In this phase,

we leverage a perturbation mechanism to protect its privacy for the weight matrix inspired by [40]. We define the weight of each edge originating from the previous layer as the "weight feature" of the nodes in the current layer. The "weight feature" of the nodes shares the same dimension as the number of nodes in the previous layer.

As an illustration, take a model of U_i that is fully connected. Given the current p_{th} layer and its input weight list $\boldsymbol{w}_{p-1}$ from the previous $p - 1_{th}$ layer, the "weight feature" of the current layer is denoted as $\boldsymbol{w}_{i,p} = \boldsymbol{w}_{p-1} \in \mathbb{R}^d$, where d is the number of nodes in the $p - 1_{th}$ layer. In this way, we convert the perturbation of the weight list into a variant of private extraction of node features on the graph. In order to tackle the computational burden of perturbations, we perform a layer-wise perturbation by extending the multi-bit mechanism in [39] for private extraction of the "weight feature" layer by layer. Assuming that each node of current p_{th} layer owns d-dimensional "weight feature" vector $\boldsymbol{w}_{i,p}$ whose elements fall into the range $[\boldsymbol{w}_{min}, \boldsymbol{w}_{max}]$, we apply the perturbation on $\boldsymbol{w}_{i,p}$ layer by layer to get the corresponding perturbed "weight feature" vector $\tilde{\boldsymbol{w}}_{i,p}$. Sharing the "weight feature" list $\tilde{\tilde{\boldsymbol{W}}}_i$ with other model owners, after all of the "weight feature" perturbations are completed. The process can be formulated as

$$\tilde{\boldsymbol{w}}_{i,p} = \boldsymbol{w}_{i,p} + MultiBit \ (\epsilon_w) \tag{9}$$

Instead of perturbing all the d features, the perturbation mechanism uniformly samples m out of d features without replacement, which with a probability formulated as

$$\text{Bernoulli} \left(\frac{1}{e^{\epsilon_w/m}+1} + \frac{\boldsymbol{w}_{i,p} - \boldsymbol{w}_{min}}{\boldsymbol{w}_{max} - \boldsymbol{w}_{min}} \cdot \frac{e^{\epsilon_w/m}-1}{e^{\epsilon_w/m}+1} \right) \tag{10}$$

For each sampled feature, a corresponding Bernoulli variable is drawn from the distribution whose parameter depends on the value of the feature and the privacy budget ϵ_w. When $m = d$, the perturbation reduces to applying the 1-bit mechanism with a privacy budget of ϵ_w/d to every single feature. The following theorem ensures that the perturbation mechanism is ϵ_w-LDP.

Theorem 2. *The weight matrix perturbation mechanism $\mathcal{M}_w$ preserves ϵ_w-local differential privacy.*

Proof. The perturbation mechanism $\mathcal{M}_w$ is based multi-bit mechanism applied on the "weight feature" vector $\boldsymbol{w}_{i,p}$. According to the proof of [39], we show that for any two weight features $\boldsymbol{w}_{i,p,1}$ and $\boldsymbol{w}_{i,p,2}$, we have $\frac{\boldsymbol{Pr}[\mathcal{M}_w(\boldsymbol{w}_{i,p,1})=\tilde{\boldsymbol{w}}_{i,p}]}{\boldsymbol{Pr}[\mathcal{M}_w(\boldsymbol{w}_{i,p,2})=\tilde{\boldsymbol{w}}_{i,p}]} \le e^{\epsilon_w}$. For any dimension $q \in \{1, 2, \ldots, d\}$, we get $\tilde{\boldsymbol{w}}_{i,p,q} \in \{-1, 0, 1\}$. When $q \notin \mathcal{S}$ with probability $1 - m/d$, the $\tilde{\boldsymbol{w}}_{i,p,q} = 0$, therefore:

$$\frac{\boldsymbol{Pr}\left[\mathcal{M}_w(\boldsymbol{w}_{i,p,1})_q = 0\right]}{\boldsymbol{Pr}\left[\mathcal{M}_w(\boldsymbol{w}_{i,p,2})_q = 0\right]} = \frac{1-m/d}{1-m/d} = 1 \le e^{\epsilon_w}, \forall\epsilon_w > 0 \tag{11}$$

Similarly, for $\tilde{\boldsymbol{w}}_{i,p} \in \{-1, 1\}$, we obtain the following inequality:

$$\frac{Pr\left[\mathcal{M}_w(\boldsymbol{w}_{i,p,1})_q \in \{-1,1\}\right]}{Pr\left[\mathcal{M}_w(\boldsymbol{w}_{i,p,1})_q \in \{-1,1\}\right]} \leq \frac{\frac{m}{d} \cdot \frac{e^{\epsilon_w/m}-1}{e^{\epsilon_w/m}+1}}{\frac{m}{d} \cdot \frac{1}{e^{\epsilon_w/m}+1}} \leq e^{\epsilon_w/m} \tag{12}$$

Therefore, with layer-by-layer iterations, we have

$$\frac{Pr\left[\mathcal{M}_w\left(\boldsymbol{w}_{i,p,1}\right) = \tilde{\boldsymbol{w}}_{i,p}\right]}{Pr\left[\mathcal{M}_w\left(\boldsymbol{w}_{i,p,2}\right) = \tilde{\boldsymbol{w}}_{i,p}\right]} = \prod_{q=1}^{d} \frac{Pr\left[\mathcal{M}_w(\boldsymbol{w}_{i,p,1})_q = \tilde{\boldsymbol{w}}_{i,p,q}\right]}{Pr\left[\mathcal{M}_w(\boldsymbol{w}_{i,p,2})_q = \tilde{\boldsymbol{w}}_{i,p,q}\right]}$$

$$= \prod_{y \mid \tilde{\boldsymbol{w}}_{i,p,y}=0} \frac{Pr\left[\mathcal{M}_w\left(\boldsymbol{w}_{i,p,1}\right)_y = 0\right]}{Pr\left[\mathcal{M}_w\left(\boldsymbol{w}_{i,p,2}\right)_y = 0\right]}$$

$$\times \prod_{z \mid \tilde{\boldsymbol{w}}_{i,p,z} \in \{-1,1\}} \frac{Pr\left[\mathcal{M}_w\left(\boldsymbol{w}_{i,p,1}\right)_z \in \{-1,1\}\right]}{Pr\left[\mathcal{M}_w\left(\boldsymbol{w}_{i,p,2}\right)_z \in \{-1,1\}\right]}$$

$$= \prod_{\tilde{\boldsymbol{w}}_{i,p,z} \in \{-1,1\}} \frac{Pr\left[\mathcal{M}_w\left(\boldsymbol{w}_{i,p,1}\right)_z \in \{-1,1\}\right]}{Pr\left[\mathcal{M}_w\left(\boldsymbol{w}_{i,p,2}\right)_z \in \{-1,1\}\right]}$$

$$\leq \prod_{\tilde{\boldsymbol{w}}_{i,p,z} \in \{-1,1\}} e^{\epsilon_w/m}$$

$$\leq e^{\epsilon_w} \tag{13}$$

Thus, we complete the proof of the weight matrix perturbation mechanism.

