

Federated Prompt-Tuning with Heterogeneous and Incomplete Multimodal Client Data

Thu Hang Phung¹ Duong M. Nguyen¹ Thanh Trung Huynh² Quoc Viet Hung Nguyen³
 Trong Nghia Hoang^{4*} Phi Le Nguyen^{1*}
¹ Institute of AI Innovation and Societal Impact, Hanoi University of Science and Technology
² EPFL ³ Griffith University ⁴ Washington State University

Abstract

This paper introduces a generalized federated prompt-tuning framework for practical scenarios where local datasets are multi-modal and exhibit different distributional patterns of missing features at the input level. The proposed framework bridges the gap between federated learning and multi-modal prompt-tuning which have traditionally focused on either uni-modal or centralized data. A key challenge in this setting arises from the lack of semantic alignment between prompt instructions that encode similar distributional patterns of missing data across different clients. To address this, our framework introduces specialized client-tuning and server-aggregation designs that simultaneously optimize, align, and aggregate prompt-tuning instructions across clients and data modalities. This allows prompt instructions to complement one another and be combined effectively. Extensive evaluations on diverse multi-modal benchmark datasets demonstrate that our work consistently outperforms state-of-the-art (SOTA) baselines.

1. Introduction

The recent emergence of exceedingly large models [3] with immense scale and versatility due to being pre-trained on large swaths of generic data has resulted in a paradigm shift in machine learning (ML). Most task-specific models are now fine-tuned versions of those large models instead of being created from de novo learning processes. For example, most state-of-the-art (SOTA) models solving specific NLP tasks are fine-tuned versions of existing large pre-trained or foundation models such as BERT [12], GPT [5], and T5 [40]. Likewise, SOTA models for various computer vision tasks have also been derived from foundation models such as the vision transformer [13] pre-trained on extensive datasets like ImageNet [11]. The standard fine-tuning approach, however, assumes centralized data which is not always possible in practical scenarios where fine-tuning data

are scattered across multiple private siloed systems.

To overcome this siloed data challenge, federated fine-tuning has recently been considered to aggregate the fine-tuned updates across private devices running their own (local) fine-tuning processes [9, 32, 47]. This approach can be further combined with an existing parameter-efficient tuning method called prompt tuning [25]. This method focuses on creating lightweight signals such as additional tokens that are prepended to the input of the pre-trained models [41, 44, 59, 60], providing them valuable contextual information to perform specific tasks more effectively. Local sets of prompts can then be aggregated using the conventional parameter aggregation procedure in FL [16, 29, 47, 50, 58].

However, such federated prompt-tuning techniques are restricted to learning with uni-modal data. In more practical scenarios, advances in sensor and device manufacturing increasingly enable the capture of multiple data modalities (e.g., text, images), leading to richer datasets for fine-tuning [6]. Leveraging such diverse multi-modal information can significantly enhance model performance, as demonstrated by the success of multi-modal foundation models like CLIP [39] and ViLT [21]. This has further motivated research on multi-modal federated learning [8, 33, 37, 52] and fine-tuning pre-trained multi-modal LLMs, i.e., multi-modal federated fine-tuning [56].

Despite these advances, multi-modal federated learning research often adopts specific local architectures which cannot be learned via fine-tuning existing pre-trained multi-modal foundation models [8, 33, 37, 52]. Existing multi-modal federated fine-tuning research on the other hand assumes that all local datasets have the same set of modalities [56]. As such, these approaches are not suitable for situations where different clients might collect data using different types of sensors. Furthermore, the multi-modal sensor data across clients might even have different missing patterns of data entries. This leads to local datasets with both intra- and inter-heterogeneities. Intra-heterogeneity occurs when individual samples within a single dataset

*Corresponding Authors

have missing entries for different subsets of features. Inter-heterogeneity arises when different clients have different distributions over their data-missing patterns.

There also exists recent work in prompt-tuning that learns a specific prompt for each subset of missing modalities [24]. Input samples having the same set of missing modalities are then augmented with the same (learnable) prompt during training, specializing it to complement for the corresponding missing modalities. This approach is, however, restricted to centralized data setting (i.e., no inter-heterogeneities) and does not generalize well to federated settings where different sets of local prompts might be biased towards different distributional patterns of missing data. A naive aggregation of these prompts might therefore collapse them into less informative prompts which are ignorant of the heterogeneities in missing-data patterns across clients, resulting in poor performance (see Fig. 1).

To bridge this gap between multi-modal prompt-tuning and federated learning, we propose a multimodal federated prompt-tuning approach that optimizes and aggregates two distinct sets of (local) prompts for each client: inter-client and intra-client prompts. Inter-client prompts breakdown and encode prototypical distribution patterns of missing input data, which can be aligned and aggregated across clients to represent shared distributional information. In contrast, intra-client prompts encode information tailored to support specific local subsets of missing modalities, independent of the distributional patterns of missing data at each client. This is achieved via the following contributions:

1. We develop a local input-adaptive fine-tuning algorithm which identifies and updates a subset of inter- and intra-client prompts that are most relevant to each individual sample. This is formulated as prompt-retrieval task which focuses on learning mapping functions that push an input towards its relevant prompts on an embedding space while pulling away from irrelevant prompts on an embedding space. This helps decompose and distill local fine-tuning knowledge into different parts of the prompt set, preventing learned instructions for different fine-tuning patterns from being overloaded on the same prompts (Section 3.2).
2. We develop a server-side model for aggregating and aligning inter-client prompts across clients. Both generation and alignment are framed as part of a clustering task, where inter-client prompts encoding similar input-level distribution patterns of missing data are identified and merged into more comprehensive and versatile prompt instructions. The empirical effectiveness of these combined prompts in reducing local training losses is used as a learning feedback to optimize this clustering model (Section 3.3).
3. We evaluate and demonstrate the effectiveness of our proposed approach through extensive experiments on two benchmark datasets, MM-IMDB [1] and UPMC Food-

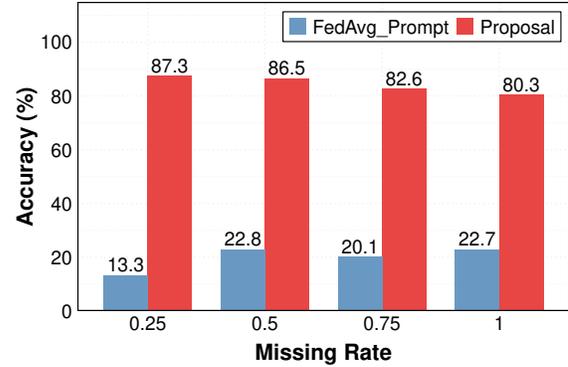


Figure 1. Our approach with prompt alignment outperforms FedAvg prompt-tuning (w/o alignment) across settings with different data missing rates.

101 [46], comparing its performance against a variety of existing multimodal federated learning methods as well as repurposed prompt-tuning baselines for federated settings. The reported results show that our method consistently and significantly outperforms the baselines across various modality-missing scenarios, achieving new SOTA performance for multi-modal federated learning (Section 4).

A bird-view of our proposed framework, Federated Prompt-Tuning with Heterogeneous and Incomplete Multimodal Client Data (FED-PRIME), is also provided in Figure 2.

2. Related Works

Multi-Modal Federated Learning. Federated learning (FL) [17, 31, 53, 54] is a collaborative learning paradigm that enables multiple devices to train a shared model while keeping local data private. This design addresses concerns over privacy, security, and inefficiency in transferring large datasets [45, 57]. While FL has been widely studied [10, 14, 15, 18, 26, 27, 30, 34–36, 42, 43], most existing methods focus on uni-modal data (e.g., image or text). However, with the rapid advancement of mobile phones and IoT devices [6], local data increasingly comes from multiple modalities. Integrating such heterogeneous signals enables richer, more robust representations [19, 48], motivating growing interest in multi-modal FL. Yet many current approaches rely on custom architectures that cannot exploit pre-trained multi-modal models. More recent studies explore federated fine-tuning of pre-trained multi-modal models, but they typically assume each client holds the same set of modalities [56], limiting applicability in realistic heterogeneous environments.

Missing Data in Multi-Modal FL. In real-world settings, local devices collect data from different types of sensors, leading to datasets with varying feature modalities (inter-heterogeneity). Moreover, sensor malfunctions or data corruption can cause missing features, with devices exhibiting distinct missing patterns (intra-heterogeneity).

Existing solutions either treat each modality independently [55, 61] or are restricted to centralized learning [7, 38, 51], both overlooking inter-modal dependencies. Recent FL methods address inter-heterogeneity [8, 37] but not intra-heterogeneity, while approaches like FedInMM [52] and FedMAC [33] tackle intra-client missing data but ignore cross-client alignment, limiting effectiveness. A related line of work, including FedMSplit [8] and CACMRN [49], handles modality-incomplete data under a personalized FL framework, where each client retains its model. This paradigm differs from standard FL, which aims to improve a single shared model, making these methods not directly comparable.

Addressing Missing Data with Prompt-Tuning. Prompt-tuning [41, 44, 59, 60] provides an effective way to adapt pre-trained models by learning prompts that add contextual information to input embeddings. Recent work has extended this idea to missing-modality scenarios by learning a prompt for each subset of missing features [20, 24], so that inputs with similar missing patterns are augmented with suitable prompts. However, these methods are developed for centralized training and do not address the additional heterogeneity introduced in FL, where prompts must be aligned and aggregated across clients. Designing effective alignment and aggregation strategies for such prompts is therefore a key challenge, as demonstrated in our framework.

3. Multi-Modal Federated Prompt-Tuning

Multimodal federated learning faces significant challenges due to missing data modalities, leading to both inter-client and intra-client heterogeneities that complicate knowledge aggregation and degrade overall model performance – see Fig. 1. To address this challenge (Section 3.1), we present a federated prompt-tuning framework based on the pre-trained Vision and Language Transformer (ViLT) model, which comprises a prompt-based client design and a server-side prompt-aggregation algorithm. The client design features a fine-tuning algorithm that learns and decomposes the fine-tuned knowledge into separate sets of inter-client and intra-client prompts specialized for fine-tuning at different input-level patterns of missing data and input-agnostic patterns of missing modalities (Section 3.2). The server-side algorithm then learns alignment between inter-client prompts specializing for similar input-level patterns of missing data and aggregates them to create more versatile prompts. This is also combined with a standard FEDAVG-based aggregation of intra-prompts specializing for input-agnostic patterns of missing modalities (Section 3.3) – see Fig. 2.

3.1. Problem Formulation

In multimodal federated learning (MMFL), we assume there are n clients and r data modalities (e.g., text, image).

Each client $t \in [n]^1$ has access to a private dataset $D_t \triangleq \{\mathbf{x}(M_{t,s}), z_{t,s}\}_{s=1}^m$, where $\mathbf{x}(M_{t,s})$ denote the observed features for the set of modalities, $M_{t,s} \subseteq [r]$ recorded for the s -th data point of the t -th client, and $z_{t,s}$ denotes the corresponding target label. The goal is to learn a multimodal model \mathbf{w}_* using all the private datasets $\{D_t\}_{t=1}^n$ without centralizing them. That is,

$$\mathbf{w}_* = \arg \min_{\mathbf{w}} \left\{ L(\mathbf{w}) \triangleq \frac{1}{n} \sum_{t=1}^n L_t(\mathbf{w}) \right\} \text{ with} \quad (1)$$

$$L_t(\mathbf{w}) \triangleq \frac{1}{n} \sum_{s=1}^m \ell \left(F(\mathbf{x}(M_{t,s}); \mathbf{w}), z_{t,s} \right)$$

where $F(\mathbf{x}(M_{t,s}); \mathbf{w})$ denotes the multimodal prediction model with parameter \mathbf{w} and $\ell(F(\mathbf{x}(M_{t,s}); \mathbf{w}), z_{t,s})$ denotes the individual loss of \mathbf{w} on the s -th data point. Previous work [8, 33, 37, 52] in MMFL has largely focused on designing $F(\mathbf{x}(M_{t,s}); \mathbf{w})$ without leveraging existing knowledge in pre-trained foundation model. Instead, we let $F(\mathbf{x}(M_{t,s}); \mathbf{w})$ be a prompt-tuned version of a pre-trained transformer-based model,

$$F(\mathbf{x}(M); \mathbf{w}) \triangleq F_c \left(F_p \left(F_e(\mathbf{x}(M)) \circ \mathbf{w}_p \right); \mathbf{w}_c \right) \quad (2)$$

where $F_c(\cdot; \mathbf{w}_c)$ is the prediction head parameterized with \mathbf{w}_c , $F_e(\mathbf{x}(M)) = F_e(\mathbf{x}_1) \circ \dots \circ F_e(\mathbf{x}_{|M|})$ where $F_e(\mathbf{x}_i)$ denote the tokenized input sequence of each unimodal feature \mathbf{x}_i in $\mathbf{x}(M) \triangleq (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{|M|})$, $F_p(\cdot)$ denotes the frozen multimodal pre-trained model – e.g., ViLT [21] – and \mathbf{w}_p denotes the set of (learnable) prompts. Here, \circ denotes the set concatenation operator and $\mathbf{w} = (\mathbf{w}_c, \mathbf{w}_p)$ denotes the full set of tuning parameters. Due to the decentralized and private nature of the local datasets $\{D_t\}_{t=1}^n$, solving Eq. (1) directly is not possible. Instead, following the federated learning [31] paradigm, its optimal solution \mathbf{w}_* is approximated via a number of iterations over two interleaving client- and server-side steps:

$$\mathbf{w}_t^{(h+1)} = \text{update} \left(L_t, \mathbf{w}_*^{(h)} \right) \forall t \in [n], \quad (3)$$

$$\mathbf{w}_*^{(h+1)} = \text{aggregate} \left(\mathbf{w}_1^{(h+1)}, \dots, \mathbf{w}_n^{(h+1)} \right), \quad (4)$$

where $\mathbf{w}_t^{(h)}$ denote the estimation of the local tuning parameter \mathbf{w}_t of client t after h iterations. For example, the update step can be a vanilla gradient descent (GD) optimization of the prompt ω and the aggregate step denotes the average of the locally optimized prompts. Such vanilla approach however does not account for both the inter- and intra-heterogeneities among local prompts (see Section 1) which ultimately risks aggregating incompatible prompts and degrade the overall performance (see Section 4). To avoid this, we will introduce novel designs for client- and server-prompts as detailed in the followings.

¹ $[n] \triangleq \{1, 2, \dots, n\}$.

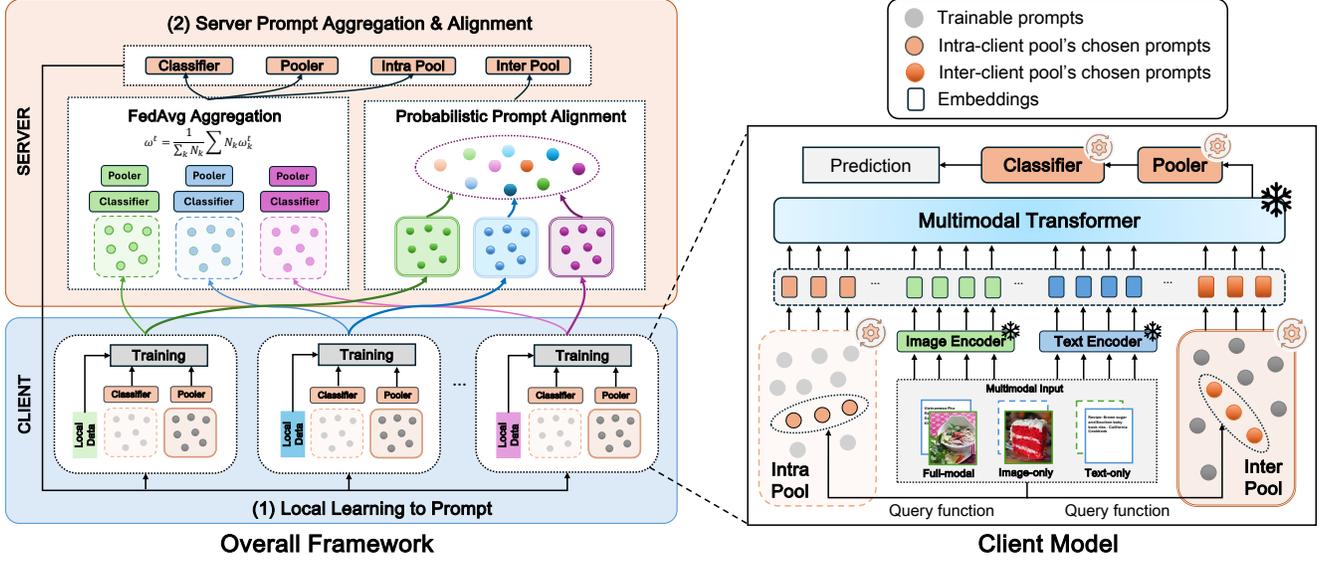


Figure 2. Overview of the proposed multi-modal federated prompt-tuning framework – see Alg. 3. Each client maintains two (learnable) sets of intra- and inter-client prompts. At the beginning of each iteration, each client performs local training via Eq. (7) and Eq. (8). Local sets of intra- and inter-client prompts are then sent to the server for aggregation – see Eq. (10)–Eq. (14).

3.2. Client Design

To enable effective prompt-aggregation across heterogeneous clients, we need to separate the fine-tuning treatment for input-level patterns of missing data and input-agnostic patterns of missing modalities. This can be achieved via further parameterizing $\mathbf{w}_p = (\mathbf{w}_p^{inter}, \mathbf{w}_p^{intra})$ with two distinct subsets of τ inter- and intra-client prompts:

$$\mathbf{w}_p^{inter} \triangleq \{\mathbf{p}_1^{inter}, \mathbf{p}_2^{inter}, \dots, \mathbf{p}_\tau^{inter}\}, \quad (5)$$

$$\mathbf{w}_p^{intra} \triangleq \{\mathbf{p}_1^{intra}, \mathbf{p}_2^{intra}, \dots, \mathbf{p}_\tau^{intra}\}, \quad (6)$$

which are learned separately. The learned inter-client prompts \mathbf{w}_p^{inter} across clients will be sent to the server which clusters them into different groups based on their specialization for similar input-level patterns of missing data – Section 3.3. Local sets of intra-client prompts \mathbf{w}_p^{intra} can be averaged via FEDAVG [31] since their specialization is for input-agnostic patterns of missing modalities. Intuitively, the aggregation mechanisms will determine how different aspects of fine-tuning knowledge are encoded across the prompt sets. For example, if knowledge regarding different input-level patterns of missing data is (incorrectly) incorporated into $\{\mathbf{p}_i^{intra}\}_{i=1}^\tau$, they will be averaged out and collapsed into less informative prompts due to the non-clustering aggregation mechanism of $\{\mathbf{p}_i^{intra}\}_{i=1}^\tau$.

As this potential loss of information will decrease the overall performance, the (learnable) prompt parameters will be updated to prevent this from happening. Otherwise, if (common) knowledge patterns regarding input-agnostic patterns of missing modalities are (incorrectly) incorporated

into $\{\mathbf{p}_i^{inter}\}_{i=1}^\tau$, their modeling bandwidth for inter-client heterogeneities will be reduced, which will also hamper the performance. This will, in turn, generate a prompt update signal to avoid this as well. Furthermore, within each set of inter- and intra-client prompts, there might also be different patterns that need to be separated and internalized into different subset of prompts.

To enable this, we develop an input-adaptive learning mechanism for both sets of prompts which identifies and updates the most relevant subset of prompts for each (multimodal) input sample. This is formulated as an information-retrieval task, learning a pair of key and query functions, $k(\mathbf{p})$ and $q(\mathbf{x}(M))$, which transform the prompt \mathbf{p} and input \mathbf{x} with the set M observed modalities, respectively, into a shared metric space. We can then define a geometric input-prompt distance, e.g., $d(\mathbf{x}(M), \mathbf{p}) = \cos(q(\mathbf{x}(M)), k(\mathbf{p}))$ which can be used to measure their contextual relevance. To avoid loading diverse fine-tuning knowledge into a single prompt, we regularize the (local) loss function to penalize it when a less relevant prompt is selected as fine-tuning instruction for an input $\mathbf{x}(M)$:

$$\mathbf{w}_t \triangleq \arg \min_{\mathbf{w}} L'_t(\mathbf{w}) \quad \text{where} \quad (7)$$

$$L'_t(\mathbf{w}) \triangleq \sum_{s=1}^m \left\{ \ell(F(\mathbf{x}(M_{t,s}); \mathbf{w}'), z_{t,s}) \right\} + \sum_{s=1}^m \left\{ r(\mathbf{x}(M_{t,s}), \mathbf{w}'_p) \right\}, \quad (8)$$

where $\ell(\mathbf{w}')$ is defined in Eq. (1) and $\mathbf{w}' = (\mathbf{w}_c, \mathbf{w}'_p)$ denote the customized set of tuning instructions for each

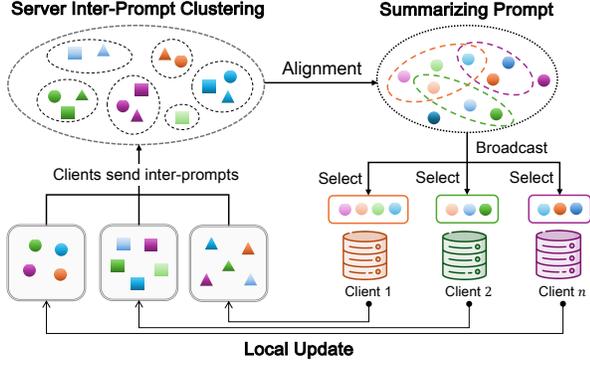


Figure 3. Workflow of the prompt-alignment algorithm. At each iteration, each client samples a subset of summarizing prompts using its query and key functions. The client then performs a local update, resulting in heterogeneous prompt sets, which are subsequently sent to the server to be clustered and summarized into new summarizing prompt sets for the next iteration.

input $\mathbf{x}(M_{t,s})$ where $\mathbf{w}'_p = (\mathbf{w}_p^{inter'}, \mathbf{w}_p^{intra'})$ such that $\mathbf{w}_p^{inter'} \subseteq \mathbf{w}_p^{inter}$ and $\mathbf{w}_p^{intra'} \subseteq \mathbf{w}_p^{intra}$ are the sets of κ most relevant prompts to $\mathbf{x}(M_{t,s})$ in \mathbf{w}_p^{inter} and \mathbf{w}_p^{intra} . The regularizer $r(\mathbf{x}(M), \mathbf{w}'_p)$ is then defined as

$$r(\mathbf{x}(M), \mathbf{w}'_p) \triangleq \sum_{\mathbf{p} \in \mathbf{w}'_p} \left\{ d(\mathbf{x}(M), \mathbf{p}) \right\}. \quad (9)$$

Intuitively, the regularizer penalizes the total distance between the (multimodal) input $\mathbf{x}(M)$ and the κ closest prompts according to the aforementioned distance metric $d(\mathbf{x}(M), \mathbf{p})$. Minimizing both the original local loss and this regularizer allows the local model to simultaneously learn to (1) internalize fine-tuning patterns of inter- and intra-heterogeneities to minimize the original loss $L_t(\mathbf{w})$; and (2) ensure that selected prompts for a given input will be close to it to reduce the penalty $r(\mathbf{x}(M), \mathbf{w}'_p)$. A pseudocode of this client design is detailed in Suppl. A.

3.3. Server Aggregation

We will now establish a principled algorithm to align and aggregate inter-client prompts specializing in similar input-level patterns of missing data². This is not straight-forward since inter-client prompts at the same position in the prompt sets across different clients might not specialize for similar patterns of missing data due to client heterogeneities: certain missing-data patterns are not observed at certain clients.

To find the correct alignment, we formulate it as a clustering optimization task. Suppose each client t communicates a set of inter-client prompts $\mathbf{w}_t^{inter} = (\mathbf{p}_t^i)_{i=1}^\tau$. Each

²This section only describes the aggregation algorithm for inter-client prompts since intra-client prompts are specialized for input-agnostic patterns of missing modalities and can be aggregated via FEDAVG [31].

potential clustering of these prompts thus defines a candidate alignment, i.e., prompts in the same cluster are considered aligned and can be aggregated. A (learnable) cluster center can then be associated with the aggregated prompt – see Fig. 3. Both the cluster assignment and centers can be found via solving the following optimization task with discrete alignment variables $\alpha_t^{pq} \in \{0, 1\}$ and continuous center variables θ_q for each cluster $q \in [n \times \tau]$:

$$\min_{\alpha, \theta, \gamma} G(\alpha, \theta, \gamma) \triangleq \sum_{t=1}^n \sum_{p=1}^\tau \sum_{q=1}^{n \times \tau} \alpha_t^{pq} \cdot c(\mathbf{p}_t^p, \theta_q; \gamma) \quad (10)$$

$$\text{s.t. } \alpha_t^{p,1} + \dots + \alpha_t^{p,n \times \tau} = 1, \quad (11)$$

where $\alpha \triangleq (\alpha_t^{p,q})$ and $\alpha_t^{p,q} \in \{0, 1\} \forall (t, p, q)$ indicates whether the p -th inter-client prompt of client t is matched to the q -th cluster. The corresponding cluster center is then optimized to have minimum total distance to all cluster members according to a (learnable) cost function $c(\mathbf{p}_t^p, \theta_q)$. The optimized cluster center can then be used as the aggregated or summarizing prompt. Minimizing Eq. (10) thus allows us to align and summarize/aggregate the inter-client prompts across clients simultaneously. Note that Eq. (11) enforces a constraint that intuitively indicates two prompts from the same clients should not be matched to the same cluster and get aggregated. This is because local prompts from the same client have already been learned to encode different fine-tuning aspects. Aggregating them will therefore collapse them into less informative ones.

In addition, some summarizing/aggregated prompts might be more generic than others and can be used to provide fine-tuning instructions to a wider range of (multimodal) inputs. The update of such prompts should be prioritized during training and to achieve that, we need a mechanism to capture this *prompt popularity*. To achieve this, we augment the optimization task in Eq. (10) above with a regularizer that penalizes alignment to *less popular* prompts,

$$R(\alpha, \zeta) \triangleq \sum_{t=1}^n \sum_{p=1}^\tau \sum_{q=1}^{n \times \tau} \left(\alpha_t^{p,q} \cdot \log U(\theta_q; \zeta) \right) \quad (12)$$

$$+ \sum_{t=1}^n \sum_{p=1}^\tau \sum_{q=1}^{n \times \tau} \left(1 - \alpha_t^{p,q} \right) \cdot \log \left(1 - U(\theta_q; \zeta) \right),$$

where $U(\theta_q; \zeta)$ is the popularity function of θ_q , which is parameterized by learnable parameter ζ . For example, $U(\theta_q; \zeta) = \sigma(g(\theta_q; \zeta)) \in (0, 1)$ ³ where $g(\theta_q; \zeta)$ is a deep neural net parameterized with ζ . Putting Eq. (13) and Eq. (10) together, we arrive at our final loss function,

$$\text{minimize}_{\alpha, \theta, \gamma, \zeta} \left\{ G(\alpha, \theta, \gamma) + R(\alpha, \zeta) \right\}, \quad (13)$$

³ $\sigma(\cdot)$ denotes the sigmoid function.

which can be optimized via alternating between (1) optimizing (θ, ζ, γ) while fixing α ; and (2) optimizing α while fixing (θ, ζ, γ) . The first optimization sub-task is straight-forward while the latter is less trivial due to the constraint as well as the discrete nature of the optimizing variables. To sidestep the technical challenge in solving for α while fixing (θ, ζ, γ) , we approximate it by optimizing for $\alpha_t = \{\alpha_t^{p,q}\}_{p,q}$ while also fixing $\alpha_{-t} = \{\alpha_{-t}^{p,q}\}_{p,q}$. This results in a minimum weighted linear sum task,

$$\min_{\alpha_t} L(\alpha_t) \triangleq \sum_{p=1}^{\tau} \sum_{q=1}^{n \times \tau} \left(\alpha_t^{p,q} \cdot v(p, q) \right) \quad \text{where} \quad (14)$$

$$v(p, q) \triangleq \left(\log \left\{ \frac{U(\theta_q; \zeta)}{1 - U(\theta_q; \zeta)} \right\} + c(\mathbf{p}_t^p, \theta_q; \gamma) \right),$$

which can be solved effectively in polynomial time using the Hungarian algorithm [23]. An overview and pseudocode of the server-side aggregation algorithm are presented in Fig. 3 and Suppl. A. A thorough evaluation of this framework is detailed next in Section 4.

4. Empirical Evaluation

This section provides a detailed empirical evaluation of our proposed framework, FED-PRIME.⁴ Our evaluation is conducted on two datasets UPMC Food-101 [46] and MM-IMDB [1] and several baselines adapting existing works in federated fine-tuning (for unimodal data) and multi-modal prompt-tuning (for centralized data) to our novel setting on decentralized multi-modal data with missing modalities and features. Detailed description of the datasets and baselines are provided in Section 4.1. Our main results and ablation studies are reported in Section 4.2.

4.1. Experiment Settings

A. Datasets. Following the experiment protocol in [24], we evaluate FED-PRIME on two multimodal datasets:

1. UPMC Food-101 [46] is an image-text classification dataset consisting of noisy image-text pairs collected from Google Image Search. Its classes are aligned with those in the larger, publicly available ETHZ Food-101 dataset [4]. For the federated setting, we select only eight most frequent classes from the original UPMC Food-101.

2. MM-IMDB [1] is a movie genre classification dataset with both image and text modalities. Since each movie can belong to multiple genres, the task is originally a multi-label classification problem, where genres are predicted using image, text, or both. For the federated setting, we use only single-label instances and select eight most frequent classes only from the original MM-IMDB.

⁴Our implementation is available at <https://github.com/hangpt01/FedPrime>.

Each dataset is first partitioned into train and test sets with the 80:20 ratio. The (heterogeneous) missing patterns on the train/test data are simulated using the following protocols.

B. Training Data Simulation. Following [24], we simulate missing data phenomenon within a dataset using a missing rate $\eta \in (0, 1)$ as detailed in 3 simulation cases below:

1. Missing Image. We drop the image features in η percent of the data and keep the remaining $(1 - \eta)$ percent complete.

2. Missing Text. We drop the text features in η percent of the data and keep the remaining $(1 - \eta)$ percent complete.

3. Missing Both. We drop the image component in $\eta/2$ percent of the data, and the text component in the other $\eta/2$ percent, leaving the remaining $(1 - \eta)$ percent complete.

The missing-simulated train set is then partitioned into n subsets for n clients via random sampling. The resulting client datasets will therefore have heterogeneous distributions over missing-data patterns.

C. Test Data Simulation. There are 5 test scenarios, which include: **(1) Full Modal** (no missing data); **(2) Image Only** (drop all text features); **(3) Text Only** (drop all image features); **(4) Miss Both** (run **Missing Both** on test data); and **(5) ~Train** (same simulation as the training data).

D. Metrics. Following prior protocol in [24], we use F1-macro and classification accuracy to evaluate the following baseline methods on MM-IMDB and UPMC Food-101.

E. Baseline Methods. Following [24], we adopt the pre-trained multimodal transformer ViLT [21] as the backbone of our client design. Each client model is a fine-tuned version of this backbone which is generated based on its own local dataset. As no prior work has addressed modality missing in multimodal fine-tuning within federated learning settings, our comparison baselines mostly feature adapted versions of existing work in federated learning and multi-modal prompt-tuning. This includes adaptations of the latest missing prompt-tuning SOTA method in [24] into the context of the classical FedAvg method as well as that of a more recent FL approach, FedMSplit [8]. These are referred to as FEDAVG-P and FEDMSPLIT-P, respectively. In addition, to underscore the importance of separating fine-tuning knowledge across intra- and inter-client prompt sets, we also include two simplified variants of FED-PRIME utilizing either inter-client prompt pooling or intra-client prompt pooling, along with their respective aggregations. We refer to these variants as FED-INTER and FED-INTRA, respectively (see more details in Supplementary B).

4.2. Experimental Results

This section provide extensive performance comparison between FED-PRIME and other baseline methods across various settings of missing data in the train and test sets, as previously stated in Section 4.1. In addition, we also provide

Table 1. Performance achieved by the baselines on the Food-101 (left) and MM-IMDB (right) datasets under various missing-data settings in the train and test sets. This includes all combinations of the 3 train scenarios and 5 test scenarios (i.e., 15 experiments) per dataset, as previously stated in Section 4.1. The best and second-best results are highlighted in **bold red**, and **blue**, respectively. Note that in the **Miss Both** training scenarios, **Test (\sim Train)** and **Test (Miss Both)** represent the same test scenario according to the definition of \sim Train. The reported result for **Test (Miss Both)** in this case is the same as that for **Test (\sim Train)**, as indicated by a dash (-).

Datasets		UPMC Food-101					MM-IMDB				
Train	Method	Test (\sim Train)	Test (Miss Both)	Test (Full Modal)	Test (Text only)	Test (Image only)	Test (\sim Train)	Test (Miss Both)	Test (Full Modal)	Test (Text only)	Test (Image only)
Miss Text	FEDAVG-P	15.71 \pm 2.27	14.90 \pm 1.57	21.56 \pm 7.81	16.91 \pm 0.69	15.36 \pm 0.43	22.42 \pm 2.27	21.89 \pm 1.34	30.78 \pm 1.65	18.40 \pm 0.06	14.53 \pm 4.56
	FEDMSPLIT-P	15.62 \pm 1.51	17.50 \pm 1.97	25.27 \pm 8.22	18.78 \pm 0.64	17.50 \pm 1.97	21.02 \pm 1.89	19.97 \pm 0.74	24.39 \pm 6.08	14.38 \pm 2.76	18.09 \pm 6.13
	FED-INTER	54.82 \pm 19.01	48.87 \pm 24.64	59.17 \pm 27.06	35.13 \pm 26.78	56.59 \pm 15.12	18.25 \pm 3.50	16.95 \pm 3.57	18.67 \pm 7.63	15.03 \pm 4.66	18.01 \pm 1.95
	FED-INTRA	61.71 \pm 17.22	48.09 \pm 19.12	62.06 \pm 26.98	22.51 \pm 5.92	62.64 \pm 11.83	13.38 \pm 1.73	12.77 \pm 0.85	12.55 \pm 1.67	11.31 \pm 0.38	14.33 \pm 1.80
	FED-PRIME	78.88 \pm 0.90	80.38 \pm 0.65	92.12 \pm 0.40	73.01 \pm 4.25	76.83 \pm 1.22	31.92 \pm 0.20	31.48 \pm 0.30	37.67 \pm 0.04	30.60 \pm 1.41	30.69 \pm 1.41
	Improv. (%)	27.82 \uparrow	64.48 \uparrow	48.44 \uparrow	107.83 \uparrow	22.65 \uparrow	42.37 \uparrow	43.81 \uparrow	22.35 \uparrow	66.30 \uparrow	69.65 \uparrow
Miss Image	FEDAVG-P	17.35 \pm 4.77	15.12 \pm 1.48	16.84 \pm 2.37	18.12 \pm 6.49	14.81 \pm 0.24	27.69 \pm 5.97	22.55 \pm 3.06	31.94 \pm 0.98	23.76 \pm 11.72	12.29 \pm 0.47
	FEDMSPLIT-P	74.16 \pm 10.56	48.88 \pm 10.26	45.64 \pm 32.43	88.65 \pm 2.17	14.81 \pm 0.90	19.11 \pm 11.33	16.61 \pm 7.22	18.19 \pm 12.55	18.12 \pm 12.30	12.81 \pm 1.25
	FED-INTER	77.96 \pm 11.62	64.62 \pm 10.22	82.08 \pm 7.75	77.69 \pm 12.35	37.56 \pm 6.49	18.79 \pm 5.23	17.93 \pm 3.60	20.56 \pm 3.56	17.67 \pm 6.79	15.47 \pm 2.51
	FED-INTRA	22.84 \pm 3.52	20.13 \pm 1.72	23.48 \pm 1.86	24.46 \pm 2.99	16.66 \pm 1.32	15.75 \pm 4.34	14.06 \pm 3.16	15.68 \pm 4.77	14.53 \pm 3.65	11.71 \pm 0.42
	FED-PRIME	90.55 \pm 0.22	79.12 \pm 0.49	92.89 \pm 0.21	90.18 \pm 0.29	54.14 \pm 2.50	36.08 \pm 0.35	31.35 \pm 0.61	38.49 \pm 0.56	36.91 \pm 0.59	18.15 \pm 0.66
	Improv. (%)	16.15 \uparrow	22.44 \uparrow	13.17 \uparrow	1.73 \uparrow	44.14 \uparrow	30.30 \uparrow	39.02 \uparrow	20.51 \uparrow	55.35 \uparrow	17.32 \uparrow
Miss Both	FEDAVG-P	14.57 \pm 1.50	-	17.17 \pm 4.37	16.40 \pm 4.05	13.24 \pm 0.32	26.45 \pm 2.63	-	33.03 \pm 2.56	24.12 \pm 11.30	20.21 \pm 1.98
	FEDMSPLIT-P	49.15 \pm 24.76	-	64.78 \pm 36.62	64.62 \pm 36.51	21.49 \pm 7.19	24.25 \pm 5.02	-	26.05 \pm 11.17	26.02 \pm 9.64	19.79 \pm 6.20
	FED-INTER	56.32 \pm 21.77	-	69.57 \pm 19.41	45.15 \pm 34.09	59.30 \pm 10.84	26.53 \pm 0.90	-	31.97 \pm 2.22	29.69 \pm 2.21	21.63 \pm 0.77
	FED-INTRA	49.28 \pm 32.87	-	56.70 \pm 37.90	43.24 \pm 34.19	49.85 \pm 25.44	11.90 \pm 0.37	-	12.47 \pm 0.45	11.46 \pm 0.33	12.83 \pm 0.92
	FED-PRIME	84.44 \pm 2.65	-	93.64 \pm 0.58	87.95 \pm 0.91	72.41 \pm 3.88	32.01 \pm 2.51	-	38.68 \pm 0.65	31.00 \pm 2.97	26.01 \pm 0.12
	Improv. (%)	49.93 \uparrow	-	34.60 \uparrow	36.10 \uparrow	22.11 \uparrow	20.66 \uparrow	-	17.11 \uparrow	4.41 \uparrow	20.25 \uparrow

(*) **Improv.** shows the relative performance improvement between our proposal and the second-best. (in percentage).

ablation studies on (1) performance robustness against various missing rate η ; (2) the ablated impact of having inter- and intra-client prompt sets; and (3) training convergence.

4.2.1. Main Results

We compare the performance of FED-PRIME against the other baselines on an extensive set of experiments across all datasets. This includes all combinations of the 3 missing simulation scenarios on train data and 5 test scenarios (with both missing and no missing data) as previously detailed in parts **B** and **C** of Section 4.1. All results are reported in Table 1 which show that FED-PRIME consistently outperforms all baselines in all train/test scenarios with various missing-data patterns. The performance on the UPMC Food-101 and MM-IMDB datasets is measured in terms of classification accuracy and F1-macro, respectively. The F1-macro performance on MM-IMDB is relatively low across all settings due to the class imbalance nature of the dataset. In each experiment, the relative performance improvement of FED-PRIME over the second best baseline is reported, which ranges between 1.73% and 107.83% for the UPMC Food-101 dataset and between 4.41% to 69.65% for the MM-IMDB dataset. Notably, FED-PRIME significantly outperforms both of its simplified variants FED-INTRA and FED-INTER, which either handle intra- or inter-heterogeneities, in all settings. This highlights the remarkable synergies and essential of handling both types of heterogeneities in multi-modal federated learning with missing data. In addition, it can also be observed that in more than half of the experiment settings, the second best base-

line is either FED-INTRA or FED-INTER which significantly outperform both FEDAVG-P and FEDMSPLIT-P in those cases. These observations complement our findings of the isolated impact of these components in our ablation studies in Section 4.2.4.

4.2.2. Impact of Missing Modalities.

Across different testing scenarios, we observe that the highest test performance is always achieved when there is no missing data in the test set as expected. Another interesting observation is that in the **Miss Text** training scenario, where approximately 70% of the text features are lost on the training data, the proposed model still performs surprisingly well on the **Text Only** test scenarios. This suggests that the multi-modal prompt alignment is effective in recovering text information, highlighting the proposed design’s adaptability. However, in the **Miss Image** training scenario, the test performance on the **Image Only** case suffers instead a significant performance drop. Upon investigation, we believe this is likely due to the VILT backbone’s lightweight image processing module, which was optimized to work in tandem with text for competitive performance, hinting at a potential text-centered nature of VILT during pre-training. This is in fact well-aligned with existing observations that co-training with text often leads to improved performance.

4.2.3. Robustness Against Various Missing Rates.

To demonstrate the robustness of FED-PRIME against different missing rates η , we run multiple comparisons of FED-PRIME, FEDAVG-P, and its centralized variant [24]

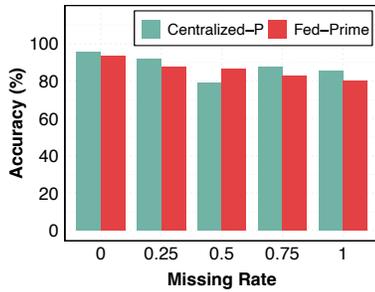


Figure 4. Performance comparison on UPMC Food-101 under **Miss Both** scenarios with various missing rates.

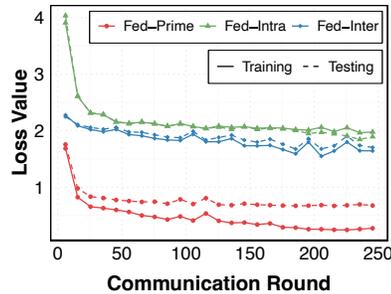


Figure 5. Plots of the train/test loss convergence on Food-101 under **Miss-Text**. FED-PRIME converges fastest.

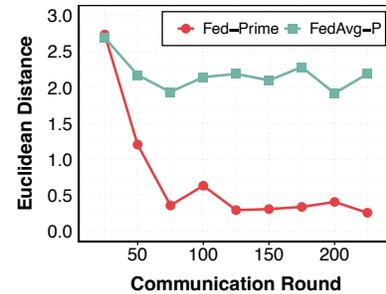


Figure 6. Convergence of Euclidean distance between aggregated prompt centroids on the UPMC Food-101 under **Miss Image** scenario.

across a wide range of $\eta \in \{0.00, 0.25, 0.50, 0.75, 1.00\}$ under the **Miss Both** training scenarios. The reported results in Fig. 1 and Fig. 4 show that our model significantly outperforms FEDAVG-P and its centralized counterpart Centralized-P. As the missing rate increases, our approach consistently maintains a competitive accuracy above 80%, closely approaching the upper-bound performance of FEDAVG-P in the centralized setting, which benefits from a centralized dataset with better-aligned modalities. These results demonstrate the robustness of our model across varying missing rates and highlight the effectiveness of our federated prompt-tuning approach, achieving performance comparable to a centralized prompt-tuning model despite having access only to decentralized local datasets with heterogeneous missing patterns.

4.2.4. Ablation Study

As explained previously in Section 4.1, FED-INTER and FED-INTRA are in variants of FED-PRIME when we drop the intra- and inter-client prompt sets from the client design, respectively. Table 1 shows that FED-PRIME consistently outperforms both FED-INTER and FED-INTRA in all test conditions with substantial gap. The lack of intra-client prompts causes FED-INTER to underperform FED-PRIME suggests that having a dedicated prompt set to capture the common and highly generalizable input-agnostic aspect of fine-tuning is essential to improve performance. Likewise, FED-INTRA, which omits the inter-client prompt set, also shows a significant performance drop. This highlights the importance of capturing and correctly aggregating input-level patterns of missing data through prompt alignment.

4.2.5. Convergence Analysis

Model Convergence. Fig. 5 shows that FED-PRIME has superior performance, achieving faster convergence and significantly lower final loss values in both training and testing compared to FED-INTER and FED-INTRA. This indicates better learning efficiency and robustness. Both FED-INTER and FED-INTRA converge more slowly, with

consistently higher loss values, especially in the early rounds. Additionally, our method is more stable across server-aggregation rounds, with minimal fluctuations. In contrast, FED-INTER and FED-INTRA exhibit significantly more variability, particularly toward the end. This highlights the effectiveness of FED-PRIME, with faster convergence and improved stability across rounds.

Prompt Convergence. Fig. 6 shows that FED-PRIME reduces the distance between the aggregated prompt centroids much faster than FEDAVG-P, which highlights the faster convergence rate of FED-PRIME over FEDAVG-P, indicating faster convergence and better prompt alignment. In contrast, FEDAVG-P converges more slowly and unstably, likely due to incorrect default prompt alignment. This again underscores the necessity of having a learnable prompt clustering mechanism to prevent incorrect prompt alignment.

5. Conclusion

Fine-tuning large pre-trained models for specific tasks has become an increasingly effective solution paradigm in machine learning. However, existing fine-tuning frameworks struggle in practical, multimodal settings where data is distributed across private silos with heterogeneous patterns of missing modalities and features. To address these challenges, this paper introduces a novel federated prompt-tuning approach that optimizes, aligns, and aggregates prompt instructions, compensating for gaps in data modalities and distributional patterns across a number of diverse and private local datasets. Rigorous evaluation on multimodal benchmark datasets demonstrates the effectiveness and robustness of our approach against various patterns of missing data in local datasets. We view this work as an initial step toward developing resilient, adaptable federated models that can fully leverage the diversity of multimodal data in real-world environments.

Acknowledgement

This research is funded by Hanoi University of Science and Technology (HUST) under grant number T2024-TD-002. Additionally, this work utilized GPU compute resource at SDSC and ACES through allocation CIS230391 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services and Support (ACCESS) program [2], which is supported by U.S. National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296.

References

- [1] John Arevalo, Thamar Solorio, Manuel Montes y Gómez, and Fabio A. González. Gated multimodal units for information fusion. *ArXiv*, abs/1702.01992, 2017. 2, 6, 3
- [2] Timothy J. Boerner, Stephen Deems, Thomas R. Furlani, Shelley L. Knuth, and John Towns. Access: Advancing innovation: Nsf’s advanced cyberinfrastructure coordination ecosystem: Services & support. In *Practice and Experience in Advanced Research Computing 2023: Computing for the Common Good*, page 173–176, New York, NY, USA, 2023. Association for Computing Machinery. 9
- [3] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, S. Buch, Dallas Card, Rodrigo Castellon, Niladri S. Chatterji, Annie S. Chen, Kathleen A. Creel, Jared Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren E. Gillespie, Karan Goel, Noah D. Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas F. Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshthe Khani, O. Khattab, Pang Wei Koh, Mark S. Krass, Ranjay Krishna, Rohith Kudithipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir P. Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avanika Narayan, Deepak Narayanan, Benjamin Newman, Allen Nie, Juan Carlos Nieves, Hamed Nilforoshan, J. F. Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Robert Reich, Hongyu Ren, Frieda Rong, Yusuf H. Roohani, Camilo Ruiz, Jack Ryan, Christopher R’e, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishna Parasuram Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei A. Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. On the opportunities and risks of foundation models. *ArXiv*, 2021. 1
- [4] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 – mining discriminative components with random forests. In *European Conference on Computer Vision*, 2014. 6
- [5] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *CoRR*, abs/2005.14165, 2020. 1
- [6] Alberto Brunete, Ernesto Gambao, Miguel Hernando, and Raquel Cedazo. Smart assistive architecture for the integration of iot devices, robotic systems, and multimodal interfaces in healthcare environments. *Sensors*, 21(6):2212, 2021. 1, 2
- [7] Jiayi Chen and Aidong Zhang. Hgmf: heterogeneous graph-based fusion for multimodal data with incompleteness. In *Proceedings of the 26th ACM International Conference on Knowledge Discovery & Data Mining*, pages 1295–1305, 2020. 3
- [8] Jiayi Chen and Aidong Zhang. 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*, pages 87–96, 2022. 1, 3, 6
- [9] Jinyu Chen, Wenchao Xu, Song Guo, Junxiao Wang, Jie Zhang, and Haozhao Wang. Fedtune: A deep dive into efficient federated fine-tuning with pre-trained transformers, 2022. 1
- [10] Yae Jee Cho, Andre Manoel, Gauri Joshi, Robert Sim, and Dimitrios Dimitriadis. Heterogeneous ensemble knowledge transfer for training large models in Federated Learning. *Computing Research Repository arXiv Preprints*, 2204.12703, 2022. 2
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. 1
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, 2019. Association for Computational Linguistics. 1
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 1
- [14] Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Personalized federated learning: A meta-learning ap-

- proach. *Computing Research Repository arXiv Preprints*, 2002.07948, 2020. 2
- [15] Ziwei Fan, Hao Ding, Anoop Deoras, and Trong Nghia Hoang. Personalized federated domain adaptation for item-to-item recommendation. In *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*, pages 560–570. PMLR, 2023. 2
- [16] Tao Guo, Song Guo, Junxiao Wang, Xueyang Tang, and Wenchao Xu. PromptFL: Let Federated Participants Cooperatively Learn Prompts Instead of Models-Federated Learning in Age of Foundation Model. *IEEE Transactions on Mobile Computing*, 2023. 1
- [17] Minh Hoang, Nghia Hoang, Bryan Kian Hsiang Low, and Carleton Kingsford. Collective model fusion for multiple black-box experts. In *Proceedings of the 36th International Conference on Machine Learning*, pages 2742–2750. PMLR, 2019. 2
- [18] Trong Nghia Hoang, Quang Minh Hoang, Kian Hsiang Low, and Jonathan P. How. Collective online learning of Gaussian processes in massive multi-agent systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7850–7857, 2019. 2
- [19] Parminder Kaur, Husanbir Singh Pannu, and Avleen Kaur Malhi. Comparative analysis on cross-modal information retrieval: A review. *Computer Science Review*, 39:100336, 2021. 2
- [20] Muhammad Uzair khattak, Hanoona Rasheed, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Maple: Multi-modal prompt learning. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 3
- [21] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pages 5583–5594. PMLR, 2021. 1, 3, 6
- [22] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123:32 – 73, 2016. 1
- [23] Harold W Kuhn. The Hungarian Method for the Assignment Problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955. 6
- [24] Yi-Lun Lee, Yi-Hsuan Tsai, Wei-Chen Chiu, and Chen-Yu Lee. Multimodal prompting with missing modalities for visual recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 3, 6, 7, 1
- [25] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic, 2021. Association for Computational Linguistics. 1
- [26] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2:429–450, 2020. 2, 4
- [27] Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zhihua Zhang. On the convergence of FedAvg on non-IID data. *Computing Research Repository arXiv Preprints*, 1907.02189v1, 2019. 2
- [28] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, 2014. 1
- [29] Wang Lu, Xixu Hu, Jindong Wang, and Xing Xie. FedCLIP: Fast Generalization and Personalization for CLIP in Federated Learning, 2023. 1
- [30] Tengfei Ma, Trong Nghia Hoang, and Jie Chen. Federated learning of models pre-trained on different features with consensus graphs. In *Uncertainty in Artificial Intelligence*, pages 1336–1346, 2023. 2
- [31] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pages 1273–1282, 2017. 2, 3, 4, 5
- [32] John Nguyen, Jianyu Wang, Kshitiz Malik, Maziar Sanjabi, and Michael Rabbat. Where to Begin? On the Impact of Pre-Training and Initialization in Federated Learning, 2023. 1
- [33] Manh Duong Nguyen, Trung Thanh Nguyen, Huy-Hieu Pham, Trong Nghia Hoang, Phi Le Nguyen, and Thanh Trung Huynh. Fedmac: Tackling partial-modality missing in federated learning with cross-modal aggregation and contrastive regularization. 2024. 1, 3
- [34] Nang Hung Nguyen, Phi Le Nguyen, Thuy Dung Nguyen, Trung Thanh Nguyen, Duc Long Nguyen, Thanh Hung Nguyen, Huy Hieu Pham, and Thao Nguyen Truong. FedDRL: Deep reinforcement learning-based adaptive aggregation for non-IID data in federated learning. In *Proceedings of the 51st International Conference on Parallel Processing*, pages 1–11, 2022. 2
- [35] Nang Hung Nguyen, Duc Long Nguyen, Trong Bang Nguyen, Thanh-Hung Nguyen, Huy Hieu Pham, Truong Thao Nguyen, and Phi Le Nguyen. CADIS: Handling cluster-skewed Non-IID data in federated learning with clustered aggregation and knowledge distilled regularization. In *Proceedings of the 23rd IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing*, pages 249–261, 2023.
- [36] Quang Ha Pham, Nang Hung Nguyen, Thanh Hung Nguyen, Huy Hieu Pham, Phi Le Nguyen, and Truong Thao Nguyen. SEM: A simple yet efficient model-agnostic local training mechanism to tackle data sparsity and scarcity in federated learning. In *Proceedings of the 11st International Symposium on Computing and Networking*, pages 120–126, 2023. 2
- [37] Thu Hang Phung, Binh P Nguyen, Thanh Hung Nguyen, Quoc Viet Hung Nguyen, Phi Le Nguyen, and Thanh Trung Huynh. A contrastive learning and graph-based approach for missing modalities in multimodal federated learning. In *Proceedings of the 2024 International Joint Conference on Neural Networks*, pages 1–8, 2024. 1, 3

- [38] Petra Poklukar, Miguel Vasco, Hang Yin, Francisco S Melo, Ana Paiva, and Danica Kragic. Geometric multimodal contrastive representation learning. In *Proceedings of the 2022 International Conference on Machine Learning*, pages 17782–17800, 2022. 3
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. *CoRR*, abs/2103.00020, 2021. 1
- [40] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020. 1
- [41] Kuniaki Saito, Kihyuk Sohn, Xiang Zhang, Chun-Liang Li, Chen-Yu Lee, Kate Saenko, and Tomas Pfister. Prefix conditioning unifies language and label supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2861–2870, 2023. 1, 3
- [42] Rachael Hwee Ling Sim, Yehong Zhang, Trong Nghia Hoang, Xinyi Xu, Bryan Kian Hsiang Low, and Patrick Jaillet. Incentives in private collaborative machine learning. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. 2
- [43] Canh T Dinh, Nguyen Tran, and Josh Nguyen. Personalized federated learning with Moreau envelopes. *Advances in Neural Information Processing Systems*, 33:21394–21405, 2020. 2
- [44] Maria Tsimpoukelli, Jacob L Menick, Serkan Cabi, S. M. Ali Eslami, Oriol Vinyals, and Felix Hill. Multimodal few-shot learning with frozen language models. In *Advances in Neural Information Processing Systems*, pages 200–212. Curran Associates, Inc., 2021. 1, 3
- [45] Feng Wang, M Cenk Gursoy, and Senem Velipasalar. Communication-efficient and privacy-preserving feature-based federated transfer learning. In *Proceedings of the 2022 IEEE Global Communications Conference*, pages 3875–3880, 2022. 2
- [46] Xin Wang, Devinder Kumar, Nicolas Thome, Matthieu Cord, and Frédéric Precioso. Recipe recognition with large multimodal food dataset. In *2015 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, pages 1–6, 2015. 2, 6, 3
- [47] Pei-Yau Weng, Minh Hoang, Lam M. Nguyen, My T. Thai, Tsui-Wei Weng, and Trong Nghia Hoang. Probabilistic federated prompt-tuning with non-IID and imbalanced data. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 1
- [48] Baochen Xiong, Xiaoshan Yang, Fan Qi, and Changsheng Xu. A unified framework for multi-modal Federated Learning. *Neurocomputing*, 480:110–118, 2022. 2
- [49] Baochen Xiong, Xiaoshan Yang, Yaguang Song, Yaowei Wang, and Changsheng Xu. Client-adaptive cross-model reconstruction network for modality-incomplete multimodal federated learning. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 1241–1249, 2023. 3
- [50] Fu-En Yang, Chien-Yi Wang, and Yu-Chiang Frank Wang. Efficient Model Personalization in Federated Learning via Client-Specific Prompt Generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 19159–19168, 2023. 1
- [51] Guan Yu, Quefeng Li, Dinggang Shen, and Yufeng Liu. Optimal sparse linear prediction for block-missing multimodality data without imputation. *Journal of the American Statistical Association*, 115(531):1406–1419, 2020. 3
- [52] Songcan Yu, Junbo Wang, Walid Hussein, and Patrick C.K. Hung. Robust multimodal federated learning for incomplete modalities. *Computer Communications*, 214:234–243, 2024. 1, 3
- [53] Mikhail Yurochkin, Mayank Agarwal, Soumya Ghosh, Kristjan Greenewald, Trong Nghia Hoang, and Yasaman Khazaeni. Bayesian nonparametric federated learning of neural networks. In *International Conference on Machine Learning*, 2019. 2
- [54] Mikhail Yurochkin, Mayank Argawal, Soumya Ghosh, Kristjan Greenewald, and Trong Nghia Hoang. Statistical model aggregation via parameter matching. In *Advances in Neural Information Processing Systems*, pages 10954–10964, 2019. 2
- [55] Changqing Zhang, Yajie Cui, Zongbo Han, Joey Tianyi Zhou, Huazhu Fu, and Qinghua Hu. Deep partial multi-view learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(5):2402–2415, 2020. 3
- [56] Jianyi Zhang, Hao Frank Yang, Ang Li, Xin Guo, Pu Wang, Haiming Wang, Yiran Chen, and Hai Li. Mllm-fl: Multimodal large language model assisted federated learning on heterogeneous and long-tailed data. *arXiv preprint arXiv:2409.06067*, 2024. 1, 2
- [57] Wei Zhang and Xiang Li. Data privacy preserving federated transfer learning in machinery fault diagnostics using prior distributions. *Structural Health Monitoring*, 21(4):1329–1344, 2022. 2
- [58] Haodong Zhao, Wei Du, Fangqi Li, Peixuan Li, and Gongshen Liu. FedPrompt: Communication-Efficient and Privacy Preserving Prompt Tuning in Federated Learning. In *Proc. ICASSP*, 2023. 1
- [59] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 3
- [60] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision (IJCV)*, 2022. 1, 3
- [61] Tongxue Zhou, Pierre Vera, Stéphane Canu, and Su Ruan. Missing data imputation via conditional generator and correlation learning for multimodal brain tumor segmentation. *Pattern Recognition Letters*, 158:125–132, 2022. 3