

Prompt Federated Learning for Weather Forecasting: Toward Foundation Models on Meteorological Data

Shengchao Chen, Guodong Long, Tao Shen and Jing Jiang

Australian Artificial Intelligence Institute, FEIT, University of Technology Sydney
 shengchao.chen.uts@gmail.com, {guodong.long, tao.shen, jing.jiang}@uts.edu.au

Abstract

To tackle the global climate challenge, it urgently needs to develop a collaborative platform for comprehensive weather forecasting on large-scale meteorological data. Despite urgency, heterogeneous meteorological sensors across countries and regions, inevitably causing multivariate heterogeneity and data exposure, become the main barrier. This paper develops a foundation model across regions capable of understanding complex meteorological data and providing weather forecasting. To relieve the data exposure concern across regions, a novel federated learning approach has been proposed to collaboratively learn a brand-new spatio-temporal Transformer-based foundation model across participants with heterogeneous meteorological data. Moreover, a novel prompt learning mechanism has been adopted to satisfy low-resourced sensors' communication and computational constraints. The effectiveness of the proposed method has been demonstrated on classical weather forecasting tasks using three meteorological datasets with multivariate time series.

1 Introduction

Climate change will significantly impact all regions; however, the specific effects will vary [Kjellstrom *et al.*, 2016]. Increasing global temperatures and melting ice will lead to alterations in sea levels, ocean currents, weather patterns, and cloud cover [Hagemann *et al.*, 2013]. To effectively tackle the challenge of global climate change, the implementation of a large-scale collaborative data-sharing platform is essential. Although this work is labor-intensive and demands a multitude of skilled experts, the utilization of machine learning techniques can enhance efficiency in addressing this problem. Nonetheless, the machine learning domain faces a challenge when attempting to employ a centralized uniform model to serve all regions due to their heterogeneity. An effective solution involves pre-training a foundational model using extensive weather data and enabling each region to fine-tune the model using a relatively small data to enhance its ability to capture local weather patterns.

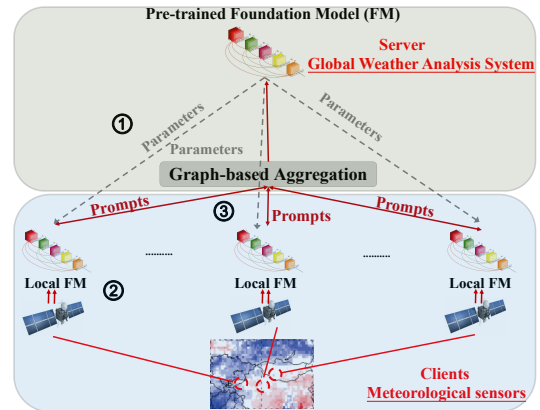


Figure 1: Our *MetePFL* for weather forecasting. i) pre-trained FM initializes the local FM; ii) local FM trains using local data; iii) the server aggregates and transmits local prompts' parameters.

Weather forecasting is a fundamental analytical task aimed at modeling the dynamic changes of weather on both a global and regional scale. Multi-sensor weather forecasting serves as a critical tool in mitigating the loss of human lives and property by providing early warnings for extreme weather events resulting from global climate change [Chattopadhyay *et al.*, 2020]. The objective of this approach is to capture potential correlations between multiple meteorological factors and the tendency of weather variations in order to gain a comprehensive understanding of specific regions. Unlike conventional time-series data, weather time-series data in meteorology are gathered from sensing devices distributed across diverse geographical locations [Campbell and Diebold, 2005].

One-step forecasting [Chen *et al.*, 2022b; Chen *et al.*, 2023], empowered by recent advancements in deep learning, have garnered considerable attention due to their high efficiency. However, the performance of these strategies is hindered by the non-stationary nature of weather changes. This limitation arises from their reliance on fixed patterns derived from prior knowledge. To address this issue and capture temporal information, other studies [Karevan and Suykens, 2020; Alléon *et al.*, 2020] suggest formulating the tasks as autoregression problems. These studies utilize preceding time step variables to predict variables at the subsequent time step [Chen and Lai, 2011]. This technique is commonly im-

plemented using RNN or Transformer models. The choice of these models is based on their superior performance in time series analysis [Shi *et al.*, 2015].

However, previous works focus on using a uniform model to serve all regions regardless of heterogeneity. In contrast, foundation model (FM) represents a novel service architecture that aims to pre-train a large model with extensive data. Subsequently, this model can be fine-tuned for specific tasks using relevant data, such as weather forecasting in a particular region. The FM has the capability to capture common knowledge shared among multiple tasks or participants. Further refinements can then enhance its alignment with the specific requirements of a given task. The FM has demonstrated remarkable success in Natural Language Processing (NLP), exemplified by ChatGPT. Notably, recent progress in foundation models has been observed across diverse domains, including ViT [Xu *et al.*, 2022], BERT [Yates *et al.*, 2021], and CLIP [Radford *et al.*, 2021].

In contrast to existing FMs, training a FM on weather forecasting tasks must tackle the following challenges. First, sharing raw data across countries/regions will not be easy. Second, transmitting and processing the continuously collected data is a challenge for low-resourced sensors or devices. Third, real-time forecasting is critically important. In summary, we need a solution to tackle data security, communication, and computation efficiency issues and provide on-device decisions independently.

