

# Mitigating Modality Quantity and Quality Imbalance in Multimodal Online Federated Learning

Heqiang Wang, Weihong Yang, Xiaoxiong Zhong, Jia Zhou,  
Fangming Liu, Weizhe Zhang, Keqin Li *Fellow, IEEE*

**Abstract**—The Internet of Things (IoT) ecosystem produces massive volumes of multimodal data from diverse sources, including sensors, cameras, and microphones. With advances in edge intelligence, IoT devices have evolved from simple data acquisition units into computationally capable nodes, enabling localized processing of heterogeneous multimodal data. This evolution necessitates distributed learning paradigms that can efficiently handle such data. Furthermore, the continuous nature of data generation and the limited storage capacity of edge devices demand an online learning framework. Multimodal Online Federated Learning (MMO-FL) has been identified as a compelling strategy to address these requirements. However, MMO-FL faces new challenges due to the inherent instability of IoT devices, which often results in modality quantity and quality imbalance (QQI) during data collection. In this study, we systematically investigate the impact of QQI within the MMO-FL framework and present a comprehensive theoretical analysis quantifying how both types of imbalance degrade learning performance. To tackle these issues, we propose the Modality Quantity and Quality Rebalanced (QQR) algorithm, a prototype learning based method designed to operate in parallel with the training process. Extensive experiments on two real-world multimodal datasets show that the proposed QQR algorithm consistently outperforms benchmarks under modality imbalance conditions with promising learning performance.

**Index terms**— Federated Learning, Multimodal Learning, Online Learning, Internet of Thing, Modality Imbalanced.

## I. INTRODUCTION

The accelerated development of the Internet of Things (IoT) [1] has resulted in an extraordinary increase in data generated by diverse interconnected devices, such as smart home systems [2], wearable health trackers [3], and industrial sensors [4]. To enable intelligent applications and services within this ecosystem, artificial intelligence, particularly machine learning and deep learning, has become an essential approach for building models from large-scale IoT data. Traditionally, model training has been conducted on centralized cloud platforms or in data centers. However, as both the volume of IoT data and the number of connected devices continue to rise, this centralized paradigm encounters scalability and efficiency bottlenecks. Sending large volumes of unprocessed data to a central server requires substantial network bandwidth and induces high communication costs, making it unsuitable for latency-critical scenarios such as autonomous driving [5] and

real-time health monitoring [6]. Moreover, sending sensitive information to the cloud raises significant privacy risks [7]. With IoT devices evolving from basic data acquisition units into capable edge nodes, there is growing potential to utilize their local computing resources to address these challenges. In this regard, federated learning (FL) [8]–[10] has emerged as an attractive distributed learning framework. FL enables distributed model training among devices without transferring raw data, offering a communication-efficient and privacy-preserving alternative to centralized training. By minimizing data transfer and maintaining data confidentiality, FL provides a scalable pathway for deploying intelligent IoT applications.

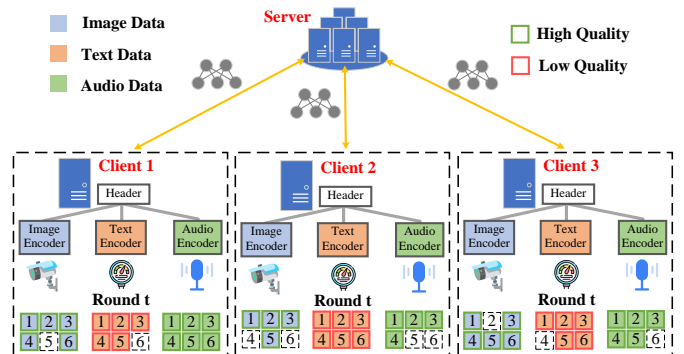


Fig. 1: IoT-Based MMO-FL with Modality Imbalance

Conventional FL frameworks for IoT have largely focused on unimodal data. In reality, however, IoT systems operate in inherently multimodal environments, where a variety of sensors capture data from different modalities [11], such as visual information from cameras, audio signals from microphones, and structured or time-series data from other sensing devices. Such multimodal inputs provide richer and more comprehensive information for downstream tasks. To leverage this diversity, multimodal federated learning (MFL) [12] has been introduced as an extension of FL, facilitating joint model training over decentralized multimodal datasets. In a typical MFL architecture, separate modality-specific encoders are employed, each designed to process a particular data type and extract discriminative features from high-dimensional raw inputs. The extracted features are then fused and passed to a shared head network, typically comprising deep neural layers and a softmax classifier, to produce the final prediction.

In conventional MFL, all training data are collected in advance, and model training proceeds in an offline manner. In contrast, real-world applications often require an online learning setting, where data are continuously acquired through-

H. Wang, W. Yang, X. Zhong, J. Zhou, F. Liu and W. Zhang are with Department of New Network, Peng Cheng Laboratory, Shenzhen, 518066, China. K. Li is with School of Computer Science, State University of New York, New Paltz, NY 12561, USA. (Corresponding Authors: Weihong Yang, Xiaoxiong Zhong.)

out the training process. Under such real-time and non-ideal conditions, variations in sensor performance and differences in modality characteristics can lead to modality imbalance during learning. In this work, modality imbalance is discussed in two primary forms: **Quantity Imbalance**: differences in sensor reliability or operation result in unequal amounts of data collected for each modality on a given client [13], and **Quality Imbalance**: variations in sensor capability cause differences in data integrity, with some modalities producing low-noise data while others suffer from higher noise levels [14]. These issues are typically absent in offline settings, where all data are pre-collected, but they become especially pronounced in online scenarios. If not addressed, modality imbalance can severely degrade learning performance and, in extreme cases, prevent model convergence.

Motivated by the above challenges, this work focuses on addressing modality quantity and quality imbalance in the Multimodal Online Federated Learning (MMO-FL) framework [15]. According to our knowledge, this is the first systematic investigation of how multimodal imbalances affect learning performance in MMO-FL, examined from both theoretical and experimental perspectives. An overview of the MMO-FL framework with modality imbalance is shown in Fig. 1. The quantity and quality of different modal data vary across clients and modalities at any given time, and these variations are inherently uncontrolled. The major contributions of our work are outlined below:

- 1) We study the problem of modality quantity and quality imbalance in MMO-FL scenarios, which primarily stems from the suboptimal conditions of data collection in online learning settings. Different sensors on clients acquire data from different modalities, and variations in their capabilities and operating frequencies result in disparities in both the quantity and quality of collected data. These imbalances ultimately impact the learning process.
- 2) We present a detailed theoretical analysis of MMO-FL in the presence of modality quantity and quality imbalance. The derived regret bound explicitly captures the performance degradation caused by each imbalance type and offers insights for subsequent algorithm design.
- 3) To address modality quantity and quality imbalance in the MMO-FL framework, we propose the Modality Quantity and Quality Rebalancing (QQR) algorithm. Built upon prototype learning, QQR provides efficient online rebalancing strategies that separately mitigate the effects of quantity and quality imbalance, thereby improving the learning performance of MMO-FL.
- 4) We evaluate the proposed QQR algorithms within the MMO-FL framework on two multimodal datasets, UCI-HAR and MVSA-Single. The experimental results indicate that QQR outperforms benchmark methods in addressing both modality quantity and quality imbalance on both datasets. In addition, the results demonstrate strong generalization capability across diverse modality combinations in modality imbalance scenarios.

The rest of this paper is structured as outlined below. Section II reviews related work. Section III outlines the framework

of the system and formulates the core problem. Section IV outlines the MMO-FL workflow under modality quantity and quality imbalance. Section V presents the regret analysis of MMO-FL and evaluates the impact of these imbalances on learning performance. Section VI presents the proposed QQR algorithm. Section VII presents experimental results demonstrating the efficacy of the proposed QQR algorithm under the MMO-FL framework. The paper is concluded in Section VIII.

## II. RELATED WORK

### A. Multimodal Federated Learning

MFL focuses on training task-specific models across clients that hold data from multiple modalities, enabling effective utilization of diverse and distributed information sources. As interest in MFL continues to grow, various algorithms have been proposed to address its unique challenges and improve model performance. One central issue is modality heterogeneity, where clients possess differing subsets of modalities, complicating model aggregation and limiting the efficiency of knowledge sharing. Several studies have investigated solutions for heterogeneous modality fusion to overcome this issue, including [16]–[18]. Another key challenge lies in selecting appropriate modalities for training under limited computational and communication resources. To address this, MPriorityFed [19] introduces an adaptive resource allocation strategy that prioritizes modality encoders based on their utility and cost, improving efficiency. In addition to the above challenges, the problem of missing modalities has attracted increasing attention [20]. Missing data may result from incomplete data acquisition [21], sensor malfunctions [22], or privacy-related restrictions [20], all of which negatively affect model performance. For instance, the MFCPL framework [23] employs cross-modal prototypes to facilitate knowledge transfer at both shared and modality-specific levels. However, most of the existing methods assume an offline learning scenario with access to a static dataset, limiting their applicability in real-world environments where data arrives sequentially and models must adapt continuously. Moreover, the missing modality problem [15] can be viewed as a specific case of the broader modality quantity imbalance addressed in this work. Unlike previous studies, we consider a more general setting where parts of a modality may be absent rather than the entire modality being unavailable. Furthermore, our approach incorporates the influence of modality quality on learning performance, a factor often overlooked in previous MFL works. Owing to differences in motivation between this work and [15], the theoretical analysis, algorithmic design, and experimental evaluation diverge significantly. The theoretical framework covers broader and more generalized scenarios, the algorithm enhances prototype utilization efficiency, and the experimental evaluation provides a more comprehensive level of detail.

### B. Online Federated Learning

Online learning processes data in a sequential manner, updating models incrementally, which makes it well-suited for scenarios involving continuous data streams and the need for

real-time adaptation [24]. Such methods are computationally efficient and remove the requirement of having the entire dataset beforehand, making them particularly advantageous in memory-constrained IoT settings. Within the FL paradigm, OFL extends these principles to distributed networks of decentralized learners [25], [26]. Unlike traditional offline FL, OFL focuses on minimizing long-term cumulative regret rather than static optimization objectives during local updates. Although research on OFL is still relatively limited, several works have made notable progress. For example, [27] proposes a communication-efficient OFL approach that achieves strong performance while reducing communication costs, and [28] introduces FedOMD, an OFL method designed for uncertain environments that can handle streaming data without assuming prior knowledge of loss distributions. While most of these studies address the horizontal FL (HFL) setting, [29] investigates the vertical FL (VFL) context, presenting an online VFL framework for cooperative spectrum sensing with sublinear regret guarantees. In the industrial domain, [4] tackles noise interference and device heterogeneity challenges in online VFL systems. However, all these efforts focus on unimodal OFL. In practical IoT applications, multimodal data is common, requiring joint learning from multiple sensor types. Recent work [15] introduces the concept of MMO-FL and addresses the problem of missing modalities. It is worth noting that missing modalities may be regarded as a specific instance of the quantity imbalance problem considered in this study, while our work investigates a more general setting that encompasses both quantity and quality imbalance.

