

A Federated Large Language Model for Long-Term Time Series Forecasting

Raed Abdel-Sater¹ and A. Ben Hamza^{1,*}

¹Concordia University, Montreal, Canada

Abstract. Long-term time series forecasting in centralized environments poses unique challenges regarding data privacy, communication overhead, and scalability. To address these challenges, we propose FedTime, a federated large language model (LLM) tailored for long-range time series prediction. Specifically, we introduce a federated pre-trained LLM with fine-tuning and alignment strategies. Prior to the learning process, we employ K-means clustering to partition edge devices or clients into distinct clusters, thereby facilitating more focused model training. We also incorporate channel independence and patching to better preserve local semantic information, ensuring that important contextual details are retained while minimizing the risk of information loss. We demonstrate the effectiveness of our FedTime model through extensive experiments on various real-world forecasting benchmarks, showcasing substantial improvements over recent approaches. In addition, we demonstrate the efficiency of FedTime in streamlining resource usage, resulting in reduced communication overhead.

1 Introduction

Long-term time series forecasting involves predicting future values based on historical data over an extended period. The primary goal is to provide reliable and accurate predictions for multiple timestamps ahead, which can be crucial for planning, decision-making, risk management, and resource allocation in various domains such as finance, energy management, transportation, and environmental monitoring.

Recent advances in deep learning, particularly within Transformer- and LLM-based models [14, 32, 4, 33, 18, 3], have showcased significant progress in long-term time series forecasting. These models capitalize on the self-attention mechanism and its variants, providing a distinct advantage in capturing and modeling long-range dependencies within sequential data. While LLMs are adept at processing discrete tokens, time series data presents a unique challenge due to its continuous nature. Also, LLMs do not inherently possess the capability to interpret time series patterns, as this knowledge is not typically included in their pre-training. In addition, centralized learning models, including LLM-based methods, often require collecting and storing vast amounts of data in a central server, raising concerns about user privacy and data security. Thus, effectively leveraging LLMs for long-term time series forecasting, while ensuring accuracy and data privacy, highlights the challenges associated with centralized learning models, emphasizing the need for federated learning approaches [9]. Federated learning is an emerging machine learning paradigm that enables

edge devices to utilize their local datasets to collaboratively train a global model with the help of a central server, while preserving data privacy by keeping the data on the respective devices. At each iteration, the server broadcasts the current global model to the devices for local training, and aggregates the local model updates from the devices to update the global model. Federated learning offers the potential to enhance long-term prediction accuracy [13]. However, certain challenges arise from skewed data distributions that may impact the training quality [12]. Another challenge in federated learning is systems heterogeneity, which arises from the variation in storage, computational, and communication capabilities influenced by differences in hardware specifications, and power constraints [9]. This leads to heightened concerns around stragglers and fault tolerance. Several strategies have been proposed to address this issue involving adaptive optimization techniques that cater to the varying capabilities of devices, optimizing resource usage based on each device's characteristics [9]. Additionally, employing techniques like asynchronous updates and gradient compression can reduce communication overhead and accommodate devices with limited connectivity or power [15]. These strategies collectively aim to enhance the efficiency and reliability of federated learning systems in heterogeneous environments. However, these techniques do not significantly enhance model accuracy in scenarios involving optimization and transfer learning, particularly in contexts where data heterogeneity arises from collective settings [27].

In this paper, we introduce a federated learning framework, dubbed FedTime, for long-term time series forecasting, while preserving data privacy. Local data are used for training individual predictive LLM models on edge devices, allowing for rapid, localized predictions while reducing the load on cloud or central servers. The cloud server plays a dual role in initializing the system and aggregating local LLMs for global updates, a key aspect of federated learning. To optimize the use of limited computational resources on edge devices, we employ parameter-efficient tuning techniques [8, 7], which are crucial in minimizing computational and communication overhead, facilitating local training. Our learning framework encompasses training procedures on local edge devices (e.g., EV charging stations) and subsequent aggregation processes on the central server side by leveraging the pre-trained Large Language Model Meta AI (LLaMA-2) [26]. This LLaMA model is a group of foundational pre-trained LLMs with various sizes (7B, 13B, 33B, and 65B parameters), offering a range of capabilities suited to different needs and computational resources. The key contributions of this paper can be summarized as follows: (i) We introduce a federated large language framework, called FedTime, allowing for collaborative model train-

* Corresponding Author. Email: hamza@ciise.concordia.ca

ing across edge devices while preserving data privacy; (ii) we design a model architecture using a two-phase fine-tuning strategy, leveraging quantized and low-ranking adaptation, and aligning the model with time series data via direct preference optimization; and (iii) we demonstrate through extensive experiments the superior performance of FedTime over strong baselines across various forecasting datasets, especially for long-range forecasts. Moreover, we empirically show that FedTime exhibits reduced communication overhead.

2 Related Work

A sizable body of research has been developed to design Transformer-based methods for long-term time series forecasting [14, 32, 28, 33, 18]. Pyraformer [14] presents a pyramidal attention module, leveraging inter-scale and intra-scale connections, which enhances model capability to capture temporal patterns. Informer [32] design a ProbSparse self-attention mechanism to improve model prediction capacity. Autoformer [28] adapts decomposition blocks and auto-correlation mechanisms to help separate the long-term trend information from predicted hidden variables. FEDformer [33] employs fast Fourier transform to decompose sequences, enhancing the extraction of long-term information from the data. PatchTST [18] segments time series data into subseries-level patches, which serve as input tokens to a Transformer-based architecture. It also employs channel-independence, where each channel contains a single univariate time series that shares the same embedding and Transformer weights across all series, allowing efficient processing of multiple time series. More recently, various pre-trained LLMs, such as the generative pre-trained Transformer (GPT) [19] and LLaMA [26] models rooted in the Transformer architecture, have emerged as powerful methods in generating high-quality outputs in tasks spanning natural language processing [25, 17, 24, 31, 30] and time series forecasting [3], leveraging their ability to discern and comprehend intricate dependencies within data. Chang *et al.* introduce an LLM-based method, called LLM4TS, by adapting the GPT model as a foundational backbone to time series data and then fine-tuning it for time series forecasting tasks. Although our work also involves a two-stage fine-tuning strategy, it differs from LLM4TS in that our FedTime method not only focuses on efficient fine-tuning, but is also tailored for federated learning setups, aiming at long-term time series forecasting. The key strength of FedTime, compared to Transformer- and LLM-based approaches, lies in its ability to allow edge devices to collaborate in model training while keeping data decentralized. This preserves data privacy, a crucial aspect, especially in sensitive domains like healthcare. In addition, FedTime leverages parameter-efficient fine-tuning techniques, enabling it to update only a small number of local model parameters. This efficiency in training and model updates is crucial for real-world applications, especially in resource-constrained edge devices such as EV charging stations.