This paper will design a novel machine-learning approach to train foundation models on weather forecasting tasks. The model will be capable of understanding and constructing the complex spatiotemporal relationship of meteorological data to provide reliable analysis support on weather forecasting and global climate challenges. Specifically, we propose a novel **Meteorological Prompt Federated Learning (MetePFL)** approach to collaboratively learn a Transformer-based foundation model (FM) across devices with multivariate time-series data (see Figure 1). The *MetePFL* only considers the model parameters exchange among devices rather than direct data sharing. Considering the low-resourced sensors' communication efficiency constraint, a brand-new prompt learning mechanism is introduced upon a pre-trained FM to comprehensively explore the correlation among weather-related variables while computing a few parameters.

Three weather forecasting datasets based on multivariate time series with multiple meteorological factors, i.e., precipitation, temperature, upstream, are leveraged to verify the effectiveness of *MetePFL*. The main contributions of this work are summarized in four-fold:

- This is the first work to explore a foundation model-based solution to enhance weather forecasting tasks towards a global scale.
- The proposed prompt federated learning approach is a novel mechanism to collaboratively learn a foundation model for the applications with many satellite sites or stations across regions.
- A spatio-temporal prompt learning mechanism has been designed to efficiently tackle multivariable time series.

- Experiments on three real datasets have demonstrated the effectiveness and superior of our proposed solution. It is worth noting that we obtain excellent performance with only 2.38% of the model's parameters trained.

2 Related Work

Weather Forecasting. Weather forecasting plays a crucial role in the global climate analysis system. Conventional forecasting methods utilize numerical weather prediction (NWP) models, which incorporate physical constraints to simulate weather phenomena [Bauer *et al.*, 2015]. However, with the emergence of data-driven approaches, weather forecasting has shifted towards approaches driven by data, such as ARIMA [Chen and Lai, 2011], SVM [Sapankevych and Sankar, 2009], and NNs [Voyant *et al.*, 2012]. While these basic models exhibit potential, they encounter difficulties in comprehending nonlinear temporal dynamics. Deep Learning methods, especially models based on Recurrent Neural Networks (RNNs), have exhibited promising outcomes in weather forecasting [Shi *et al.*, 2015]. Recently, Transformers have demonstrated superior performance compared to RNN-based models in time series analysis [Bojesomo *et al.*, 2021], thereby gaining popularity for weather-related tasks [Chen *et al.*, 2023]. However, these models neglect the data exposure concerns when utilizing multi-sensor data for real-world tasks and modeling spatio-temporal correlations.

Foundation Model. A fully supervised learning paradigm needs a large scale of data. The foundation model provides a practical solution for scenario-specific tasks, aiming to pre-train a model using extensive prior knowledge. The FM has widespread applications in NLP [Yates *et al.*, 2021] and CV [Xu *et al.*, 2022; Radford *et al.*, 2021], providing an effective cross-task learning strategy. For example, ChatGPT can be used as a baseline model for researchers to fine-tune it to achieve more accurate responses for downstream tasks.

Federated Learning. Federated learning (FL) is a new learning paradigm to embody the collaborative training of models without requiring data exposure from each participant (e.g., meteorological sensors) [McMahan *et al.*, 2017; Long *et al.*, 2021; Long *et al.*, 2020; Jiang *et al.*, 2020]. Vanilla FL suffer from the heterogeneity of data and devices. Personalized FL aims to solve the problem by multiple techniques [Tan *et al.*, 2022a; Chen *et al.*, 2022a; Wang *et al.*, 2022; Ma *et al.*, 2022; Tan *et al.*, 2021; Zhang *et al.*, 2023; Li *et al.*, 2023; Long *et al.*, 2023; Li *et al.*, 2019; Li *et al.*, 2020; Gao *et al.*, 2022] via training better personalized model for each client. However, these methods are not suitable for weather forecasting due to all parameters must be considered during communication so that hinder the real-time forecasting. pre-train-based strategy can mitigate the problem [Tan *et al.*, 2022b] but can not explore the spatio-temporal correlations. Different from the above methods, this paper focuses on establishing a high-efficiency FL approach that provides analytical support to across-regional weather forecasting systems with heterogeneous meteorological data.

Prompt Learning. Prompt learning as a lightweight mechanism is widely used in NLP [Li and Liang, 2021; Liu *et al.*,

2021], which requires fewer parameters and is more adaptive than fine-tuned pre-trained models by represented by several prompt tuning strategies in different applications [Zhou *et al.*, 2022a]. Different from language data, understanding multidimensional correlation among multivariate data in the weather forecasting task is critical. However, the key point is often ignored by the federated prompt learning method [Guo *et al.*, 2022]. The paper introduces a novel prompt mechanism within the FL framework based on pre-trained FM to explore the temporal dynamics and the potential correlation among clients while computing only a few parameters.

3 Problem Formulation

Given N clients that possess individual local private datasets D , each client has a multivariate time series denoted as $\mathbf{X}_i \in \mathbb{R}^{m \times n}$. In this notation, each sample at a specific time step t is represented as $\mathbf{x}_t \in \mathbb{R}^{1 \times n}$. Weather forecasting using multivariate time series can be defined as the process of utilizing historical values of all variables for a duration of P periods to predict the values of a specific variable in the future over Q periods, can be defined below:

$$[\mathbf{x}_{t-P}, \mathbf{x}_{t-P+1}, \dots, \mathbf{x}_t] \xrightarrow{f} [\mathbf{x}'_{t+1}, \mathbf{x}'_{t+2}, \dots, \mathbf{x}'_{t+Q}], \quad (1)$$

where f is a learning system, and $\mathbf{x}'_t \in \mathbb{R}^{1 \times 1}$ is the value of the variable to be forecasting at the t -th time step. Valida FL system aims to minimize the average loss of the global model w on all clients' local dataset:

$$F(w) := \arg \min_{w_1, w_2, \dots, w_N} \sum_{k=1}^N \frac{n_k}{n} F_k(w_k), \quad (2)$$

where n_k is the number of samples hold by the k -th client. n is the number of samples held by all clients. $F_k(w_k)$ denotes the local objective function of k -th client. The distinguishing factor is that each client possesses a unique pattern, and sensors are deployed in specific locations, resulting in a statistically heterogeneous environment. To address this challenge, PFL is typically modeled as a bi-level optimization problem.