### C. Modality Quality and Quantity Imbalance

Multimodal learning often faces modality quality imbalance, where certain modalities provide high-quality signals, while others are degraded by noise or carry less informative content. To address this issue, work [30] proposed a quality-aware multimodal fusion framework grounded in uncertainty estimation, achieving provably robust fusion performance under varying data quality conditions. Similarly, a comprehensive survey by [14] highlights quality-varying multimodal data as a central challenge and categorizes existing approaches that tackle noise, missing modalities, and uneven signal quality. Recent studies have continued to explore inter-modal quality imbalance. For instance, [31] and [32] propose targeted methods to mitigate quality disparities across modalities, thereby improving overall multimodal learning performance. In parallel, growing attention has been given to modality quantity imbalance, where certain modalities are sparsely observed or underrepresented. [33] examine this issue within the MFL context, identifying modal bias, a phenomenon in which clients possessing only unimodal data exert disproportionate influence on the global model. To address this, they introduce a balanced modality selection strategy to regulate client and modality participation, along with leveraging global prototypes for performance enhancement. Building on this work, their follow-up method, FedCMI [34], introduces cross-modal infiltration, which transfers knowledge from strong to weak modalities via projector modules and class-wise temperature scaling. While

previous studies have addressed modality quality and quantity imbalance individually, as noted in the aforementioned works, these approaches are predominantly based on offline learning and do not account for the dynamic characteristics of online learning scenarios. Moreover, prior research has not examined the impact of modality quality and quantity imbalance simultaneously, which introduces additional complexity and imposes stricter requirements on algorithm design.

## III. SYSTEM MODEL

Before presenting the details of the system model, we provide a summary of the key notations in Table I.

TABLE I: Key Notations

Symbol	Semantics
$K$	The number of clients
$M$	The number of modalities
$T$	The number of global rounds
$N$	The number of data samples collected by client
$E$	The number of local iterations
$C$	The number of classes
$\theta^m$	The modality encoder
$\theta^0$	The head encoder
$\eta$	The learning rate
$Z$	The feature extractor
$\Theta$	The overall model
$P_k^{t,m}$	The data availability status
$Q_k^{t,m}$	The modality quality status
$D_k^t$	The local training dataset
$G_k^{t,\tau}$	The gradient of the local overall model
$v_k^{t,m}$	The local prototype
$v_c^{t,m}$	The instantaneous global prototype
$\bar{v}_c^{t,m}$	The cumulative global prototype
$\bar{Y}^t$	The cumulative global prototype collection
$\mathcal{L}_{PCE}^m$	The prototype cross entropy
$\mathcal{L}_{PLR}$	The prototypical quality rebalancing loss
$\lambda_p$	The intra-round quantity imbalance ratio
$\lambda_r$	The inter-round quantity imbalance ratio
$\delta_r$	The intra-round quality imbalance ratio

Consider an IoT-enabled smart factory comprising a cloud server and  $K$  workstations serving as clients. Each workstation is equipped with multiple sensors to monitor factory conditions across various modalities at different locations, including visual, acoustic, and environmental sensors (e.g., temperature and humidity). With real-time data streaming from these sensors, the goal is to collaboratively train a unified global model using multimodal data segmented across various clients. This configuration forms the basic setting for the MMO-FL problem. During factory operation, sensors at each client continuously gather new data over time, and the timeline is divided into discrete intervals denoted as  $t = 1, 2, \dots, T$ . For simplicity, each interval is treated as a global training round.

In each global round, every client  $k \in \mathcal{K}$  independently acquires data for each modality using its corresponding sensors. Let  $N$  denote the target number of samples each client aims to collect per round. In practice, however, variations in sensor reliability and environmental factors can lead to modality quantity and quality imbalance, where some modalities fail to produce complete or valid samples. As a result, the actual local training dataset of client  $k$  at round  $t$  is defined as:

$$D_k^t = \left( \left[ \left( X_k^{t,1}, P_k^{t,1} \right), \dots, \left( X_k^{t,M}, P_k^{t,M} \right); Q_k^t; Y_k^t \right] \right)$$

$$= \left\{ \left\{ \left[ \left( x_{k,n}^{t,m}, p_{k,n}^{t,m} \right); y_{k,n}^t \right] \right\}_{n=1}^N; q_k^{t,m} \right\}_{m=1}^M \quad (1)$$

Here,  $x_{k,n}^{t,m}$  represents the  $m^{\text{th}}$  modality data of the  $n^{\text{th}}$  sample in client  $k$  collected at global round  $t$ , and  $y_{k,n}^t$  denotes the corresponding label. The binary indicator  $p_{k,n}^{t,m} \in \{0, 1\}$  specifies whether modality  $m$  of the  $n^{\text{th}}$  sample is available: a value of 1 indicates successful data collection, whereas 0 indicates that the modality is missing. Similarly, the binary indicator  $q_k^{t,m} \in \{0, 1\}$  reflects the quality status of modality  $m$ , which can be determined by monitoring the operating condition of the client's sensor. A value of 1 indicates that the collected modality data is of normal quality, while 0 signifies low quality. The overall dataset aggregated across all clients at global round  $t$  is defined as  $\mathcal{D}^t = \sum_{k=1}^K \mathcal{D}_k^t$ . Due to the variability in  $P_k^{t,m}$  and  $Q_k^t$  across clients and rounds, the local datasets are inherently heterogeneous, leading to modality quality and quantity imbalance during training.

In the MFL, the global model trained collaboratively by the clients and aggregated on the server consists of two main components: the modality encoders  $\theta^1, \dots, \theta^M$ , which extract feature representations from raw inputs, and the head encoder  $\theta^0$ , which integrates these features to generate the final output. The architecture of each modality encoder is tailored to its specific data type. For client  $k$  at round  $t$ , the feature vector for the  $m^{\text{th}}$  modality, extracted by modality encoder  $\theta^m$ , is denoted as  $Z_k^{t,m} = \theta^m(X_k^{t,m})$ . The header encoder  $\theta^0$  then combines the outputs  $Z_k^{t,m}$  from all modality encoders to accomplish the learning objective. Based on these definitions, the loss function for the aggregated training dataset at round  $t$  is formulated as follows:

$$F_t(\Theta, \mathcal{D}^t) = \frac{1}{K} \sum_{k=1}^K f_t \left( \theta^0 \left( Z_k^{t,1}, \dots, Z_k^{t,M} \right), Y_k^t \right) \quad (2)$$

Here,  $\Theta = \{\theta^0, \theta^1, \dots, \theta^M\}$  denotes the overall model. The term  $\theta^0 \left( Z_k^{t,1}, \dots, Z_k^{t,M} \right)$  represents the predicted labels generated by the head encoder, while  $f_t$  is the loss function measuring the difference between the predictions and the ground-truth labels. Notably, only samples containing all modalities can be used to update the client's loss function. This constraint can significantly hinder learning performance, as a large portion of the additional data, available only in a subset of modalities, is excluded from the model update process.

Since the training process relies on continuously collected real-time data rather than a fixed dataset, adopting an online learning approach is essential. Let the global model at each round be denoted as  $\Theta^1, \dots, \Theta^T$ . The learning regret,  $\text{Reg}_T$  is defined to measure the discrepancy between the cumulative loss incurred by the learner and that of the best fixed model chosen in hindsight. Specifically:

$$\text{Reg}_T = \sum_{t=1}^T F_t(\Theta^t; \mathcal{D}^t) - \sum_{t=1}^T F_t(\Theta^*; \mathcal{D}^t) \quad (3)$$

Here,  $\Theta^* = \arg \min_{\Theta} \sum_{t=1}^T F_t(\Theta; \mathcal{D}^t)$  represents the optimal fixed model chosen in hindsight. The objective is to minimize the learning regret, which is equivalent to minimizing the total loss accumulated across all rounds. Notably, if the regret grows

sublinearly with respect to  $T$ , the online learning algorithm can asymptotically reduce the training loss over time.

The fixed optimal strategy in hindsight refers to a policy that could be devised by a centralized entity with complete prior knowledge of all per-round loss functions. In our setting, achieving such a strategy would require access to future information, including all upcoming data collections across rounds, which is inherently unpredictable due to its stochastic nature. Consequently, the complete loss functions are unknown at the outset and evolve dynamically over time. In this context, regret serves as a theoretical measure to quantify the disparity in performance between the proposed algorithm and the ideal but unattainable optimal strategy. For experimental validation, the proposed online algorithm is evaluated with practical criteria, such as test accuracy.

#### IV. THE WORKFLOW OF MMO-FL WITH MODALITY QUANTITY AND QUALITY IMBALANCE

This section outlines the overall workflow of MMO-FL and examines the challenges posed by modality quantity and quality imbalance during the learning process. MMO-FL extends the HFL framework, in which each client independently collects local data samples and the server combines these updates to form the global model. Under modality imbalance, however, MMO-FL faces both quantitative and qualitative disparities among different modalities, which can disrupt the training process and substantially impair learning performance. In each global round  $t \in \mathcal{T}$ , where  $\mathcal{T} = \{0, 1, 2, \dots, T-1\}$ , all clients perform a specified number of local training iterations, denoted by the parameter  $E$ . The index  $\tau = 0, 1, 2, \dots, E$  tracks these local iterations. Each global round  $t$  then proceeds through the following steps.