3 Method

In this section, we formulate the time series forecasting problem and provide an approach overview. Then, we present the main building blocks of the proposed federated learning framework.

3.1 Problem Statement and Approach Overview

Problem Statement. Time series forecasting refers to the process of predicting future values over a period of time using historical data. Let $\mathbf{X}_{1:L} = (\mathbf{x}_1, \dots, \mathbf{x}_L)^T \in \mathbb{R}^{L \times M}$ be a history sequence of L

multivariate time series, where for any time step t , each row $\mathbf{x}_t = (x_{t1}, \dots, x_{tM}) \in \mathbb{R}^{1 \times M}$ is a multivariate vector consisting of M variables or channels.

Given a history sequence $\mathbf{X}_{1:L}$ with look-back window L , the goal of multivariate time series forecasting is to predict a sequence $\mathbf{X}_{L+1:L+T} = (\mathbf{x}_{L+1}, \dots, \mathbf{x}_{L+T})^T \in \mathbb{R}^{T \times M}$ for the future T timesteps. To this end, we introduce a federated pre-trained large language model by leveraging channel independence and patching [18], as well as supervised and downstream fine-tuning strategies [3].

Approach Overview. The proposed federated learning framework comprises two main components: local training operations conducted on the client side and aggregation operations carried out on the server side. These components collaborate to facilitate efficient training. Prior to the learning process, we apply K-means clustering as a pre-processing step to group similar edge devices or clients (i.e., based on cluster size and performance). Clustering not only helps reduce the impact of biased prediction [22], but also ensures that sensitive data remains localized within clusters, reducing the risk of data exposure during model training. Then, we employ a pre-trained LLaMA model [26] as the foundational backbone to leverage the inherent capabilities of LLMs in capturing data patterns, thereby facilitating its application to time series forecasting within our federated learning framework. To improve the global model convergence rate, we use federated averaging with adaptive learning rates using Adam optimizer [16], which dynamically adjusts the learning rates for each local device based on the model updates. For each cluster, a global model is produced by aggregating the model updates from the cluster members. We also employ parameter-efficient fine-tuning techniques [7], which offer an efficient way to adapt pre-trained LLMs to diverse downstream tasks by focusing on updating only a small number of model parameters, and hence transmit updates faster, leading to quicker convergence and reduced communication overhead. While our proposed approach builds upon these existing techniques, its technical novelties lie in how these techniques are integrated and adapted to address specific challenges in long-term time series forecasting. In addition to proposing a novel two-phase fine-tuning strategy that leverages quantized and low-ranking adaptation techniques, along with direct preference optimization, to align the model with time series data effectively, our federated learning framework streamlines resource usage by minimizing computational and communication overhead, thereby improving the scalability and efficiency of the training process.

3.2 Proposed Learning Framework

We present federated learning framework, designed to overcome the classical challenges within a centralized training environment, such as network congestion, low bandwidth, latency, and privacy. The proposed federated architecture is built on the LLaMA structure, enabling parallel feature learning across various time series forecasting benchmarks. Figure 1 depicts the main building blocks of our learning framework, including time series data normalization, patching and positional encodings, LLM encoder and robust fine-tuning strategies. Our two-phase approach consists of supervised fine-tuning, followed by forecasting fine-tuning.

Channel Independence. Channel independence refers to a strategy where the features or variables of a multivariate time series are treated separately [18], emphasizing independence rather than mixing or amalgamating information from different channels or variables. This approach aims to retain the distinct characteristics of each

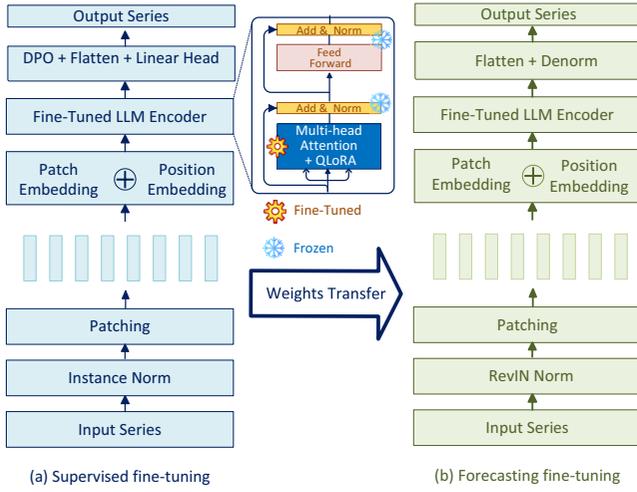


Figure 1. FedTime Architecture. (a) Supervised fine-tuning with direct preference optimization (DPO) integration for backbone model alignment. (b) Forecasting fine-tuning using reversible instance normalization (RevIN).

variable without merging them together. Specifically, the input time series $\mathbf{X}_{1:L} = (\mathbf{x}_1, \dots, \mathbf{x}_L)^\top$ is split into M univariate series $\mathbf{x}^{(i)} = (x_1^{(i)}, \dots, x_L^{(i)})^\top \in \mathbb{R}^L$, where $\mathbf{x}^{(i)}$ is the i th column of $\mathbf{X}_{1:L}$. Each of these univariate series is fed into the model backbone. Our proposed federated large language model takes $\mathbf{x}^{(i)}$ as input and returns a T -dimensional vector of predictions $\hat{\mathbf{x}}^{(i)} = (\hat{x}_{L+1}^{(i)}, \dots, \hat{x}_{L+T}^{(i)})^\top$.

Normalization. In the supervised fine-tuning phase, we employ instance normalization before patching to normalize each data instance $\mathbf{x}^{(i)}$ with zero mean and unit standard deviation, which are then added back to the output prediction. In the subsequent forecasting fine-tuning phase, each input series is normalized using the reversible instance normalization (RevIN) technique [10], which addresses challenges related to shifts in data distributions over time. RevIN consists of two main steps: normalization and denormalization. In the first step, the input undergoes normalization to standardize its distribution in terms of mean and variance. After the model generates output sequences, RevIN reverses the normalization process by denormalizing these outputs.

Patching. We adopt the patching strategy from [18], which involves dividing the sequential data into smaller patches. Specifically, each input series $\mathbf{x}^{(i)} \in \mathbb{R}^L$ is partitioned into patches, resulting in a sequence of patches $\mathbf{X}_p^{(i)} \in \mathbb{R}^{N \times P}$, where P is the patch length and N is the resulting number of patches, which also serves as the input sequence length for the Transformer encoder. The benefits of patching include better preservation of local semantic information in patches, computational and memory efficiency, and an expanded historical view.

