

# FedMultimodal: A Benchmark For Multimodal Federated Learning

Tiantian Feng  
University of Southern California  
Los Angeles, CA, USA  
tiantianf@usc.edu

Rajat Hebbar  
University of Southern California  
Los Angeles, CA, USA  
rajatheb@usc.edu

Mi Zhang  
The Ohio State University  
Columbus, OH, USA  
mizhang.1@osu.edu

Digbalay Bose  
University of Southern California  
Los Angeles, CA, USA  
dbose@usc.edu

Anil Ramakrishna  
Amazon Alexa AI  
Los Angeles, CA, USA  
aniramak@amazon.com

Salman Avestimehr  
University of Southern California  
Los Angeles, CA, USA  
avestime@usc.edu

Tuo Zhang  
University of Southern California  
Los Angeles, CA, USA  
tuozhang@usc.edu

Rahul Gupta  
Amazon Alexa AI  
Boston, MA, USA  
gupra@amazon.com

Shrikanth Narayanan  
University of Southern California  
Los Angeles, CA, USA  
shri@sipi.usc.edu

## ABSTRACT

Over the past few years, Federated Learning (FL) has become an emerging machine learning technique to tackle data privacy challenges through collaborative training. In the Federated Learning algorithm, the clients submit a locally trained model, and the server aggregates these parameters until convergence. Despite significant efforts that have been made to FL in fields like computer vision, audio, and natural language processing, the FL applications utilizing multimodal data streams remain largely unexplored. It is known that multimodal learning has broad real-world applications in emotion recognition, healthcare, multimedia, and social media, while user privacy persists as a critical concern. Specifically, there are no existing FL benchmarks targeting multimodal applications or related tasks. In order to facilitate the research in multimodal FL, we introduce FedMultimodal, the first FL benchmark for multimodal learning covering five representative multimodal applications from ten commonly used datasets with a total of eight unique modalities. FedMultimodal offers a systematic FL pipeline, enabling end-to-end modeling framework ranging from data partition and feature extraction to FL benchmark algorithms and model evaluation. Unlike existing FL benchmarks, FedMultimodal provides a standardized approach to assess the robustness of FL against three common data corruptions in real-life multimodal applications: missing modalities, missing labels, and erroneous labels. We hope that FedMultimodal can accelerate numerous future research directions, including designing multimodal FL algorithms toward extreme data heterogeneity, robustness multimodal FL, and efficient multimodal FL. The datasets and benchmark results can be accessed at: <https://github.com/usc-sail/fed-multimodal>.

## CCS CONCEPTS

• **Computing methodologies** → **Distributed computing methodologies**.

## KEYWORDS

Federated Learning, Multimodal Learning, Multimodal Benchmark

### ACM Reference Format:

Tiantian Feng, Digbalay Bose, Tuo Zhang, Rajat Hebbar, Anil Ramakrishna, Rahul Gupta, Mi Zhang, Salman Avestimehr, and Shrikanth Narayanan. 2023. FedMultimodal: A Benchmark For Multimodal Federated Learning. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23)*, August 6–10, 2023, Long Beach, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3580305.3599825>

## 1 INTRODUCTION

With rapid advances in machine learning (ML) [39] in the past decade, modern mobile devices and wearable sensors [7, 52] have revolutionized applications and services in industries ranging from entertainment and transportation to healthcare and defense, significantly changing how people live, work, and interact with each other. These intelligent sensing devices, equipped with sensors of multiple modalities, can capture diverse information about a user, including but not limited to physiological, emotional, and rich spatiotemporal contextual information [4, 6, 17, 22, 54]. These data records are typically transmitted to remote servers for centralized training of the ML models. However, the collection of human-centered data raises significant concerns about compromising user privacy due to association with sensitive environments and contexts, ranging from homes, workplaces, and business meetings to hospitals and schools [56]. Therefore, it is critical to ensure that modern ML systems can protect user privacy by preventing any unauthorized access to data [19, 49].

In response to this, ML practitioners have developed Federated Learning as an alternative paradigm to build models, without the need to transfer user data from the edge devices [37]. Unlike centralized training, models are trained locally using locally stored



This work is licensed under a Creative Commons Attribution-ShareAlike International 4.0 License.

**Table 1: Experimental considerations between FedMultimodal and existing multimodal FL studies.**

	Missing Modalities	Erroneous Labels	Missing Labels
SSCL [59]	✗	✗	✓
MMFed [74]	✗	✗	✗
FedMsplit [10]	✓	✗	✗
CreamFL [77]	✓	✗	✗
<b>FedMultimodal (Ours)</b>	✓	✓	✓

data, and updated parameters are transmitted to the server instead of raw data. FL allows clients to train a model collaboratively without sharing their local data, making it one of the most emerging privacy-enhancing learning algorithms in ML research [33].

Previous works in FL have primarily focused on designing robust and efficient algorithms for federated model training. *FedAvg* [46] was the earliest FL optimization algorithm to train the model in the distribution mechanism. In *FedAvg*, each client executes local model updates before submitting the updates to the server. Even though *FedAvg* offers possibilities for deploying FL in the wild, it often encounters slow convergence as a consequence of gradient drifting from data heterogeneity. As such, researchers have proposed algorithms such as stochastic Controlled Averaging Algorithm (SCAFFOLD) [35] and FedProx [40] to minimize the impact of gradient drift for heterogeneous data. For example, SCAFFOLD accelerates the training speed through control variates which prevent the client gradients from drifting away from the global optima. Similarly, [57] introduced adaptive optimization algorithms, FedOpt, that allow server optimization through momentum.

To facilitate FL research in more diverse problem domains, a number of FL benchmarks have been developed in the past few years. For example, LEAF [8] was the earliest FL benchmark which includes multiple FL training tasks on 5 datasets covering various computer vision and NLP tasks. FedML [28], besides providing an open-source library and a platform for federated learning deployment, it includes multiple FL benchmarks on computer vision and health [29], data mining [27], IoT [80], and NLP [43]. More recently, [38] announced a multi-domain FL benchmark called FedScale. FedScale included implementations with 20 realistic FL datasets mainly in computer vision and natural language processing applications. [14] introduced an FL simulation tool named FLUTE, which covers the application of CV, NLP, and audio tasks. Meanwhile, [79] presented an audio-centric federated learning framework, FedAudio, which focused on speech emotion recognition, keyword spotting, and audio event classification. Further, FLamby [66] is a recently proposed FL benchmark for a wide range of healthcare applications such as identifying lung nodules and predicting death risks. FederatedScope [73] also incorporates various benchmarks for federated learning in CV, NLP, and data mining [71].

**Existing Multimodal Federated Learning Works:** While existing FL benchmarks largely focus on unimodal applications such as computer vision (CV), natural language processing (NLP), and speech recognition, a significant number of real-world applications are associated with multimodal data streams. As listed in Table 1, [59] was one of the earliest to investigate FL using multi-sensory

data. They proposed a self-supervised learning approach called Scalogram-signal Correspondence Learning (SSCL) to learn robust multi-modal representations in FL. More recently, [74] designed a multi-modal FL framework named MMFed using the cross-attention mechanism. Moreover, [10] proposed an FL framework called FedM-Split that targeted the issue of missing modalities in the multimodal setup. CreamFL [77] provides a multi-modal FL framework using contrastive representation-level ensemble to learn a larger server model from heterogeneous clients across multi-modalities. However, existing multimodal FL frameworks perform their evaluation using their defined experimental setups, thus making it challenging for researchers to compare their methods with existing state-of-the-art fairly and effectively.