1) **Client - Online Data Collect:** At the start of each global round  $t$ , client  $k$  aims to collect a new training dataset  $\mathcal{D}_k^t$  for that round. However, due to the inherent instability of IoT sensors, the collected data may exhibit modality quantity and quality imbalances. Data collection is conducted separately for each modality, producing paired feature-label sub-datasets that capture both quantitative and qualitative characteristics. To differentiate between the effects of quantity imbalance and quality imbalance, we define two types of local datasets to represent these conditions. The dataset only affected by **modality quantity imbalance** is denoted as  $\mathcal{D}_k^{t,-}$  and is defined as:

$$\mathcal{D}_k^{t,-} = \left( \left[ \left( X_k^{t,1}, \hat{P}_k^{t,1} \right), \dots, \left( X_k^{t,M}, \hat{P}_k^{t,M} \right); Q_k^t, Y_k^t \right] \right)$$

Here,  $\hat{P}_k^{t,m}$  represents the actual data availability status for modality  $m$ , which differs from the ideal collection state  $P_k^{t,m}$ . In the ideal case, all entries of  $P_k^{t,m}$  are equal to 1. In contrast,  $\hat{P}_k^{t,m}$  may contain zeros, indicating missing modality data for certain samples. To further characterize the impact of **modality quality imbalance**, we further define the dataset  $\tilde{\mathcal{D}}_k^{t,-}$  as follow:

$$\tilde{\mathcal{D}}_k^{t,-} = \left( \left[ \left( \tilde{X}_k^{t,1}, \hat{P}_k^{t,1} \right), \dots, \left( \tilde{X}_k^{t,M}, \hat{P}_k^{t,M} \right); \hat{Q}_k^t, Y_k^t \right] \right)$$

In this context,  $\hat{Q}_k^t = \{q_k^{t,m} \mid m \in \mathcal{M}\}$  indicates the actual quality condition of the data acquired by client  $k$  at global  $t$ . If

all elements in the set are equal to 1, it means that the collected data is of normal quality, which implies  $\tilde{X}_k^{t,m} = X_k^{t,m}$ . Conversely, if any  $q_k^{t,m}$  is equal to 0, it reflects that the corresponding data is of low quality and is represented as  $\tilde{X}_k^{t,m}$ . Such online fluctuations in data quantity and quality do not occur in conventional offline multimodal FL settings. In our work, the fluctuations in data quality and quantity we refer to primarily arise from the nature of real-time data arrival, where both the quantity and quality of data from a single client can vary significantly across different rounds. This highlights the inherent volatility of online settings. In contrast, while traditional offline scenarios may exhibit heterogeneity across clients, they generally do not face temporal fluctuations in data quality or quantity from the same client. As a result, offline approaches can often mitigate such issues through preprocessing and denoising before training. However, in online learning, such fluctuations must be addressed dynamically, necessitating the design of online algorithms to continuously adapt to these fluctuations. Although training can still proceed with the current imbalanced dataset, these conditions inevitably affect the training process and learning performance. Therefore, dedicated strategies are necessary to handle such imbalances effectively and ensure robust model training.

2) **Client - Local Model Update:** After data collection, each client  $k$  starts its own local training using the current global model  $\Theta^t$ , received from the server. This model serves as the starting point for updating a new local model based on the collected, real-world, non-ideal training dataset  $\tilde{\mathcal{D}}_k^{t,-}$ . Each client performs  $E$  iterations of online gradient descent (OGD) [29] using the full training dataset. The update process is formulated as follows:

$$\begin{aligned}\tilde{\Theta}_k^{t,0} &= \Theta^t \\ \tilde{\Theta}_k^{t,\tau+1} &= \tilde{\Theta}_k^{t,\tau} - \eta \tilde{\mathbf{G}}_k^{t,\tau,-}, \quad \forall \tau = 0, \dots, E-1 \\ \tilde{\Theta}_k^{t,E} &= \tilde{\Theta}_k^{t,E}\end{aligned}\quad (4)$$

Here,  $\tilde{\mathbf{G}}_k^{t,\tau,-} = \nabla F_t(\tilde{\Theta}_k^{t,\tau}, \tilde{\mathcal{D}}_k^{t,-})$  denotes the gradient computed on the current local dataset  $\tilde{\mathcal{D}}_k^{t,-}$ , which contains all modalities but is subject to both quantity and quality imbalance. The term  $\mathbf{G}_k^{t,\tau,-} = \nabla F_t(\tilde{\Theta}_k^{t,\tau}, \mathcal{D}_k^{t,-})$  represents the gradient influenced only by modality quantity imbalance, while  $\tilde{\mathbf{G}}_k^{t,\tau} = \nabla F_t(\tilde{\Theta}_k^{t,\tau}, \tilde{\mathcal{D}}_k^t)$  corresponds to the gradient affected solely by modality quality imbalance. In all cases,  $\eta$  denotes the learning rate. Since  $\Theta_k^{t,0}$  is initialized from the final model obtained in the previous global round, no local training has been performed at this stage. Consequently, it is not affected by modality imbalance, yielding  $\Theta_k^{t,0} = \tilde{\Theta}_k^{t,0}$ . For notational consistency with the subsequent derivations, we use  $\tilde{\Theta}_k^{t,0}$  here.

It should be noted that modality quantity imbalance causes some samples to contain complete modality information, while others include only partial modality data. At the same time, modality quality imbalance introduces discrepancies between certain modality representations and their true underlying values. Although these imbalances may degrade model performance, the training process remains viable, and model updates can still be carried out.

3) **Client - Local Model Upload:** Each client will upload the corresponding local model  $\tilde{\Theta}_k^{t+1}$  to the server after finishing the  $E$  iterations of local model update.

4) **Server - Global Model Update:** The global model is updated at the server through the aggregation of local model updates collected from clients, as expressed by the following equation:

$$\Theta^{t+1} = \frac{1}{K} \sum_{k \in \mathcal{K}} \tilde{\Theta}_k^{t+1} \quad (5)$$

The server then transmits the newly updated global model to all clients for the next training round. This process is repeated until the predefined number of global rounds is completed in the online learning setting.

From the preceding MMO-FL workflow under modality quantity and quality imbalance, it is clear that training can still proceed using the non-ideal local dataset  $\tilde{\mathcal{D}}_k^{t,-}$ . However, the resulting performance is likely to be suboptimal due to missing modalities and degraded data quality. A natural approach is to augment or correct  $\tilde{\mathcal{D}}_k^{t,-}$ , but this is often impractical because complete modality information is rarely available. To address this challenge, we propose an online compensation mechanism that operates concurrently with the training process to mitigate the negative effects of modality imbalance. Before detailing the algorithm, we present a thorough theoretical analysis quantifying the impact of modality quantity and quality imbalance on the regret bound within the MMO-FL framework.

## V. THEORETICAL ANALYSIS

In this section, we present a comprehensive regret analysis of the proposed MMO-FL algorithm, highlighting several additional challenges in the theoretical study. First, the online learning paradigm requires assessing the long-term cumulative loss  $Reg_T$  rather than focusing solely on convergence at a specific round. Second, due to the multimodal nature of the problem, the feature models corresponding to each modality must be examined individually, making it unsuitable to treat the local model as a single unified entity. Third, appropriate assumptions and configurations are necessary to account for both modality quantity and quality imbalance. These factors call for new proof techniques and supplementary assumptions to accurately model the problem's complexity. The proof is conducted in three stages: (1) We derive the regret bound for the case where the number of local iterations satisfies  $E > 1$  without considering modality quantity and quality imbalance. (2) We extend the analysis to account for modality quality imbalance under the same local iteration setting. (3) We will further explore the regret bound when  $E > 1$  while accounting for the impact of both modality quantity and quality imbalance.

### A. MMO-FL with local iterations $E > 1$ and without modality quantity and quality imbalance

To facilitate our analysis, we first introduce several additional definitions. After applying some basic transformations to the global model update equation above, we derive an alternative form of the global model update equation for the

case where the local iteration is  $E > 1$  and no modality quantity and quality imbalance happen, as follows:

$$\Theta^{t+1,0} = \Theta^{t,0} - \frac{\eta}{K} \sum_{k=1}^K \sum_{\tau=0}^{E-1} \mathbf{G}_k^{t,\tau} \quad (6)$$

where  $\mathbf{G}_k^{t,\tau}$  denotes the gradient of the local overall model for client  $k$  across all  $M$  modalities for local round  $t$  and local iteration  $\tau$ , which equals:

$$\mathbf{G}_k^{t,\tau} = \left[ \left( \mathbf{G}_k^{t,\tau,0} \right)^\top, \dots, \left( \mathbf{G}_k^{t,\tau,m} \right)^\top, \dots, \left( \mathbf{G}_k^{t,\tau,M} \right)^\top \right]^\top \quad (7)$$

Subsequently, we will introduce the assumptions that are standard for analyzing online convex optimization, as referenced in [4], [29]. Some assumptions are defined at the modality level, tailored to the requirements of the multimodal parameter formulation.

**Assumption 1.** For any  $\mathcal{D}^t$ , the loss function  $F_t(\Theta; \mathcal{D}^t)$  is convex with respect to  $\Theta$  and differentiable.

**Assumption 2.** The loss function is  $L$ -Lipschitz continuous, the partial derivatives for each modality satisfies:  $\|\nabla_m F_t(\Theta)\|^2 \leq L^2$ .

**Assumption 3.** The partial derivatives, corresponding to the consistent loss function, fulfills the following condition:

$$\left\| \mathbf{G}_k^{t,\tau'} - \mathbf{G}_k^{t,\tau} \right\| \leq \varphi \left\| \Theta_k^{t,\tau'} - \Theta_k^{t,\tau} \right\|$$

where  $\tau'$  and  $\tau$  indicate that they correspond to different local iterations.

For the purposes of subsequent theoretical analysis, we consider a  $D$ -dimensional vector for each modality in both the overall gradient and model. We define an arbitrary vector element  $d \in [1, D]$  in overall gradient for modality  $m$  as  $\mathbf{G}_{k,d}^m$ , and similarly, the arbitrary vector element  $d \in [1, D]$  in overall model for modality  $m$  is denoted as  $\Theta_{k,d}^m$ . Including the head encoder  $m = 0$ , each overall gradient and model consists of a total of  $(M + 1)D$  vector elements.

**Assumption 4.** The arbitrary vector element  $d$  in the overall model  $\Theta_{k,d}^m$  for any modality  $m$  is bounded by:  $\left| \Theta_{k,d}^m \right| \leq \sigma$ .

Assumption 1 guarantees the convexity of the objective function, allowing us to apply convex optimization properties. Assumption 2 places a bound on the magnitude of the loss function's partial derivatives at the modality level. Assumption 3 restricts the variation in these partial derivatives to a specific range, corresponding to the changes in the model across two different local iterations under the same loss function, thereby reflecting the smoothness property. Finally, Assumption 4 defines the allowable range for each vector element of the overall model. Then, we derive the following Theorem 1.