$$F(v; w) := \arg \min_{\{w_1, w_2, \dots, w_N\}, \{v_1, v_2, \dots, v_N\}} \sum_{k=1}^N G_k(v_k, w),$$

$$i.e. \quad G_k(v_k, w) = \frac{n_k}{n} F_k(v_k) + \lambda R(v_k, w), \quad (3)$$

where each client hold a personalized model parameterized by v_i , w denotes the global model. $R(\cdot)$ is the regularization term to control model update, via avoiding the local model updating be far away to the optimal global model.

4 Meteorological Prompt Federated Learning

The framework of *MetePFL* is depicted in Figure 1. In contrast to conventional Federated Learning (FL) where random global parameters are broadcasted to each client, *MetePFL* employs a fixed FM, thereby reducing computation costs and improving performance without requiring extensive back-propagation. During each round, only the prompt parameters of the clients are taken into consideration. The *MetePFL* framework consists of the Spatial-Temporal Prompt (STP) and the optimization process.

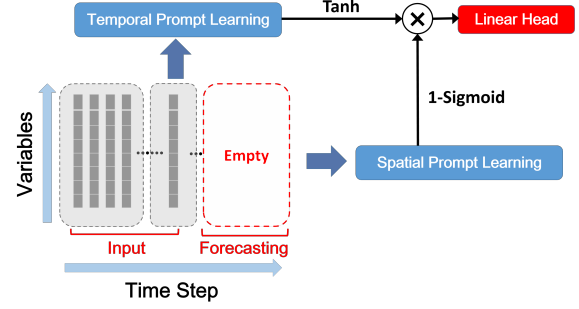


Figure 2: Schematic of Spatial-Temporal Prompt Learning.

4.1 Spatial-Temporal Prompt

The Spatial-Temporal Prompt (STP) shown in Figure 2 can be divided into Temporal prompt learning (TPL) and Spatial prompt learning (SPL).

Temporal prompt learning (TPL). To effectively understand the temporal dynamics under multivariate interactions, we use a multi-step incremental learning mechanism for learning the prompt parameters along the time dimension. The TPL comprise four phases, is shown in Figure 3.

Given a time-series $\mathbf{X} \in \mathbb{R}^{m \times n}$, the m and n represent the number of time steps and variables, respectively. Define an initial step l , the four temporal prompt format: (1) initial prompt $\hat{\mathbf{P}}$; (2) Temporal prompt-I: \mathbf{P}_1 ; (3) Temporal prompt-II: $\mathbf{P}_2 \in \mathbb{R}^{2l \times n}$; (4) Temporal prompt-III: $\mathbf{P}_3 \in \mathbb{R}^{2l \times n}$.

The first p steps are fixed as $Seq \in \mathbb{R}^{p \times n}$. In addition, the data within the p to k steps are spliced with the $\hat{\mathbf{P}} \in \mathbb{R}^{(k-p+l) \times n}$ to generate \mathbf{P}_1 by

$$\|\mathbf{X}^{(p \sim k)}, \hat{\mathbf{P}}\| \rightarrow \mathbf{P}_{T,1}, \mathbf{P}_{T,1} \in \mathbb{R}^{(2k-2p+l) \times n}, \quad (4)$$

where $\|\|\|$ denote the *concat* operation. To enhance the accuracy of forecasting performance, we adopt a two-stage objective in the initial learning phase: " $\mathbf{X}^{(0 \sim p)} \xrightarrow{1} \mathbf{X}^{(0 \sim k)} \xrightarrow{2} \mathbf{X}^{(0 \sim k+l)}$ ". During this phase, we encourage the FM to enhance prompt parameters associated with prior knowledge in $\mathbf{X}^{(p \sim k)}$ in order to establish a foundation for the subsequent stage. In the next stage, the FM learns the remaining $\mathbf{P}_{T,1}$ parameters in $\mathbf{X}^{(k \sim k+l)}$. This approach enables the FM to develop a comprehensive understanding of the relationships among variables across different time periods, rather than establishing rigid back-and-forth associations. The formulation of the initial learning phase is:

$$\mathbf{R}_1 = \text{FM}([\text{Seq}, \text{embed}(\mathbf{P}_{T,1})]), \quad (5)$$

where the $\text{embed}(x)$ represents the learnable position encoder, and the FM is the pre-trained foundation model.