Patch and Position Embeddings. Patch embeddings are obtained by mapping each patch in the sequence of patches to D dimensions with a trainable linear projection $\mathbf{W}_p \in \mathbb{R}^{P \times D}$. Position embeddings are added to patch embeddings via a learnable position encoding matrix $\mathbf{W}_{\text{pos}} \in \mathbb{R}^{N \times D}$ to retain positional information of the patches. Summing patch and position embeddings yields an embedding $\mathbf{X}_d^{(i)} \in \mathbb{R}^{N \times D}$ given by

$$\mathbf{X}_d^{(i)} = \mathbf{X}_p^{(i)} \mathbf{W}_p + \mathbf{W}_{\text{pos}}, \quad (1)$$

which serves as input to the multi-head self-attention mechanism of the Transformer encoder.

LLM Encoder. We employ the lightweight LLaMA2-7B model as

our LLM encoder. This LLaMA model is fundamentally built on the Transformer architecture, distinguished by its 7 billion parameters, and benefits from a comprehensive training on an expansive corpus containing 2 trillion tokens. At their core, Transformers rely on self-attention, which operates on the matrix $\mathbf{X}_d^{(i)}$ of tokens (i.e., patch embeddings with added positional encodings) by first linearly projecting it into a query matrix $\mathbf{Q}^{(i)}$, a key matrix $\mathbf{K}^{(i)}$ and a value matrix $\mathbf{V}^{(i)}$ as follows:

$$\mathbf{Q}^{(i)} = \mathbf{X}_d^{(i)} \mathbf{W}_q, \quad \mathbf{K}^{(i)} = \mathbf{X}_d^{(i)} \mathbf{W}_k, \quad \mathbf{V}^{(i)} = \mathbf{X}_d^{(i)} \mathbf{W}_v, \quad (2)$$

where $\mathbf{W}_q \in \mathbb{R}^{D \times d_k}$, $\mathbf{W}_k \in \mathbb{R}^{D \times d_k}$ and $\mathbf{W}_v \in \mathbb{R}^{D \times d_k}$ are learnable weight matrices, and d_k is the projection dimension. Then, the self-attention (SA) output is a $N \times d_k$ matrix defined as the weighted sum of the value vectors for each token

$$\text{SA}(\mathbf{X}_d^{(i)}) = \text{softmax} \left(\frac{\mathbf{Q}^{(i)} \mathbf{K}^{(i)\top}}{\sqrt{d_k}} \right) \mathbf{V}^{(i)}, \quad (3)$$

where the weights are attention scores computed as the dot product between the query and key vectors, scaled by the square root of the projection dimension, and followed by softmax applied row-wise.

In the multi-head attention mechanism (MSA), multiple self-attention operations are performed in parallel, each with its own set of learnable weight matrices. For simplicity, we assume $d_k = D/h$, where h is the number of attention heads. The outputs of h heads are concatenated and linearly transformed using a learnable weight matrix $\mathbf{W}_o \in \mathbb{R}^{D \times D}$ to obtain the MSA output:

$$\text{MSA}(\mathbf{X}_d^{(i)}) = \text{Concat}(\mathbf{Y}_1^{(i)}, \dots, \mathbf{Y}_h^{(i)}) \mathbf{W}_o \in \mathbb{R}^{N \times D}, \quad (4)$$

where $\mathbf{Y}_j^{(i)} = \text{SA}_j(\mathbf{X}_d^{(i)}) \in \mathbb{R}^{N \times \frac{D}{h}}$ is the output of the j -th attention head.

As shown in Figure 1, the Transformer encoder architecture is comprised of a series of alternating layers, each containing an MSA mechanism and a feedforward neural network (FFN) block with SwiGLU activation function [23]. Before each block, layer normalization is applied, and after each block, residual connections are employed. Hence, the resulting output of the Transformer encoder is $\mathbf{Z}_d^{(i)} \in \mathbb{R}^{N \times D}$, which subsequently undergoes a flatten layer with linear head to produce the model predictions $\hat{\mathbf{x}}^{(i)} = (\hat{x}_{L+1}^{(i)}, \dots, \hat{x}_{L+T}^{(i)})^\top$.

Model Fine-Tuning. Parameter efficient fine-tuning (PEFT) methods offer an efficient way to adapt pre-trained LLMs to diverse downstream tasks without the need to fine-tune the entire set of model parameters. To fine-tune our pre-trained LLM, we employ quantized and low-rank adaptation (QLoRA) [7], a PEFT technique that aims to minimize the number of parameters updated during fine-tuning while preserving or improving performance on downstream tasks. QLoRA builds upon LoRA [8], which freezes the weight matrices of the linear projection layers within the self-attention mechanism of the Transformer encoder and represents each of their update matrices as a product of two trainable low-rank matrices. This low-rank matrix decomposition significantly reduces the number of trainable parameters, making the model more efficient for downstream tasks. In addition to low-rank decomposition, QLoRA also incorporates quantization techniques that reduce the precision of the weight values, thereby further decreasing the model's memory footprint and computational requirements without substantially sacrificing performance. With QLoRA, only 1.2% of the model's parameters are considered trainable, whereas using LoRA increases this percentage to 1.5% [3].

Model Alignment. Supervised fine-tuning is used as a fundamental precursor for subsequent forecasting fine-tuning of the LLaMA model. To enhance the alignment of our model with time series data after the supervised fine-tuning phase, we integrate direct preference optimization (DPO) [20] into our fine-tuning process. DPO, which involves the analysis of 10K comparison data pairs from the Ultra-Feedback dataset [5], is applied post-supervised fine-tuning to enhance model performance for specific data points. It optimizes its objective function, producing the optimal policy to an implicit reward function that is the best fit to the preference data. We use DPO to capture any change of variables, ensuring a more effective adaptation of the LLaMA model to the intricacies of time series forecasting. Finally, we transfer the updated weights of the backbone model to the forecasting fine-tuning phase.

Loss Function. The parameters of our FedTime model are learned by minimizing the training objective, which is the mean squared error (MSE) defined as

$$\mathcal{L} = \frac{1}{MT} \sum_{i=1}^M \sum_{\tau=L+1}^{L+T} \|\mathbf{x}_\tau^{(i)} - \hat{\mathbf{x}}_\tau^{(i)}\|^2, \quad (5)$$

where $\mathbf{x}_\tau^{(i)}$ and $\hat{\mathbf{x}}_\tau^{(i)}$ are the ground-truth and prediction, respectively, $\tau \in \{L+1, \dots, L+T\}$, L is the look-back window, T is the prediction length/horizon (i.e., future timesteps), and M is the number of time series variables (i.e., number of features or channels).

Algorithm. The main algorithmic steps of our federated learning framework are summarized in Algorithm 1. Prior to the learning process, K-means clustering is used as a pre-processing step to group edge devices into clusters. In practice, K-means clustering may occur in scenarios where a large number of edge devices are deployed in geographically dispersed locations, such as smart city infrastructures, industrial IoT deployments, or distributed sensor networks.