**Theorem 1.** Under Assumption 1-4, MMO-FL with local iterations  $E > 1$  and excluding the impact of modality quantity and quality imbalance, yields the following regret bound:

$$Reg_T = \sum_{t=1}^T \sum_{k=1}^K \mathbb{E}_t [F_t(\Theta^{t,0}; \mathcal{D}_k^t)] - \sum_{t=1}^T \sum_{k=1}^K F_t(\Theta^*; \mathcal{D}_k^t)$$

$$\leq \frac{K \|\Theta^{1,0} - \Theta^*\|^2}{2\eta E} + \frac{\eta T K E (M + 1) L^2}{2} + 2\eta T D E K (M + 1)^2 \varphi \sigma L$$

*Proof.* The detail proof can be found in Appendix B.  $\square$

According to Theorem 1, by setting  $\eta = \mathcal{O}(1/\sqrt{T})$ , the MMO-FL can achieve a sublinear regret rate  $\mathcal{O}(\sqrt{T})$ . A sublinear regret bound means that the average regret per round, calculated as the total regret divided by the number of rounds, approaches zero as the number of rounds increases indefinitely. In the following analysis, we will consider the effect of modality quality and quantity imbalance.

### B. MMO-FL with local iterations $E > 1$ and with modality quality imbalance

In this section, we extend the proof to the case where the number of local iterations satisfies  $E > 1$  while accounting for the impact of modality quality imbalance. Using a similar transformation approach as before, we obtain the global model update equation as follows:

$$\Theta^{t+1,0} = \Theta^{t,0} - \frac{\eta}{K} \sum_{k=1}^K \sum_{\tau=0}^{E-1} \tilde{\mathbf{G}}_k^{t,\tau} \quad (8)$$

We introduce an additional assumption to bound the influence of low-quality data on gradient computation. In particular, we assume that the reduction in data quality results from noise introduced during the data collection process.

**Assumption 5.** For any arbitrary component  $d$  of the overall gradient corresponding to modality  $m$ , the difference between the gradients computed with and without the influence of noise is bounded by a finite range:

$$\left| \tilde{\mathbf{G}}_{k,d}^{t,\tau,m} - \mathbf{G}_{k,d}^{t,\tau,m} \right| \leq \rho_m^t \quad (9)$$

where  $\rho_m^t$  denotes the maximum deviation induced by noise in modality  $m$  at round  $t$ .

Assumption 5 ensures that the deviation in each component of the gradient vector remains bounded within a specified range when comparing gradients computed with and without noise. The parameter  $\rho$  quantifies the extent of the noise impact, where a larger value of  $\rho$  indicates a higher degree of modality quality imbalance. Based on this assumption, we derive Theorem 2, which incorporates the effect of modality quality imbalance into the analysis.

**Theorem 2.** Under Assumption 1-5, MMO-FL with local iterations  $E > 1$  and including the impact of modality quality imbalance, yields the following regret bound:

$$\begin{aligned} Reg_T &= \sum_{t=1}^T \sum_{k=1}^K \mathbb{E}_t [F_t(\Theta^{t,0}; \tilde{\mathcal{D}}_k^t)] - \sum_{t=1}^T \sum_{k=1}^K F_t(\Theta^*; \tilde{\mathcal{D}}_k^t) \\ &\leq \frac{K \|\Theta^{1,0} - \Theta^*\|^2}{2\eta E} + 2T\eta D E K (M + 1)^2 \sigma \varphi L \\ &\quad + 2(M + 1) D K \sigma \sum_{t=1}^T \rho_{\max}^t + \eta T K E (M + 1) L^2 \end{aligned}$$

$$+ \eta KE(M+1)^2 D^2 \sum_{t=1}^T (\rho_{\max}^t)^2$$

Where  $\rho_{\max}^t = \max_{m \in \mathcal{M}} \rho_m^t$ .

*Proof.* The detail proof can be found in Appendix D.  $\square$

Based on Theorem 2, by choosing the learning rate as  $\eta = \mathcal{O}(1/\sqrt{T})$ , MMO-FL with the impact of modality quality imbalance can achieve a regret bound of  $\mathcal{O}(\sqrt{T} + \sum_{t=1}^T \rho_{\max}^t + \frac{1}{\sqrt{T}} \sum_{t=1}^T (\rho_{\max}^t)^2)$  over  $T$  time rounds. Our analysis reveals that the term  $\mathcal{O}(\sum_{t=1}^T \rho_{\max}^t + \frac{1}{\sqrt{T}} \sum_{t=1}^T (\rho_{\max}^t)^2)$ , which captures the impact of low-quality data, is essential in determining whether sublinear regret is achievable. When low-quality data is involved, the behavior of this term becomes unpredictable and cannot be ignored. Consequently, this additional error term must be retained in the overall regret bound, making it a generalized result that accounts for quality imbalance. To mitigate the effect of low-quality, we aim to propose an algorithm that encourages a gradual reduction of  $\rho_{\max}^t$  throughout the training process. This ensures that as  $T \rightarrow \infty$ , the average regret  $\frac{\text{Reg}_T}{T}$  approaches zero. To support this theoretical claim, we will introduce additional theoretical analysis in the end of section VI that illustrate how  $\rho_{\max}^t$  evolves when our proposed algorithm is applied. To further our analysis, we now incorporate the influence of modality quantity imbalance, which introduces the following case.

### C. MMO-FL with local iterations $E > 1$ and with modality quantity and quality imbalance

In this section, we extend the proof to the case where the local iteration  $E > 1$  and considering the effect with both modality quality and quantity imbalance. Following a similar transformations approach, we derive the global model update equation as follows:

$$\Theta^{t+1,0} = \Theta^{t,0} - \frac{\eta}{K} \sum_{k=1}^K \sum_{\tau=0}^{E-1} \tilde{\mathbf{g}}_k^{t,\tau,-} \quad (10)$$

To further characterize the impact of modality quantity imbalance, we introduce an additional assumption that bounds the deviation between the gradient computed on individual data samples and the corresponding batch gradient.

**Assumption 6.** For any arbitrary component  $d$  of the overall gradient corresponding to modality  $m$ , the deviation between the gradient computed from a single data sample and the corresponding batch gradient is bounded within a finite range:

$$\left| g_{k,d}^{t,\tau,m} - \mathbf{G}_{k,d}^{t,\tau,m} \right| \leq \gamma_m \quad (11)$$

Here,  $g_{k,d}^{t,\tau,m}$  denotes the gradient computed from a single data sample, while  $\mathbf{G}_{k,d}^{t,\tau,m} = \frac{1}{N} \sum_{n=1}^N g_{k,d}^{t,\tau,m}$  represents the corresponding batch gradient averaged over  $N$  samples.

Assumption 6 ensures a bounded deviation between the gradient computed from individual data samples and the average gradient over the batch. This allows us to quantify the impact of modality quantity imbalance. Then we derive

Theorem 3, which integrates the effects of both modality quality and quantity imbalance into the regret bound.

**Theorem 3.** Under Assumption 1-6, MMO-FL with local iterations  $E > 1$  and including the impact of modality quantity and quality imbalance, yields the following regret bound:

$$\begin{aligned} \text{Reg}_T &= \sum_{t=1}^T \sum_{k=1}^K \mathbb{E}_t \left[ F_t(\Theta^{t,0}; \tilde{\mathcal{D}}_k^{t,-}) \right] - \sum_{t=1}^T \sum_{k=1}^K F_t(\Theta^*; \tilde{\mathcal{D}}_k^{t,-}) \\ &\leq \frac{K \|\Theta^{1,0} - \Theta^*\|^2}{2\eta E} + \frac{3}{2} \eta T K E (M+1) L^2 \\ &\quad + 6\eta T K E (M+1)^2 D^2 \gamma_{\max}^2 \left( \frac{N - N_{\min}}{N} \right)^2 \\ &\quad + \frac{3}{2} \eta K E (M+1)^2 D^2 \sum_{t=1}^T (\rho_{\max}^t)^2 + 2\eta T D E K (M+1)^2 \sigma \varphi L \\ &\quad + 2T D K (M+1) \sigma \gamma_{\max} \left( \frac{N - N_{\min}}{N} \right) \\ &\quad + 4D K (M+1) \sigma \sum_{t=1}^T \rho_{\max}^t \end{aligned} \quad (12)$$

Where  $\gamma_{\max} = \max_{m \in \mathcal{M}} \gamma_m$  and  $N_{\min} = \min_{m,t,k} N_k^{t,m}$ .

*Proof.* The detail proof can be found in Appendix F.  $\square$

Based on Theorem 3, by choosing the learning rate as  $\eta = \mathcal{O}(1/\sqrt{T})$ , MMO-FL with the impact of modality quality and quantity imbalance can achieve a regret bound of  $\mathcal{O}(\sqrt{T} + \sum_{t=1}^T \rho_{\max}^t + \frac{1}{\sqrt{T}} \sum_{t=1}^T (\rho_{\max}^t)^2 + T \gamma_{\max} \left( \frac{N - N_{\min}}{N} \right))$  over  $T$  time rounds. In addition to the term  $\mathcal{O}(\sum_{t=1}^T \rho_{\max}^t + \frac{1}{\sqrt{T}} \sum_{t=1}^T (\rho_{\max}^t)^2)$  that appeared in Theorem 2, the regret bound is further influenced by the term  $T \gamma_{\max} \left( \frac{N - N_{\min}}{N} \right)$ , which captures the effect of modality quantity imbalance. To achieve a sublinear regret bound, two strategies can be considered: reducing  $\gamma_{\max}$  or increasing  $N_{\min}$  to approach  $N$ . Since  $\gamma_{\max}$  is typically determined by the data and not controllable, a more practical approach is to compensate for missing data to raise  $N_{\min}$  toward  $N$ . This observation motivates the design of the algorithm proposed in the subsequent section. The theoretical analysis above provides clear guidance for algorithm design. In the following section, we introduce the rebalancing algorithm, which leverages prototype learning to address both types of modality imbalance.

## VI. MODAL QUANTITY AND QUALITY REBALANCING ALGORITHM

Building upon the theoretical analysis presented in the previous section, it is evident that both modality quality imbalance and modality quantity imbalance have a substantial impact on the regret bound and ultimately degrade learning performance. To mitigate these effects, dedicated algorithmic solutions are necessary. In this section, we introduce the Modality Quantity and Quality Rebalancing (QQR) algorithm, which leverages prototype learning techniques [35] to enable effective and efficient rebalancing of modality quality and quantity during the learning process. Several prototype-based approaches have been explored in multimodal federated learning. For example,

FedProto [36] is designed to alleviate client heterogeneity by exchanging and aggregating class prototypes rather than gradients, using global prototypes to regularize local training and thereby improving convergence and generalization under non-IID conditions. In contrast, FedPCL [37] introduces a lightweight federated learning framework that enables clients to collaboratively integrate multiple fixed pre-trained models through prototype sharing and prototype-wise contrastive learning, which supports the construction of personalized representations with low communication overhead. Although both methods rely on prototypes, they target challenges that are fundamentally different from those addressed in our work. The proposed QQR algorithm is detailed in Algorithm 1 and comprises three main components: Online Global Prototype Construction (OGPC), Prototypical Quantity Rebalancing (PNR), and Prototypical Quality Rebalancing (PLR). The following subsections provide a detailed exposition of all components.

---

**Algorithm 1** QQR Algorithm
 

---

```

1: for  $t = 1, 2, \dots, T - 1$  do
2:   Client Side:
3:   for  $k = 1, 2, \dots, K$  do
4:     if Client satisfies the condition  $\Phi_g$  then
5:       Calculates the local prototype  $v_{k,c}^{t,m}$  via Eq. 13.
6:       Integrates and uploads the  $v_{k,c}^{t,m}$  to server.
7:     end if
8:     Downloads and utilizes the collection  $\bar{\mathcal{V}}^t$ .
9:     Prototypical Quantity Rebalancing:
10:    Builds the  $\tilde{Z}_k^{t,m}$  via Eq. 18 and Eq. 19.
11:    Prototypical Quality Rebalancing:
12:    Utilizes the PLR loss  $\mathcal{L}_{PLR}$  in Eq. 22 for updating.
13:   end for
14:   Server Side:
15:   Collects the local prototype  $v_{k,c}^{t,m}$  from clients.
16:   Calculates the  $v_c^{t,m}$  via Eq. 15.
17:   Updates the  $\bar{v}_c^{t,m}$  via Eq. 16.
18:   Aggregates and distributes the  $\bar{\mathcal{V}}^t$  to all clients.
19: end for

```

---

### A. Online Global Prototype Construction

The global prototype construction process is designed to maintain long-term global prototypes, which are made available to clients as needed during the quantity and quality rebalancing stages. To progressively approximate the true class representations, these prototypes are updated at each global round  $t$ . Let  $c \in \{1, \dots, C\}$  denote the class labels. The local prototype is defined as the average of the feature representations extracted by the modality encoder from the available samples of a given class that meet the normal quality criterion. This can be formally expressed as follows:

$$v_{k,c}^{t,m} = \frac{1}{\sum_c P_{k,c}^{t,m}} \sum_{n \in \mathcal{N}(\Phi_g)} \theta^m(x_{k,n}^{t,m}) \quad (13)$$

$$\text{where } \Phi_g = \left\{ p_{k,n}^{t,m} = 1, q_k^{t,m} = 1, y_{k,n}^t = c \right\} \quad (14)$$

To ensure fairness among prototypes for different modalities, we introduce two additional principles. First, local prototypes

derived from various modalities should conform to a unified structure. This ensures that the features extracted by the modality encoders, denoted as  $\theta^m(x_{k,n}^{t,m})$ , remain consistent across modalities, even if the original data formats differ. By passing data through their respective encoders, the resulting features can be structurally aligned in a shared representation space. Second, to maintain comparability and training stability, local prototypes across modalities should be normalized, ensuring their magnitudes are consistent throughout the learning process.

At the end of each training round, every client generates local prototypes for each modality and transmits them to the server. Upon receiving these local prototypes, the server constructs an instantaneous global prototype for modality  $m$  and class  $c$  at global round  $t$  using the following formulation:

$$v_c^{t,m} = \frac{1}{K} \sum_{k \in \mathcal{K}} v_{k,c}^{t,m} \quad (15)$$

Given the nature of online learning, data collected at each global round may be biased and may not cover all class categories. According to the law of large numbers [38], with an increasing number of samples, the sample mean approaches the actual population mean. To tackle this problem, it is essential to maintain a cumulative global prototype that continuously captures the semantic representation of each class for every modality across training rounds. The cumulative global prototype for modality  $m$  and class  $c$  at global round  $t$  is defined as follows:

$$\bar{v}_c^{t,m} = \frac{(t-1)\bar{v}_c^{t-1,m} + v_c^{t,m}}{t} \quad (16)$$

The cumulative global prototype  $\bar{v}_c^{t,m}$  is continuously updated and maintained on the server, this will serve as a key component for subsequent quantity and quality rebalancing procedures. The server retains the full set of cumulative global prototypes, organized by modality and class, denoted as  $\bar{\mathcal{V}}^t$ , and defined as follows:

$$\bar{\mathcal{V}}^t = \begin{bmatrix} \bar{v}_1^{t,1} & \dots & \bar{v}_c^{t,1} & \dots & \bar{v}_C^{t,1} \\ \dots & \dots & \dots & \dots & \dots \\ \bar{v}_1^{t,m} & \dots & \bar{v}_c^{t,m} & \dots & \bar{v}_C^{t,m} \\ \dots & \dots & \dots & \dots & \dots \\ \bar{v}_1^{t,M} & \dots & \bar{v}_c^{t,M} & \dots & \bar{v}_C^{t,M} \end{bmatrix} \quad (17)$$

This collection is continuously updated over successive training rounds, maintained on the server, and shared with clients as needed. In the following sections, we describe how clients utilize the cumulative global prototype collection to perform modality quantity rebalancing and modality quality rebalancing, respectively.

### B. Prototypical Quantity Rebalancing

In this phase, each client executes the PNR algorithm. The core objective of PNR is to address modality quantity imbalance by compensating for discrepancies in sample counts across modalities through prototype-based augmentation. The overall process of the PNR algorithm is illustrated in Step 2 of Fig. 2. Upon receiving the cumulative global prototype collection from the server, each client identifies its own

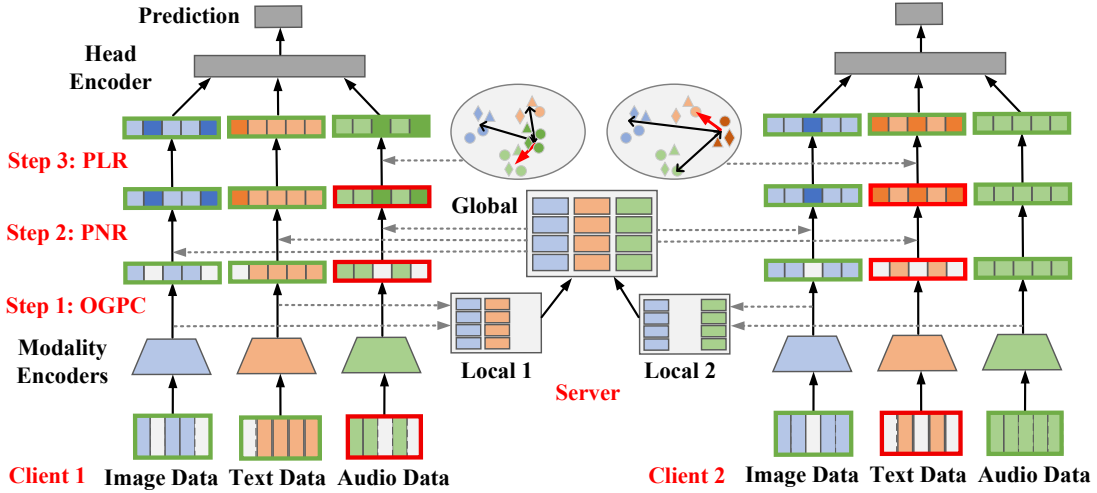


Fig. 2: Illustration of QQR: The QQR framework comprises three main components. First, OGPC generates local prototypes at each client and aggregates them at server to produce cumulative global prototypes. Second, PNR leverages the cumulative global prototype collection to substitute missing feature embeddings, thereby addressing modality quantity imbalance. Finally, PLR utilizes both the cumulative global prototype collection and the relative relationships among modalities to correct modality quality imbalance. The PNR and PLR procedures can be executed either independently or concurrently.

modality imbalance scenario and generates virtual feature representations to substitute for missing features in missing data samples. This process is described in detail below:

$$\tilde{Z}_k^{t,m} = [\tilde{z}_{k,1}^{t,m}, \dots, \tilde{z}_{k,n}^{t,m}, \dots, \tilde{z}_{k,N}^{t,m}] \quad (18)$$

where  $\tilde{z}_{k,n}^{t,m}$  denotes the corresponding feature representations for data sample  $n$  of modality  $m$  at client  $k$  during round  $t$ . The value of  $\tilde{z}_{k,n}^{t,m}$  is determined based on the availability status of the sample, as outlined below:

$$\tilde{z}_{k,n}^{t,m} = \begin{cases} \theta^m(x_{k,n}^{t,m}) & \text{if } p_{k,n}^{t,m} = 1 \\ \bar{v}_{c(k,n)}^{t,m} & \text{if } p_{k,n}^{t,m} = 0 \end{cases} \quad (19)$$

In this context,  $c(k, n)$  denotes the class label associated with the  $n$ -th data sample of client  $k$ . The predicted label is then obtained by applying the head encoder  $\theta^0$  to the aggregated feature representations across all modalities:

$$\theta^0 \left( Z_k^{t,1}, \dots, \tilde{Z}_k^{t,m}, \dots, Z_k^{t,M} \right) \quad (20)$$

After this step, the standard MMO-FL training procedure, as outlined in the previous section, is carried out. This approach effectively mitigates the impact of modality quantity imbalance by aligning with the theoretical objective of increasing  $N_{min}$ , as discussed in our earlier analysis.

### C. Prototypical Quality Rebalancing

In this phase, the client will execute the PLR algorithm. The core idea of the PLR algorithm is to alleviate modality quality imbalance by refining the representations of lower-quality modalities. This is achieved by aligning these representations toward their expected semantic positions, guided by the relative distances among modality-specific prototypes. This prototype-based alignment facilitates effective quality correction across modalities. The overall mechanism of PLR is illustrated in Step 3 of Fig. 2. Inspired by prior work [31], we

introduce the Prototype Cross Entropy (PCE), a loss function that captures the discrepancy between the feature representation of a given modality and its corresponding prototype in each global round. The PCE is formally defined as follows:

$$\mathcal{L}_{PCE}^m(f_t) = \mathbb{E} \left[ -\log \frac{\exp(-d(\tilde{z}_{k,n}^{t,m}, \bar{v}_{c(k,n)}^{t,m}))}{\sum_{c=1}^C \exp(-d(\tilde{z}_{k,n}^{t,m}, \bar{v}_c^{t,m}))} \right] \quad (21)$$

Here,  $d(\cdot, \cdot)$  denotes the distance function, which is instantiated as the Euclidean distance. The Prototypical Quality Rebalancing loss, denoted by  $\mathcal{L}_{PLR}$ , is defined as a weighted combination of the standard cross-entropy (CE) loss and PCE loss. It is given as follows:

$$\mathcal{L}_{PLR} = \mathcal{L}_{CE} + \beta(1 - q_k^{t,m})\mathcal{L}_{PCE}^m \quad (22)$$

In this formulation,  $\beta$  is a hyper-parameter to control the degree of modulation. The indicator variable  $q_k^{t,m}$  represents the quality status of modality  $m$  for client  $k$  at round  $t$ . If the modality is identified as low quality,  $q_k^{t,m} = 0$ , the loss function incorporates the PCE term to reduce the impact of the degraded data. Otherwise, when  $q_k^{t,m} = 1$ , no adjustment is applied. Each client dynamically adjusts its training loss based on the quality status of its local modalities.