Splicing discrete prompts to form a complete series is challenging because their parameters cannot match the FM simultaneously. To establish the relationship between two prompts, a dual-correct strategy is adopted in the second learning phase. This strategy encourages the FM to correct $\mathbf{P}_{T,1}$ based on the previous $\mathbf{P}_{T,1}$ while learning $\mathbf{P}_{T,2}$. The objectives of the second learning phase are " $\mathbf{X}^{(0 \sim k+l)} \rightarrow \mathbf{X}^{(0 \sim k+3l)}$ ". This process can be formulated as mapping the input data

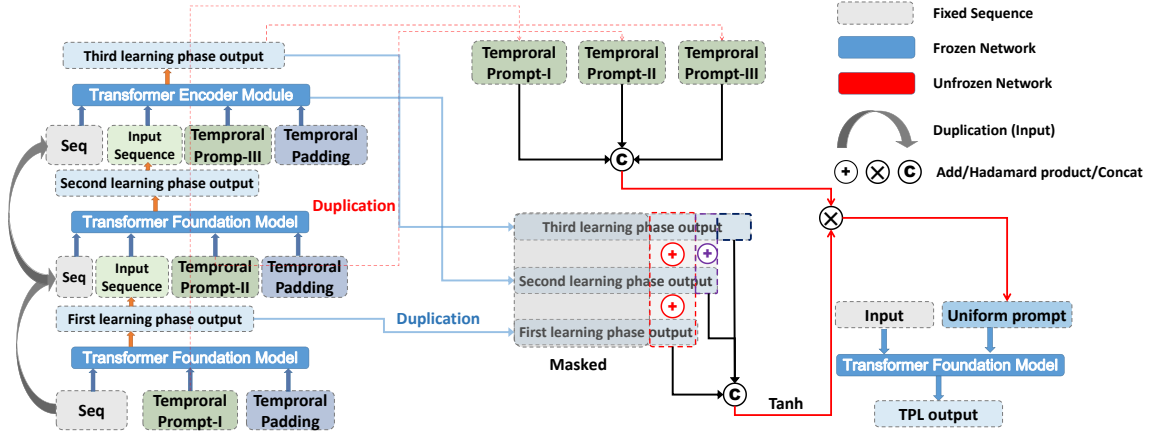


Figure 3: The learning strategy of Temporal Prompt Learning, which consists of four different learning phases.

from $\mathbf{X}^{(0 \sim k+l)}$ to $\mathbf{X}^{(0 \sim k+3l)}$ during the second phase.

$$\begin{aligned} \|\mathbf{P}_{T,1}, \mathbf{P}_{T,2}\| &\rightarrow \mathbf{P}_{T,2}, \mathbf{P}_{T,2} \in \mathbb{R}^{3l \times n}, \\ \mathbf{R}_2 &= \text{FM}(\|\text{Seq}_2, \mathbf{R}_1^{p/2 \sim p}, \text{embed}(\mathbf{R}_1^{p \sim p+l}), \text{embed}(\mathbf{P}_{T,2})\|). \end{aligned} \quad (6)$$

The objective of the third learning phase can be expressed as “ $\mathbf{X}^{(0 \sim k+3l)} \rightarrow \mathbf{X}^{(0 \sim k+5l)}$ ”, for further correcting these previously learned prompts to improving the continuity of several prompts and providing smooth transitions between them. The third learning phase can be formulated as:

$$\begin{aligned} \|\mathbf{P}_{T,2}, \mathbf{P}_{T,3}\| &\rightarrow \mathbf{P}_{T,3}, \mathbf{P}_{T,3} \in \mathbb{R}^{4l \times n}, \\ \mathbf{R}_3 &= \text{FM}(\|\text{Seq}_3, \hat{\mathbf{P}} * \mathbf{R}_2^{p \sim (p+l)}, \text{embed}(\mathbf{P}_{T,2}), \text{embed}(\mathbf{P}_{T,3})\|). \end{aligned} \quad (7)$$

To prevent the prompt parameters from being overly biased toward expressing short-term over long-term dependence, the final stage of learning intends to uniformly adjust these parameters. The corresponding values from the three previous learning phases are concatenated along the time dimension and then multiplied by the uniform prompt $\mathbf{P}_{T,w}$. The final learning can be expressed as follows.

$$\begin{aligned} \|\mathbf{P}_{T,1}, \mathbf{P}_{T,2}, \mathbf{P}_{T,3}\| &\rightarrow \mathbf{P}_{T,w}, \\ \|\mathbf{P}_{T,1} + \mathbf{P}_{T,2} + \mathbf{P}_{T,3}, \mathbf{P}_{T,2} + \mathbf{P}_{T,3}, \mathbf{P}_{T,3}\| &\rightarrow \hat{\mathbf{P}}_{T,w}, \\ \mathbf{R}_T &= \text{FM}(\|\text{Seq}, \frac{1}{m-k} \sum_{i=k}^m [\text{tanh}(\hat{\mathbf{P}}_{T,w}) * \mathbf{P}_{T,w}]\|). \end{aligned} \quad (8)$$

Spatial Prompt Learning (SPL). We regard multiple meteorological factors within a specific space. This enables the establishment of correlations between these factors from a spatial perspective on the local client. The SPL are shown in Figure 4.

The trainable parameters that serve as spatial prompts can be represented as $\mathbf{P}_S \in \mathbb{R}^{(m-k) \times 1}$, where $m-k$ represents the length of the forecasting period. Prior to the initial learning phase, the first p hours of the data are considered fixed and denoted as $\text{Seq} \in \mathbb{R}^{k \times n}$. Subsequently, Seq are combined with the initial spatial prompts along the temporal dimension, and any gaps in the spatial dimension are filled with zero-valued parameters, resulting in $\hat{\mathbf{P}}_S \in \mathbb{R}^{(m-k) \times (n-1)}$.