The FedTime algorithm initializes a global model, distributes it to selected edge devices for local training, aggregates the updated models from the participating devices at the central server, and iterates this process for multiple rounds, allowing the global model to gradually improve by aggregating insights from diverse edge devices while preserving data privacy. By training cluster-specific models, FedTime can capture patterns specific to different subsets of edge devices. Moreover, instead of transmitting updates from each individual edge device to the central server, FedTime aggregates updates from devices within the same cluster. This helps reduce the overall communication overhead by consolidating updates from devices with similar data distributions.

4 Experiments

In this section, we present experimental evaluation of the proposed federated learning approach by comparing it with recent state-of-the-art methods.

4.1 Experimental Setup

Datasets. We conduct a comprehensive evaluation of our FedTime model on various benchmark datasets: Weather¹, Traffic², Electricity³, and four ETT datasets (ETTh1, ETTh2, ETTm1, ETTm2) [28]. Dataset statistics are summarized in Table 1: (1) **Weather** is recorded

¹ <https://www.bgc-jena.mpg.de/wetter>

² <http://pems.dot.ca.gov>

³ <https://archive.ics.uci.edu/ml/datasets/ElectricityLoadDiagrams20112014>

Algorithm 1 FedTime: Federated LLM Algorithm

- 1: **Input:** Global model parameters θ ; number of edge devices S ; number of clusters K ; local datasets $\mathcal{D}_1, \dots, \mathcal{D}_S$
- 2: **Output:** Vector of aggregated global cluster-specific models θ_K
- 3: Apply K-means algorithm to split devices into K clusters $\mathcal{C} = \{c_1, c_2, \dots, c_K\}$ based on similar attributes.
- 4: Initialize global cluster model parameters θ_K
- 5: **for** each cluster c in \mathcal{C} **do**
- 6: **for** each round $r = 1, 2, \dots$ **do**
- 7: **while** model θ_K has not converged yet **do**
- 8: **for** each edge device $s \in \{1, \dots, S\}$ **in parallel do**
- 9: $\theta_K \leftarrow \text{UpdateDevice}(\mathcal{D}_s, \theta_c)$
- 10: **end for**
- 11: **end while**
- 12: Server aggregates final global models from cluster-specific models using weighted averaging

$$\theta_c = \frac{\sum_s w_{s,c} \theta_K}{\sum_s w_{s,c}}$$

where $w_{s,c}$ is the weight of device s in cluster c .

- 13: **end for**
- 14: Send global cluster-specific model θ_c to local device s
- 15: **end for**
- 16: **UpdateDevice**(\mathcal{D}_s, θ)
- 17: **for** each cluster c in \mathcal{C} **do**
- 18: **for** each local device s in cluster c **do**
- 19: Receive global model parameters θ from the server
- 20: Train local model using local data \mathcal{D}_s
- 21: Update local model parameters

$$\theta \leftarrow \theta - \eta \nabla \mathcal{L}(\theta; \mathcal{D}_s)$$

where η is the learning rate and $\mathcal{L}(\theta; \mathcal{D}_s)$ is the MSE local loss.

- 22: Send local model parameters θ to the server
- 23: **end for**
- 24: **end for**
- 25: **Return:** θ_K

at 10-minute intervals for the entire year of 2020, and contains 21 meteorological indicators, such as air temperature and humidity. (2) **Traffic** is a collection of hourly data from California Department of Transportation, and describes the road occupancy rates measured by different sensors on San Francisco Bay area freeways. (3) **Electricity** contains the hourly electricity consumption of 321 customers from 2012 to 2014. (4) **ETT** contains data collected from electricity transformers, including load and oil temperature that are recorded every 1 hour for ETTh datasets and every 15 minutes for ETTm datasets between July 2016 and July 2018.

Table 1. Summary statistics of benchmark datasets.

| Dataset | Features | Timesteps | Granularity |
|-------------|----------|-----------|-------------|
| Weather | 21 | 52,696 | 10 min |
| Traffic | 862 | 17,544 | 1 hour |
| Electricity | 321 | 26,304 | 1 hour |
| ETTh1 | 7 | 17,420 | 1 hour |
| ETTh2 | 7 | 17,420 | 1 hour |
| ETTm1 | 7 | 69,680 | 15 min |
| ETTm2 | 7 | 69,680 | 15 min |

Baseline Methods. We evaluate the performance of our FedTime model against recent state-of-the-art methods, including Pyraformer [14], Informer [32], Autoformer [4], FEDformer [33], DLinear [29], N-HiTS [2], PatchTST [18], TiDE [6], and LLM4TS [3].

Evaluation Metrics. We use the mean squared error (MSE) and mean absolute error (MAE) as metrics for evaluating both long-term time series forecasting. MSE measures the average squared difference between the predicted and actual values. A lower MSE indicates that the predictions are closer to the actual values on average. On the other hand, MAE measures the average absolute difference between the predicted and actual values. Similar to MSE, a lower MAE indicates better predictive accuracy, with predictions closer to the actual values. We report the performance at various prediction horizons and look-back windows.

Implementation Details. In our federated learning setup, we assume the presence of 555 eligible edge devices for training. All experiments are carried out on a GPU server with an AMD EPYC 7502 processor, an NVIDIA RTX A6000 48GB GDDR6 and 512GB memory. The algorithms are implemented in PyTorch and PySyft to simulate a federated learning environment. For hyper-parameters optimization, we apply a grid search on a batch size of range {128, 256, 512, 1024}, a learning rate of range {0.01, 0.001, 0.0001}, and a local training round of range {40, 80, 200}. Using hyper-parameter tuning, the batch size and learning rate are set to 512 and 0.001, respectively. In all experiments, we split the datasets into 80% training and 20% testing. To update QLoRA parameters, we employ FedAdam [21], an adaptive optimization method that helps improve the convergence of federated learning.

4.2 Results and Analysis

Long-Term Time Series Forecasting. We evaluate the forecasting performance of our FedTime model against strong baseline methods on seven forecasting benchmark datasets, and the results are reported in Table 2. For Transformer-based baselines, the values of the MSE and MAE error metrics are taken from [18]. The best results are shown in bold, and the second best results are underlined. As can be seen, FedTime consistently outperforms the baselines on most datasets, particularly for long-term prediction horizons. In terms of MSE, Table 2 reveals that FedTime performs better than LLM4TS in 24 out of 28 evaluations spanning seven datasets and four forecast horizons, improving upon this best performing baseline by relative error reductions of 15.56% and 20% on the Traffic and Electricity datasets, respectively, for $T = 720$, while maintaining a smaller number of learnable parameters compared to LLM4TS. Similarly, our method outperforms LLM4TS in 25 out of 28 MAE evaluations, yielding relative error reductions of 10.98% and 7.36% on the ETTm1 and ETTm2 datasets, respectively, for $T = 720$. Moreover, FedTime performs better or on par with the best performing LLM4TS baseline on the Weather dataset across all the prediction horizons. It is worth pointing out that MAE is more interpretable than MSE since it directly represents the average magnitude of prediction errors. A low MAE implies that FedTime provides accurate estimates of forecasting time series, thus demonstrating the effectiveness of the our framework.

In Figure 2, we report the visual results in terms of MSE on six forecasting datasets for various values of the look-back window length L ranging from 24 to 720. As can be seen, FedTime exhibits consistent and superior performance, especially for long look-

back windows. On the Traffic dataset, it is evident that FedTime holds a competitive edge, consistently outperforming Transformer-based models such as Informer, Autoformer and FEDformer. On the Weather dataset, FedTime stands out, especially between the look-back window lengths of 192 and 336. On the ETTh1 and ETTh2 datasets, FedTime’s performance peaks within the intermediate look-back window span of 48 to 192. Overall, FedTime yields superior performance over the baselines across the various forecasting datasets, especially for long-range look-back scenarios.

Federated Performance. In contrast to the centralized training of the LLaMA model, federated learning adopts a decentralized methodology, allowing for localized model updates and only transmitting the model weights, thus enhancing data privacy. In Figure 3, the training and testing learning curves show that our federated LLM model converges approximately $3\times$ faster than its centralized counterpart under analogous experimental conditions. We see that while the centralized LLaMA model faltered in convergence even post 200 epochs, the federated FedTime model adeptly minimized the loss function in merely 70 epochs. Also, the federated model’s uniform training trajectory, which is largely attributed to the incorporation of a cohesive cluster of edge devices, diverges sharply from the inconsistency seen in the centralized counterpart.

For the sake of federated comparison, we implemented a federated version of PatchTST [18], which we refer to as Fed-PatchTST. We evaluate the performance of FedTime against federated stacked LSTM (FSLSTM) [1] and Fed-PatchTST, and the results are reported in Table 3, which shows that FedTime consistently outperforms these federated baselines across all the forecasting benchmark datasets. This better performance is particularly notable for longer prediction horizons, highlighting the robustness and adaptability of FedTime to long-term time series forecasting. For instance, the percentage gap and percentage difference between FedTime and Fed-PatchTST, in terms of MAE, on the Electricity dataset are 39.24% and 32.8%, respectively, indicating a significant improvement in prediction accuracy with our model.

4.3 Communication Overhead

Evaluating the communication overhead is crucial in assessing the feasibility of a model’s deployment in federated environments. We use the real-world ACN dataset [11], which is an adaptive charging network dataset containing information about an electric vehicle (EV) single charging session at two distinct sites: California Institute of Technology (Caltech 540 stations) and Jet Propulsion Laboratory (JPL 400 stations). An exploratory data analysis of the dataset shows that the data distribution has an increasing trend at both Caltech and JPL sites. These upward trends in the distribution of energy demand at both sites underscores the need for enhanced capacity to accommodate these surges and mitigate escalated demand charges. The ACN dataset consists of 1.5 million charging sessions with a window time of 19 months. Both Caltech and JPL sites show a repetitive usage pattern with much higher utilization during weekdays than on weekends, as shown in Figure 4, which depicts the distribution of the energy delivered.

One of the key metrics of communication overhead is the data volume transferred, measured in Megabytes. As shown in Figure 5, our FedTime model yields reduced communication overhead compared to the baseline methods. This reduction is even more pronounced when considering metrics like message count and communication time. A key factor contributing to FedTime’s efficient performance is the strategic engagement of edge devices in the learning process. This

Table 2. Long-term time series forecasting results for various prediction lengths $T \in \{96, 192, 336, 720\}$. The best results are in **bold**, while the second-best are underlined.

| Method | Autoformer | | FEDformer | | DLinear | | N-HiTS | | PatchTST/64 | | TiDE | | LLM4TS | | FedTime | | |
|-------------|------------|-------|-----------|-------|---------|-------|--------|-------|-------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | MSE | MAE | |
| Weather | 96 | 0.249 | 0.329 | 0.238 | 0.314 | 0.176 | 0.237 | 0.158 | 0.195 | 0.149 | 0.198 | 0.166 | 0.222 | 0.147 | 0.196 | 0.150 | 0.194 |
| | 192 | 0.325 | 0.370 | 0.275 | 0.329 | 0.220 | 0.282 | 0.211 | 0.247 | 0.194 | 0.241 | 0.209 | 0.263 | <u>0.191</u> | <u>0.238</u> | 0.180 | 0.217 |
| | 336 | 0.351 | 0.391 | 0.339 | 0.377 | 0.265 | 0.319 | 0.274 | 0.300 | 0.245 | 0.282 | 0.254 | 0.301 | <u>0.241</u> | <u>0.227</u> | 0.209 | 0.227 |
| | 720 | 0.415 | 0.426 | 0.389 | 0.409 | 0.323 | 0.362 | 0.401 | 0.413 | 0.314 | 0.334 | 0.313 | 0.340 | 0.313 | 0.329 | <u>0.318</u> | <u>0.335</u> |
| Traffic | 96 | 0.597 | 0.371 | 0.576 | 0.359 | 0.410 | 0.282 | 0.402 | 0.282 | 0.360 | 0.249 | <u>0.336</u> | 0.253 | 0.372 | 0.259 | 0.328 | 0.215 |
| | 192 | 0.607 | 0.382 | 0.610 | 0.380 | 0.423 | 0.287 | 0.420 | 0.297 | 0.379 | <u>0.256</u> | <u>0.346</u> | 0.257 | 0.391 | 0.265 | 0.339 | 0.233 |
| | 336 | 0.623 | 0.387 | 0.608 | 0.375 | 0.436 | 0.296 | 0.448 | 0.313 | 0.392 | 0.264 | <u>0.355</u> | 0.260 | 0.405 | <u>0.275</u> | 0.347 | 0.288 |
| | 720 | 0.639 | 0.395 | 0.621 | 0.375 | 0.466 | 0.315 | 0.539 | 0.353 | 0.432 | 0.286 | <u>0.386</u> | <u>0.273</u> | 0.437 | 0.292 | 0.369 | 0.239 |
| Electricity | 96 | 0.196 | 0.304 | 0.186 | 0.311 | 0.140 | 0.237 | 0.147 | 0.249 | 0.129 | 0.222 | 0.132 | 0.229 | <u>0.128</u> | <u>0.223</u> | 0.118 | 0.158 |
| | 192 | 0.211 | 0.324 | 0.197 | 0.311 | 0.153 | 0.249 | 0.167 | 0.269 | 0.147 | <u>0.240</u> | <u>0.147</u> | 0.243 | <u>0.146</u> | 0.242 | 0.133 | 0.315 |
| | 336 | 0.214 | 0.327 | 0.213 | 0.314 | 0.169 | 0.267 | 0.186 | 0.290 | 0.163 | <u>0.259</u> | <u>0.161</u> | 0.261 | 0.163 | <u>0.258</u> | 0.148 | 0.255 |
| | 720 | 0.236 | 0.342 | 0.233 | 0.344 | 0.203 | 0.301 | 0.243 | 0.340 | 0.197 | 0.290 | <u>0.196</u> | 0.294 | 0.220 | <u>0.292</u> | 0.176 | 0.288 |
| ETTh1 | 96 | 0.435 | 0.446 | 0.376 | 0.415 | 0.375 | 0.399 | 0.378 | 0.436 | 0.379 | 0.401 | <u>0.375</u> | 0.398 | 0.371 | 0.394 | 0.392 | 0.402 |
| | 192 | 0.456 | 0.457 | 0.423 | 0.446 | 0.405 | 0.420 | 0.427 | 0.436 | 0.413 | 0.429 | 0.412 | 0.422 | 0.403 | <u>0.412</u> | <u>0.404</u> | 0.360 |
| | 336 | 0.486 | 0.487 | 0.444 | 0.462 | 0.439 | 0.443 | 0.458 | 0.484 | 0.435 | 0.436 | 0.435 | 0.433 | 0.428 | <u>0.422</u> | <u>0.425</u> | 0.421 |
| | 720 | 0.515 | 0.517 | 0.469 | 0.492 | 0.472 | 0.490 | 0.472 | 0.551 | <u>0.446</u> | 0.464 | 0.454 | 0.465 | <u>0.422</u> | <u>0.444</u> | 0.414 | 0.436 |
| ETTh2 | 96 | 0.332 | 0.368 | 0.332 | 0.374 | 0.289 | 0.353 | 0.274 | 0.345 | 0.274 | 0.337 | 0.270 | 0.336 | <u>0.269</u> | <u>0.332</u> | 0.249 | 0.303 |
| | 192 | 0.426 | 0.434 | 0.407 | 0.446 | 0.383 | 0.418 | 0.353 | 0.401 | 0.338 | <u>0.376</u> | 0.332 | 0.380 | <u>0.328</u> | <u>0.377</u> | 0.318 | 0.343 |
| | 336 | 0.477 | 0.479 | 0.400 | 0.447 | 0.448 | 0.465 | 0.382 | 0.425 | 0.363 | 0.397 | 0.360 | 0.407 | <u>0.353</u> | <u>0.396</u> | 0.344 | 0.368 |
| | 720 | 0.453 | 0.490 | 0.412 | 0.469 | 0.605 | 0.551 | 0.625 | 0.557 | 0.393 | 0.430 | 0.419 | 0.451 | <u>0.383</u> | <u>0.425</u> | 0.376 | 0.428 |
| ETTm1 | 96 | 0.510 | 0.492 | 0.326 | 0.390 | 0.299 | 0.343 | 0.302 | 0.350 | 0.293 | 0.346 | 0.306 | 0.349 | <u>0.285</u> | <u>0.343</u> | 0.283 | 0.339 |
| | 192 | 0.514 | 0.495 | 0.365 | 0.415 | 0.335 | 0.365 | 0.347 | 0.383 | 0.333 | 0.370 | 0.335 | <u>0.366</u> | <u>0.324</u> | <u>0.366</u> | 0.301 | 0.338 |
| | 336 | 0.510 | 0.492 | 0.392 | 0.425 | 0.369 | 0.386 | 0.369 | 0.402 | 0.369 | 0.392 | 0.364 | <u>0.384</u> | <u>0.353</u> | 0.385 | 0.319 | 0.365 |
| | 720 | 0.527 | 0.493 | 0.446 | 0.458 | 0.425 | 0.421 | 0.431 | 0.441 | 0.416 | 0.420 | 0.413 | <u>0.413</u> | <u>0.408</u> | 0.419 | 0.328 | 0.373 |
| ETTm2 | 96 | 0.205 | 0.293 | 0.180 | 0.271 | 0.167 | 0.260 | 0.176 | 0.255 | 0.166 | 0.256 | <u>0.161</u> | <u>0.251</u> | 0.165 | 0.254 | 0.148 | 0.227 |
| | 192 | 0.278 | 0.336 | 0.252 | 0.318 | 0.224 | 0.303 | 0.245 | 0.305 | 0.223 | 0.296 | <u>0.215</u> | <u>0.289</u> | 0.225 | 0.292 | 0.197 | 0.261 |
| | 336 | 0.343 | 0.379 | 0.324 | 0.364 | 0.281 | 0.342 | 0.295 | 0.346 | 0.274 | 0.329 | <u>0.267</u> | <u>0.326</u> | 0.268 | 0.326 | 0.219 | 0.279 |
| | 720 | 0.414 | 0.419 | 0.410 | 0.420 | 0.397 | 0.421 | 0.401 | 0.413 | 0.362 | 0.385 | <u>0.352</u> | 0.383 | 0.356 | <u>0.380</u> | 0.305 | 0.352 |

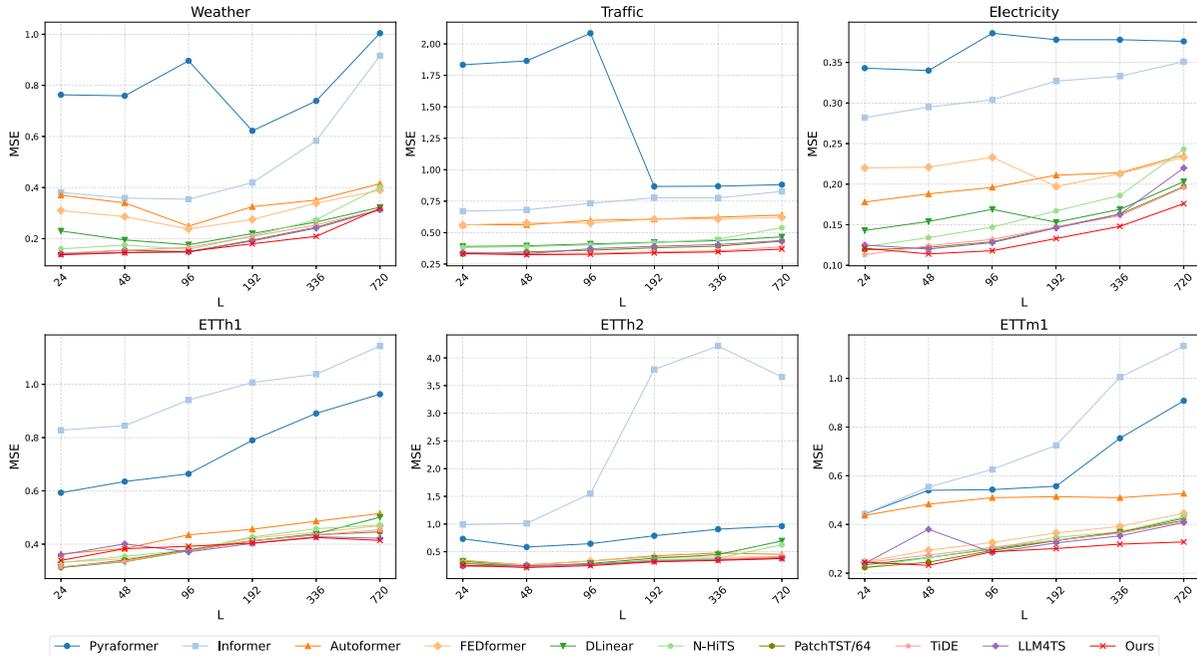


Figure 2. Long-term forecasting performance with varying look-back window lengths $L \in \{24, 48, 96, 192, 336, 720\}$. We set the prediction horizon to $T = 720$.

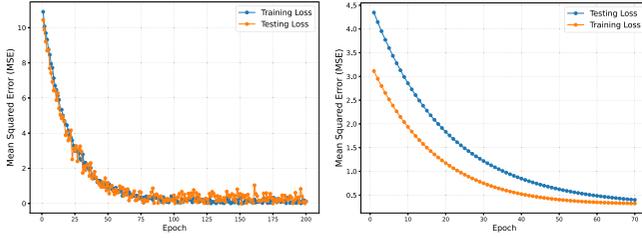


Figure 3. Training and testing learning curves of the centralized LLaMA model (left) and FedTime (right).

Table 3. Performance comparison of FedTime and federated models across benchmark datasets with prediction length $T = 720$.

| Method | Fed-PatchTST | | FSLSTM | | FedTime | |
|-------------|--------------|-------|--------|-------|--------------|--------------|
| | MSE | MAE | MSE | MAE | MSE | MAE |
| Weather | 0.363 | 0.495 | 0.451 | 0.421 | 0.318 | 0.335 |
| Traffic | 0.388 | 0.272 | 0.450 | 0.386 | 0.369 | 0.239 |
| Electricity | 0.231 | 0.401 | 0.310 | 0.466 | 0.176 | 0.288 |
| ETTh1 | 0.456 | 0.520 | 0.522 | 0.621 | 0.414 | 0.436 |
| ETTh2 | 0.402 | 0.521 | 0.452 | 0.533 | 0.376 | 0.428 |
| ETTh1 | 0.340 | 0.421 | 0.352 | 0.483 | 0.328 | 0.373 |
| ETTh2 | 0.335 | 0.384 | 0.401 | 0.495 | 0.305 | 0.352 |

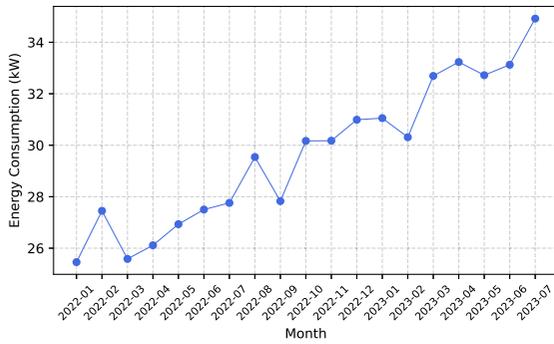


Figure 4. ACN energy delivery distribution based on disconnection time.

decentralized approach, leveraging the inherent distributed nature of IoT-enabled devices, facilitates a system where updating model weights incurs a significantly lower overhead than transferring entire datasets. Such an optimized model structure and design strategy make FedTime notably efficient in comparison with centralized architectures such as Informer, Pyraformer, Autoformer, and LLM4TS. Moreover, Figure 5 alludes to the efficiency of models with optimized weight structures using PEFT techniques.

4.4 Ablation Study

In Figure 6, we evaluate the performance of various FedTime variants on the ACN dataset for Caltech site. In the following, we introduce these variants and analyze their respective performances:

- **FedTime without Clustering:** This model aligns relatively well with the actual energy consumption in general up to the 20-hour mark. Yet, between the 20- and 40-hour intervals, the model tends to underestimate consumption. Between 60 to 80 hours, the model’s predictions come closer to the actual data, but discrepancies are noticeable toward the end of the 100-hour period. The curve’s fluctuations are likely due to the mix of devices from different-sized clusters participating in this specific training round.

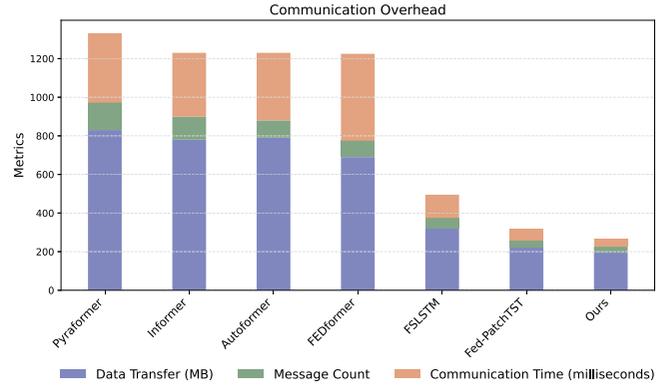


Figure 5. Communication overhead comparison between FedTime and baseline methods.

- **FedTime without PEFT:** The performance without parameter efficient fine-tuning generally tracks the consumption trend. Yet, certain intervals display noticeable over-estimations, and there are slight variances from the actual energy consumption, especially during peak and trough periods.
- **FedTime with Clustering+PEFT:** This variant manifests a notable performance over the 100-hour duration. Distinctly evident is its aptitude in closely mirroring the authentic consumption, especially in regions characterized by pronounced fluctuations, such as the 20-40 and 60-80 hour intervals. The superior performance of this model makes it the best choice for energy consumption prediction over the scrutinized 100-hour interval at Caltech site.

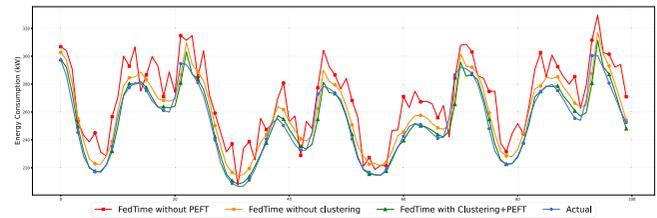


Figure 6. Comparison of predicted and actual energy consumption on Caltech site during weekdays (100-hour timestep) using FedTime variants. The horizontal axis represents the different time periods in hours.

5 Conclusion

We introduced FedTime, a federated large language model for long-term time series forecasting. FedTime leverages federated learning, ensuring decentralized model training while upholding data privacy across edge devices. Our empirical evaluations underscore FedTime’s superiority over both centralized and federated baselines in long-term time series forecasting, demonstrating significant performance enhancements across diverse datasets. Moreover, FedTime addresses communication overhead challenges by optimizing the transmission of model updates, thereby minimizing resource consumption while enhancing forecasting accuracy. In summary, FedTime offers a versatile and efficient solution that strikes a balance between privacy preservation, communication efficiency, and forecasting accuracy in decentralized settings. Looking ahead, our future research endeavors will explore synergies between federated frameworks and other distributed systems, with a keen focus on integrating blockchain technology.