**Our Contributions:** In this work, we introduce FedMultimodal, a FL benchmark for multimodal applications. We summarize our key contributions as follows:

- FedMultimodal includes **ten representative datasets** covering **five diverse application scenarios** – emotion recognition, multimodal action recognition, human activity recognition, healthcare, and social media – that are well aligned with FL. We present systematic benchmark results on the above datasets to facilitate researchers to fairly compare their algorithms.
- To help the community accurately compare performance and ensure **reproducibility**, FedMultimodal presents an **open-source end-to-end FL simulation framework** and includes capabilities to perform data partitioning, feature processing, and multimodal training. FedMultimodal offers support for several popular FL optimizers including FedAvg [46], FedProx [40], FedrS [41], SCAFFOLD [35], and FedOpt [57], and provide flexibility that allows users to customized the trainers on the included datasets. The source codes and user guides are available at <https://github.com/usc-sail/fed-multimodal>.
- In addition to ensuring accessibility and reproducibility, the benchmark provides a **robustness assessment module** that allows researchers to simulate challenges uniquely tied to multimodal FL applications in real-world scenarios. As listed in Table 1, previous works on multimodal FL provide limited assessments of the robustness under real-world settings. Specifically, FedMultimodal emulates **missing modalities**, **missing labels**, and **erroneous labels** on top of the provided datasets to simulate scenarios when deploying FL systems in real-world settings. This is a crucial difference and a unique contribution of FedMultimodal compared to existing FL literature.

## 2 MULTIMODAL DATASETS AND TASKS

Table 2 provides an overview of the 10 datasets included in FedMultimodal. These 10 multimodal datasets cover five diverse tasks – Emotion Recognition, Multimedia Action Recognition, Human Activity Recognition, Healthcare, and Social Media classification. One important reason we select these 10 datasets is that they are publicly available, thus ensuring ease of accessibility and reproducibility. In this section, we provide a brief overview of each included dataset and the corresponding tasks.

**Table 2: Overview of the 10 datasets included in FedMultimodal.**

Task	Dataset	Partition	Client Num.	Modalities	Features	Metirc	Validation Protocol	Total Instance
ER	MELD	Natural	86	Audio, Text	MFCCs, MobileBert	UAR	Pre-defined	9,718
	CREMA-D	Natural	72	Audio, Video	MFCCs, MobileNetV2		5-Fold	4,798
MAR	UCF101	Synthetic	100	Audio, Video	MFCCs, MobileNetV2	Top1 Acc	Pre-defined	6,837
	MiT10	Synthetic	200	Audio, Video	MFCCs, MobileNetV2			41.6K
	MiT51	Synthetic	2000	Audio, Video	MFCCs, MobileNetV2			157.6K
HAR	UCI-HAR	Synthetic	105	Acc, Gyro	Raw	F1	Pre-defined	8,979
	KU-HAR	Natural	66	Acc, Gyro	Raw		5-Fold	10.3K
Health	PTB-XL	Natural	34	I-AVF, V1-V6	Raw	F1	Pre-defined	21.7K
SM	Hateful-Memes	Synthetic	50	Image, Text	MobileNetV2, MobileBert	AUC	Pre-defined	10.0K
	CrisisMMD	Synthetic	100		MobileNetV2, MobileBert	F1	Pre-defined	18.1K

## 2.1 Emotion Recognition (ER)

Emotion recognition (ER) has broad applicability of ER in virtual assistant-based tasks, human behavior analysis, and AI-assisted education, making it a valuable research topic in FL [21, 67]. FedMultimodal benchmark incorporates two widely used datasets in this category: MELD and CREMA-D.

**MELD** is a multiparty dialog dataset [55] containing over 9k utterances with audio and transcripts data from the Friends TV series. Due to the imbalanced label distribution in the dataset, we keep 4 emotions with the most samples i.e., neutral, happy, sad, and angry.

**CREMA-D** has 7,442 audio-visual clips recorded by 91 actors [9]. Each speaker was instructed to utter 12 sentences emulating 6 emotions: neutral, happy, anger, disgust, fear, and sad.

## 2.2 Multimodal Action Recognition (MAR)

The task of MAR consists of classifying a video into action categories based on underlying visual and audio modalities. In FedMultimodal, we include two well-known MAR testbeds: UCF101 and Moments in Time (MiT).

**UCF101** dataset [63] consists of 13,320 web videos with 101 sport-based action labels. However, data associated with only 51 labels are presented with video and audio information, resulting in less than 7,000 videos for the experiments. The duration of the videos ranges from several seconds to over 20 seconds. We subsample the video at the frame rate of 1Hz to reduce the computation overhead.

**Moments in Time (MiT)** is a large-scale MAR ( 1 million) dataset [50] with short (3 seconds) videos with overall list of 339 action labels. It is worth noting that MiT is a challenging dataset, with state-of-the-art top-1 accuracy close to 35% [58]. Given the inherent difficulty of this task, we tackle the easier classification problem by creating partitions of data with fewer distinct labels. We create two sub-datasets, MiT10 and MiT51, from the original MiT dataset. MiT10 and MiT51 contain videos of the 10 and the 51 most frequent labels. Similar to the UCF101 setting, we subsample the video every 10 frames to accommodate the computing constraints in FL.

## 2.3 Human Activity Recognition (HAR)

HAR identifies human actions based on wearable data such as accelerometers and gyroscopes. Due to the nature of its wearable-friendly attribute, it has become a prevalent research topic in FL

[61]. FedMultimodal provides the implementation on two HAR datasets: UCI-HAR and KU-HAR. In our experiments, we treat the accelerometer and gyroscope data as two different modalities.

**UCI-HAR** dataset [3] consists of smartphone sensors (Accelerometer and Gyroscope) data from 30 subjects (19-48 yrs old) performing six daily activities: walking, walking upstairs, walking downstairs, sitting, standing, laying. The participants wear smartphones on their waists during the collection phase. The accelerometer and gyroscope data are sampled at 50Hz.

**KU-HAR** is a recent human activity recognition dataset [62] collected with a group of 90 participants (75 male and 15 female) on 18 different activities. Instead of evaluating the 18 activities, we decided to keep 6 activities existing in the UCI-HAR dataset while adding jumping and running activities.

## 2.4 Healthcare

Healthcare ML applications have made immense progress in a range of domains e.g., heart-disease classification over the last decade. FedMultimodal explores the problem of ECG classification based on the PTB-XL [69] dataset. **PTB-XL** [69] includes over 20,000 clinical 12-lead ECG recordings from 18,885 patients for a multi-label classification task. There are 5 classes describing ECG diagnosis, including normal ECG, myocardial infarction, ST/T change, conduction disturbance, and hypertrophy. As suggested by [64], we use the ECG data provided at the sampling frequency of 100 Hz. We separate the readings from electrodes I, II, III, AVL, AVR, AVF, and V1-V6 as two modalities suggested by [2].