B. Decentralized Federated Graph Matching

By performing information extraction and perturbation on all pre-trained models, the model owners can obtain privacy-preserving information that can be shared for cross-client model fusion. The goal is to find a suitable method that can achieve optimal fusion results without relying on the training data or revealing the original details of the models. One naive approach is to directly average the weights of corresponding positions in the weight space, similar to the FedAvg algorithm [42]. However, it's important to note that FedAvg still requires the participation of training data from each client during the aggregation process. While the idea of aggregation is similar, the proposed method aims to achieve model fusion without relying on the explicit use of training data. Moreover, recent popular studies [7], [9], [10] have demonstrated the limitations of the above idea and utilized matching and alignment of network neurons for model fusion and effectively improved the accuracy of the fusion model. To address model privacy concerns, we propose a novel approach using decentralized federated graph matching. By leveraging a small amount of data for network inference and recording activation values, we preserve the privacy of the original training data while obtaining valuable information for fusion. This enables us to balance model performance and privacy preservation, facilitating effective model fusion with minimal data involvement.

In this process, the model owner U_i shares the perturbed graph features $\tilde{\boldsymbol{X}}_i$ and $\tilde{\boldsymbol{W}}_i$ to the U_j instead of directly sharing the true values of the model weights. The U_j takes the received perturbed activation and weight of the model from U_i to calculate the similarity with the neurons of the model from U_j layer by layer. The similarity is denoted as an affinity

matrix in graph matching, which can be used to solve the permutation matrix according to Eq.(1), to adjust the ordering of neurons, and then to fuse the model after the alignment of neurons. In our methodology, we evaluate the affinity of neuron activations and weights based on the *Gaussian Kernel* applied by [9], respectively. Then we merge these two matrices $\tilde{\boldsymbol{K}_w}$ and $\tilde{\boldsymbol{K}_{act}}$ together after normalization, as follows

$$\tilde{\boldsymbol{K}}_{[a,b,f,k]} = \tilde{\boldsymbol{K}}_{\boldsymbol{w}[a,b,f,k]} + \langle \tilde{\boldsymbol{K}_{act}[a,f]} \rangle_{s=(1,1)} \tag{14}$$

Since the dimensions of the two matrices are different, we use a sliding window-like scheme $\langle \cdot \rangle$, with $strides = (1,1)$, to merge them by adding them together and reducing the influence of the magnitude between the indicators by means of layer-wise normalization to ensure the efficiency of the merged affinity matrix.

Given the merged affinity matrix $\tilde{\boldsymbol{K}}_{[a,b,f,k]}$, the model alignment's target is to maximize the following objective function:

$$\max_{\boldsymbol{P}} \sum_{a=0}^{d_\Sigma-1} \sum_{b=0}^{d_\Sigma-1} \boldsymbol{R}_{[a,b]} \boldsymbol{P}_{[a,b]} \text{ s.t. constraints in Eq. (1).} \tag{15}$$

where $\boldsymbol{R}_{[a,b]} = \sum_{f=0}^{d_\Sigma-1} \sum_{k=0}^{d_\Sigma-1} \tilde{\boldsymbol{K}}_{[a,b,f,k]} \boldsymbol{P}_{[f,k]}$. The constrained optimization problem described above is a linear assignment problem that usually is solved using the Hungarian algorithm [43] or the Sinkhorn algorithm [44] with a relaxed projection.

C. Private Neuron Alignment and Model Fusion

The problem of Eq.(15) is an NP-hard and a special case of *Lawler's Quadratic Assignment Problem* whose memory cost can be $O((d_\Sigma)^4)$ and difficult to avoid. To alleviate the computational complexity, we refer to the formulation of [9] and obtain the private permutation matrix $\tilde{\boldsymbol{P}}$ with decomposed diagonal matrices in a differentially private manner. Subsequently, we solved the objective function by the "*Spectral Graph Matching*" solver [45], a method that obtains the eigenvectors of the input affinity matrix by power iteration. Thus, we can obtain the aligned model by permuting the node ordering using $\tilde{\boldsymbol{P}}$.

Private Model Fusion. If we have two models with aligned neural networks, the challenge lies in effectively integrating them while maintaining privacy. One straightforward approach is to apply the LDP mechanism and aggregate the perturbed weights. However, this approach may lead to a cumulative degradation of model accuracy over time. Therefore, we design a *Perturbation-Filter Adapter(PFA)* based on Gaussian randomization and search for a trade-off between privacy and accuracy for the exchange of model weights. The PFA consists of two parts, *Randomized Perturbation Unit (RPU)* and *Smoothing Filter Unit (SFU)*. Before sharing the weight of the aligned neural networks, the RPU performs perturbation on the weight $\boldsymbol{W}\boldsymbol{p}_i$, after the neurons have been permuted, by adding randomized Gaussian noise, that is

$$\widetilde{\boldsymbol{W}}_i \leftarrow \tilde{\boldsymbol{W}}\boldsymbol{p}_i = \boldsymbol{W}\boldsymbol{p}_i + Gaussian \ (\epsilon_f, \delta) \tag{16}$$

where $\epsilon_f, \delta > 0$ and δ is typically small. According to definition in [25], [26], the PRU satisfies the (ϵ_f, δ)-LDP. In, particular, (ϵ_f, δ)-LDP is also called approximate LDP and the case of $\delta = 0$ is called pure LDP.

To reduce the performance degradation of model fusion, we use SFU to standardize the noise-added results to ensure the effectiveness of the model. Since the noise follows the Gaussian distribution, then we transform it through the corresponding probability density function. We compute the *Gaussian Error Function* which denoted as $erf(\cdot)$, and could be standardized as following

$$\boldsymbol{W}_i = \max\left\{0, erf\left(\frac{\widetilde{\boldsymbol{W}_i} - \mu}{\sigma \cdot \sqrt{2}}\right)\right\} \tag{17}$$

where μ and δ denote the mean and standard deviation of the weights. Ultimately, after performing the above operations, we fuse the two models in the weight space by averaging them according to different scales α and save the model with the best performance.