The first learning phase comprises:

$$\begin{aligned} \|\mathbf{P}_S, \hat{\mathbf{P}}_S\| &\rightarrow \mathbf{P}_{S,1}, \\ \mathbf{P}'_{S,1} &= \mathbf{P}_{S,1} * \text{FM}(\|\text{Seq}, \mathbf{S}_1\|). \end{aligned} \quad (9)$$

In learning spatial prompts for the i^{th} variable, the i^{th} learning phase of SPL can be formulated as follow:

$$\begin{aligned} \|\mathbf{P}'_{S,i-1}, \hat{\mathbf{P}}_S \in \mathbb{R}^{(m-k) \times (n-i)}\| &\rightarrow \mathbf{P}_{S,i}, \\ \mathbf{P}'_{S,i} &= \mathbf{P}_{S,i} * \text{FM}(\|\text{Seq}, \mathbf{P}_{S,i}\|). \end{aligned} \quad (10)$$

One advantage of SPL over auto-regression is its continuous correction of learned prompt parameters from the previous iteration. This correction utilizes all the previously predicted values, leading to improved forecast accuracy. Moreover, SPL improves the model’s perception of correlations among multiple variables by considering adjacent variables together. Through multiple iterations of regression, SPL enhances the model’s ability to account for these spatial correlations.

To prevent isolating the TPL and SPL effects within the STP, we combine the result of TPL and SPL empirically using a Gate operation. The output of the STP is formulated as $\mathbf{R} = [(1 - \text{sigmoid}(\mathbf{R}_S)) * \text{tanh}(\mathbf{R}_T)] * \mathbf{W}$.

4.2 Optimization of MetePFL

The optimization objective of *MetePFL* is

$$\arg \min_{\mathbf{P}} \sum_{k=1}^N \frac{n_k}{n} [F_k(\{\mathbf{P}_k\}) + \lambda R(\{\mathbf{P}_k\}, \{\mathbf{P}^g\})] + \tau \mathcal{G}(\mathbf{A}) \quad (11)$$

where $\{\mathbf{P}_k\}$ is prompts including $\mathbf{P}_S, \mathbf{P}_T$, $\{\mathbf{P}^g\}$ is the global prompt parameters, $\mathcal{G}(\mathbf{A})$ is a regularization term used to represent the correlation among client according the adjacent matrix \mathbf{A} . specifically, the term $F_k(\{\mathbf{P}\}) + \lambda R(\{\mathbf{P}\}, \{\mathbf{P}^g\})$ can be formulated as $\mathcal{L}_k = \mathcal{L}_{\text{MSE}} + \lambda \|\{\mathbf{P}_k\} - \{\mathbf{P}^g\}\|_2^2$. The optimization objective on i -th client as:

$$\begin{aligned} \arg \min_{\mathbf{A}} \sum_{k=1}^N \lambda R(\{\mathbf{P}_k\}, \{\mathbf{P}^g\}) + \tau \mathcal{G}(\mathbf{A}), \\ \text{s.t. } \mathbf{P}^s \in \arg \min_{\{\mathbf{P}_k\}} \mathbf{A}_{j,i} S(\{\mathbf{P}_j\}, \{\mathbf{P}_i\}), \\ \text{i.e. } \mathbf{P}^g = G(\{\mathbf{P}_1^s\}, \{\mathbf{P}_2^s\}, \dots, \{\mathbf{P}_N^s\}), \end{aligned} \quad (12)$$

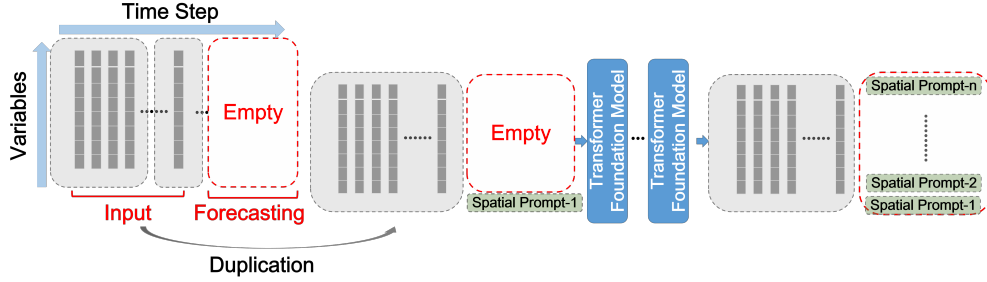


Figure 4: The learning strategy of Spatial Prompt Learning, the empty is a location in data where variables need forecasting.

Algorithm 1 *MetePFL* algorithm.

```

Initialized  $P_T, P_S$ .
for each communication round  $t = 0, 1, 2, \dots, T$  do
  Local model initialize:
  for each client  $i = 0, 1, 2, \dots, N$  in parallel do
     $0 \leftarrow P_T, P_S$ 
  Local model update:
  for each client  $i = 0, 1, 2, \dots, N$  in parallel do
    Update  $P_T, P_S$  for  $e$  local steps:
    Train  $P_T, P_S$  with loss function  $\mathcal{L}$  in Eq.(11)
  end for
  Each selected client sends  $P_T, P_S$  to the server
  Aggregation:
   $A \leftarrow \text{GraphGenerator}(\{P_1\}, \{P_2\}, \dots, \{P_N\})$ 
   $A' \leftarrow \text{GAT}(A)$ 
  Update  $P$  for  $r$  steps  $\text{GCN}(A', \{P_i\}_{i=1}^N)$ :
   $\{P_i^s\}_{i=1}^N = A' \{P_i\}_{i=1}^N$ 
   $\{P_i^s\}_{i=1}^N = \alpha \{P_i^s\}_{i=1}^N + (1 - \alpha) \{P_i\}_{i=1}^N$ 
  Get Global Prompts  $\{P^g\} \leftarrow G(\{P_1^s\}, \{P_2^s\}, \dots, \{P_N^s\})$ 
end for