## 2.5 Social Media (SM)

Social media has become an increasingly important tool during disasters and emergencies for people to track the latest updates in the area, especially the impact (e.g., property damage, injury, and death) of the disaster, as well as urgent needs for help. However, the widespread adoption of social media has also drawn significant concerns about spreading misinformation, thus urging the need to identify and mitigate this misleading and harmful content. To accelerate FL research in this domain, FedMultimodal incorporates two social-media-based multimodal datasets related to hateful content and crisis information classification.

**Hateful Memes** dataset was released from the 2020 Hateful Memes Challenge [36] that focus on detecting hateful speech in memes.




Figure 1: The overall architecture of the end-to-end multimodal federated learning framework included in FedMultimodal.

The database includes 10,000 multimodal data with image and text pairs with binary classes.

**CrisisMMD** [1] comprises 18.1k tweets containing both paired visual and textual information. It collects relevant tweets from seven prominent natural disasters such as California Wildfires (2017). One of the purposes of the dataset is to identify the impact of the disaster like utility damage and injured or dead people.

### 3 END-TO-END MULTIMODAL FEDERATED LEARNING FRAMEWORK

To benchmark the performance of the multimodal datasets described in Section 2 as well as to support future research in the area of multimodal federated learning, we have built an end-to-end multimodal federated learning research framework. Figure 1<sup>1</sup> illustrates the overall architecture of the framework. As shown, our framework covers the complete pipeline of multimodal federated learning, which includes six key components: (1) non-IID data partitioning, (2) feature processing, (3) multimodal models, (4) fusion schemes, (5) federated optimizers, and (6) real-world noise factor emulator. In particular, one key difference between FedMultimodal and existing multimodal FL literature is that FedMultimodal takes the real-world noise factors into consideration and examines model robustness to three real-world noise factors: missing modalities, missing labels, and erroneous labels. In this section, we describe each of the six key components in detail.

#### 3.1 Non-IID Data Partitioning

The non-IID data partitioning is a fundamental step in emulating FL experiments. The first partition scheme is through the unique client identifier. For example, speech-related datasets, like CREMA-D and MELD, comprise speech-text or speech-visual data organized by speaker IDs. Hence, it is natural to use speaker IDs to partition the client data in FL, creating authentic non-IID data distributions. Similarly, we consider partitions in datasets like KU-HAR and PTB-XL comprise data with based on participant IDs and clinical site IDs, respectively. On the other hand, other multimodal datasets used in this paper, including MAR and SM datasets, do not have such realistic client partitions thus requiring ML practitioners to synthesize non-IID data distributions. Following prior works, we partition these datasets using Dirichlet distribution with  $\alpha \in \{0.1, 5.0\}$  to control the level of data heterogeneity, where  $\alpha = 0.1$  and  $\alpha = 5.0$  represents high heterogeneity and low heterogeneity, respectively. Although the original UCI-HAR datasets consist of data partitioned by participants, each participant performed the same amount of activities from each category, making the label distribution IID in UCI-HAR. Hence, we increase the heterogeneity of data distribution by dividing each participant’s data using the Dirichlet distribution.

#### 3.2 Feature Processing

Instead of training the model using raw input data like images and texts from scratch, FedMultimodal leverages well-established pre-trained models as backbone networks to extract features for the training downstream models. Unlike the centralized training paradigm, feature processing in federated learning benchmarks demands considerations in computation efficiency and feasibility.

<sup>1</sup>Figure 1 uses image sources from <https://openmoji.org/>




Figure 2: The architecture of the basic model.

Mainly, the selected feature needs to align with the computation capabilities available on the edge computing devices. For example, it is unrealistic to assume that edge devices could load and run large transformer-based [68] models for inference purposes without sacrificing system performance. Hence, we focus on implementing mobile-friendly feature extraction pipelines in FedMultimodal which are listed below, targeting swift computation, efficient storage, and ease of deployment.

- **Visual:** For the visual data, our benchmark supports **MobileNetV2** [30] and **MobileViT** [47] as the embedding network to extract latent presentations. The complete MobileNetV2 and MobileViT have 4.3M and 2.7M parameters, respectively, making them practical visual feature backbones in FL. Due to space constraints, we report benchmark results with MobileNetV2 in this paper.
- **Text:** FedMultimodal integrates both **MobileBERT** [65] and **DistillBERT** [60] to extract representations from textual data. MobileBERT uses a bottleneck structure to reduce the parameter size from 340M to 25M when compared to BERT [13], while DistillBERT applies a knowledge distillation process that decreases the BERT model to 66M parameters. We decide to benchmark with the MobileBERT feature backbone given the page constraints.
- **Audio:** FedMultimodal uses Mel-frequency cepstral coefficients (MFCCs) due to their widespread usage in the state-of-the-art speech recognition models like Wav2Vec 2.0 [11].
- **Other Modalities:** We use the raw data with the remaining modalities in the FedMultimodal. These data streams include accelerometer, gyroscope, and ECG readings.

### 3.3 Multimodal Models

**Model Design Principles.** Compared to remote servers, edge computing nodes are more appropriate for lightweight computing tasks due to constraints in computation resources, storage capabilities, battery capacities, and communication bandwidths. When designing ML models for resource-constrained devices, a significant design consideration is to reduce the number of parameters in edge ML models, thus reducing memory and execution latency. Such models can either be the backbone feature extraction models or application-specific prediction models. One major design principle of FedMultimodal is to study lightweight but effective solutions for multimodal FL learning instead of training models with multi-million parameters.

Table 3: Hyperparameters for training FL models with multimodal data using FedAvg.

Multimodal Dataset	Client Sample Rate	Learning Rate	Training Rounds
MELD	10%	0.01	200
CREMA-D	10%	0.05	200
UCF101	10%	0.05	200
MiT10	5%	0.05	300
MiT51	5%	0.05	300
UCI-HAR	10%	0.05	200
KU-HAR	10%	0.05	200
PTB-XL	25%	0.05	200
Hateful-Memes	25%	0.05	200
CrisisMMD	10%	0.05	200

**Model Architecture.** With this design principle in mind, we construct ML models mainly based on the 1D Conv-RNN/MLP architecture. Even though the transformer-based model has achieved SOTA performance in diverse applications, these models typically include millions or even billions of parameters, making them impractical to use in FL settings as a result of massive computations, memory usage, and battery consumption during back-propagation [70].

An example model architecture is presented in Figure 2. Specifically, the multimodal model in FedMultimodal includes an encoder, a modality-fusion block, and a downstream classifier. The encoder part follows either Conv+RNN architecture or RNN-only architecture. The encoder that adopts Conv+RNN architecture takes the input of audio, accelerometer, gyroscope, and ECG information, otherwise uses RNN-only architecture. Following encoder modules, FedMultimodal uses a late-fusion mechanism to combine modality-specific representations into a multimodal representation. The multimodal representation is then fed through 2 dense layers for downstream predictions.