The steps described above form the core components of the proposed QQR algorithm. In the following, we present additional considerations and potential extensions aimed at enhancing the algorithm's effectiveness.

**Discussion on  $\rho_{\max}^t$ :** After applying the PLR algorithm to address modality quality imbalance, we proceed with a theoretical analysis to characterize the evolution of  $\rho_{\max}^t$  over time under the influence of the proposed PLR algorithm. This leads to the following proposition:

**Proposition 1.** *Under the previously stated assumptions, the upper bound of the gradient noise, denoted as  $\rho_{\max}^t$ , is expected to decrease monotonically with the progression of global round  $t$ , as:  $\mathbb{E}[\rho_{\max}^{t+1}] < \mathbb{E}[\rho_{\max}^t]$ .*

*Proof.* The detail proof can be found in Appendix G.  $\square$

The above proposition provides theoretical validation for the effectiveness of the proposed algorithm and enhances the depth of our theoretical analysis.

**Communication Efficiency Strategy:** Within the QQR algorithm, the cumulative global prototype functions as an approximate representation that captures the most salient and discriminative features of the original data. However, high numerical precision may not be necessary for this purpose. To further reduce communication costs, quantization techniques can be employed when transmitting the local prototypes. In our experimental evaluation, we assess the impact of this strategy on learning performance under both modality quantity and quality imbalance scenarios.

## VII. EXPERIMENT

In this section, we will experimentally evaluate the performance of the MMO-FL algorithm under conditions of modality quantity and quality imbalance. All experiments were carried out on a workstation running Ubuntu 18.04, featuring an Intel Core i7-10700KF CPU at 3.8 GHz and an NVIDIA GeForce RTX 3070 GPU. The detailed experimental settings are described below.

### A. Datasets

To simulate the MMO-FL setting in Internet of Things scenarios, we utilize two real-world multimodal datasets, namely UCI-HAR and MVSA-Single. These datasets differ in their modality types, as the UCI-HAR dataset consists of sensor signal data, while the MVSA-Single dataset includes paired image and text content. These modality differences naturally lead to variations in experimental results. A detailed description of each dataset is provided below.

**UCI-HAR:** The UCI-HAR dataset is a well-established benchmark for human activity recognition. It contains a total of 10,299 samples collected from 30 individuals (average age 24), each performing one of six predefined activities: walking, walking downstairs, walking upstairs, standing, sitting, and lying down. Data were recorded using smartphone-based inertial sensors, including a 3-axis accelerometer and gyroscope, capturing motion signals in three dimensions. The sensors sampled at a rate of 50 Hz, producing 128 readings per axis within each time window. This dataset is used in our experiments to investigate sensor-based human activity recognition using 3D motion data.

**MVSA-Single:** The MVSA-Single dataset is a benchmark resource widely used in multimodal sentiment analysis, emphasizing the joint modeling of textual and visual information from social media posts. It contains 5,129 samples, each comprising an image paired with accompanying text. Every image-text pair is annotated with one sentiment categories: Positive, Neutral, or Negative, representing the overall emotional expression conveyed through the combined modalities.

### B. Online Data Generation

In the context of online learning, the training dataset must evolve over time, with new data acquired at the beginning of each global round. To ensure that the model receives adequate data for effective training, an initial dataset is gathered at the outset of the learning process. Given the differences in size and characteristics between two datasets, we adopt dataset-specific strategies for simulating online data generation in each case.

**UCI-HAR:** For the UCI-HAR dataset, the training setup includes five clients. Each client is initially allocated 2000 samples, drawn from a Dirichlet distribution with ratio  $\alpha$  to simulate non-IID long-term data sources. During each global round, clients maintain an evolving local dataset consisting of 500 samples. To mimic online data collection, 20 new samples are added from the long-term pool at every round, while the oldest 20 samples are simultaneously discarded. This continuous replacement strategy ensures that local data remains dynamic.

**MVSA-Single:** For the MVSA-Single dataset, the training setup consists of five clients. Each client is initially provided with 1500 samples drawn using a Dirichlet distribution with ratio  $\alpha$  to reflect a non-IID long-term data distribution. During training, each client maintains an evolving online dataset of 800 samples. At each global round, 20 new samples are added from the long-term data pool, and the oldest 20 samples are removed, ensuring a constant dataset size. This streaming update strategy supports the dynamic nature of online learning by continuously refreshing the local data.

### C. Imbalance Simulation

**Quantity Imbalance Simulation:** To evaluate the impact of modality quantity imbalance, we introduce two parameters:  $\lambda_p$ , named intra-round quantity imbalance ratio, which denotes the proportion of missing data relative to the total dataset, and  $\lambda_r$ , named inter-round quantity imbalance ratio, which represents the proportion of global rounds affected by quantity imbalance relative to the total number of rounds. A higher  $\lambda_p$  indicates a greater extent of data loss during the collection process, while a higher  $\lambda_r$  reflects a more frequent occurrence of quantity imbalance throughout the training period. To isolate this effect and eliminate confounding factors, the experimental design includes the following constraints. First, we assume that all clients experience the same value of  $\lambda_p$  across all global rounds. Second, given that the datasets used in the experiments consist of two modalities, we restrict each instance of missing data to affect only one modality per modality pair. In the following experiments, we systematically examine the impact of varying degrees of modality quantity imbalance on learning performance by controlling the value of  $\lambda_p$  and  $\lambda_r$ .

**Quality Imbalance Simulation:** To evaluate the impact of modality quality imbalance, we introduce the parameter inter-round quality imbalance ratio  $\delta_r$  to represent the proportion of rounds in which the client collects low-quality data out of the total number of learning rounds. A higher value of  $\delta_r$  indicates that a client collects low-quality data in a greater number of global rounds. To control variables and isolate the effect of modality quality imbalance, the experimental

design follows two rules. First, all clients are assigned the same  $\delta_r$  value across training rounds, and second since the dataset contains only two modalities, at most one modality is permitted to collect low-quality data in any given round. In the subsequent experiments, we systematically investigate how varying degrees of modality quality imbalance affect learning performance by adjusting the value of  $\delta_r$ .

#### D. Model Details

In the following, we detail the model architectures and key parameters used in our experiments, presented separately for the two datasets.

**UCI-HAR:** The dataset includes two distinct modalities: accelerometer and gyroscope signals, each requiring a dedicated encoder architecture. For the accelerometer data, we use a CNN-based model as the encoder, featuring five convolutional layers ending with a fully connected layer. For the gyroscope data, we use an LSTM-based model with one LSTM layer and one fully connected layer. Both encoders produce 128-dimensional feature vectors. A shared classification head, consisting of two fully connected layers, is used to generate the final predictions. The learning rate is 0.08, with a decay factor of 0.95 until it reaches 0.001. The batch size is 128.

**MVSA-Single:** The dataset includes two distinct modalities: text and image data, which require separate encoder networks tailored to their data types. For the image data, a four-layer CNN is utilized, with its final layer adjusted to yield a 128-dimensional embedding. For text data, a two-layer LSTM network is employed, also configured to output a 128-dimensional feature vector. The modality-specific features are subsequently processed by a shared header model consisting of two fully connected layers. The learning rate initialized at 0.01, which exponentially decays by a factor of 0.99 until it reaches a minimum of 0.001. The batch size used throughout is 128.

#### E. Benchmarks

In our experimental evaluation, we consider several baseline methods to benchmark the performance of our proposed approach. Since the challenges of modality quality and quantity imbalance in the context of MMO-FL remain unexplored, there are no well-established baselines available for direct comparison. To bridge this gap, we incorporate both optimistic and pessimistic bounds on learning performance. A detailed explanation of each benchmark is provided below.

For validating the utility of the PNR algorithm, we consider the following three benchmarks for comparison.

**Full Collection (FC).** In this scenario, each client is able to successfully collect complete modality information for all data samples, resulting in no quantity imbalance. Consequently, the entire dataset can be fully utilized for training, serving as an upper bound for achievable learning performance.

**Incomplete Subset (IS).** In this scenario, each client collects modality-specific data based on its actual operating conditions. Due to variations in sensor performance, some data samples may contain only partial modality information, leading to modality quantity imbalance. As a result, training can only be performed on the subset of data that contains all

modalities after modality alignment, which serves as a lower bound for evaluating learning performance.

**Zero Padding (ZP).** In this scenario, to maximize data utilization based on the incomplete set, we apply zero-padding to compensate for the missing modality components. This allows training to be conducted using the available modality information from all collected data samples.

**Average Imputation (AI).** In this scenario, to handle missing modality components in incomplete samples, we adopt Average Imputation as a widely used baseline method [39]. Specifically, the missing entries in a given modality are filled with the empirical mean value computed from the available samples of that modality, which enables all data points to participate in training without introducing modality quantity imbalance. This imputation strategy maintains both dataset size and computational efficiency by replacing absent values with a stable statistical estimate.

To assess the effectiveness of the PLR algorithm, we compare it against the following two benchmarks.

**Pristine Quality (PQ).** In this scenario, all modality-specific sensors on the client operate reliably, consistently collecting normal quality data without exhibiting temporal variations in data quality.


**Bare Quality (BQ).** In this scenario, certain modality-specific sensors on the client produce low-quality data due to hardware limitations or malfunctions. Furthermore, this degraded data is not corrected or enhanced through any compensation methods. Consequently, the client proceeds with local model training using the degraded input.

**Median Aggregation (MA).** In this scenario, to mitigate the adverse impact of modality quality imbalance on global learning, we adopt Median Aggregation as a robust aggregation strategy [40]. Specifically, in each global round, the server updates the global model by taking the coordinate-wise median of the uploaded local model parameters instead of their average, thereby suppressing the influence of clients whose updates are severely distorted by low-quality modality data.

#### F. Simulation Results

In this section, we assess the proposed MMO-FL framework under modality quantity and quality imbalance scenarios. To enable controlled analysis and prevent variable cross-interference, we first examine the effect of quantity imbalance and assess the performance of the PNR algorithm. We then investigate the quality imbalance problem and evaluate the PLR algorithm. Next, we perform ablation studies to analyze performance under different parameter settings, considering the combined effects of quantity and quality imbalance. Finally, we assess the performance impact and communication efficiency achieved through quantized prototype uploads. The results represent averages computed over 10 random runs.


**Performance Comparison (Quantity Imbalance).** We begin by evaluating the learning performance with the impact of quantity imbalance within the MMO-FL scenario. The test accuracy comparison between proposed PNR and benchmarks are shown in Fig. 3(a) and Fig. 3(b) based on UCI-HAR dataset with configuration  $[\lambda_r = 0.5, \lambda_p = 0.9, \alpha = 1]$  and



(a) Test Accuracy with UCI-HAR (b) Test Accuracy with MVSA-Single

Fig. 3: Performance comparison of proposed algorithm and benchmarks with quantity imbalance.

MVSA-Single dataset with  $[\lambda_r = 0.5, \lambda_p = 0.9, \alpha = 5]$ , respectively. Based on the above figures, we get several key findings. First, the FC setting achieved the highest performance because it includes all data and does not experience quantity imbalance. In contrast, the IS setting yielded the lowest performance, indicating that quantity imbalance without mitigation can significantly degrade learning. Second, both the ZP and AI methods improve performance relative to IS, indicating that filling missing inputs with zeros or average values is more effective than discarding incomplete samples during training. Among the two, AI tends to achieve slightly better results than ZP because its imputed values fall within a more reasonable range, thereby introducing smaller distortions in gradient updates. Finally, the proposed PNR algorithm consistently outperforms both ZP and AI, since substituting missing entries with cumulative global prototypes yields stronger cross-modal alignment and leads to superior learning performance. The above findings are consistent across both datasets.




(a) Test Accuracy with UCI-HAR (b) Test Accuracy with MVSA-Single


Fig. 4: Performance comparison of proposed algorithm and benchmarks with quality imbalance.

**Performance Comparison (Quality Imbalance).** In this part, we will evaluate the learning performance with the impact of quality imbalance within the MMO-FL scenario. To emulate low-quality data conditions, we inject high-intensity additive white Gaussian noise (AWGN) with a SNR of 10 dB into the original dataset. The test accuracy comparison between proposed PLR and benchmarks are shown in Fig. 4(a) and Fig. 4(b) based on UCI-HAR dataset with configuration  $[\delta_r = 0.5, \beta = 0.5, \alpha = 1]$  and MVSA-Single dataset with  $[\delta_r = 0.5, \beta = 0.5, \alpha = 5]$ , respectively. Based on the simulation results, we can summarize the following findings. First, a comparison between PQ and BQ shows that adding substantial noise to multimodal data produces low-quality inputs that markedly degrade learning performance. Second, MA leverages median-based updates instead of mean-based ones, offering increased robustness against updates influenced by

low-quality data and thereby enhancing learning performance. However, its overall performance remains inferior to that of PLR. Finally, the proposed PLR algorithm effectively reduces the negative impact of low-quality data, with performance improvements becoming more pronounced as the number of training rounds increases. This improvement is primarily attributed to the PCE loss, which reduces the disparity between the feature extractor of low-quality data and cumulative global prototypes, thereby gradually diminishing the adverse effects of quality imbalance.



(a) Test Accuracy with UCI-HAR (Quantity Imbalance) (b) Test Accuracy with MVSA-Single (Quantity Imbalance)



(c) Test Accuracy with UCI-HAR (Quality Imbalance) (d) Test Accuracy with MVSA-Single (Quality Imbalance)

Fig. 5: Performance comparison under online vs. offline prototype updates.

**Performance Comparison (Online Updates).** In this part, we will evaluate the learning performance with the impact of online and offline prototype update strategies within the MMO-FL scenario. To evaluate the effectiveness of online prototype updates, we simulate two offline strategies: Offline Quantity Rebalancing (OFNR) and Offline Quality Rebalancing (OFLR). In both cases, the prototype is initialized only once using a static dataset during the first round and remains unchanged throughout the training process. This fixed prototype is applied consistently during the subsequent modality quantity and quality rebalancing procedures. We compare the performance of these offline strategies against our online prototype update approach. The test accuracy results are presented in Fig. 5(a)(c) and Fig. 5(b)(d), based on the UCI-HAR and MVSA-Single datasets, respectively, following the previously described experimental configuration. Across both datasets and under both types of imbalance, models using offline prototype updates consistently lag behind those employing online updates in terms of test accuracy. This emphasizes the necessity of developing algorithms tailored for online settings.

Based on the above performance comparisons, which confirm the effectiveness of the proposed OGPC, PNR and PLR algorithms under the MMO-FL setting, we next conduct a

series of ablation studies to evaluate the specific impact of key parameters on learning performance.




Fig. 6: Performance comparison of proposed algorithm and benchmarks with different  $\lambda_p$  ratio.

**Impact of Ratio  $\lambda_p$  (Quantity Imbalance).** In this part, we will explore the effects with the intra-round quantity imbalance ratio  $\lambda_p$  on learning performance. The results for the UCI-HAR and MVSA-Single datasets are presented in Fig. 6(a) and Fig. 6(b), respectively. Several key observations can be made: First, as  $\lambda_p$  increases, learning performance improves because more data is available and fewer quantity imbalances need to be compensated by the PNR algorithm. In contrast, when  $\lambda_p$  decreases substantially (e.g.,  $\lambda_p = 0.5$ ), the MVSA-Single dataset shows a pronounced drop in performance. This decline results from severe data loss, where even with compensation, achieving the desired results becomes difficult. The degree of degradation varies across datasets.




Fig. 7: Performance comparison of proposed algorithm and benchmarks with different  $\lambda_r$  ratio.

**Impact of Ratio  $\lambda_r$  (Quantity Imbalance).** In this part, we will explore the effects with the inter-round quantity imbalance ratio  $\lambda_r$  on learning performance. The results for the UCI-HAR and MVSA-Single datasets are presented in Fig. 7(a) and Fig. 7(b), respectively. We have the following findings: First, smaller values of  $\lambda_r$  yield better learning performance because the proportion of rounds with quantity imbalance is lower, allowing more rounds to collect complete modal data. When  $\lambda_r$  reaches 0.7, both datasets exhibit a significant drop in learning performance. This decline is due to the high frequency of quantity imbalance, where even with compensation, sustaining strong learning performance remains difficult. This phenomenon is consistently observed across both datasets.

**Impact of Ratio  $\delta_r$  (Quality Imbalance).** In this part, we will explore the effects with the inter-round quality imbalance ratio  $\delta_r$  on learning performance. The results for the UCI-HAR and MVSA-Single datasets are presented in Fig. 8(a) and Fig. 8(b), respectively. We observe that learning performance decreases as  $\delta_r$  increases. When  $\delta_r$  is smaller, fewer data




Fig. 8: Performance comparison of proposed algorithm and benchmarks with different  $\delta_r$  ratio.

samples are affected by quality imbalance. This not only increases the number of updates to the cumulative global prototypes but also reduces the number of rounds in which the PCE is used to mitigate quality imbalance. Together, these factors contribute to improved learning performance.

Next, we assess the efficiency of the proposed PNR and PLR algorithms by applying quantized upload strategies to further reduce communication overhead and enhance their practicality. We evaluate performance under quantity imbalance and quality imbalance scenarios separately.




Fig. 9: Performance evaluation of the proposed algorithm under varying quantization levels and quantity imbalance. (c) Test Accuracy with UCI-HAR by communication cost. (d) Test Accuracy with MVSA-Single by communication cost.

samples are affected by quality imbalance. This not only increases the number of updates to the cumulative global prototypes but also reduces the number of rounds in which the PCE is used to mitigate quality imbalance. Together, these factors contribute to improved learning performance.

**Impact of Quantized Upload (Quantity Imbalance).** In this part, we examine the effect of applying quantization to local prototypes during the PNR upload process as a means of further reducing communication overhead. We employ a uniform scalar quantizer, where the parameter  $b$  represents the number of bits per component used for compression. Without quantization,  $b = 32$ , corresponding to full-precision transmission. Applying quantization with  $b$  bits results in  $2^b$  quantization levels. We evaluate the performance impact for  $b \in [2, 4]$ . The results, as illustrated in Fig. 9(a) and Fig. 9(b) by round, and Fig. 9(c) and Fig. 9(d) by communication cost, illustrate the trade-off between model performance and communication efficiency. As expected, larger  $b$  values yield better performance but incur higher communication costs.

Notably, prototypes remain effective at low precision, and quantization leads to only a minor performance drop while substantially reducing communication overhead. This confirms that PNR can be efficiently integrated into the MMO-FL framework without imposing a significant communication burden. Similarly, the PLR results, presented in Fig. 10(a) and Fig. 10(b) by round, and Fig. 10(c) and Fig. 10(d) by communication cost, exhibit the same trade-off pattern. As with PNR, low-precision quantization in PLR achieves substantial communication savings with minimal performance degradation, leading to the same simulation conclusion.

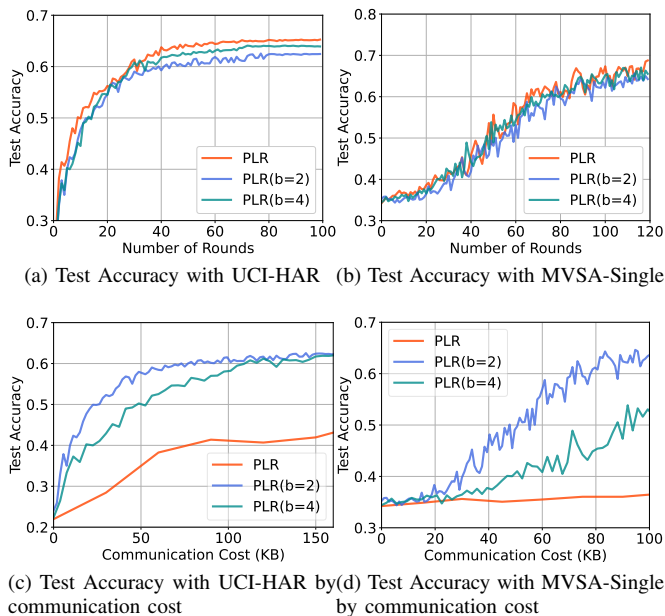


Fig. 10: Performance evaluation of the proposed algorithm under varying quantization levels and quality imbalance.

## VIII. CONCLUSION

In this work, we investigate the MMO-FL framework, which supports distributed multimodal learning with online data collection on edge devices. We focus on non-ideal conditions where modality quantity and quality imbalance arise due to variations in device performance during data acquisition. To address these challenges, we propose the QQR algorithms, designed to mitigate quantitative and qualitative modality imbalances, respectively. The proposed approach is supported by comprehensive theoretical analysis and validated through extensive experimental evaluations. For future work, we aim to advance the theoretical foundations of MMO-FL by exploring additional challenges to further enrich research in this domain. From a theory perspective, data arrives in a streaming fashion, and concept drift frequently occurs due to factors such as sensor recalibration, operational shifts, or evolving cross-modal semantics. In such cases, simple round-wise prototype averaging is insufficient, as it fails to distinguish between transient noise and genuine semantic changes, potentially leading to prototype contamination and degraded performance. This highlights the necessity of robust, advanced prototype update mechanisms specifically designed for online scenarios. From a practical perspective, we plan to develop

a real-world testbed leveraging IoT edge devices for actual data collection and distributed and online model training. This deployment will help uncover additional issues related to modality quantity and quality imbalance that may not be evident in simulation environments, thereby enabling further refinement and optimization of the proposed QQR algorithms.

## REFERENCES

- [1] L. Atzori, A. Iera, and G. Morabito, "The internet of things: A survey," *Computer networks*, vol. 54, no. 15, pp. 2787–2805, 2010.
- [2] H. Wang, J. Bian, and J. Xu, "On the local cache update rules in streaming federated learning," *IEEE Internet of Things Journal*, vol. 11, no. 6, pp. 10 808–10 816, 2023.
- [3] T. Wu, F. Wu, C. Qiu, J.-M. Redouté, and M. R. Yuce, "A rigid-flex wearable health monitoring sensor patch for iot-connected healthcare applications," *IEEE Internet of Things Journal*, vol. 7, no. 8, pp. 6932–6945, 2020.
- [4] H. Wang, X. Zhong, K. Liu, F. Liu, and W. Zhang, "Denosing and adaptive online vertical federated learning for sequential multi-sensor data in industrial internet of things," *IEEE Transactions on Mobile Computing*, 2025.
- [5] J. Ma, L. Li, and C. Xu, "Autors: Environment-dependent real-time scheduling for end-to-end autonomous driving," *IEEE Transactions on Parallel and Distributed Systems*, vol. 34, no. 12, pp. 3238–3252, 2023.
- [6] F. Al-Turjman and S. Alturjman, "Context-sensitive access in industrial internet of things (iiot) healthcare applications," *IEEE Transactions on Industrial Informatics*, vol. 14, no. 6, pp. 2736–2744, 2018.
- [7] T.-H. Wang, P.-N. Chen, and Y.-C. Huang, "Probabilistic byzantine attack on federated learning," *IEEE Transactions on Signal Processing*, 2025.
- [8] W. Y. B. Lim, N. C. Luong, D. T. Hoang, Y. Jiao, Y.-C. Liang, Q. Yang, D. Niyato, and C. Miao, "Federated learning in mobile edge networks: A comprehensive survey," *IEEE Communications Surveys and Tutorials*, vol. 22, no. 3, pp. 2031–2063, 2020.
- [9] J. Bian, L. Wang, K. Yang, C. Shen, and J. Xu, "Accelerating hybrid federated learning convergence under partial participation," *IEEE Transactions on Signal Processing*, vol. 72, pp. 3258–3271, 2024.
- [10] S. Salgia, T. Gabay, Q. Zhao, and K. Cohen, "A communication-efficient adaptive algorithm for federated learning under cumulative regret," *IEEE Transactions on Signal Processing*, vol. 72, pp. 735–743, 2024.
- [11] Z. Huang, X. Xu, J. Ni, H. Zhu, and C. Wang, "Multimodal representation learning for recommendation in internet of things," *IEEE Internet of Things Journal*, vol. 6, no. 6, pp. 10 675–10 685, 2019.
- [12] L. Che, J. Wang, Y. Zhou, and F. Ma, "Multimodal federated learning: A survey," *Sensors*, vol. 23, no. 15, p. 6986, 2023.
- [13] L. P. Le, T. Nguyen, M. A. Riegler, P. Halvorsen, and B. T. Nguyen, "Multimodal missing data in healthcare: A comprehensive review and future directions," *Computer Science Review*, vol. 56, p. 100720, 2025.
- [14] Q. Zhang, Y. Wei, Z. Han, H. Fu, X. Peng, C. Deng, Q. Hu, C. Xu, J. Wen, D. Hu *et al.*, "Multimodal fusion on low-quality data: A comprehensive survey," *arXiv preprint arXiv:2404.18947*, 2024.
- [15] H. Wang, X. Liu, X. Zhong, L. Chen, F. Liu, and W. Zhang, "Multimodal online federated learning with modality missing in internet of things," *IEEE Transactions on Mobile Computing*, 2025.
- [16] X. Ouyang, Z. Xie, H. Fu, S. Cheng, L. Pan, N. Ling, G. Xing, J. Zhou, and J. Huang, "Harmony: Heterogeneous multi-modal federated learning through disentangled model training," in *Proceedings of the 21st Annual International Conference on Mobile Systems, Applications and Services*, 2023, pp. 530–543.
- [17] H. Chen, Y. Zhang, D. Krompass, J. Gu, and V. Tresp, "Feddat: An approach for foundation model finetuning in multi-modal heterogeneous federated learning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 10, 2024, pp. 11 285–11 293.
- [18] J. Chen and A. 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*, 2022, pp. 87–96.
- [19] J. Bian, L. Wang, and J. Xu, "Prioritizing modalities: Flexible importance scheduling in federated multimodal learning," *arXiv preprint arXiv:2408.06549*, 2024.
- [20] M. Ma, J. Ren, L. Zhao, S. Tulyakov, C. Wu, and X. Peng, "Smil: Multimodal learning with severely missing modality," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 3, 2021, pp. 2302–2310.

- [21] C. Zhang, X. Chu, L. Ma, Y. Zhu, Y. Wang, J. Wang, and J. Zhao, "M3care: Learning with missing modalities in multimodal healthcare data," in *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2022, pp. 2418–2428.
- [22] H. Maheshwari, Y.-C. Liu, and Z. Kira, "Missing modality robustness in semi-supervised multi-modal semantic segmentation," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2024, pp. 1020–1030.
- [23] H. Q. Le, C. M. Thwal, Y. Qiao, Y. L. Tun, M. N. Nguyen, E.-N. Huh, and C. S. Hong, "Cross-modal prototype based multimodal federated learning under severely missing modality," *Information Fusion*, p. 103219, 2025.
- [24] W. Rang, D. Yang, D. Cheng, and Y. Wang, "Data life aware model updating strategy for stream-based online deep learning," *IEEE Transactions on Parallel and Distributed Systems*, vol. 32, no. 10, pp. 2571–2581, 2021.
- [25] S. Hong and J. Chae, "Communication-efficient randomized algorithm for multi-kernel online federated learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 12, pp. 9872–9886, 2021.
- [26] Y. Lin, W. Li, J. Song, and X. Li, "Online federated reproduced gradient descent with time-varying global optima," *IEEE Transactions on Signal Processing*, 2025.
- [27] D. Kwon, J. Park, and S. Hong, "Tighter regret analysis and optimization of online federated learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [28] A. Mitra, H. Hassani, and G. J. Pappas, "Online federated learning," in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 4083–4090.
- [29] H. Wang and J. Xu, "Online vertical federated learning for cooperative spectrum sensing," *IEEE Transactions on Cognitive Communications and Networking*, 2024.
- [30] Q. Zhang, H. Wu, C. Zhang, Q. Hu, H. Fu, J. T. Zhou, and X. Peng, "Provable dynamic fusion for low-quality multimodal data," in *International conference on machine learning*. PMLR, 2023, pp. 41 753–41 769.
- [31] Y. Fan, W. Xu, H. Wang, J. Wang, and S. Guo, "Pmr: Prototypical modal rebalance for multimodal learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 20 029–20 038.
- [32] Y. Wei, S. Li, R. Feng, and D. Hu, "Diagnosing and re-learning for balanced multimodal learning," in *European Conference on Computer Vision*. Springer, 2024, pp. 71–86.
- [33] Y. Fan, W. Xu, H. Wang, F. Huo, J. Chen, and S. Guo, "Overcome modal bias in multi-modal federated learning via balanced modality selection," in *European Conference on Computer Vision*. Springer, 2024, pp. 178–195.
- [34] Y. Fan, W. Xu, H. Wang, J. Zhu, and S. Guo, "Balanced multi-modal federated learning via cross-modal infiltration," *arXiv preprint arXiv:2401.00894*, 2023.
- [35] M. Yang, J. Xu, W. Ding, and Y. Liu, "Fedhap: federated hashing with global prototypes for cross-silo retrieval," *IEEE Transactions on Parallel and Distributed Systems*, vol. 35, no. 4, pp. 592–603, 2023.
- [36] Y. Tan, G. Long, L. Liu, T. Zhou, Q. Lu, J. Jiang, and C. Zhang, "Fedproto: Federated prototype learning across heterogeneous clients," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 36, no. 8, 2022, pp. 8432–8440.
- [37] Y. Tan, G. Long, J. Ma, L. Liu, T. Zhou, and J. Jiang, "Federated learning from pre-trained models: A contrastive learning approach," *Advances in neural information processing systems*, vol. 35, pp. 19 332–19 344, 2022.
- [38] P.-L. Hsu and H. Robbins, "Complete convergence and the law of large numbers," *Proceedings of the national academy of sciences*, vol. 33, no. 2, pp. 25–31, 1947.
- [39] C. Zhang, Y. Cui, Z. Han, J. T. Zhou, H. Fu, and Q. Hu, "Deep partial multi-view learning," *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 5, pp. 2402–2415, 2020.
- [40] M. Fang, S. Nabavirazavi, Z. Liu, W. Sun, S. S. Iyengar, and H. Yang, "Do we really need to design new byzantine-robust aggregation rules," *Network and Distributed System Security Symposium (NDSS)*, 2025.