```

where the $A \in \{0, 1\}$, $S(\{P_j\}, \{P_i\})$ is the similarity of prompt parameters of client i and client j measured by cosine or distance, $G(\cdot)$ is the average operation. During the optimization of *MetePFL*, each client update their model via solving the local objective function as \mathcal{L}_k after them receive the fixed foundation model at first. Then each client upload their prompts P rather than complete model to the server that conduct graph-based aggregation, which significantly reduce the communication overhead while exploring the potential correlations among clients. The aggregation including two steps: Graph Attention Network (GAT) [Veličković *et al.*, 2017]-based graph structure learning that explore the dynamic correlations among clients and Graph Convolution Network (GCN) [Kipf and Welling, 2016] that utilized to parameters reconstruction using learned adjacent matrix A and the prompt parameters uploaded by clients. The GCN automatically updates the parameters of each node by aggregating the models of its neighbors in the graph.

5 Experiments

Baselines. We compare our *MetePFL* with STGCN [Yu *et al.*, 2017], LSTM [Graves, 2012], ConvLSTM [Shi *et al.*, 2015], Transformer [Zerveas *et al.*, 2021], Informer [Zhou *et al.*, 2021], Autoformer [Wu *et al.*, 2021], and Fedformer [Zhou *et al.*, 2022b]. The LSTM-based models have

four layers. The Transformer consists of an eight-layer Encoder, while the Informer, Autoformer, and Fedformer consist of two encoders and a decoder. The Transformer models are trained by data-centric and FL setting, respectively.

Datasets. We compiled three multivariate time series datasets from NASA¹, Average Precipitation (AvePRE), Surface Temperature (SurTEMP), and Surface Upstream (SurUPS) collected by 88, 525, and 238 devices, respectively. All three datasets cover the hour-by-hour variability of 12 different weather-related meteorological variables.

Experimental Setups. Models’ input and output dimensions are $C = 12$ and $C = 1$, respectively. The dataset is split into training, validation, and testing in a 6:2:2 ratio. For pre-training, we use 2/3 of the training set (i.e., 50% of the entire dataset) for training and 1/6 as the validation set (i.e., 10% of the complete dataset) based on the above partition, following the pre-training strategy from Zerveas *et al.* [Zerveas *et al.*, 2021]. For fine-tuning and prompt learning, the last 1/6 of the training set is used for training, while the validation and test sets remain unchanged. In the federated training process, we set k to control the number of clients participating in training per round, and we use $k = 0.1$ and $k = 0.2$ in the experiments. The forecasting uses 12 time steps in history ($\mathcal{P}=12$, i.e., the past twelve hours) to predict 15 time steps in the future ($\mathcal{Q}=15$, i.e., fifty hours in the future) with a time window length of 27 h. We set l to 3, considering the validity time and trigger threshold of weather events. All models were trained on an NVIDIA Tesla V100 GPU using an initial learning rate of $1e^{-3}$ and a batch size of 128, ADAM [Kingma and Ba, 2014] as the optimizer. The algorithm used in the FL of the Transformer-based network is FedAvg. The communication round was set to 20, $\alpha = 0.5$, and the early stopping strategy was applied. Mean absolute error (MAE) and root mean absolute error (RMSE) as evaluation metrics. The code is available at <https://github.com/shengchaochen82/MetePFL>.

5.1 Overall Comparison

Table 1 presents a performance comparison between our *MetePFL* and baselines. Results indicate the superiority of the Transformer-based model over the LSTM-based model and STGCN when trained centrally across all three datasets. However, when utilized as a FM within the FL framework, the performance of the Transformer-based model is dimin-

¹<https://disc.gsfc.nasa.gov/>

Model	Temporal Encoding	AvePRE		SurTEMP		SurUPS	
		MAE	RMSE	MAE	RMSE	MAE	RMSE
STGCN	None	0.331	0.815	0.196	0.257	0.298	0.399
ConvLSTM	None	0.305	0.730	0.198	0.266	0.311	0.416
LSTM	None	0.326	0.781	0.212	0.274	0.335	0.431
Transformer	Learnable	0.301	0.714	0.236	0.313	0.369	0.475
	Fixed	0.332	0.744	0.239	0.320	0.351	0.449
FEDformer	Learnable	0.239	0.547	0.165	0.214	0.205	0.271
	Fixed	0.248	0.564	0.165	0.216	0.201	0.264
Autoformer	Learnable	0.271	0.589	0.167	0.235	0.201	0.265
	Fixed	0.269	0.590	0.176	0.228	0.212	0.279
Informer	Learnable	0.213	0.543	0.191	0.245	0.251	0.330
	Fixed	0.216	0.547	0.193	0.251	0.240	0.311
Fed-Transformer	Learnable	0.454/0.402	0.927/0.892	0.780/0.684	0.910/0.793	0.621/0.522	0.769/0.640
Fed-FEDformer	Learnable	0.397/0.372	0.791/0.726	0.684/0.530	0.822/0.680	0.612/0.512	0.754/0.647
Fed-Autoformer	Learnable	0.425/0.349	0.784/0.724	0.742/0.627	0.924/0.765	0.578/0.503	0.715/0.602
Fed-Informer	Learnable	0.385/0.361	0.865/0.768	<u>0.647/0.513</u>	<u>0.790/0.656</u>	0.605/0.543	0.737/0.724
PromptFL-Transformer	Learnable	0.427/0.389	0.828/0.786	0.683/0.612	0.824/0.741	0.603/0.519	0.766/0.641
<i>MetePFL*</i>	Learnable	0.389/0.376	0.631/0.626	0.592/0.522	0.611/0.597	0.584/0.485	0.721/0.610
<i>MetePFL</i>	Learnable	0.378/0.342	0.628/0.605	0.556/0.542	0.601/0.569	0.521/0.460	0.642/0.589

Table 1: Performance comparison of *MetePFL* with baselines, the first eight models are trained from scratch using full parameters. Fed-Transformer refers to training the Transformer model from scratch in a federated learning (FL) setting, the last two models employ FL-based prompt learning methods with a pre-trained Transformer as the FM, the symbol * indicates a FedAvg-based implementation, underline means the optimal in FL full parameters training, **Bold** means the optimal in FL prompt learning strategy.