### 3.4 Fusion Schemes

In this work, we present two basic fusion approaches: *concatenation-based fusion* and *attention-based fusion*. In the concatenation-based fusion, the average pooling operation is first performed on the GRU output. After that, we concatenate the pooling embeddings to form the multimodal embedding. On the other hand, the attention-based fusion concatenates the temporal output from each modality without the average pooling step. We apply an attention mechanism similar to hierarchical attention [76]. Given the concatenated multimodal data  $h$ , the attention procedures are as follows:

$$u = \tanh(Wh+b); a = \text{softmax}(u^T c)$$

$$v = \sum_i a_i h_i$$

The concatenated multimodal data  $h$  is first fed through a one-layer MLP to get representation  $u$ . We then use a context vector  $c$  to obtain a normalized importance score through a softmax function. After that, we compute the final multimodal embedding  $v$  as a weighted sum of  $h$  based on the weights  $a$ . Here, we can further implement a multi-head attention mechanism by having multiple  $c$ . We would also stress that this attention mechanism is lightweight, thus making it realistic to deploy on a variety of edge devices. In addition, the attention mechanism allows us to mask the missing

**Table 4: Benchmarking performance. Text colors in red and blue denote the best performance using *Concatanation-based Fusion* and *Attention-based Fusion*, respectively. † indicates the best performance score of the corresponding dataset.**

	ML Datasets	$\alpha$	Metric	Concatanation-based Fusion				Attention-based Fusion			
				FedAvg	FedProx	FedRS	FedOpt	FedAvg	FedProx	FedRS	FedOpt
<b>Natural Partition</b>	MELD	-	UAR	48.08	48.47	49.21	<b>50.66</b>	54.37	54.67	53.82	<b>55.37</b> †
	CREMA-D		UAR	61.52	61.64	<b>62.17</b>	61.14	61.66	62.03	60.41	<b>62.66</b> †
	KU-HAR		F1	61.56	61.13	61.26	<b>68.82</b>	61.78	61.78	62.04	<b>71.41</b> †
	PTB-XL		F1	61.87	61.00	-	<b>62.83</b> †	61.88	61.71	-	<b>62.42</b>
<b>Synthetic Partition</b>	UCF101	5.0	Acc	67.98	67.98	67.98	<b>74.38</b>	75.13	74.51	75.27	<b>75.89</b> †
	MiT10		Acc	55.39	55.39	55.29	<b>55.47</b>	57.10	<b>57.93</b> †	56.82	57.25
	MiT51		Acc	28.96	28.62	27.67	<b>35.01</b>	33.90	34.46	33.74	<b>35.62</b> †
	Hateful-Memes		AUC	58.23	<b>59.90</b> †	57.98	<b>58.01</b>	57.83	59.09	56.67	<b>59.51</b>
	CrisisMMD		F1	43.67	43.37	<b>44.30</b> †	43.44	39.11	39.36	<b>41.01</b>	38.74
	UCI-HAR		F1	78.78	78.07	77.87	<b>84.83</b>	77.75	77.38	76.82	<b>85.17</b> †
<b>Synthetic Partition</b>	UCF101	0.1	Acc	64.57	64.55	61.17	<b>74.17</b>	74.53	74.71	73.27	<b>75.05</b> †
	MiT10		Acc	44.84	50.03	45.92	<b>50.10</b>	42.96	45.47	46.23	<b>50.76</b> †
	MiT51		Acc	28.63	27.98	27.92	<b>33.46</b>	32.41	32.55	31.99	<b>35.35</b> †
	Hateful-Memes		AUC	51.02	<b>59.86</b>	51.28	<b>58.08</b>	49.68	59.44	49.80	<b>60.51</b> †
	CrisisMMD		F1	9.90	10.65	9.28	<b>26.82</b>	8.49	25.31	10.12	<b>27.59</b> †
	UCI-HAR		F1	77.50	77.34	73.68	<b>78.97</b>	76.66	76.58	68.65	<b>79.80</b> †

modalities in the computation, providing a simple yet effective solution to train FL models with missing modalities.

### 3.5 Federated Optimizers

First, most existing FL training algorithms are validated in unimodal settings, and their efficacy on multimodal tasks remains unexplored. As a result, FedMultimodal is suited to several popular FL algorithms, including FedAvg [46], FedProx [40], FedRS [41], and FedOpt [57]. Particularly, FedOpt holds state-of-the-art performance across multiple unimodal applications [57]. One objective of FedMultimodal is to provide comprehensive evaluations across various FL algorithms.

### 3.6 Real-world Noise Factor Emulator

Prior literature (see Table 1) on multimodal FL provides little or no assessment of their robustness in real-life settings. In order to provide a comprehensive evaluation of multimodal FL models toward safe and robust deployment, FedMultimodal enables the emulation of missing modalities, missing labels, and erroneous labels for real-world multimodal FL.

**Missing Modality.** In practice, data sources, whether they are microphones, cameras, mobile hardware, or medical electrodes, are prone to data imperfections or complete data losses (e.g., missing modalities) caused by firmware malfunctions, network disconnections, or sensor damages [20]. Hence, it is critical to design an emulator module to synthesize the cases of missing modalities for some clients. FedMultimodal provides the simulations of missing modalities as suggested by [10], where the availability of each modality follows a Bernoulli distribution. We set an equal missing rate  $q$  for each modality in the following experiments.

**Missing Labels.** Not only can the multimodal FL encounter data imperfection challenges, it can also suffer from missing label problems. Surprisingly, most prior works have made the ideal assumption that

the data stored on edge devices are fully annotated with ground-truth labels. However, in a more realistic real-world FL setting, we argue that only a portion of the data can come with labels. As such, FedMultimodal allows the missing label simulation to assess the risk of decreased robustness.

**Erroneous Labels.** In addition to missing labels, real-world FL implementations encounter label noise as a result of bias, skill differences, and labeling errors from the annotators. Inspired by [79], we apply a label error generation process described in [51]. In summary, the erroneous labels are generated using a transition matrix  $Q$ , where  $Q_{i,j} = \mathbb{P}(y = j | y = i)$  indicates the chance of ground-truth label  $i$  being erroneously annotated with label  $j$ .

## 4 EXPERIMENTS AND DISCUSSION

### 4.1 Experimental Details

**Setup.** We adopt the RNN-only model architecture to the video and text modalities while utilizing the Conv-RNN model architecture in other modalities. Specifically, the model with the convolutional module consists of 3 convolution layers with the number of filters in {16, 32, 64} and the filter kernel size of  $5 \times 5$ . Moreover, we set the hidden layer size of RNN as 128. We choose ReLU as the activation function and the dropout rate as 0.2. The number of attention heads is 6 in all experiments. We fixed the batch size for all datasets to 16 and the local epoch to 1 for all experiments.

Additionally, we set the training epochs as 200 for all datasets except the MiT sub-datasets. However, the total training epoch is 300 in MiT10 and MiT51 as these 2 datasets contain more data than the other datasets. Hyperparameter details such as the learning and client sampling rates for each dataset are listed in Table 3. Due to the limited number of clients in the Hateful Memes dataset and PTB-xl datasets, we apply a higher client sample rate in these 2 datasets. In the FedOpt algorithm, we search the server learning rate in a range from  $10^{-3}$  to  $2.5 \times 10^{-3}$ . Meanwhile, the proximal term ranges from  $10^{-2}$  to  $10^0$  in the FedProx algorithm.




Figure 3: Performance comparisons between multimodal and unimodal learning under FL settings.




Figure 4: Relative performance changes under different missing modality rates.

**Evaluation Metrics.** We follow established practices from the literature while conducting evaluations on each dataset. Specifically, evaluation metrics (e.g., F1) and validation protocols (e.g., predefined splits) are two fundamental components to ensure comparability with past (and future) works. With datasets that provide a pre-defined partition for training/validation/testing, we repeat the experiments 5 times using different seeds. We perform 5-fold cross-validation on datasets without such pre-defined experimenting rules. We provide details about our evaluation methods in Table 2.