In the model fusion phase, the U_i fuses its model $\boldsymbol{W}_i$ with the model $\boldsymbol{W}_j$ at varying scales, allowing users to choose the fusion performance according to their specific needs. The fusion process is represented as follows

$$\boldsymbol{W}_{\mathcal{F}} = \frac{1}{2}\left(\alpha \times \boldsymbol{W}_i + (1 - \alpha) \times \boldsymbol{W}_j\right) \tag{18}$$

The securely fused model can then be reused and shared with other users for access and utilization.

Theorem 3. *The PFA mechanism $\mathcal{M}_f$ preserves (ϵ_f, δ)-local differential privacy.*

Proof. Due to the page limitation, we refer to the detailed proof of Theorem 3 to [26] and [25]. The mechanism $\mathcal{M}_f$ achieves (ϵ_f, δ)-LDP with the probability at least $1 - \delta$. In [26], the noise with standard deviation to each weight with l_2-sensitivity *sen* and was given by $\delta = \sqrt{2\ln\frac{2}{\delta}} \cdot \frac{sen}{\epsilon}$. Likewise, for delta in [25], it was given by $\delta = \sqrt{2\ln\frac{1.25}{\delta}} \cdot \frac{sen}{\epsilon}$. Especially, the noise we added following a Gaussian distribution with mean 0 and variance δ^2.

Corollary. *The randomized mechanisms $\mathcal{M}_x$, $\mathcal{M}_w$ and $\mathcal{M}_f$ jointly ensure $(\epsilon_x + \epsilon_w + \epsilon_f)$-local differential privacy.*

Therefore, the following process will also provide local differential privacy for the entire multi-model fusion and pre-trained model pool applications, referring to the composition theorem in [25].

V. EXPERIMENTS

We conduct a comprehensive set of experiments to evaluate the privacy-utility performance of the proposed method Priv-Fusion. Our investigation addresses various parameter settings that may impact its effectiveness, and aims to answer the following research questions:

- **RQ1:** How does the performance of PrivFusion compare to conventional methods that have no privacy guarantees?
- **RQ2:** Is PrivFusion's performance controllable and does it achieve the desired utility-privacy trade-off under different privacy budgets?

- **RQ3:** Which factor, perturbing activation values or perturbing weights, plays a more critical role in the effectiveness of model fusion when adjusting the privacy budget?
- **RQ4:** How does the hybrid differentially private randomized mechanism perform? Does it offer advantages over single mechanisms?

A. Experimental settings

1) Datasets: In this section, we draw inspiration from two existing model fusion methods, namely OTFusion [7] and GAMF [9]. We evaluate our proposed approach using the widely used image classification datasets, MNIST and CIFAR-10. The data augmentation settings employed in our experiments are consistent with those specified in OTFusion. For more comprehensive information about the datasets and augmentation settings, we recommend referring to their respective open-source repositories.

We explore two different data partition settings for local models/users: 1) Homogeneous data partition, where each client acquires data from an independent and identically distributed (IID) source, and the data distribution across classes is nearly equal among users. 2) Heterogeneous data partition, where each user acquires data from a non-IID source, resulting in variations in the data distribution across users. To simulate a heterogeneous data split, we followed the approach used in previous works such as [7]. Specifically, for the MNIST digit classification task, we designed Model A to possess a unique capability of recognizing a specific "personalized" label, which in this case was digit 4. Model B, on the other hand, comprised the majority of the remaining training set (excluding label 4), accounting for four-fifths of the data, while Model A held the remaining one-fifth. By incorporating these distinct data partition settings, we aim to evaluate the performance and robustness of our proposed method in scenarios where clients have varying access to different classes of data.

2) Backbone Models and Training Settings: The graph-based model fusion approach we employed in our study has a memory allocation requirement for initialization, resulting in significant overhead. To accommodate these constraints within our experimental setup, we conducted evaluations using relatively modest network architectures. For the MNIST dataset, we selected a fully connected network comprising five layers, along with a small convolutional neural network (CNN) consisting of two convolutional layers and one fully connected layer. As for the CIFAR-10 dataset, we employed a classic CNN architecture with two convolutional layers and three fully connected layers. The training process involved training both models on the MNIST dataset for 20 epochs, while the CIFAR-10 training was conducted for a total of 200 epochs, with all other training settings kept consistent.

By employing this standardized approach, we ensured fair and consistent comparisons between the compared methods when evaluating the two models across different datasets. While it is important to note that the scalability of our findings to larger-scale networks may be limited, our study provides valuable insights into the functionality and effectiveness of the proposed technique. Furthermore, **the PrivFusion can be**

TABLE I
COMPARISON OF TWO PRE-TRAINED MODELS FUSION ON TWO PUBLIC DATASETS

Dataset	Data Partition	Metric	Individual Models [Acc(%) , ...]	PE Acc(%)	VA Acc(%)	OTFusion Acc(%)	GAMF Acc(%)	GIT RE-BASIN Acc(%)	PrivFusion (ours) Acc(%)
MNIST	Homogeneous	best	[93.06, 93.17]	**93.41**	92.85	93.13	93.17	93.11	93.01
		top3-Avg		**93.39**	92.23	92.99	93.17	93.05	92.90
	Heterogeneous	best	[77.63, 91.11]	90.74	90.91	90.64	**91.11**	77.63	91.10
		top3-Avg		90.25	90.08	89.79	**91.10**	77.46	91.09
CIFAR-10	Homogeneous	best	[55.75, 56.49]	**56.80**	52.00	55.49	56.49	55.75	56.56
		top3-Avg		**56.75**	47.11	53.81	56.49	53.98	56.31
	Heterogeneous	best	[44.06, 64.05]	63.65	62.80	62.68	64.08	46.43	**64.30**
		top3-Avg		62.59	54.71	59.75	64.06	46.11	**64.07**

easily scaled to larger networks with increased computational resources, offering the potential for future research in the context of larger-scale networks.