FM	Algorithm	AvePRE		SurTEMP		SurUPS	
		MAE	RMSE	MAE	RMSE	MAE	RMSE
Transformer*	FedAtt	0.507/0.467	0.836/0.823	0.978/0.947	1.279/1.186	0.705/0.686	0.828/0.820
	FedProx	0.567/0.531	0.845/0.827	0.922/0.901	1.141/1.102	0.688/0.672	0.814/0.810
	Scaffold	0.567/0.536	0.833/0.811	0.930/0.899	1.232/1.200	0.697/0.676	0.817/0.808
	FedAvg	0.611/0.591	0.823/0.810	0.998/0.896	1.118/1.115	0.706/0.699	0.832/0.821
Transformer	<i>MetePFL</i> (FedAtt)	0.383/0.357	0.735/0.618	0.576/ 0.520	0.603/0.575	0.511/0.482	0.642/0.610
	<i>MetePFL</i> (FedProx)	0.399/0.385	0.691/0.633	0.564/0.542	0.686/0.667	0.556/0.512	0.702/0.651
	<i>MetePFL</i> (Scaffold)	0.385/0.353	0.755/0.627	0.602/0.531	0.727/0.600	0.560/0.512	0.719/0.645
	<i>MetePFL</i> (FedAvg)	0.389/0.376	0.631/0.626	0.592/0.522	0.611/0.597	0.584/0.485	0.721/0.610
	<i>MetePFL</i>	0.378/0.342	0.628/0.605	0.556/0.542	0.601/0.569	0.521/0.460	0.642/0.589

Table 2: Performance comparison between fine-tuned Transformer and *MetePFL* with different FL algorithm, * implying that the model applies fine-tuning strategy, the *MetePFL* (·) means that the implementation based on other FL algorithm, **Bold** means the optimal results.

ished compared to the trained FM. This reduction can be attributed to the heterogeneity of weather data collected from multiple sensors. Notably, while the Transformer may be less effective than other similar models in centralized training, it demonstrates a significant performance advantage over Fed-FEDformer, Fed-Autoformer, and Fed-Informer when employed as an FM within the *MetePFL*. This observation suggests that STP enhances the Transformer’s capability to comprehend spatiotemporal data. Furthermore, our reliable FM and aggregation algorithms (Transformer and FedAvg) surpass PromptFL [Guo *et al.*, 2022]. This outcome provides additional validation of the effectiveness and superiority of our proposed *MetePFL*.

To evaluate the effectiveness and superiority of *MetePFL* over fine-tuning, we implement different *MetePFL* version based on four FL algorithms: FedAtt [Jiang *et al.*, 2020], Fed-Prox [Sahu *et al.*, 2018], Scaffold [Karimireddy *et al.*, 2020], and FedAvg [McMahan *et al.*, 2017]. We maintained the prompt setting and compared the results against fine-tuning. The results are presented in Table 2. Our experiments reveal

the following findings: (1) *MetePFL* outperforms the fine-tuning method, highlighting the sensitivity of data correlation to the STP approach; (2) the graph-based aggregation used in *MetePFL* surpasses other algorithms, indicating its effectiveness in mitigating the negative impact of Non-IID.

5.2 Framework Applicability

To determine the applicability of *MetePFL*, we replaced its FM with pre-trained Informer and Autoformer using the same pre-training strategy. Additionally, we used fine-tuning and PromptFL as reference strategies. The results as shown in Table 3. Our proposed *MetePFL* remains valid for other Transformer FMs and significantly outperforms fine-tuning and PromptFL. By comparing the performance of *MetePFL* under different FL algorithm, we demonstrate the effectiveness of the graph-based aggregation once again.

5.3 Parameter Utilization

Table 4 compares parameter utilization for different federated prompt learning strategies. PromptFL and *MetePFL* are sig-