## 4.2 Overall Performance

We first report the comparisons between two fusion mechanisms (**attention-based fusion** and **concatenation-based fusion**) in Table 4. From the results, we can find that the attention-based fusion mechanism outperforms concatenation-based fusion in the majority of the datasets. Specifically, the attention-based fusion mechanism leads to better performances in most high data heterogeneity conditions, but it underperforms the concatenation-based fusion in two synthetic datasets with  $\alpha = 5.0$  (low data heterogeneity). Moreover, we observe that the FedOpt algorithm consistently yields better baselines compared to other FL algorithms with a few exceptions in low data heterogeneity conditions. However, we would like to highlight that, in practice, FedOpt requires additional hyperparameter tuning on the server learning rate to reach the best performance. Overall, these results imply that the fusion mechanism is a critical factor impacting multimodal model performance in data heterogeneous FL.

Moreover, HAR tasks are associated with the highest performance scores, suggesting the simplicity of this learning task. In contrast, classification results on the Hateful Memes dataset and CrisisMMD dataset imply that social media classification is a challenging task using FL, with the best model performance on the CrisisMMD dataset below 30%. A plausible explanation is that the pre-trained models that we rely on are not generalized to social media data, generating image and textual features that are unrepresentative of downstream learning models. On the other hand, performances on the MiT51 dataset demonstrate similar findings pointed out from [50], validating that MiT is a challenging dataset. However, reducing the number of labels indeed simplifies the predicting task, resulting in moderate model performance on MiT10.

## 4.3 Uni-modality vs. Multi-modalities

One fundamental research question centering around multimodal learning is its performance compared to unimodal models. For example, the previous multimodal benchmark [42], with an emphasis on centralized learning setup, demonstrates that unimodal learning could yield similar performance to multimodal models with fewer parameters. Similar to MultiBench, FedMultimodal provides the unimodal FL to compare with multimodal FL baselines. We summarize the benchmark comparisons between unimodal FL and multimodal FL in Figure 3. The comparisons use datasets with natural partitions or high data heterogeneity partitions. Overall, we observe that unimodal learning provides competitive performance compared with the multimodal FL benchmarks, complying with centralized benchmark results reported in [42]. Nevertheless, in




Figure 5: Relative performance changes under different label missing rates.




Figure 6: Relative performance changes under different erroneous label rates.

most scenarios, multimodal learning still outperforms unimodal learning, whereas the performance gap between multimodal and unimodal FL is within 5% in the majority of the datasets.

#### 4.4 Impact of Missing Modalities

As described in earlier sections, a unique challenge associated with multimodal learning is dealing with scenarios of missing modalities [10, 77]. In this section, we benchmark our selected datasets with different rates of missing modalities. In this experiment, we assume that the availability of each modality follows a Bernoulli distribution with a missing rate of  $q$ . Following the experiment protocol presented by [10], we set a uniform missing rate of  $q$  for each modality, where  $q \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$ . As described in the multimodal model section, attention-based fusion allows us to train the model even with missing data through masking. To train the model with the missing entries, we fill the missing data with 0 [53] while masking out the corresponding data points in calculating attention scores.

We present the relative model performance changes at different missing modality rates in Figure 4. From the graph, we find that the relative performance changes with missing rates below 30% are not substantial, suggesting that a small amount of missing modalities in deployment does not impact the final model performance in multimodal FL. Furthermore, we observe that the model performance starts to decline substantially at the missing rate of 50%. Surprisingly, we observe that half of the models have relative performance decreases that are under 10%, suggesting that the provided baseline models that use attention-based fusion still learn useful information in these cases. However, we find that the missing modality introduces a significantly larger impact on CrisisMMD data and the

HAR applications compared to the other multimodal applications we evaluated. This suggests that future FL HAR research should carefully consider missing modalities as a part of the evaluation.

#### 4.5 Impact of Missing Labels

Missing labels is a widely presented challenge in FL. In this section, we perform evaluations of missing label conditions using the FedMultimodal benchmark. Similar to the missing modality experiment, we assume that the availability of each label follows a Bernoulli distribution with a missing rate  $l$ . We apply the FedMultimodal benchmark to emulate the missing label rate  $l \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$ . The goal of this experiment is to quantify the impact of missing labels on the overall performance of benchmarks, hence we do not integrate any mitigation methodologies, such as semi-supervised learning or self-supervised learning in our experiments.

Results on relative model performance changes at different label missing ratios using the FedMultimodal framework are presented in Figure 5. Overall, we can observe that missing labels have a reduced impact on the model performance when compared to the missing modality scenario. For instance, the model performance suffers less than 10% relative performance decreases in the majority of the datasets with the exception of KU-HAR datasets. When the missing label ratio is below 50%, we observe minor performance decreases in all datasets. Surprisingly, we identify that CrisisMMD yields worse performance at 30% than at 10% missing label ratio. We conjecture that the reason behind this result might be attributed to the data labeling quality of the Hateful-Memes dataset. For example, determining whether the content is hateful or not can be very subjective, and such subjectiveness could hurt labeling quality.






Figure 7: Relative performance changes with 30% data corrupted (missing modalities vs. missing labels vs. erroneous labels).

#### 4.6 Impact of Erroneous Labels

Besides missing modalities and missing labels, erroneous labels frequently exist in FL [75]. In this section, we report our benchmark performance with erroneous labels. Similar to previous experiments, we search the erroneous label ratio  $l \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$ , where  $l$  represents the amount of data with erroneous labels. Similar to the experiment setup in [79], our benchmark defines the sparsity of erroneous label transition matrix  $Q$  as 0.4. The error sparsity specifies the possible number of unique labels  $k$  that one label can be wrongly annotated with, with a small sparsity error rate corresponding to a larger  $k$ .

The complete results of the relative model performance changes at different levels of erroneous label ratios are shown in Figure 6. Compared to the missing modalities experiment, the erroneous label condition leads to substantially larger performance decreases. For example, more than half of the datasets have the relative performance decreases above 10% at the erroneous label ratio of 30%. Moreover, a 20% performance drop can be identified from these datasets when the erroneous label ratio reaches 50%. To compare the impact of data corruption conditions in FL, we plot the relative performance changes with different data corruption conditions at the data corrupted ratio of 30% in Figure 7. We can observe that performances of multimodal FL are more susceptible to label noises than missing modalities or missing labels. Based on these observations, our future benchmark directions also include implementations of backdoor attacks and mitigation in FedMultimodal.

### 5 LIMITATIONS AND FUTURE WORK

**Scale of Datasets and Models.** The dataset selection criteria of FedMultimodal ensures that the chosen datasets are representative and diversified across different dimensions such as application scenarios, data size, and number of clients. In addition, FedMultimodal only includes ML models that align with the use cases of FL, taking into account the computational limitations of edge devices. We acknowledge that FedMultimodal currently does not cover several promising multimodal applications, such as medical image analysis, autonomous driving, and virtual reality, and the range of the supported models is limited. We will continuously update FedMultimodal to support new tasks such as Ego4D [26], as well as newer feature extraction models.

**Scale of Modality Fusion Schemes.** Currently, FedMultimodal includes two basic approaches for modality fusion: concatenation and attention. Modality fusion under FL remains an open problem,

and our objective is to draw attention to the need for developing more advanced modality fusion schemes under FL [32, 45].