3) Compared Methods: We compare the model fusion performance of the proposed PrivFusion against the following ones:

- **Prediction Ensemble (PE)**. Prediction ensembling involves retaining all models and averaging their predictions or output layer scores, representing the optimal performance that can be achieved by consolidating into a single model, albeit an unattainable one.
- **Vanilla Averaging (VA)** is a simple but efficient algorithm that denotes the direct averaging of parameters. This approach shares a similar concept with the classic method **FedAvg** [42] in the context of federated learning. Both methods aim to aggregate model parameters from different clients to obtain a global model that represents the collective knowledge of all participants.
- **OTFusion** [7]. As previously mentioned, published at *NeurIPS 2020*, the model fusion problem is formulated as a linear assignment problem and solved using the Wasserstein barycenter.
- **GAMF** [9]. As aforementioned, published at *ICML 2022*, it formulates the model fusion problem as a graph-matching task to maximize weight similarity and iteratively updates matchings in a consistency-maintaining manner.
- **GIT RE-BASIN** [10]. A prominent work, published at *ICLR 2023*, which merges two models in weight space by introducing three algorithms that permute the units of one model to align with a reference model.

Due to the differences in the settings between our method and federated learning, we do not include comparisons with related works in the field of federated learning. In terms of privacy protection mechanisms, we have selected several independent mechanisms to compare with our proposed hybrid mechanism.

- **Laplace Mechanism** [41]. The Laplace mechanism is commonly used for count queries, sum queries, and other low-sensitivity queries. It is particularly suitable for discrete and bounded data domains.
- **Gaussian Mechanism** [25]. The Gaussian mechanism is commonly employed for query tasks that involve

continuous data domains. It is particularly suitable for high-sensitivity queries such as range queries and mean queries.
- **MultiBit Mechanism** [39]. The MultiBit mechanism, introduced at *CCS 2021*, extends the 1-bit mechanism [46] for collecting multidimensional features. It has been shown to outperform the widely used mechanism in terms of performance.

4) Experimental Setup: Our experiments were conducted on a computing platform comprising two Intel Xeon Silver 4214R processors, each with 12 cores and a total of 192GB of RAM, and two GTX 3090 GPUs. Following the experimental phase, we proceeded to conduct model fusion using various proportions for comparative analysis. The fusion involved combining the two models in different ratios from $\{0, 0.1, 0.2, ..., 1\}$, and their performance was assessed using the accuracy metric.

To assess the impact of different fusion ratios on overall performance, we conducted a comparative analysis between the average results of the top-3 models and the best-performing model. This approach provided valuable insights into the effectiveness of diverse fusion strategies in enhancing accuracy. By considering both the average results of the top-3 models, which represent the stability of the fused models and the best-performing model, which represents the optimal performance available for users to choose from, we gained a comprehensive understanding of the influence of fusion ratios on performance enhancement.

B. Performance Evaluation

RQ1: PrivFusion demonstrates high utility accompanied by acceptable performance loss. Table I presents a comprehensive performance analysis of different methods on the MNIST and CIFAR-10 datasets. Two data partitioning approaches, homogeneous and heterogeneous, are considered for evaluation. In this study, our proposed method focuses on striking a balance between privacy protection and utility. Across the MNIST dataset, the PrivFusion, with all privacy budgets consistently set to 1, achieves performance that is comparable to the optimal method under both data partitioning approaches. This indicates that our approach effectively manages the trade-off between preserving privacy and maintaining model accuracy. Specifically, when using the homogeneous

TABLE II
COMPARISON OF DIFFERENT PRIVACY BUDGETS' EFFECT

Dataset	Data Partition	Individual Models [Acc(%) of each model, ...]	Privacy Budget			Acc(%)	
			ϵ_w	ϵ_a	ϵ_f	top3-Avg	best
MNIST	Homogeneous	[93.06, 93.17]	0.01	0.01	0.1	87.50 ± 0.4	92.27
			0.01	0.1	0.1	$\mathbf{90.48 \pm 0.2}$	**92.66**
			0.1	0.01	0.1	88.29 ± 0.3	92.41
			0.01	0.01	0.01	19.97 ± 0.8	27.55
			0.01	0.1	0.01	$\mathbf{24.58 \pm 0.5}$	33.12
			0.1	0.01	0.01	22.32 ± 0.7	**44.02**
	Heterogeneous	[77.63, 91.11]	0.01	0.01	0.1	89.41 ± 0.2	90.91
			0.01	0.1	0.1	$\mathbf{89.67 \pm 0.1}$	90.83
			0.1	0.01	0.1	89.47 ± 0.3	**91.02**
			0.01	0.01	0.01	37.80 ± 0.7	41.86
			0.01	0.1	0.01	$\mathbf{46.49 \pm 0.4}$	66.00
			0.1	0.01	0.01	43.61 ± 0.5	**72.78**
CIFAR-10	Homogeneous	[55.75, 56.49]	0.01	0.01	0.1	44.11 ± 0.3	52.05
			0.01	0.1	0.1	$\mathbf{47.15 \pm 0.3}$	**54.04**
			0.1	0.01	0.1	45.02 ± 0.5	53.83
			0.01	0.01	0.01	14.75 ± 0.8	16.22
			0.01	0.1	0.01	$\mathbf{15.39 \pm 0.5}$	**18.54**
			0.1	0.01	0.01	14.84 ± 0.4	17.47
	Heterogeneous	[44.06, 64.05]	0.01	0.01	0.1	51.96 ± 0.4	62.48
			0.01	0.1	0.1	$\mathbf{55.13 \pm 0.2}$	**63.63**
			0.1	0.01	0.1	52.66 ± 0.6	63.57
			0.01	0.01	0.01	14.85 ± 0.7	20.01
			0.01	0.1	0.01	$\mathbf{17.45 \pm 0.5}$	**23.02**
			0.1	0.01	0.01	16.77 ± 0.7	22.06

data partition, our method demonstrates competitive results compared to PE, VA, GAMF, OTFusion and GIT RE-BASIN. Moreover, when considering the average of the top 3 model accuracy and the best model accuracy, the PrivFusion performs comparably to the optimal method (highlighted in bold) under both data partitioning approaches.

Moving to the CIFAR-10 dataset, the PrivFusion delivers impressive results by achieving high model accuracy while preserving privacy. In the Homogeneous data partition scenario, our method outperforms VA and OTFusion, while achieving comparable results to PE, GIT RE-BASIN and GAMF. The PrivFusion consistently exhibits competitive performance on the CIPAF-10 compared to all other methods being evaluated. Especially in the case of heterogeneous data partitioning, PrivFusion achieves the best performance, which may be attributed to the unpredictable contribution of randomized perturbations to the models. These findings underscore the practical potential of our approach in privacy-preserving machine learning applications.

RQ2: The perturbation of PrivFusion is effective and controllable, allowing for the achievement of a utility-privacy trade-off by adjusting the privacy budget. Fig. 3 illustrates the influence of adjusting the privacy budget on the accuracy of the fused model, highlighting the controllable nature of this adjustment. By manipulating the privacy budget, the degree of perturbation applied to the model can be controlled, leading to different levels of performance degradation. When the privacy budget is set to 1, indicating minimal perturbation, the fused model achieves the highest accuracy among different privacy budget settings. This implies that with a larger privacy budget
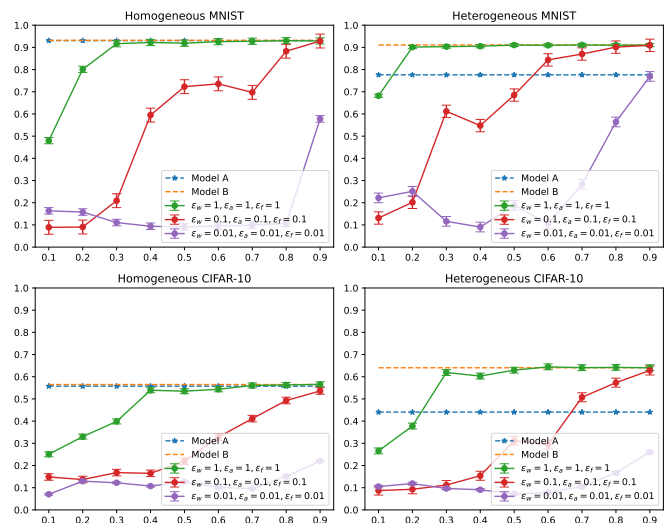

Fig. 3. The effect of the privacy budget on the performance of PrivFusion.

and smaller perturbations, the model retains a higher level of accuracy. As the privacy budget decreases to 0.1, indicating increased perturbations, there is a slight reduction in model accuracy compared to the budget of 1. However, the observed performance degradation remains relatively small, suggesting that the model can tolerate a certain level of perturbation while maintaining reasonable accuracy. Further reducing the privacy budget to 0.01, representing even larger perturbations, leads to a more pronounced decline in model accuracy. The increased perturbation significantly impacts the model's performance,

resulting in a noticeable reduction in accuracy.

These findings highlight the fundamental trade-off between privacy preservation and model performance. By adjusting the privacy budget, practitioners can exert control over the level of perturbation applied to the model, allowing them to influence the delicate balance between maintaining privacy and preserving model accuracy. This control over the privacy-performance trade-off provides valuable flexibility for practitioners in privacy-preserving machine learning scenarios.

RQ3: The perturbation of activations has a greater influence compared to weights on fused model performance. Table II illustrates the influence of adjusting the privacy budget, which corresponds to varying levels of perturbation strength, on the perturbation of model weights(ϵ_w), activation values(ϵ_a), and the private fusion(ϵ_f) process. The results highlight the importance of the privacy budget in determining the extent of perturbation and its impact on the final results. Specifically, the fused model can get the optimal performance (highlighted in bold with their standard deviation) when $\epsilon_a > \epsilon_w$, and the analysis reveals that perturbing activation values have a greater impact on the final results compared to perturbing model weights. This suggests that privacy-preserving techniques should focus more on perturbing activation values while minimizing the perturbation on model weights to achieve better results. Additionally, it is observed that the perturbation strength in fusion, the ϵ_f which setting to 0.1 and 0.01, has the greatest impact on the accuracy of the fused model. This parameter is especially crucial when performing model exchange and sharing. Furthermore, the analysis shows that perturbing activation values and model weights mainly affect the performance of neuron similarity and matching mapping. This suggests that perturbing these factors may not significantly impact the overall performance of the model.

In conclusion, the results demonstrate that adjusting the privacy budget plays a crucial role in privacy-preserving machine learning. Focusing on perturbing activation values and minimizing the perturbation on model weights can lead to improved results. Moreover, the perturbation strength in fusion is a crucial parameter that significantly impacts the accuracy of the fused model. The findings can inform the development of privacy-preserving techniques for machine learning and contribute to the advancement of the field.

RQ4: The hybrid differentially private randomized mechanism is more effective compared to a single mechanism. In Table III, we present the results which validate the superiority of the hybrid mechanism in model fusion performance. The comparative methods include the Non-Private mechanism, Laplace mechanism, Gaussian mechanism, MultiBit mechanism, and our proposed hybrid mechanism with Perturbation-Filter Adapter (PFA). The non-private mechanism serves as the upper bound for model fusion performance, while other methods exhibit performance degradation to varying degrees. Each method employs a specific perturbation mechanism in different stages of the fusion process. The results unequivocally demonstrate that the hybrid mechanism outperforms the other methods in terms of model accuracy.

For different components, it is recommended to apply

TABLE III
COMPARISON OF DIFFERENT PERTURBATION MECHANISMS EFFECT

Perturbation Mechanism	Acc(%)	
	top3-Avg	best
Non-Private	93.17	93.17
Laplace	87.14 ± 0.3	87.77
Gaussian	85.29 ± 0.5	87.31
MultiBit	90.87 ± 0.1	91.83
Hybrid (ours)	88.81 ± 0.5	89.65
Hybrid + PFA (ours)	**92.90 ± 0.2**	**93.01**

different perturbation mechanisms. Specifically, **the Laplace mechanism with lower sensitivity is suitable for perturbing activation values, as it provides strong differential privacy guarantees and enhances security. On the other hand, the MultiBit mechanism is effective for perturbing weights, as it ensures both utility and tolerance for model parameter count and precision bits. Lastly, for the final private model fusion, the Gaussian mechanism based on a uniform distribution with a more relaxed privacy guarantee can be utilized. This hybrid mechanism achieves a balanced trade-off between privacy preservation and model utility.** These findings align with the theoretical analysis, which rigorously demonstrates that the hybrid mechanism satisfies the constraints and requirements of local differential privacy.

Furthermore, integrating PFA with the hybrid mechanism enhances its performance even further. The experimental results corroborate the effectiveness of this integration, revealing improved model accuracy while ensuring privacy preservation. Notably, the performance of the hybrid mechanism with PFA closely approaches that of the non-private model fusion, indicating a minimal loss in model accuracy due to privacy protection. In summary, the analysis of Table III unequivocally confirms the superiority of the hybrid mechanism in model perturbation. It effectively addresses various perturbation tasks and strikes a delicate balance between privacy and utility, as evidenced by its superior model accuracy compared to single Laplace, Gaussian, and MultiBit mechanisms. The incorporation of the Perturbation-Filter Adapter (PFA) serves to augment the performance of the hybrid mechanism, bringing its fusion model closer to the lossless model fusion without privacy protection.

C. Utility and Privacy Analysis

Utility Analysis. The utility is primarily demonstrated through the acceptance of model fusion results and the reduced reliance on extensive participation of raw data in the PrivFusion framework. However, it is important to consider the impact of the graph matching process on space requirements and overall usability. In the model fusion process, graph matching incurs significant spatial overhead due to the construction of the graph structure and the calculations involved in affinity calculation. This overhead can potentially affect the usability of the PrivFusion framework. However, extensive experimentation and evaluation have shown that PrivFusion achieves satisfactory model fusion results across various settings. Despite the space overhead, PrivFusion manages to strike a balance

TABLE IV
THE PERFORMANCE OF REAL-WORLD APPLICATION TASKS WITH DIFFERENT METRICS.

	Models	Privacy Budget $[\epsilon_w, \epsilon_a, \epsilon_f]$	Acc	ma_F1	w_F1	ma_Rec	w_Rec	ma_Prec	w_Prec
unbalanced	Model A	-	77.05%	0.7677	0.772	77.39%	77.05%	76.69%	77.87%
	Model B	-	77.30%	0.7653	0.7725	76.42%	77.30%	76.67%	77.22%
	Fused Model	-	76.91%	0.7573	0.7663	75.34%	76.91%	76.59%	76.79%
	PrivFusion(ours)	[0.01, 0.1, 1]	**77.10%**	**0.7605**	**0.769**	**75.72%**	**77.10%**	**76.67%**	**76.96%**
		[0.01, 0.1, 0.1]	76.57%	0.7539	0.7629	75.01%	76.57%	76.20%	76.43%
		[0.01, 0.1, 0.01]	63.65%	0.6243	0.6357	62.38%	63.65%	62.49%	63.50%
sd2t2	Model A	-	61.94%	0.6191	0.6173	64.22%	61.94%	64.54%	66.50%
	Model B	-	79.45%	0.7905	0.7954	79.36%	79.45%	78.88%	79.76%
	Fused Model	-	80.14%	0.7936	0.8004	79.12%	80.14%	79.72%	80.04%
	PrivFusion(ours)	[0.01, 0.1, 1]	**80.14%**	**0.7941**	**0.8006**	**79.22%**	**80.14%**	**79.77%**	**80.05%**
		[0.01, 0.1, 0.1]	79.40%	0.7862	0.7931	78.40%	79.40%	78.93%	79.30%
		[0.01, 0.1, 0.01]	59.59%	0.5887	0.5796	63.61%	59.59%	66.27%	68.87%
non-IID	Model A	-	95.50%	0.9549	0.955	95.47%	95.50%	95.53%	95.51%
	Model B	-	87.94%	0.8788	0.8797	88.21%	87.94%	87.80%	88.26%
	Fused Model	-	53.39%	0.5294	0.5328	52.95%	53.39%	52.98%	53.23%
	PrivFusion(ours)	[0.01, 0.1, 1]	**53.48%**	**0.5227**	**0.5284**	**52.58%**	**53.48%**	**52.73%**	**52.96%**
		[0.01, 0.1, 0.1]	51.49%	0.5134	0.5154	51.36%	51.40%	51.35%	51.62%
		[0.01, 0.1, 0.01]	46.27%	0.3164	0.2928	50.00%	46.27%	23.14%	21.41%

between privacy preservation and model utility. It successfully preserves privacy without compromising the overall effectiveness of model fusion. Moreover, PrivFusion reduces the reliance on large amounts of original data for model fusion. This is particularly beneficial in scenarios where access to extensive raw data is limited or impractical. By leveraging the shared knowledge and representations in the pre-trained models, the PrivFusion effectively combines models without requiring excessive participation of raw data, thus enhancing usability and reducing data dependency.

In terms of resource consumption, the non-private model fusion has an average processing time of 0.9s, while PrivFusion exhibits a slightly higher average processing time of 1.3s. PrivFusion requires more storage space due to its graph matching process. However, it offers a trade-off by sacrificing some space to achieve faster model updates. This makes PrivFusion suitable for applications that require frequent model updates and prioritize responsiveness over storage efficiency. The increased time overhead in PrivFusion mainly arises from computing activation values during inference and incorporating noise at different stages. Compared to the number of communication rounds typically involved in federated learning, our approach substantially reduces the overall communication overhead.

Privacy Analysis. The privacy analysis for PrivFusion focuses primarily on the treatment of model parameters, which undergo standardized processing. Unlike raw data, model parameters have lower sensitivity, and the level of noise introduced during privacy-preserving mechanisms can be controlled. To ensure the privacy of individual parties' models, PrivFusion employs local differential privacy. By incorporating randomized mechanisms like the hybrid-differentially private approach, PrivFusion introduces decentralized federated graph matching with the privacy-preserving affinity calculation process. This noise injection effectively safeguards the privacy

of individual model parameters while enabling efficient model fusion. To preserve privacy, PrivFusion avoids direct sharing of raw data and instead capitalizes on the shared knowledge and representations within pre-trained models. This collaborative approach allows parties to benefit from model fusion without disclosing their proprietary data or compromising their privacy. Moreover, the standardized processing of model parameters guarantees that sensitive information remains protected throughout the fusion process.

In contrast, previous researches such as [46], [47] typically evaluate their systems by testing different values of ϵ and selecting the minimum value that produces satisfactory results for deployment. For example, Microsoft uses $\epsilon = 1$ for collecting telemetry data from Windows users [46], while Apple's choice of ϵ in iOS and macOS ranges from 2 to 8 [47]. However, in our experiments, we have observed that our model remains remarkably robust even when subjected to small ϵ values below 1. This resilience suggests that our model consistently performs well under stringent privacy constraints, reinforcing its effectiveness in privacy-preserving scenarios.

D. Real-World Applications

To further investigate the effectiveness of our proposed method PrivFusion in real-world applications, we collected 10,217 color fundus images for diagnosing age-related macular degeneration (shown in Fig. 4) from 17 specialized ophthalmology hospitals in 15 regions. The dataset was split into training, validation, and test sets with an 0.8:0.05:0.15 ratio, respectively. We trained two models with 10 and 50 epochs to fuse in order to simulate the underfitting scenario (sd2t2). To investigate the influence of class imbalance on model fusion performance, 80% of normal data was randomly removed from the dataset, and one model was trained by removing 70% of patient data from the same dataset. These two trained models were then fused (unbalanced). In federated learning, different
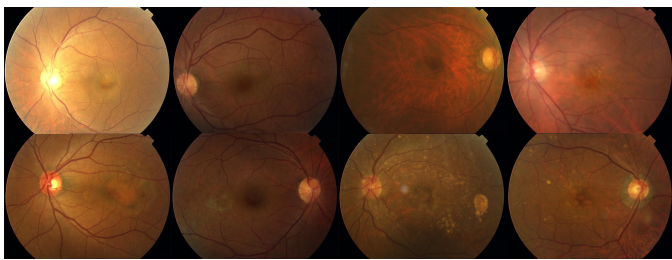
Fig. 4. Typical color fundus images in our age-related macular degeneration dataset.

participating parties' data often results in non-IID situations, so we categorized the data by region into two parties, each with data from 7 and 8 regions, and visualized the data in the form of Fig 4. The figures demonstrated significant statistical differences between the datasets from the two region collections. To assess the effectiveness of our method under non-IID conditions, we trained a model separately with data from each participating party and conducted a model fusion experiment (non-IID).
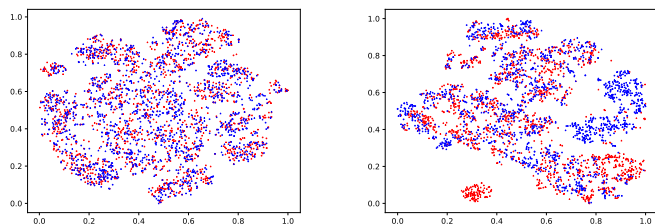

Fig. 5. The homogeneous and heterogeneous distribution of real application dataset.

Fig. 5 presents the true distribution of the training dataset, where the left plot represents an IID scenario, while the right plot portrays a non-IID scenario. The red and blue colors indicate distinct datasets, with circles and triangles representing different classes. In the left plot, both datasets exhibit a high degree of similarity, indicating a significant overlap in their distributions. Consequently, when integrating or combining models, we can expect improved performance since the datasets possess comparable characteristics. Conversely, the right plot demonstrates substantial dissimilarity and dispersion in the distributions of the two datasets. This discrepancy suggests the presence of diverse patterns and distinct variations between the datasets. Integrating models trained on such disparate datasets presents a greater challenge. The disparities in distribution can introduce inconsistencies in learned representations and decision boundaries, thereby impeding effective model fusion and optimal performance. Hence, the analysis of Fig. 5 underscores the significance of considering the distribution of training data when performing model fusion or combination. When datasets share similar distributions, fusion processes are likely to yield superior outcomes. Conversely, when significant distributional differences exist, fusion tasks become more intricate, necessitating meticulous consideration and potentially specialized techniques to address the disparities.

In Table IV, we present the experimental results under non-homogeneous, underfitting, and non-IID settings. We evaluated the model fusion using metrics such as accuracy (Acc), macro F1-score (ma_F1), weighted F1-score (w_F1), macro recall (ma_Rec), weighted recall (w_Rec), macro precision (ma_Prec), and weighted precision (w_Prec). The results demonstrate that PrivFusion achieves satisfactory performance across all these metrics. PrivFusion effectively addresses the challenges associated with non-homogeneous data, underfitting, and non-IID distribution. It successfully combines models from multiple parties while preserving privacy and maintaining good utility. The measured metrics consistently indicate high performance in terms of accuracy, F1 score, recall, and precision.

Moreover, PrivFusion allows users to adjust the privacy budget, enabling a trade-off between utility and privacy. By varying the privacy budget, users can fine-tune the level of privacy protection while ensuring acceptable utility. The results in Table IV confirm the effectiveness of PrivFusion in model fusion under diverse and challenging data settings. It achieves a balance between utility and privacy, facilitating secure collaboration and model fusion among multiple parties with satisfactory performance across various evaluation metrics. These findings highlight the practicality and effectiveness of PrivFusion in real-world scenarios.

VI. CONCLUSION

Decentralized model fusion has emerged as a promising approach for aggregating models while preserving privacy in collaborative settings. In this paper, we present PrivFusion, a novel architecture for privacy-preserving decentralized model fusion. PrivFusion addresses the challenges of sharing trained models under local differential privacy constraints, ensuring data confidentiality and secure collaboration. Our method leverages a graph-based representation of the neural network and employs randomized mechanisms for decentralized federated graph matching to protect the privacy of neuron alignment during model fusion. We analyze the impact of different privacy budgets on model fusion and highlight the significant contribution of activation values to the fused model's performance. Moreover, our findings demonstrate the advantages of adopting a hybrid LDP approach, which combines the benefits of LDP mechanisms for preserving privacy during the fusion process. Experimental results on two image datasets and real-world healthcare applications validate the effectiveness of PrivFusion in preserving privacy while maintaining model utility.

In conclusion, PrivFusion offers a valuable solution for privacy-preserving decentralized model fusion, with potential applications in various domains. Future research can explore the scalability of PrivFusion and extend its use to large-scale pre-trained models in collaborative scenarios. Additionally, investigating noise-adaptive perturbation techniques is another avenue for future research. The advancements made in this paper contribute to the field of privacy-preserving decentralized model fusion, enabling secure and collaborative data analysis.

REFERENCES

[1] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill *et al.*, "On the opportunities and risks of foundation models," *arXiv preprint arXiv:2108.07258*, 2021.

[2] HuggingFace, "Huggingface," https://huggingface.co.

[3] Alibaba, "Modelscope," https://modelscope.cn.

[4] J. Howard and S. Ruder, "Universal language model fine-tuning for text classification," *arXiv preprint arXiv:1801.06146*, 2018.

[5] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," in *International conference on machine learning*. PMLR, 2015, pp. 1180–1189.

[6] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.

[7] S. P. Singh and M. Jaggi, "Model fusion via optimal transport," *Advances in Neural Information Processing Systems*, vol. 33, pp. 22 045–22 055, 2020.

[8] D. C. Nguyen, Q.-V. Pham, P. N. Pathirana, M. Ding, A. Seneviratne, Z. Lin, O. Dobre, and W.-J. Hwang, "Federated learning for smart healthcare: A survey," *ACM Computing Surveys (CSUR)*, vol. 55, no. 3, pp. 1–37, 2022.

[9] C. Liu, C. Lou, R. Wang, A. Y. Xi, L. Shen, and J. Yan, "Deep neural network fusion via graph matching with applications to model ensemble and federated learning," in *International Conference on Machine Learning*. PMLR, 2022, pp. 13 857–13 869.

[10] S. K. Ainsworth, J. Hayase, and S. Srinivasa, "Git re-basin: Merging models modulo permutation symmetries," *arXiv preprint arXiv:2209.04836*, 2022.

[11] A. K. Akash, S. Li, and N. G. Trillos, "Wasserstein barycenter-based model fusion and linear mode connectivity of neural networks," 2022.

[12] X. Jin, X. Ren, D. Preotiuc-Pietro, and P. Cheng, "Dataless knowledge fusion by merging weights of language models," 2022.

[13] L. Breiman, "Bagging predictors," *Machine learning*, vol. 24, pp. 123–140, 1996.

[14] D. H. Wolpert, "Stacked generalization," *Neural networks*, vol. 5, no. 2, pp. 241–259, 1992.

[15] Y. Freund and R. E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," *Journal of computer and system sciences*, vol. 55, no. 1, pp. 119–139, 1997.

[16] L. Zhang, L. Shen, L. Ding, D. Tao, and L.-Y. Duan, "Fine-tuning global model via data-free knowledge distillation for non-iid federated learning," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 10 164–10 173.

[17] H. Wang, M. Yurochkin, Y. Sun, D. S. Papailiopoulos, and Y. Khazaeni, "Federated learning with matched averaging," in *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia*, 2020.

[18] X.-C. Li, Y.-C. Xu, S. Song, B. Li, Y. Li, Y. Shao, and D.-C. Zhan, "Federated learning with position-aware neurons," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 10 072–10 081.

[19] F. Tramèr, F. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, "Stealing machine learning models via prediction apis." in *USENIX security symposium*, vol. 16, 2016, pp. 601–618.

[20] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill *et al.*, "On the opportunities and risks of foundation models," *arXiv preprint arXiv:2108.07258*, 2021.

[21] 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*, 2015, pp. 1322–1333.

[22] L. Zhu, Z. Liu, and S. Han, "Deep leakage from gradients," *Advances in neural information processing systems*, vol. 32, 2019.

[23] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in *2017 IEEE symposium on security and privacy (SP)*. IEEE, 2017, pp. 3–18.

[24] N. Carlini, S. Chien, M. Nasr, S. Song, A. Terzis, and F. Tramer, "Membership inference attacks from first principles," in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022, pp. 1897–1914.

[25] C. Dwork, A. Roth *et al.*, "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.

[26] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, "Our data, ourselves: Privacy via distributed noise generation," in *Advances in Cryptology - EUROCRYPT 2006*, S. Vaudenay, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 486–503.

[27] C. Gentry, *A fully homomorphic encryption scheme*. Stanford university, 2009.

[28] W. Wang, Y. Hu, L. Chen, X. Huang, and B. Sunar, "Exploring the feasibility of fully homomorphic encryption," *IEEE Transactions on Computers*, vol. 64, no. 3, pp. 698–706, 2013.

[29] Y. Lindell, "Secure multiparty computation for privacy preserving data mining," in *Encyclopedia of Data Warehousing and Mining*. IGI global, 2005, pp. 1005–1009.

[30] P. Mohassel and Y. Zhang, "Secureml: A system for scalable privacy-preserving machine learning," in *2017 IEEE symposium on security and privacy (SP)*. IEEE, 2017, pp. 19–38.

[31] S. Truex, N. Baracaldo, A. Anwar, T. Steinke, H. Ludwig, R. Zhang, and Y. Zhou, "A hybrid approach to privacy-preserving federated learning," in *Proceedings of the 12th ACM workshop on artificial intelligence and security*, 2019, pp. 1–11.

[32] Y. Qu, S. Yu, L. Gao, W. Zhou, and S. Peng, "A hybrid privacy protection scheme in cyber-physical social networks," *IEEE Transactions on Computational Social Systems*, vol. 5, no. 3, pp. 773–784, 2018.

[33] Ú. Erlingsson, V. Pihur, and A. Korolova, "Rappor: Randomized aggregatable privacy-preserving ordinal response," in *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, 2014, pp. 1054–1067.

[34] M. Yang, L. Lyu, J. Zhao, T. Zhu, and K.-Y. Lam, "Local differential privacy and its applications: A comprehensive survey," *arXiv preprint arXiv:2008.03686*, 2020.

[35] C. Wei, S. Ji, C. Liu, W. Chen, and T. Wang, "Asgldp: collecting and generating decentralized attributed graphs with local differential privacy," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 3239–3254, 2020.

[36] S. P. Kasiviswanathan, K. Nissim, S. Raskhodnikova, and A. D. Smith, "Analyzing graphs with node differential privacy." in *TCC*, vol. 13. Springer, 2013, pp. 457–476.

[37] P. Liu, Y. Xu, Q. Jiang, Y. Tang, Y. Guo, L.-e. Wang, and X. Li, "Local differential privacy for social network publishing," *Neurocomputing*, vol. 391, pp. 273–279, 2020.

[38] C. Wu, F. Wu, Y. Cao, Y. Huang, and X. Xie, "Fedgnn: Federated graph neural network for privacy-preserving recommendation," *arXiv preprint arXiv:2102.04925*, 2021.

[39] S. Sajadmanesh and D. Gatica-Perez, "Locally private graph neural networks," in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 2130–2145. [Online]. Available: https://doi.org/10.1145/3460120.3484565

[40] W. Lin, B. Li, and C. Wang, "Towards private learning on decentralized graphs with local differential privacy," *IEEE Transactions on Information Forensics and Security*, vol. 17, pp. 2936–2946, 2022.

[41] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*. Springer, 2006, pp. 265–284.

[42] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.

[43] H. W. Kuhn, "The hungarian method for the assignment problem," *Naval Research Logistics Quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/nav.3800020109

[44] M. Cuturi, "Sinkhorn distances: Lightspeed computation of optimal transport," in *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2*, ser. NIPS'13. Red Hook, NY, USA: Curran Associates Inc., 2013, p. 2292–2300.

[45] M. Leordeanu and M. Hebert, "A spectral technique for correspondence problems using pairwise constraints," in *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, vol. 2. IEEE, 2005, pp. 1482–1489.

[46] B. Ding, J. Kulkarni, and S. Yekhanin, "Collecting telemetry data privately," *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[47] J. Tang, A. Korolova, X. Bai, X. Wang, and X. Wang, "Privacy loss in apple's implementation of differential privacy on macos 10.12," *CoRR*, vol. abs/1709.02753, 2017. [Online]. Available: http://arxiv.org/abs/1709.02753