Acknowledgements

This work was supported in part by the Discovery Grants Program of the Natural Sciences and Engineering Research Council of Canada under grant RGPIN-2024-04291.

References

- [1] R. Abdel-Sater and A. B. Hamza. A federated learning approach to anomaly detection in smart buildings. *ACM Transactions on Internet of Things*, 2(4):1–23, 2021.
- [2] C. Challu, K. G. Olivares, B. N. Oreshkin, F. G. Ramirez, M. M. Canseco, and A. Dubrawski. N-HITS: Neural hierarchical interpolation for time series forecasting. In *Proc. AAAI Conference on Artificial Intelligence*, pages 6989–6997, 2023.
- [3] C. Chang, W.-C. Peng, and T.-F. Chen. LLM4TS: Two-stage fine-tuning for time-series forecasting with pre-trained LLMs. *arXiv preprint arXiv:2308.08469*, 2023.
- [4] M. Chen, H. Peng, J. Fu, and H. Ling. Autoformer: Searching transformers for visual recognition. In *Proc. IEEE International Conference on Computer Vision*, pages 12270–12280, 2021.
- [5] G. Cui, L. Yuan, N. Ding, G. Yao, W. Zhu, Y. Ni, G. Xie, Z. Liu, and M. Sun. UltraFeedback: Boosting language models with high-quality feedback. In *arXiv preprint arXiv:2310.01377*, 2023.
- [6] A. Das, W. Kong, A. Leach, R. Sen, and R. Yu. Long-term forecasting with TiDE: Time-series dense encoder. *Transactions on Machine Learning Research*, 2023.
- [7] T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer. QLoRA: Efficient finetuning of quantized LLMs. In *Advances in Neural Information Processing Systems*, 2023.
- [8] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022.
- [9] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*.
- [10] T. Kim, J. Kim, Y. Tae, C. Park, J.-H. Choi, and J. Choo. Reversible instance normalization for accurate time-series forecasting against distribution shift. In *International Conference on Learning Representations*, 2022.
- [11] Z. J. Lee, T. Li, and S. H. Low. ACN-Data: Analysis and applications of an open EV charging dataset. In *Proc. International Conference on Future Energy Systems*, 2019.
- [12] Q. Li, Z. Wen, Z. Wu, S. Hu, N. Wang, Y. Li, X. Liu, and B. He. A survey on federated learning systems: Vision, hype and reality for data privacy and protection. *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [13] Q. Li, Y. Diao, Q. Chen, and B. He. Federated learning on non-IID data silos: An experimental study. In *Proc. IEEE International Conference on Data Engineering*, pages 965–978, 2022.
- [14] S. Liu, H. Yu, C. Liao, J. Li, W. Lin, A. X. Liu, and S. Dustdar. Pyraformer: Low-complexity pyramidal attention for long-range time series modeling and forecasting. In *International Conference on Learning Representations*, 2021.
- [15] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Proc. International Conference on Artificial Intelligence and Statistics*, pages 1273–1282, 2017.
- [16] J. Mills, J. Hu, and G. Min. Communication-efficient federated learning for wireless edge intelligence in IoT. *IEEE Internet of Things Journal*, 7(7):5986–5994, 2019.
- [17] S. Murugesan and A. K. Cherukuri. The rise of generative artificial intelligence and its impact on education: The promises and perils. *Computer*, 56(5):116–121, 2023.
- [18] Y. Nie, N. H. Nguyen, P. Sinthong, and J. Kalagnanam. A time series is worth 64 words: long-term forecasting with Transformers. In *International Conference on Learning Representations*, 2023.
- [19] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray, et al. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, pages 27730–27744, 2022.
- [20] R. Rafailov, A. Sharma, E. Mitchell, S. Ermon, C. D. Manning, and C. Finn. Direct preference optimization: Your language model is secretly a reward model. In *Advances in Neural Information Processing Systems*, 2023.
- [21] S. J. Reddi, Z. Charles, M. Zaheer, Z. Garrett, K. Rush, J. Konečný, S. Kumar, and H. B. McMahan. Adaptive federated optimization. In *International Conference on Learning Representations*, 2021.
- [22] Y. M. Saputra, D. T. Hoang, D. N. Nguyen, E. Dutkiewicz, M. D. Mueck, and S. Srikanteswara. Energy demand prediction with federated learning for electric vehicle networks. In *Proc. IEEE Global Communications Conference*, 2019.
- [23] N. Shazeer. GLU variants improve Transformer. *arXiv preprint arXiv:2002.05202*, 2020.
- [24] R. Taori, I. Gulrajani, T. Zhang, Y. Dubois, X. Li, C. Guestrin, P. Liang, and T. B. Hashimoto. Stanford Alpaca: An instruction-following LLaMA model, 2023.
- [25] R. Thoppilan, D. De Freitas, J. Hall, N. Shazeer, A. Kulshreshtha, H.-T. Cheng, A. Jin, T. Bos, L. Baker, Y. Du, et al. LaMDA: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*, 2022.
- [26] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, et al. LLaMA: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [27] J. Wen, Z. Zhang, Y. Lan, Z. Cui, J. Cai, and W. Zhang. A survey on federated learning: challenges and applications. *International Journal of Machine Learning and Cybernetics*, 14(2):513–535, 2023.
- [28] H. Wu, J. Xu, J. Wang, and M. Long. Autoformer: Decomposition Transformers with auto-correlation for long-term series forecasting. In *Advances in Neural Information Processing Systems*, pages 22419–22430, 2021.
- [29] A. Zeng, M. Chen, L. Zhang, and Q. Xu. Are Transformers effective for time series forecasting? In *Proc. AAAI conference on Artificial Intelligence*, pages 11121–11128, 2023.
- [30] Z. Zhang, A. Zhang, M. Li, H. Zhao, G. Karypis, and A. Smola. Multimodal chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923*, 2023.
- [31] L. Zheng, W.-L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. Xing, et al. Judging LLM-as-a-judge with MT-bench and chatbot arena. In *Advances in Neural Information Processing Systems*, 2023.
- [32] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, and W. Zhang. Informer: Beyond efficient Transformer for long sequence time-series forecasting. In *Proc. AAAI Conference on Artificial Intelligence*, pages 11106–11115, 2021.
- [33] T. Zhou, Z. Ma, Q. Wen, X. Wang, L. Sun, and R. Jin. FEDformer: Frequency enhanced decomposed Transformer for long-term series forecasting. In *Proc. International Conference on Machine Learning*, pages 27268–27286, 2022.