**Data Heterogeneity.** As discussed in previous sections, addressing the data heterogeneity challenge is critical in FL. While many FL studies focus their experiments on the unimodal setup, there is a lack of extensive research on tackling data heterogeneity in multimodal FL. To address this gap, the FedMultimodal benchmark provides opportunities to facilitate fundamental research in this direction. In the future, it is of further interest to explore knowledge-transfer learning approaches as suggested in [12, 31, 44], within the context of multimodal FL.

**Label Scarcity.** One major challenge for FL is the lack of qualitative labels. FedMultimodal enables researchers to efficiently perform experiments on multimodal FL with missing labels by providing the ability to emulate experimental conditions with missing labels. We hope FedMultimodal brings unique benefits for ML practitioners to develop self-supervised learning [15, 78, 84] and semi-supervised learning [34, 81, 82] algorithms under multimodal FL.

**Privacy Leakage.** While sharing model updates is considered to be more private than sharing raw data, recent works have revealed that FL can still be susceptible to privacy attacks. These attacks include (but are not limited to) membership inference attacks [48], reconstruction attacks [25, 83], attribute inference attacks [18] and label inference attacks [24]. Consequently, an emerging research direction for expanding FedMultimodal is to explore the privacy leakages in multimodal FL. Apart from identifying privacy risks associated with multimodal FL, it is also crucial to investigate privacy-enhancing techniques, such as differential privacy [16, 23, 72] and secure aggregation [5] as promising areas of research within the scope of FedMultimodal to mitigate privacy attacks.

### 6 CONCLUSION

In this paper, we presented a new framework for multimodal federated learning, named FedMultimodal, which enables federated learning in multimodal applications. We further established a reproducible benchmark of results for 5 multimodal FL applications covering 10 datasets for future comparisons. We also benchmarked results on model robustness to missing modalities, missing labels, and noisy labels in each of these tasks.

### 7 ACKNOWLEDGEMENT

This work is in part supported by USC Amazon Center, as well as research gift awards from Intel, Meta, and Konica Minolta.

## REFERENCES

- [1] Firoz Alam, Ferda Ofli, and Muhammad Imran. 2018. Crisismmd: Multimodal twitter datasets from natural disasters. In *Twelfth international AAAI conference on web and social media*.
- [2] Erick A Perez Alday, Annie Gu, Amit J Shah, Chad Robichaux, An-Kwok Ian Wong, Chengyu Liu, Feifei Liu, Ali Bahrami Rad, Andoni Elola, Salman Seyedi, et al. 2020. Classification of 12-lead ecgs: the physionet/computing in cardiology challenge 2020. *Physiological measurement* 41, 12 (2020), 124003.
- [3] Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra Perez, and Jorge Luis Reyes Ortiz. 2013. A public domain dataset for human activity recognition using smartphones. In *Proceedings of the 21th international European symposium on artificial neural networks, computational intelligence and machine learning*. 437–442.
- [4] Burçin Becerik-Gerber, Gale M. Lucas, Ashrant Aryal, Mohamad Awada, Mario Bergés, Sarah Billington, Olga Boric-Lubecke, Ali Ghahramani, Arsalan Heydari, Christoph Höelscher, Farrokh Jazizadeh, Azam Khan, Jared Langevin, Ruying Liu, Frederick Marks, Matthew Louis Mauriello, Elizabeth L. Murnane, Haeyoung Noh, Marco Pritoni, Shawn C Roll, Davide Schaumann, Mir Hasan Seyedrezaei, John Ellor Taylor, Jie Zhao, and Runhe Zhu. 2022. The field of human building interaction for convergent research and innovation for intelligent built environments. *Scientific Reports* 12 (2022).
- [5] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. 2017. Practical secure aggregation for privacy-preserving machine learning. In *proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 1175–1191.
- [6] Brandon M Booth, Tiantian Feng, Abhishek Jangalwa, and Shrikanth S Narayanan. 2019. Toward robust interpretable human movement pattern analysis in a workplace setting. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 7630–7634.
- [7] Brandon M Booth, Karel Mundnich, Tiantian Feng, Amrutha Nadarajan, Tiago H Falk, Jennifer L Villatte, Emilio Ferrara, and Shrikanth Narayanan. 2019. Multimodal human and environmental sensing for longitudinal behavioral studies in naturalistic settings: Framework for sensor selection, deployment, and management. *Journal of medical Internet research* 21, 8 (2019), e12832.
- [8] Sebastian Caldas, Sai Meher Karthik Duddu, Peter Wu, Tian Li, Jakub Konečný, H Brendan McMahan, Virginia Smith, and Ameet Talwalkar. 2018. Leaf: A benchmark for federated settings. *arXiv preprint arXiv:1812.01097* (2018).
- [9] Houwei Cao, David G Cooper, Michael K Keutmann, Ruben C Gur, Ani Nenkova, and Ragini Verma. 2014. Crema-d: Crowd-sourced emotional multimodal actors dataset. *IEEE transactions on affective computing* 5, 4 (2014), 377–390.
- [10] Jiayi Chen and Aidong Zhang. 2022. 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*. 87–96.
- [11] Li-Wei Chen and Alexander Rudnicky. 2021. Exploring Wav2vec 2.0 fine-tuning for improved speech emotion recognition. *arXiv preprint arXiv:2110.06309* (2021).
- [12] Yae Jee Cho, Andre Manoel, Gauri Joshi, Robert Sim, and Dimitrios Dimitriadis. 2022. Heterogeneous ensemble knowledge transfer for training large models in federated learning. *arXiv preprint arXiv:2204.12703* (2022).
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *ArXiv abs/1810.04805* (2019).
- [14] Dimitrios Dimitriadis, Mirian Hipolito Garcia, Daniel Madrigal Diaz, Andre Manoel, and Robert Sim. 2022. Flute: A scalable, extensible framework for high-performance federated learning simulations. *arXiv preprint arXiv:2203.13789* (2022).
- [15] Nanqing Dong and Irina Voiculescu. 2021. Federated contrastive learning for decentralized unlabeled medical images. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part III 24*. Springer, 378–387.
- [16] Cynthia Dwork. 2006. Differential privacy. In *Automata, Languages and Programming: 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10–14, 2006, Proceedings, Part II 33*. Springer, 1–12.
- [17] Tiantian Feng, Brandon M Booth, Brooke Baldwin-Rodríguez, Felipe Osorno, and Shrikanth Narayanan. 2021. A multimodal analysis of physical activity, sleep, and work shift in nurses with wearable sensor data. *Scientific reports* 11, 1 (2021), 8693.
- [18] Tiantian Feng, Hanieh Hashemi, Rajat Hebbar, Murali Annavaram, and Shrikanth S Narayanan. 2021. Attribute inference attack of speech emotion recognition in federated learning settings. *arXiv preprint arXiv:2112.13416* (2021).
- [19] Tiantian Feng, Rajat Hebbar, Nicholas Mehlman, Xuan Shi, Aditya Kommineni, and Shrikanth Narayanan. 2023. A Review of Speech-centric Trustworthy Machine Learning: Privacy, Safety, and Fairness. *APSIPA Transactions on Signal and Information Processing* 12, 3 (2023). <https://doi.org/10.1561/116.00000084>
- [20] Tiantian Feng and Shrikanth Narayanan. 2019. Imputing missing data in large-scale multivariate biomedical wearable recordings using bidirectional recurrent neural networks with temporal activation regularization. In *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2529–2534.
- [21] Tiantian Feng and Shrikanth Narayanan. 2022. Semi-FedSER: Semi-supervised Learning for Speech Emotion Recognition On Federated Learning using Multi-view Pseudo-Labeling. *arXiv preprint arXiv:2203.08810* (2022).
- [22] Tiantian Feng and Shrikanth S Narayanan. 2019. Discovering optimal variable-length time series motifs in large-scale wearable recordings of human bi-behavioral signals. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 7615–7619.
- [23] Tiantian Feng, Raghuvveer Peri, and Shrikanth Narayanan. 2022. User-Level Differential Privacy against Attribute Inference Attack of Speech Emotion Recognition on Federated Learning. In *Proc. Interspeech 2022*. 5055–5059. <https://doi.org/10.21437/Interspeech.2022-10060>
- [24] Chong Fu, Xuhong Zhang, Shouling Ji, Jinyin Chen, Jingzheng Wu, Shanqing Guo, Jun Zhou, Alex X Liu, and Ting Wang. 2022. Label inference attacks against vertical federated learning. In *31st USENIX Security Symposium (USENIX Security 22)*. 1397–1414.
- [25] Jiahui Geng, Yongli Mou, Feifei Li, Qing Li, Oya Beyan, Stefan Decker, and Chunming Rong. 2021. Towards General Deep Leakage in Federated Learning. *arXiv preprint arXiv:2110.09074* (2021).
- [26] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. 2022. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 18995–19012.
- [27] Chaoyang He, Keshav Balasubramanian, Emir Ceyani, Yu Rong, Peilin Zhao, Junzhou Huang, Murali Annavaram, and Salman Avestimehr. 2021. FedGraphNN: A Federated Learning System and Benchmark for Graph Neural Networks. *ArXiv abs/2104.07145* (2021).
- [28] Chaoyang He, Songze Li, Jinhyun So, Mi Zhang, Hongyi Wang, Xiaoyang Wang, Praneeth Vepakomma, Abhishek Singh, Hang Qiu, Li Shen, Peilin Zhao, Yan Kang, Yang Liu, Ramesh Raskar, Qiang Yang, Murali Annavaram, and Salman Avestimehr. 2020. FedML: A Research Library and Benchmark for Federated Machine Learning. *arXiv preprint arXiv:2007.13518* (2020).
- [29] Chaoyang He, Alay Dilipbhai Shah, Zhenheng Tang, Di Fan1Adarshan Naiynar Sivashunmugam, Keerti Bhogaraju, Mita Shimpi, Li Shen, Xiaowen Chu, Mahdi Soltanolkotabi, and Salman Avestimehr. 2021. Fedcv: a federated learning framework for diverse computer vision tasks. *arXiv preprint arXiv:2111.11066* (2021).
- [30] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Veyand, Marco Andreetto, and Hartwig Adam. 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017).
- [31] Sohei Itahara, Takayuki Nishio, Yusuke Koda, Masahiro Morikura, and Koji Yamamoto. 2021. Distillation-based semi-supervised federated learning for communication-efficient collaborative training with non-iid private data. *IEEE Transactions on Mobile Computing* 22, 1 (2021), 191–205.
- [32] Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. 2021. Perceiver: General perception with iterative attention. In *International conference on machine learning*. PMLR, 4651–4664.
- [33] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Benis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. 2021. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning* 14, 1–2 (2021), 1–210.
- [34] Yan Kang, Yang Liu, and Xinle Liang. 2022. Fedcvt: Semi-supervised vertical federated learning with cross-view training. *ACM Transactions on Intelligent Systems and Technology (TIST)* 13, 4 (2022), 1–16.
- [35] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. 2020. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*. PMLR, 5132–5143.
- [36] Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. 2020. The hateful memes challenge: Detecting hate speech in multimodal memes. *Advances in Neural Information Processing Systems* 33 (2020), 2611–2624.
- [37] Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. 2016. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492* (2016).
- [38] Fan Lai, Yinwei Dai, Xiangfeng Zhu, Harsha V Madhyastha, and Mosharaf Chowdhury. 2021. FedScale: Benchmarking model and system performance of federated learning. In *Proceedings of the First Workshop on Systems Challenges in Reliable and Secure Federated Learning*. 1–3.
- [39] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436–444.
- [40] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. 2020. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems* 2 (2020), 429–450.