Model	Strategy	AvePRE		SurTEMP		SurUPS	
Fed-Informer	Fine-tuning (FedAvg)	0.397/0.391	0.776/0.759	0.734/0.713	0.874/0.864	0.669/0.631	0.824/0.781
	Fine-tuning (FedAtt)	0.403/0.378	0.780/0.724	0.950/0.899	1.057/1.00	0.686/0.675	0.807/0.793
	PromptFL	0.407/0.361	0.742/0.727	0.722/0.700	0.867/0.832	0.671/0.658	0.786/0.753
	<i>MetePFL</i> (FedAvg)	0.382/ 0.357	0.631/0.618	0.701 /0.692	0.840/0.812	0.698/0.650	0.796/ 0.725
	<i>MetePFL</i> (FedAtt)	0.392/0.387	0.776/0.758	0.754/0.696	0.905/0.839	0.669/ 0.644	0.772/0.738
	<i>MetePFL</i>	0.363 /0.358	0.630 / 0.601	0.713/ 0.677	0.838 / 0.800	0.652 / 0.644	0.755 /0.739
Fed-Autoformer	Fine-tuning (FedAvg)	0.378/0.373	0.761/0.758	0.693/0.682	0.848/0.834	0.610/0.523	0.767/0.721
	Fine-tuning (FedAtt)	0.372/0.355	0.761/0.754	0.706/0.683	0.848/0.839	0.598/0.551	0.742/0.698
	PromptFL	0.355 /0.348	0.759/0.753	0.672/0.659	0.826/0.807	0.584/0.543	0.724/0.678
	<i>MetePFL</i> (FedAvg)	0.364/0.348	0.780/0.731	0.674/ 0.630	0.820/0.781	0.564/0.520	0.736/0.656
	<i>MetePFL</i> (FedAtt)	0.372/0.341	0.762/0.754	0.689/0.641	0.824/0.766	0.549/0.525	0.717/0.650
	<i>MetePFL</i>	0.355 / 0.334	0.750 / 0.719	0.666 / 0.630	0.814 / 0.750	0.547 / 0.516	0.712 / 0.649

Table 3: Performance comparison of pre-trained Fed-Informer and Fed-Autoformer with different learning strategies, *MetePFL* (FedAvg) and *MetePFL* (FedAtt) implies the FedAvg- and FedAtt-based implementations, respectively, **Bold** means the optimal performance.

Strategy	# of Total Param	# of Training Param	# of Participation Param
FL-Regular	3,229,857	3,229,857	100%
FL-Fine Tuning	3,288,886	109,854	30.37%
PromptFL	3,250,867	71,835	2.22%
<i>MetePFL</i> (Ours)	3,258,547	77,595	2.38%

Table 4: Comparison of parameter utilization of *MetePFL*.

nificantly more advantageous than regular training (train from scratch) and fine-tuning, with parameter utilization of 2.22% and 2.38%, respectively - nearly 28% lower than fine-tuning. While PromptFL is better than *MetePFL* in parameter utilization (-0.16%), it performs nearly 20% worse than *MetePFL* (see Table 1). Despite considering only 2.38% parameters, *MetePFL* achieves excellent performance and significantly improves inter-device communication efficiency.

5.4 Ablation Study

To evaluate the effectiveness of STP, we conducted ablation studies under general scenarios rather than under the FL framework. The results are presented in Table 5. The Encoder-only and Decoder-only Transformer underwent pre-training with one-step forecasting and auto-regressive mechanisms similar to previous experiments. The Encoder-only Transformer trained with STP outperformed the Decoder-only Transformer, indicating the superior learning mechanism of the proposed STP over conventional auto-regression. Moreover, the performance gap between the pre-trained Encoder with and without STP confirms the efficacy of the proposed STP in general scenarios.

Model	Strategy	MAE	RMSE
Encoder-only	Train from scratch	0.332	0.744
	Train from scratch & STP	0.300	0.689
	Pre-train	0.295	0.668
	Pre-train & STP	0.267	0.571
Decoder-only	Train from scratch	0.304	0.724

Table 5: Comparison of Encoder/Decoder-only Transformer under different strategies in general scenarios based on AvePRE.

The effectiveness of SPL, Gate, and their advantage over position-aware embedding (PE) was verified using the

TPL	SPL	Gate	PE	MAE	RMSE
w	w	w/o	w/o	0.301	0.608
w	w	w	w/o	0.267	0.571
w/o	w	w/o	w/o	0.284	0.580
w	w/o	w/o	w/o	0.299	0.613
w/o	w/o	w/o	w	0.284	0.617

Table 6: Ablation results of the proposed STP in Encoder-only Transformer based on the AvePRE in general scenarios.

Encoder-only Transformer. The results are presented in Table 6. SPL demonstrates a greater improvement compared to PE. Additionally, SPL enhances the model’s performance, while the model without Gate experiences a significant decline. TPL does not perform as effectively as desired in generic scenarios. Consequently, we conducted ablation experiments in the FL setting, as shown in Table 7. The results indicate that both TPL and SPL enhance the model’s forecasting performance. In conclusion, the effectiveness of the Gate operation is demonstrated, and both TPL and SPL can improve the model’s performance in the FL setting.

TPL	SPL	Gate	MAE	RMSE
w	w	w	0.378	0.628
w/o	w	w/o	0.407	0.722
w	w/o	w/o	0.415	0.740

Table 7: Ablations on FL setting.

6 Conclusion

This paper proposes a novel machine learning approach to train foundation models for weather forecasting tasks, capable of capturing the spatiotemporal relationships of meteorological data based on multivariate time series. To enhance the performance while keeping data secure and reducing communication overhead, we introduce a prompt learning mechanism based on the fixed foundation model within the FL framework. Additionally, we utilize a graph-based approach to mitigate the impact of data heterogeneity on model effectiveness. Extensive experiments on three real-world weather datasets confirm the effectiveness of our proposed method.

References

- [Alléon *et al.*, 2020] Antoine Alléon, Grégoire Jauvion, Boris Quennehen, and David Lissmyr. Plumenet: Large-scale air quality forecasting using a convolutional lstm network. *arXiv preprint arXiv:2006.09204*, 2020.
- [Bauer *et al.*, 2015] Peter Bauer, Alan Thorpe, and Gilbert Brunet. The quiet revolution of numerical weather prediction. *Nature*, 525(7567):47–55, 2015.
- [Bojesomo *et al.*, 2021] Alabi Bojesomo, Hasan Al-Marzouqi, Panos Liatsis, Gao Cong, and Maya Ramanath. