

Towards Multi-modal Transformers in Federated Learning

Guangyu Sun¹, Matias Mendieta¹, Aritra Dutta², Xin Li², and Chen Chen¹

¹ Center for Research in Computer Vision, University of Central Florida, FL, USA

² Department of Mathematics, University of Central Florida, FL, USA

{guangyu, matias.mendieta, aritra.dutta, xin.li}@ucf.edu;
chen.chen@crcv.ucf.edu

Abstract. Multi-modal transformers mark significant progress in different domains, but privacy concerns on high-quality data hinder their further improvement. Federated learning (FL) has emerged as a promising privacy-preserving paradigm for training models without direct access to the raw data held by different clients. Despite its potential, a considerable research direction regarding the unpaired uni-modal clients and the transformer architecture in FL remains unexplored. To fill this gap, this paper explores a transfer multi-modal federated learning (MFL) scenario within the vision-language domain, where clients possess data of various modalities distributed across different datasets. We systematically evaluate the performance of existing methods when a transformer architecture is utilized and introduce a novel framework called Federated modality complementary and collaboration (FedCola) by addressing the in-modality and cross-modality gaps among clients. Through extensive experiments across various FL settings, FedCola demonstrates superior performance over previous approaches, offering new perspectives on future federated training of multi-modal transformers. Code is available at <https://github.com/imguangyu/FedCola>.

Keywords: Federated Learning · Multi-modal Learning · Transformer

1 Introduction

Multi-modal transformers have led to remarkable advancements across a spectrum of downstream tasks [2, 22, 42]. Nevertheless, training these models demands voluminous and high-quality data [15, 31]. Although high-quality training datasets contain a wealth of information, acquiring such data is resource-intensive. Moreover, many of these datasets are often closely guarded by various entities and fall under diverse privacy regulations, resulting in *data silos* [21]. These silos pose remarkable challenges to further improvements in the field, as they prevent the leveraging of isolated datasets to enhance model training and performance.

In response to this challenge, **federated learning** (FL) [35] proposes a novel training paradigm typically designed to circumvent the need for direct access to




Fig. 1: The **transfer multi-modal federated learning** setting in the vision-language domain. The clients possess data of various modalities distributed across different datasets and different local training objectives. The server aims to collaboratively train a *multi-modal transformer* with the data from all clients.

raw data. FL enables a decentralized approach to model training. Specifically, a central server cyclically disseminates a *global model* to a selected pool of clients for *local training*. After the local training, the server receives the model updates from the clients, averages them with specific weights as the *global aggregation*, and applies the aggregated updates to the global model. The process continues until convergence. FL allows collaborative training on heterogeneous data spread over geographically remote clients while simultaneously preserving their data privacy and offers a promising solution to the issue of data silos by facilitating the collaborative improvement of models.

Current FL research predominantly addresses uni-modal scenarios and explores different aspects, such as local training [28, 37], personalization [3, 27, 34, 46, 48], aggregation [18, 50], and initialization [5, 39, 49]. Nonetheless, in domains with rich multi-modal data, such as healthcare and the Internet of Things (IoTs) [1, 53, 59], **multi-modal federated learning** (MFL) are gaining attention [4]. Studies have explored horizontal MFL in healthcare, which enables training clients that share identical input modalities (*e.g.*, all clients have text and image) [1, 6, 53], and vertical MFL in the IoTs, which enables clients that possess a single modality from the same users with multi-modal data (*e.g.*, some clients may have text-only data, while some have their corresponding image data) [59]. However, operating explicitly in either of the settings excludes other uni-modal clients with unpaired data from participating in the MFL process. Unpaired uni-modal data has shown the capability to improve the multi-modal model under a centralized setting [2, 44]. Therefore, transferring knowledge from those uni-modal clients can further extend the scope of training data in FL, providing an important impact for MFL. But only a few studies have explored this transfer MFL setting [33, 56] shown in Fig. 1.

Although promising, there are two key challenges in the transfer MFL due to the different modality data and training objectives across clients; see Fig. 1. First, restrictive data accessibility of the uni-modal clients. Multi-modal clients can directly gain cross-modal knowledge from local training data, but uni-modal clients have never accessed the other modalities. Consequently, the client mod-

els become biased toward the local modality, leading to a *cross-modality gap*. Meanwhile, even within the same modality, multi-modal and uni-modal models are trained with different objectives, leading to an *in-modality gap*.

To tackle these challenges, we propose a novel framework for the transfer MFL in both local training and global aggregation, leveraging the unified design of transformers across different modalities. During the local training, we propose a *mixture strategy* to leverage the modality-complementary information, together with a weight compression trick for efficient communication. During the global aggregation, we propose an *aggregation with disaggregation* strategy to aggregate general knowledge while maintaining specific knowledge for each local objective. With these collaboration techniques for both local training and global aggregation, we build and propose a novel framework, **Federated modality Complementary and collaboration** (FedCola), providing insights for FL on multi-modal transformers, and highlight the main contributions of this paper:

- (i) **Transfer multi-modal FL (§3)**. To the best of our knowledge, this paper is the first to explore transformers in transfer multi-modal FL, filling the gap between centralized and federated training of multi-modal transformers.
- (ii) **FedCola (§4)**. We propose a novel framework, FedCola, to enable advanced collaboration between different modalities with only the parameters in both local training and global aggregation.
- (iii) **Cross-modality knowledge-sharing (§4.2)**. We illustrate that cross-modality knowledge-sharing can be achieved without directly accessing the data but only with the model weights and provide new insights on convergence conditions of collaboration and quantified contributions.
- (iv) **Empirical study (§5)**. We conduct extensive experiments on real-world datasets under different FL settings, domain gaps, and numbers of participating clients, demonstrating the effectiveness of the proposed framework.

2 Related Work

Federated learning has emerged as a powerful approach for decentralized, privacy-preserving machine learning. In practice, heterogeneity in different aspects exists among clients, including statistical heterogeneity and system heterogeneity. Most current FL focuses on addressing the statistical heterogeneity for clients with data from the same domain, the same modality, but different distributions [23, 26, 28, 37, 61]. Meanwhile, the other line of work focuses on enabling edge devices with limited resources to address the system heterogeneity [8, 16, 19, 32, 36]. Meanwhile, FL can also be categorized by the focus of the product, which can either be the performance of the global model on more generalized test data or the performance of the client models on their local data. These two categories lead to two separate directions of FL: global and personalized FL [10, 25, 30, 34, 38, 46]. In this paper, we explore and address the heterogeneity in the modality level to obtain a generalized multi-modal transformer, which falls into the scope of global FL without system heterogeneity.

Multi-modal FL. Modality collaboration, or multi-modal fusion, allows different modalities to train and learn collaboratively by sharing high-level common knowledge and has recently been introduced in FL [14, 53, 59]. The scope of multi-modality in FL is limited to a handful of works, and the potential is largely unexplored. For instance, earlier works consider a homogeneous model for each modality [53, 59] and further involve clients with missing modality [1, 6]; recently, [9] took a step towards a larger server model training by knowledge transfer of diverse smaller client models, and [56] use knowledge-transfer between uni and multi-modal clients with heterogeneous data and learn a larger global model. We note that some multi-modal frameworks are restrictive, barring the participation of uni-modal clients [53]; [56] lifts this barrier but still relies on the auxiliary from a public dataset shared with all clients. In this paper, we focus on transferring the knowledge without any public data but with the multi-modal capability of transformers, further reducing the restrictions.

Vision transformers in FL. Although vision transformers have become the de facto architecture in centralized computer vision tasks, their exploration under FL remains limited. Qu *et al.* [41] evaluate the performance of vision transformers under global FL settings and showcase their robustness against data heterogeneity. Meanwhile, Sun *et al.* [46] and Li *et al.* [25] study the personalization of vision transformers in personalized FL. Further, research with other aspects of vision transformers (*e.g.*, pre-trained transformers) have also gained attention [5, 39, 47, 60]. With a different perspective in this paper, we focus on the multi-modal capability and unified architecture of the transformer, paving the path for further large federated multi-modal transformers.

3 Transfer MFL

We consider a transfer multi-modal FL setting in the vision-language domain with a total of N clients. That is, we have N_v image clients, N_l text clients, and N_{vl} image-text clients, with $N_v + N_l + N_{vl} = N$. Clients with the same modality own non-IID local training data with the same local objective. Uni-modal clients conduct classification tasks, while multi-modal clients conduct cross-modal retrieval tasks as their local objectives. The server owns a hold-off multi-modal test set for evaluation.

At the beginning of the FL process, the server initializes global models, $w^{(0)} = (w^{(0,v)}, w^{(0,vl_v)}, w^{(0,vl_l)}, w^{(0,l)})^\top$ for each modality combination. The entire process runs for T global communication rounds. In each round, t , for each modality type M with a sample rate r_M , a total of $r_M N_M$ is selected to perform the local training. Note that if all modality combinations share a fixed uniform sample rate, then $r_M = r$. Each selected client, i , downloads the global models, $w^{(t)}$, performs training for E local epochs, and sends the local update, $\nabla w_i^{(t+1)} = w_i^{(t+1)} - w^{(t)}$ to the server.

The server, first, aggregates the updates for each modality combination separately with a uni-modal aggregation algorithm. When FedAvg is used, the global

updates for modality type M is

$$\nabla \bar{w}^{(t+1,M)} = \sum_{i=1}^{r_M N_M} \frac{n_i}{n_M} \nabla w_i^{(t+1)},$$

where $n_M = \sum_{i=1, \mathcal{M}(i)=M}^{r_M N_M} n_i$ is the total training samples on selected clients and \mathcal{M} is a mapping that maps the client index i to its corresponding modality type. Additionally, in our setting, we perform further collaboration across all modality types to get collaborated updates $\nabla w^{(t+1)}$ as

$$\nabla w^{(t+1)} = \Omega \nabla \bar{w}^{(t+1)},$$

where Ω is the collaboration matrix for all parameters in all models w . Finally, we update the global models by $w^{(t+1)} = w^{(t)} + \nabla w^{(t+1)}$ that goes in for the next round. After T rounds, we evaluate $w^{(T,vl)}$ with the multi-modal test data on the server. Detailed discussion on the general transfer MFL framework and convergence guarantee in our setting are given in the supplementary material.

4 Method

In this section, we introduce our framework and outline two methods to mitigate cross-modality and in-modality gaps.

4.1 The framework

To address the cross-modality and in-modality gaps, a key goal is to unify the learned representation into a shared feature space [56]. Under centralized training, those gaps can be addressed by joint training in a multi-task fashion with all the data in different modalities. However, parameters are the only knowledge carrier in FL, due to the privacy requirements of local clients, preventing any direct joint training across client data and tasks. Instead, we leverage the unifying architecture of transformers across modalities and develop a novel framework with new insights into local training and global aggregation; see Fig. 2. Specifically, we split the model into three parts: (i) embedding layers, (ii) transformer blocks, and (iii) task head. Among different clients, the transformer blocks are in a unified architecture, while the embedding layers and task heads are specified based on their input modalities and local training objectives. We propose a novel complementary local training pipeline on the uni-modal clients (§4.2) while exploring the collaborative aggregation on the server (§4.3).

4.2 Complementary local training against cross-modality gap

Local training on uni-modal clients is inherently limited to utilizing only the local data available on the respective clients, thus preventing information access from other modalities. This limitation becomes particularly pronounced when employing disparate model architectures for encoding images and texts. Specifically, images cannot be encoded using the model deployed on text clients, and

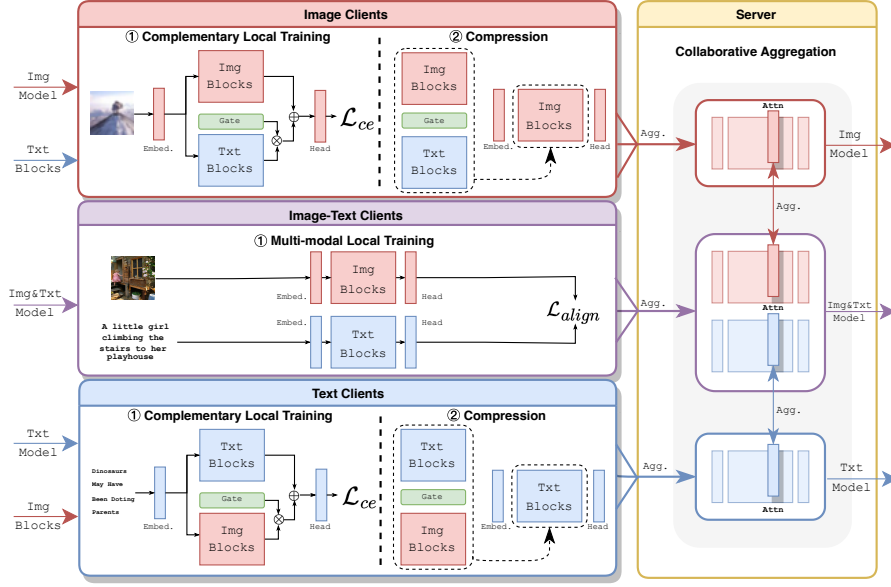


Fig. 2: Overview of our proposed framework, **FedCola**. In each round of FL, uni-modal clients download the global model for their modality along with the transformer blocks from the other modalities and perform complementary local training to address the *cross-modality gap*, while multi-modal clients perform standard multi-modal local training. Then, all clients send their local updates to the server for uni-aggregation and collaborative aggregation to address the *in-modality gap*.

vice versa. Consequently, the valuable knowledge encapsulated within the text modality, which could complement the local data, remains untapped by image clients relying solely on model weights derived from other modalities. With the adoption of transformer blocks possessing a unified architecture, tokenized local data can now be encoded using blocks originating from other modalities. This cross-modal encoding capability enables the design of paradigms to integrate signals from diverse modalities into local training processes.

As recent progress on the **Mixture of Experts** (MoE) [2, 20, 58], each layer from different expert models can be mixed by weights generated by a router, providing a new perspective for parameter collaborations. In contrast to the original MoE formulation, wherein all expert models are considered equally significant and the router learns to weigh them based on input characteristics, the scenario of local training on uni-modal clients presents a more straightforward approach. Given that all input data on uni-modal clients pertains to the local modality, there exists a clear emphasis on prioritizing local information, with models from other modalities serving as auxiliary components.

To signify local modality, we employ a learnable parameter initialized to 0 as the **gate**, g . This parameter enables adaptive learning, determining how complementary knowledge from other modalities contributes to the local modality.

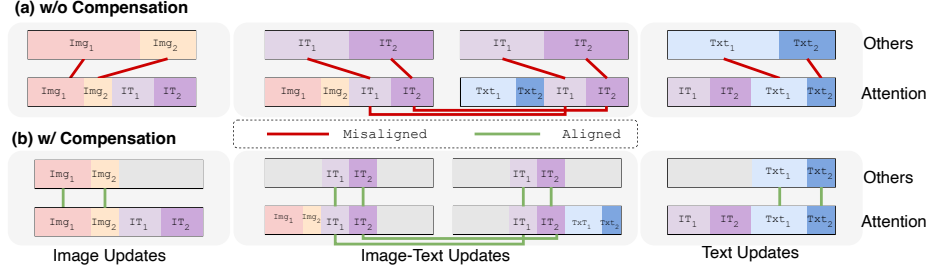


Fig. 3: Illustration of the self-attention (Attention) and other layers updates with and without the proposed compensation scheme. The width of the block indicates the aggregation coefficient on that client. Without the compensation, layer-level misalignment happens between self-attention and other layers, while modality-level misalignment happens between updates of each modality on multi-modal updates. With compensation, both misalignment is fixed.

Formally, the output of the layer with the gate g when takes \mathbf{x} as the input can be formulated as $\mathbf{W}_{\text{local}}\mathbf{x} + g\mathbf{W}_{\text{out}}\mathbf{x}$, where $\mathbf{W}_{\text{local}}$ and \mathbf{W}_{out} are the local and out model weights respectively. In this way, uni-modal clients can learn modality-complementary knowledge during local training, closing the cross-modality gap. To conduct complementary local training, we additionally download the transformer blocks from the other complementary modality and add gates for all linear layers within the transformer blocks, as shown in Fig. 2 (Left), encompassing self-attention layers and MLPs for the complementary cross-modal contribution.

However, if we directly send the updates of the modified local model, the upload communication cost will be almost doubled due to the involvement of more parameters. To remedy this, we compute the equivalent weights for each layer after the local training as $\mathbf{W}_{\text{local}} + g\mathbf{W}_{\text{out}}$ as a compression trick since they are linear, so the total uploading size is reduced to the size of the original model. Regarding the download size, we further explore options to reduce the required model weights from the complementary modality in § 6.1. In this way, we provide the flexibility of the trade-off between performance and communication costs.

4.3 Collaborative aggregation against in-modality gap

Global aggregation in FL presents both opportunities and challenges. On the one hand, aggregating model weights from diverse clients promises to enhance generalizability, while on the other, the inherent heterogeneity among clients poses significant obstacles to effective aggregation [13, 35, 54]. While aggregation within the same modality constitutes a uni-aggregation problem, extensively studied within uni-modal FL, our focus shifts towards the collaboration dynamics among different modality/task combinations.

A given modality may be represented across multiple client types, *e.g.*, image-only clients and hybrid image-text clients both use image data. The image models within these clients are trained on data from the same modality but with distinct

local objectives. Consequently, the parameters of image models within these clients encapsulate the broader domain-specific knowledge inherent to image data and the intricacies of their respective local objectives. Given that the transformer blocks adhere to a uniform architecture, direct aggregation of their model weights is feasible. However, such a uniform aggregation approach risks diluting the specific knowledge pertinent to individual local objectives. To remedy this, we propose an *aggregation with disaggregation* strategy that captures overarching modality knowledge while safeguarding task-specific insights.

Inspired by transformers in centralized training paradigms [2, 51], we observe that self-attention layers encode inter-token relationships, embodying more generalized knowledge. Conversely, the subsequent MLPs, serve to adapt this generalized knowledge to the local objectives, thus encapsulating domain-specific knowledge. This insight motivates us to develop the selective aggregation of updates exclusively from *self-attention layers* of transformer blocks within the same modality following the aggregation of updates for each modality combination of clients. Concurrently, we advocate for the disaggregation of other components, constituting what we term *in-modal collaboration*. Formally, the in-modal collaboration matrix Ω (with compatible dimensions of the identity matrix, I) is given as:

$$\Omega_{i,j}^{\text{attn}} = \begin{cases} \frac{n_i}{n_i+n_j}I & \mathcal{M}(i) = \mathcal{M}(j), \\ 0 & \text{otherwise,} \end{cases}; \quad \Omega_{i,j}^{\text{others}} = I.$$

Note that $\Omega_{i,j}^{\text{attn}}$ combines the self-attention layers and $\Omega_{i,j}^{\text{others}}$ combines the MLPs. However, such an aggregation will cause misalignment between the model weights with and without collaboration. Intuitively, the updates from a single client are coherent across different layers since they are directly trained. Scaling the updates for specific layers will break the coherence and lead to a misalignment in the global updates, as shown in Fig. 3(a). Quantitatively, during the initial uni-aggregation within the same modality type, the updates from client i are equally applied to each layer of the global model with the coefficient of $\frac{n_i}{n_{\mathcal{M}(i)}}$. For the collaborative aggregation, the updates to the self-attention layers are scaled to $\frac{n_i}{n_{\mathcal{M}(i)}+n_{v1}}$, leading to *layer-level misalignment*. Meanwhile, for the updates on the image-text model, self-attention layers in it are scaled with $\frac{n_{v1}}{n_v+n_{v1}}$ and $\frac{n_{v1}}{n_l+n_{v1}}$ separately, leading to a *modality-level misalignment*.

To fix such a misalignment, we propose a compensation scheme to align the updates for different layers. Specifically, we scale the updates to the non-collaborated layers with the same coefficient as the self-attention layers. With compatible dimensions of the identity matrix, I , and the matrix of all zeros, 0 , the compensation scheme is formulated as

$$\Omega_{\text{comp}} = \begin{pmatrix} \frac{n_v}{n_v+n_{v1}}I & \frac{n_{v1}}{n_v+n_{v1}}I & 0 & 0 \\ \frac{n_v}{n_v+n_{v1}+n_l}I & \frac{n_{v1}}{n_v+n_{v1}+n_l}I & 0 & 0 \\ 0 & 0 & \frac{n_{v1}}{n_v+n_{v1}+n_l}I & \frac{n_l}{n_v+n_{v1}+n_l}I \\ 0 & 0 & \frac{n_l}{n_{v1}+n_l}I & \frac{n_l}{n_{v1}+n_l}I \end{pmatrix}.$$

Such a compensation scheme will ensure that the updates to the different parts of the transformer blocks are aligned, as shown in Fig. 3(b). With in-modal collaboration, both uni-modal and multi-modal clients can gain more generalized multi-modal knowledge by aggregating from more training data. Meanwhile, the specific knowledge for their local objective is maintained by the aggregation with disaggregation strategy. Consequently, the in-modality gap is mitigated by the collaborative aggregation.

5 Experiments

Datasets. We follow the previous work’s setting to select the datasets for each modality type [56]. We use CIFAR-100 [24] as the image data for image clients and the AG NEWS [57] dataset as the text data for text clients. We evaluate the performance on two different multi-modal datasets, Flickr30k [40] and COCO Captions [7]. We sample 10,000 images along with 50,000 corresponding captions from Flickr30k, noted as Flickr10k, and 50,000 image-text pairs from COCO Captions. To further evaluate the impact of the domain gap between uni-modal and multi-modal clients, we use two medical domain datasets to provide a larger domain gap. Specifically, we use OrganCMNIST [55] as the medical image dataset and the Medical Abstract [43] dataset as the medical text dataset. Classification tasks are performed on the uni-modal clients trained with cross-entropy loss, while cross-modal retrieval tasks are performed on multi-modal clients trained with contrastive loss [17].

FL settings. We choose a practical FL setting as our default setting and further provide two more challenging scenarios with more heterogeneity and less participation. Specifically, we set a number of the image, text, and multi-modal clients as 12, 12, 8 in our default setting. We partition the training data to corresponding clients following non-IID Dirichlet distributions with $\alpha = 0.5$ without overlapping. In each round, $r = 0.25$ of the clients on each modality type will train 5 local epochs and participate in the aggregation in a total of $T = 30$ rounds. For the setting with more heterogeneity, we partition the data following a Dirichlet distribution with a smaller $\alpha = 0.1$, and for the setting with less participation, we use a lower participation rate of $r = 0.125$.

Model architecture. We employ an ImageNet pre-trained ViT-Small [12] as the transformer blocks. Images are embedded with a patch embedding layer with a patch size of 16, and texts are embedded with a BERT tokenizer [11]. The model architecture will be the embedding layers, transformer blocks, and the classification head in the uni-modal clients, while two separate embedding layers, transformer blocks, and retrieval heads in the multi-modal clients. All the training images are resized to 224×224 , and texts are tokenized with a maximum length of 40.

Comparison methods. Considering there is no previous work focusing on the same setting, we compare our proposed method with the following methods that can apply to the transfer multi-modal FL setting:

Table 1: Performance of (a) uni-modal methods, (b) multi-modal methods, and (c) our proposed method under different FL settings. Arrows mean setting changes compared to Default. The evaluation metric is the sum of top-1 **recalls** following previous settings.

Method		Default		More Heterogeneity		Less Participation	
		$\alpha = 0.5, r = 0.25$		$\alpha = 0.1 \downarrow, r = 0.25$		$\alpha = 0.5, r = 0.125 \downarrow$	
		FlickR	COCO	FlickR	COCO	FlickR	COCO
(a)	FedAvg	81.08	95.42	81.70	95.32	64.82	83.91
	FedProx	78.55	95.16	76.33	95.62	63.33	79.88
(b)	CreamFL	74.83	95.26	80.00	91.41	66.85	78.97
	FedIoT	85.51	98.40	83.28	95.89	61.94	80.65
(c)	FedCola	91.96	105.10	91.82	100.83	88.85	102.30

Method		More Image		More Text		Fewer Image-Text	
		$r = (0.33 \uparrow, 0.25, 0.25)$		$r = (0.25, 0.33 \uparrow, 0.25)$		$r = (0.25, 0.25, 0.125 \downarrow)$	
		FlickR	COCO	FlickR	COCO	FlickR	COCO
(a)	FedAvg	78.28	97.28	79.62	96.69	61.12	75.10
	FedProx	79.25	95.39	77.59	94.96	59.19	73.08
(b)	CreamFL	80.31	93.65	80.75	92.81	57.12	74.69
	FedIoT	82.74	95.04	78.02	97.04	60.34	76.14
(c)	FedCola	91.24	104.22	90.10	100.96	85.68	94.40

- (i) **Uni-modal methods:** We compare our proposed method with FedAvg [35] and FedProx [28] as the most widely applied baselines for uni-modal FL.
- (ii) **Adapted multi-modal methods:** Considering methods for horizontal and vertical multi-modal FL are infeasible under the transfer MFL setting, only a few multi-modal FL methods are applicable. We choose CreamFL [56], which is the state-of-the-art knowledge-distillation-based method for multi-modal FL, and FedIoT [59], which extends FedAvg with a multi-modal aggregation strategy, as the comparison methods. Implementation details of the adaptation are provided in the supplementary material.

Evaluation Metrics. Following a related setting in previous multi-modal FL [56], we evaluate the sum of the *Top-1 Recalls* on the image-to-text and text-to-image retrieval under settings with 1k and 5k test images, for which higher is better.

5.1 Evaluation under different FL settings

As shown in Table 1, we first evaluate the performance of all methods under the default setting. Consistent with previous exploration on transformers in FL [41], FedAvg is a strong baseline when transformers are applied as the backbone. Meanwhile, FedProx, which has shown effectiveness with CNNs, is less effective than FedAvg with transformers. Similarly, CreamFL and FedIoT show a performance gain compared with FedAvg only under certain settings. However, our proposed method, FedCola, consistently surpasses all other methods across




Fig. 4: Relative performance of each multi-modal method compared to FedAvg under different domain gaps

Table 2: Impact of each proposed module. CA: collaborative aggregation; CP: compensation scheme in CA; CL: complementary local training.

CA	CP	CL	R@1 _{sum}
✗	✗	✗	81.08
✓	✗	✗	88.70
✓	✓	✗	90.09
✓	✓	✓	91.96

diverse settings and datasets. FedCola’s incorporation of complementary modalities in local training enables sharing knowledge embedded within model weights across varied clients in FL. Simultaneously, through collaborative aggregation, the generalizability of multi-modal models is bolstered by leveraging collaboration with uni-modal counterparts. We explore each component further in § 6.1. This approach concurrently diminishes the cross-modality gap among distinct uni-modal clients by aggregating them with multi-modal models exhibiting superior alignment, thereby yielding substantial performance improvements.

We then evaluate the performance of all methods under the settings with more heterogeneity and less participation. As shown in Table 1 (Top), our proposed method, FedCola, still outperforms all other methods and does not undergo a significant performance drop. This result demonstrates the robustness of our proposed method in handling various FL settings.

Furthermore, the number of selected clients in each modality type will affect the ratio in the aggregation. Considering the fact that uni-modal data is usually more accessible than multi-modal data, we also evaluate the performance of all methods under the setting with more participating image clients, more participating text clients, and fewer participating multi-modal clients, as shown in Table 1 (Bottom).

We notice that FedCola consistently outperforms all other methods under all settings, demonstrating the robustness of our proposed method in handling imbalanced aggregation with a consistent performance gain compared to the FedAvg baseline.

5.2 Evaluation under different domain gaps

Considering different datasets are used for different modality types, it is important to evaluate the impact of domain gaps. Anchored by the multi-modal dataset, we further evaluate the robustness of the multi-modal performance when larger domain gaps are introduced. Specifically, we use the OrganCMNIST [55] dataset and the Medical Abstracts [43] dataset to replace the image and text datasets on uni-modal clients.

Table 3: Communication and performance trade-off with different complementary local learning strategies. Co-Layer indicates the collaborated layers.

Co-Layer	Trainable	Comm. Cost (MB)	R@1 _{sum}
None	-	208.81	90.09
Blocks	✓	371.26	91.96
Blocks	✗	371.26	91.22
MLP	✓	316.98	91.32
Attention	✓	<u>262.95</u>	<u>91.73</u>




Fig. 5: Performance of different collaborative aggregation strategies

To quantify the domain gap, we use a well-used foundation model, CLIP [42], to extract the averaged feature embeddings for each dataset. The domain gap g between the embeddings of the datasets e_1 and e_2 is then computed by the norm of the difference between the embeddings as $g = \|e_1 - e_2\|_2$. The total domain gaps in a setting will be the sum of gaps in both modalities.

We choose four settings to evaluate the performance of all methods under different domain gaps: 1) the default setting, where $g_1 = 10.11$, 2) the setting with more text gap by replacing AG NEWS with Medical Abstracts, where $g_2 = 10.84$, 3) the setting with more image gap by replacing CIFAR-100 with OrganCMNIST, where $g_3 = 11.64$, and 4) the setting with more gaps on both modalities, where $g_4 = 12.36$.

We investigate the performance of multi-modal methods compared to the FedAvg baseline relatively under these settings, and the result is shown in Fig. 4. We notice that FedCola consistently outperforms all other methods under all settings, demonstrating the robustness of our proposed method in handling various domain gaps. However, larger domain gaps do have a negative impact on the performance of FedCola, and we leave addressing the challenge of large domain gaps as future work. Similarly, FedIoT struggles the most in the setting with more gaps since its collaboration is based on the capability of multi-modal clients. When the domain gap is large, the multi-modal clients are not able to provide effective information to the uni-modal clients, leading to a performance drop. CreamFL, on the other hand, shows a more stable performance under different domain gaps since their collaboration is based on the public dataset, which is less affected by the domain gap despite being unable to consistently outperform the uni-modal FedAvg baseline.

6 Discussion

6.1 Ablation study

Impact of each module. As shown in Table 2, we observe that when the in-modal collaboration is applied, the performance is improved to 88.70 by the bet-

Table 4: Performance on each test dataset.

Method	CIFAR-100	AG NEWS	Flickr
FedAvg	89.20	87.01	81.08
FedProx	88.79	85.75	78.55
CreamFL	88.81	83.99	74.83
FedIoT	87.68	83.95	85.51
FedCola	89.23	87.07	91.96

Table 5: Scaling-up capability with more uni-modal datasets.

# Datasets	0	1	2
$R@1_{\text{sum}}$	81.08	91.96	93.25

ter generalizability gained from the collaborated self-attention layers. By further addressing the misalignment between layers and modalities, the performance is further improved to 90.09. Finally, with the cross-modal collaboration on the local training, each uni-modal model learns from the other modalities, improving the performance to 91.96.

Designs for complementary local training. For the out-modality model weights, they can be either trainable or frozen during the local training. As shown in Table 3, the trainable weights from the other modalities are more effective than the frozen out-modality weights, indicating that the collaboration is more effective when out-modality weights are updated during the local training. Meanwhile, considering introducing such out-modality weights will increase the downloading cost, so we propose alternative strategies, which are more communication-efficient. We notice that although leveraging the entire transformer blocks is more effective for complementary local training, performing it with the self-attention layers only provides a better trade-off between performance and communication cost.

Designs on collaborative aggregation. During the collaborative aggregation, we choose to perform the collaboration on the self-attention layers between the same modalities (*i.e.*, in-modal collaboration). We also study the alternative strategy for collaborative aggregation. The collaboration can be performed on MLPs or the entire transformer blocks, and the collaboration range can be extended to all models instead of the same modalities as the model architectures are the same (*i.e.*, all collaboration). As shown in Fig. 5, the self-attention layers are the most effective for collaboration, supporting that the self-attention layers carry more general and collaborative knowledge. Meanwhile, the collaboration between the same modalities is more effective than the collaboration between all models, indicating that the cross-modality gap is more significant than the in-modality gap. We leave the exploration of the cross-modal collaboration on the global aggregation as future work.

6.2 Fairness analysis

Unlike centralized training, where all data is owned by one data center aiming at one objective [13], the data in FL are from different clients for different local

tasks. We further study the performance on uni-modal tasks and quantify the contributions of each type of client to the performance gain. Such an analysis provides insights into the fairness of the collaboration and an initial guideline for profit sharing among different clients [29, 45].

Performance on uni-modal tasks. We evaluate the performance of each method on the uni-modal tasks. As shown in Table 4, we observe that FedCola maintains the performance on the uni-modal tasks and improves the multi-modal performance, while the other methods sacrifice the uni-modal performance.

Quantifying the contributions of uni-modal clients. Quantifying the contribution is a significant problem in federated learning regarding fairness. We leverage the Shapley value [52] to investigate the contribution of each type of uni-modal client. The Shapley value is a solution to the problem of fair distribution of the total gains generated by the coalition of players. It is the average marginal contribution of a player to all possible coalitions. In our case, the coalition is the combination of different modality clients. We consider all four possible coalitions of the participation of uni-modal clients: 1) No uni-modal clients, 2) only image clients, 3) only text clients, and 4) both image and text clients. For each coalition, we evaluate the performance on Flickr and assume the probability of each coalition is equal. The Shapley value of image clients and text clients is 4.74 and 6.14, respectively, indicating that the text clients contribute more to the performance gain in the default setting. Consequently, the profit sharing should be more favorable to the text clients.

6.3 Scaling-up capability

To verify the scaling-up capability, we study the trend of the performance on multi-modal evaluation when more uni-modal datasets (*i.e.*, CIFAR-100, AG NEWS, OrganCMNIST, and Medical Abstracts) are involved. As shown in Table 5, the performance increases when more generalized data is leveraged in FL, supporting the scaling-up capability of the proposed framework.

7 Conclusion

In this work, we present FedCola, a novel framework tailored for transfer MFL focussing on transformers. By integrating complementary local training and collaborative aggregation techniques that solely depend on model parameters and dispense the need for public data, FedCola demonstrates improved performance over previous approaches under diverse FL scenarios, including different domain gaps and numbers of participating clients. Our framework proves its effectiveness through extensive numerical experiments and offers new insights into the dynamics of collaboration, including conditions for convergence and the specific contributions of different client types. We envision that our contributions will encourage further investigation into multi-modal federated training paradigms, particularly to train large transformer architectures.

Acknowledgement

This work is partially supported by the NSF/Intel Partnership on MLWiNS under Grant No. 2003198.

References

1. Bao, G., Zhang, Q., Miao, D., Gong, Z., Hu, L.: Multimodal federated learning with missing modality via prototype mask and contrast. arXiv preprint arXiv:2312.13508 (2023)
2. Bao, H., Wang, W., Dong, L., Liu, Q., Mohammed, O.K., Aggarwal, K., Som, S., Piao, S., Wei, F.: Vlmo: Unified vision-language pre-training with mixture-of-modality-experts. *Advances in Neural Information Processing Systems* **35**, 32897–32912 (2022)
3. Bergou, E.H., Burlachenko, K.P., Dutta, A., Richtárik, P.: Personalized federated learning with communication compression. *Transactions on Machine Learning Research* (2023)
4. Che, L., Wang, J., Zhou, Y., Ma, F.: Multimodal federated learning: A survey. *Sensors* **23**(15), 6986 (2023)
5. Chen, H.Y., Tu, C.H., Li, Z., Shen, H.W., Chao, W.L.: On the importance and applicability of pre-training for federated learning. In: *The Eleventh International Conference on Learning Representations* (2023)
6. Chen, J., Zhang, A.: Fedmsplit: Correlation-adaptive federated multi-task learning across multimodal split networks. In: *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. pp. 87–96 (2022)
7. Chen, X., Fang, H., Lin, T.Y., Vedantam, R., Gupta, S., Dollár, P., Zitnick, C.L.: Microsoft coco captions: Data collection and evaluation server. arXiv preprint arXiv:1504.00325 (2015)
8. Cheng, S., Wu, J., Xiao, Y., Liu, Y., Liu, Y.: FedGEMS: Federated learning of larger server models via selective knowledge fusion (2022)
9. Cho, Y.J., Manoel, A., Joshi, G., Sim, R., Dimitriadis, D.: Heterogeneous ensemble knowledge transfer for training large models in federated learning. In: *International Joint Conference on Artificial Intelligence* (2022)
10. Deng, Y., Kamani, M.M., Mahdavi, M.: Adaptive personalized federated learning. arXiv preprint arXiv:2003.13461 (2020)
11. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding pp. 4171–4186 (2019)
12. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *International Conference on Learning Representations* (2020)
13. Dutta, A., Bergou, E.H., Abdelmoniem, A.M., Ho, C.Y., Sahu, A.N., Canini, M., Kalnis, P.: On the discrepancy between the theoretical analysis and practical implementations of compressed communication for distributed deep learning. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 34, pp. 3817–3824 (2020)
14. Feng, T., Bose, D., Zhang, T., Hebbar, R., Ramakrishna, A., Gupta, R., Zhang, M., Avestimehr, S., Narayanan, S.: Fedmultimodal: A benchmark for multimodal federated learning. In: *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. p. 4035–4045 (2023)

15. Gunasekar, S., Zhang, Y., Aneja, J., Mendes, C.C.T., Giorno, A.D., Gopi, S., Javaheripi, M., Kauffmann, P., de Rosa, G., Saarikivi, O., Salim, A., Shah, S., Behl, H.S., Wang, X., Bubeck, S., Eldan, R., Kalai, A.T., Lee, Y.T., Li, Y.: Textbooks are all you need (2023)
16. He, C., Annavaram, M., Avestimehr, S.: Group knowledge transfer: Federated learning of large CNNs at the edge. *Advances in Neural Information Processing Systems* **33**, 14068–14080 (2020)
17. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9729–9738 (2020)
18. Hsu, H., Qi, H., Brown, M.: Measuring the effects of non-identical data distribution for federated visual classification (2019)
19. Huang, H., Zhuang, W., Chen, C., Lyu, L.: Fedmef: Towards memory-efficient federated dynamic pruning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 27548–27557 (June 2024)
20. Jiang, A.Q., Sablayrolles, A., Roux, A., Mensch, A., Savary, B., Bamford, C., Chaplot, D.S., de las Casas, D., Hanna, E.B., Bressand, F., Lengyel, G., Bour, G., Lample, G., Lavaud, L.R., Saulnier, L., Lachaux, M.A., Stock, P., Subramanian, S., Yang, S., Antoniak, S., Scao, T.L., Gervet, T., Lavril, T., Wang, T., Lacroix, T., Sayed, W.E.: Mixtral of experts (2024)
21. Kairouz, P., McMahan, H.B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A.N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al.: Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning* **14**(1–2), 1–210 (2021)
22. Kang, W., Liu, G., Shah, M., Yan, Y.: Segvg: Transferring object bounding box to segmentation for visual grounding (2024)
23. Karimireddy, S.P., Kale, S., Mohri, M., Reddi, S., Stich, S., Suresh, A.T.: Scaffold: Stochastic controlled averaging for federated learning. In: *International conference on machine learning*. pp. 5132–5143. PMLR (2020)
24. Krizhevsky, A.: Learning multiple layers of features from tiny images pp. 32–33 (2009)
25. Li, H., Cai, Z., Wang, J., Tang, J., Ding, W., Lin, C.T., Shi, Y.: Fedtp: Federated learning by transformer personalization. *IEEE Transactions on Neural Networks and Learning Systems* (2023)
26. Li, Q., He, B., Song, D.: Model-contrastive federated learning. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10713–10722 (2021)
27. Li, T., Hu, S., Beirami, A., Smith, V.: Ditto: Fair and robust federated learning through personalization. In: *International Conference on Machine Learning*. pp. 6357–6368 (2021)
28. Li, T., Sahu, A.K., Zaheer, M., Sanjabi, M., Talwalkar, A., Smith, V.: Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems* **2**, 429–450 (2020)
29. Li, T., Sanjabi, M., Beirami, A., Smith, V.: Fair resource allocation in federated learning. In: *International Conference on Learning Representations* (2020)
30. Li, X., Jiang, M., Zhang, X., Kamp, M., Dou, Q.: FedBN: Federated learning on non-IID features via local batch normalization. In: *International Conference on Learning Representations* (2021)
31. Li, Y., Bubeck, S., Eldan, R., Giorno, A.D., Gunasekar, S., Lee, Y.T.: Textbooks are all you need ii: phi-1.5 technical report (2023)

32. Lin, T., Kong, L., Stich, S.U., Jaggi, M.: Ensemble distillation for robust model fusion in federated learning. *Advances in Neural Information Processing Systems* **33**, 2351–2363 (2020)
33. Liu, F., Wu, X., Ge, S., Fan, W., Zou, Y.: Federated learning for vision-and-language grounding problems. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 34, pp. 11572–11579 (2020)
34. Luo, J., Mendieta, M., Chen, C., Wu, S.: Pgfed: Personalize each client’s global objective for federated learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3946–3956 (October 2023)
35. McMahan, H.B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: *Artificial intelligence and statistics*. pp. 1273–1282 (2017)
36. Mendieta, M., Sun, G., Chen, C.: Navigating heterogeneity and privacy in one-shot federated learning with diffusion models (2024)
37. Mendieta, M., Yang, T., Wang, P., Lee, M., Ding, Z., Chen, C.: Local learning matters: Rethinking data heterogeneity in federated learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8397–8406 (2022)
38. Mortaheb, M., Vahapoglu, C., Ulukus, S.: Fedgradnorm: Personalized federated gradient-normalized multi-task learning. In: *2022 IEEE 23rd International Workshop on Signal Processing Advances in Wireless Communication (SPAWC)*. pp. 1–5. IEEE (2022)
39. Nguyen, J., Wang, J., Malik, K., Sanjabi, M., Rabbat, M.: Where to begin? on the impact of pre-training and initialization in federated learning. In: *Workshop on Federated Learning: Recent Advances and New Challenges (in Conjunction with NeurIPS 2022)* (2022)
40. Plummer, B.A., Wang, L., Cervantes, C.M., Caicedo, J.C., Hockenmaier, J., Lazebnik, S.: Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2641–2649 (2015)
41. Qu, L., Zhou, Y., Liang, P.P., Xia, Y., Wang, F., Adeli, E., Fei-Fei, L., Rubin, D.: Rethinking architecture design for tackling data heterogeneity in federated learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10061–10071 (2022)
42. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International conference on machine learning*. pp. 8748–8763 (2021)
43. Schopf, T., Braun, D., Matthes, F.: Evaluating unsupervised text classification: Zero-shot and similarity-based approaches. In: *Proceedings of the 2022 6th International Conference on Natural Language Processing and Information Retrieval*. p. 6–15. NLPir ’22, Association for Computing Machinery (2023)
44. Singh, A., Hu, R., Goswami, V., Couairon, G., Galuba, W., Rohrbach, M., Kiela, D.: Flava: A foundational language and vision alignment model. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 15638–15650 (2022)
45. Song, T., Tong, Y., Wei, S.: Profit allocation for federated learning. In: *2019 IEEE International Conference on Big Data (Big Data)*. pp. 2577–2586. IEEE (2019)
46. Sun, G., Mendieta, M., Luo, J., Wu, S., Chen, C.: Fedperfix: Towards partial model personalization of vision transformers in federated learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4988–4998 (2023)

47. Sun, G., Mendieta, M., Yang, T., Chen, C.: Conquering the communication constraints to enable large pre-trained models in federated learning. *arXiv* (2022)
48. Tan, A.Z., Yu, H., Cui, L., Yang, Q.: Towards personalized federated learning. *IEEE Transactions on Neural Networks and Learning Systems* **34**(12), 9587–9603 (2022)
49. Tan, Y., Long, G., Ma, J., Liu, L., Zhou, T., Jiang, J.: Federated learning from pre-trained models: A contrastive learning approach. *Advances in Neural Information Processing Systems* **35**, 19332–19344 (2022)
50. Wang, H., Yurochkin, M., Sun, Y., Papailiopoulos, D., Khazaeni, Y.: Federated learning with matched averaging. In: *International Conference on Learning Representations* (2020)
51. Wang, W., Bao, H., Dong, L., Bjorck, J., Peng, Z., Liu, Q., Aggarwal, K., Mohammed, O.K., Singhal, S., Som, S., Wei, F.: Image as a foreign language: Beit pretraining for vision and vision-language tasks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 19175–19186 (2023)
52. Winter, E.: The shapley value. *Handbook of game theory with economic applications* **3**, 2025–2054 (2002)
53. Xiong, B., Yang, X., Qi, F., Xu, C.: A unified framework for multi-modal federated learning. *Neurocomputing* **480**, 110–118 (2022)
54. Xu, H., Kostopoulou, K., Dutta, A., Li, X., Ntoulas, A., Kalnis, P.: Deepreduce: A sparse-tensor communication framework for federated deep learning. *Advances in Neural Information Processing Systems* **34**, 21150–21163 (2021)
55. Yang, J., Shi, R., Wei, D., Liu, Z., Zhao, L., Ke, B., Pfister, H., Ni, B.: Medmnist v2- a large-scale lightweight benchmark for 2d and 3d biomedical image classification. *Scientific Data* **10**(1), 41 (2023)
56. Yu, Q., Liu, Y., Wang, Y., Xu, K., Liu, J.: Multimodal federated learning via contrastive representation ensemble. In: *The Eleventh International Conference on Learning Representations* (2022)
57. Zhang, X., Zhao, J., LeCun, Y.: Character-level convolutional networks for text classification. *Advances in neural information processing systems* **28** (2015)
58. Zhang, Y., Ding, X., Gong, K., Ge, Y., Shan, Y., Yue, X.: Multimodal pathway: Improve transformers with irrelevant data from other modalities. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6108–6117 (2024)
59. Zhao, Y., Barnaghi, P., Haddadi, H.: Multimodal federated learning on iot data. In: *2022 IEEE/ACM Seventh International Conference on Internet-of-Things Design and Implementation (IoTDI)*. pp. 43–54 (2022)
60. Zhuang, W., Chen, C., Lyu, L.: When foundation model meets federated learning: Motivations, challenges, and future directions (2024)
61. Zhuang, W., Lyu, L.: Fedwon: Triumphant multi-domain federated learning without normalization. In: *The Twelfth International Conference on Learning Representations* (2024)