- [41] Xin-Chun Li and De-Chuan Zhan. 2021. Fedrs: Federated learning with restricted softmax for label distribution non-iid data. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 995–1005.
- [42] Paul Pu Liang, Yiwei Lyu, Xiang Fan, Zetian Wu, Yun Cheng, Jason Wu, Leslie Chen, Peter Wu, Michelle A Lee, Yuke Zhu, et al. 2021. Multibench: Multiscale benchmarks for multimodal representation learning. *arXiv preprint arXiv:2107.07502* (2021).
- [43] Bill Yuchen Lin, Chaoyang He, Zihang Zeng, Hulin Wang, Yufen Huang, Mahdi Soltanolkotabi, Xiang Ren, and Salman Avestimehr. 2021. Fednlp: Benchmarking federated learning methods for natural language processing tasks. *arXiv preprint arXiv:2104.08815* (2021).
- [44] Tao Lin, Lingjing Kong, Sebastian U Stich, and Martin Jaggi. 2020. Ensemble distillation for robust model fusion in federated learning. *Advances in Neural Information Processing Systems* 33 (2020), 2351–2363.
- [45] Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. 2016. Hierarchical question-image co-attention for visual question answering. *Advances in neural information processing systems* 29 (2016).
- [46] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agueri y Arcas. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*. PMLR, 1273–1282.
- [47] Sachin Mehta and Mohammad Rastegari. 2021. Mobilevit: light-weight, general-purpose, and mobile-friendly vision transformer. *arXiv preprint arXiv:2110.02178* (2021).
- [48] Luca Melis, Congzheng Song, Emiliano De Cristofaro, and Vitaly Shmatikov. 2019. Exploiting unintended feature leakage in collaborative learning. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 691–706.
- [49] Fatemehsadat Mireshtghallah, Mohammadkazem Taram, Praneeth Vepakomma, Abhishek Singh, Ramesh Kaskar, and Hadi Esmailzadeh. 2020. Privacy in deep learning: A survey. *arXiv preprint arXiv:2004.12254* (2020).
- [50] Mathew Monfort, Alex Andonian, Bolei Zhou, Kandan Ramakrishnan, Sarah Adel Bargal, Tom Yan, Lisa Brown, Quanfu Fan, Dan Gutfreund, Carl Vondrick, et al. 2019. Moments in time dataset: one million videos for event understanding. *IEEE transactions on pattern analysis and machine intelligence* 42, 2 (2019), 502–508.
- [51] Curtis Northcutt, Lu Jiang, and Isaac Chuang. 2021. Confident learning: Estimating uncertainty in dataset labels. *Journal of Artificial Intelligence Research* 70 (2021), 1373–1411.
- [52] Alexandros Pantelopoulos and Nikolaos G Bourbakis. 2009. A survey on wearable sensor-based systems for health monitoring and prognosis. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 40, 1 (2009), 1–12.
- [53] Srinivas Parthasarathy and Shiva Sundaram. 2020. Training strategies to handle missing modalities for audio-visual expression recognition. In *Companion Publication of the 2020 International Conference on Multimodal Interaction*. 400–404.
- [54] Shyamal Patel, Hyung Park, Paolo Bonato, Leighton Chan, and Mary Rodgers. 2012. A review of wearable sensors and systems with application in rehabilitation. *Journal of neuroengineering and rehabilitation* 9, 1 (2012), 1–17.
- [55] Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2018. Meld: A multimodal multi-party dataset for emotion recognition in conversations. *arXiv preprint arXiv:1810.02508* (2018).
- [56] Andrew Raji, Animikh Ghosh, Santosh Kumar, and Mani Srivastava. 2011. Privacy risks emerging from the adoption of innocuous wearable sensors in the mobile environment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 11–20.
- [57] Sashank Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and H Brendan McMahan. 2020. Adaptive federated optimization. *arXiv preprint arXiv:2003.00295* (2020).
- [58] Michael S Ryoo, AJ Piergiovanni, Mingxing Tan, and Anelia Angelova. 2019. AssembleNet: Searching for multi-stream neural connectivity in video architectures. *arXiv preprint arXiv:1905.13209* (2019).
- [59] Aaqib Saeed, Flora D Salim, Tanir Ozcelebi, and Johan Lukkien. 2020. Federated self-supervised learning of multisensor representations for embedded intelligence. *IEEE Internet of Things Journal* 8, 2 (2020), 1030–1040.
- [60] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108* (2019).
- [61] EK Sannara, Francois Portet, Philippe Lalanda, and VEGA German. 2021. A federated learning aggregation algorithm for pervasive computing: Evaluation and comparison. In *2021 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 1–10.
- [62] Niloy Sikder and Abdullah-Al Nahid. 2021. KU-HAR: An open dataset for heterogeneous human activity recognition. *Pattern Recognition Letters* 146 (2021), 46–54.
- [63] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. 2012. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402* (2012).
- [64] Nils Strodthoff, Patrick Wagner, Tobias Schaeffter, and Wojciech Samek. 2020. Deep learning for ECG analysis: Benchmarks and insights from PTB-XL. *IEEE Journal of Biomedical and Health Informatics* 25, 5 (2020), 1519–1528.
- [65] Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020. Mobilebert: a compact task-agnostic bert for resource-limited devices. *arXiv preprint arXiv:2004.02984* (2020).
- [66] Jean Ogier du Terrail, Samy-Safwan Ayed, Edwige Cyffers, Felix Grimberg, Chaoyang He, Regis Loeb, Paul Mangold, Tanguy Marchand, Othmane Marfoq, Erum Mushtaq, et al. 2022. Flamby: Datasets and Benchmarks for Cross-Silo Federated Learning in Realistic Healthcare Settings. *arXiv preprint arXiv:2210.04620* (2022).
- [67] Vasileios Tsouvalas, Tanir Ozcelebi, and Nirvana Meratnia. 2022. Privacy-preserving Speech Emotion Recognition through Semi-Supervised Federated Learning. In *2022 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops)*. IEEE, 359–364.
- [68] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>
- [69] Patrick Wagner, Nils Strodthoff, Ralf-Dieter Boussejot, Dieter Kreiseler, Fatima I Lunze, Wojciech Samek, and Tobias Schaeffter. 2020. PTB-XL, a large publicly available electrocardiography dataset. *Scientific data* 7, 1 (2020), 1–15.
- [70] Meng Wang, Weijie Fu, Xiangnan He, Shijie Hao, and Xindong Wu. 2020. A survey on large-scale machine learning. *IEEE Transactions on Knowledge and Data Engineering* (2020).
- [71] Zhen Wang, Weirui Kuang, Yuexiang Xie, Liuyi Yao, Yaliang Li, Bolin Ding, and Jingren Zhou. 2022. FederatedScope-GNN: Towards a Unified, Comprehensive and Efficient Package for Federated Graph Learning. *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (2022).
- [72] Kang Wei, Jun Li, Ming Ding, Chuan Ma, Howard H Yang, Farhad Farokhi, Shi Jin, Tony QS Quek, and H Vincent Poor. 2020. Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security* 15 (2020), 3454–3469.
- [73] Yuexiang Xie, Zhen Wang, Daoyuan Chen, Dawei Gao, Liuyi Yao, Weirui Kuang, Yaliang Li, Bolin Ding, and Jingren Zhou. 2022. FederatedScope: A Comprehensive and Flexible Federated Learning Platform via Message Passing. *ArXiv abs/2204.05011* (2022).
- [74] Baochen Xiong, Xiaoshan Yang, Fan Qi, and Changsheng Xu. 2022. A unified framework for multi-modal federated learning. *Neurocomputing* 480 (2022), 110–118.
- [75] Duygu Yaldiz, Tuo Zhang, and Salman Avestimehr. 2023. Secure Federated Learning against Model Poisoning Attacks via Client Filtering. *ArXiv abs/2304.00160* (2023).
- [76] Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies*. 1480–1489.
- [77] Qiying Yu, Yimu Wang, Ke Xu, Yang Liu, and Jingjing Liu. 2023. Multimodal Federated Learning via Contrastive Representation Ensemble. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=Hnk1WRMAYqg>
- [78] Fengda Zhang, Kun Kuang, Zhaoyang You, Tao Shen, Jun Xiao, Yin Zhang, Chao Wu, Yueting Zhuang, and Xiaolin Li. 2020. Federated unsupervised representation learning. *arXiv preprint arXiv:2010.08982* (2020).
- [79] Tuo Zhang, Tiantian Feng, Samiul Alam, Sunwoo Lee, Mi Zhang, Shrikanth S Narayanan, and Salman Avestimehr. 2022. FedAudio: A Federated Learning Benchmark for Audio Tasks. *arXiv preprint arXiv:2210.15707* (2022).
- [80] Tuo Zhang, Lei Gao, Chaoyang He, Mi Zhang, Bhaskar Krishnamachari, and Salman Avestimehr. 2021. Federated Learning for the Internet of Things: Applications, Challenges, and Opportunities. *IEEE Internet of Things Magazine* 5 (2021), 24–29.
- [81] Zhengming Zhang, Yaoqing Yang, Zhewei Yao, Yujun Yan, Joseph E Gonzalez, Kannan Ramchandran, and Michael W Mahoney. 2021. Improving semi-supervised federated learning by reducing the gradient diversity of models. In *2021 IEEE International Conference on Big Data (Big Data)*. IEEE, 1214–1225.
- [82] Yuchen Zhao, Hanyang Liu, Honglin Li, Payam Barnaghi, and Hamed Haddadi. 2020. Semi-supervised federated learning for activity recognition. *arXiv preprint arXiv:2011.00851* (2020).
- [83] Ligeng Zhu and Song Han. 2020. Deep leakage from gradients. In *Federated learning*. Springer, 17–31.
- [84] Weiming Zhuang, Xin Gan, Yonggang Wen, Shuai Zhang, and Shuai Yi. 2021. Collaborative unsupervised visual representation learning from decentralized data. In *Proceedings of the IEEE/CVF international conference on computer vision*. 4912–4921.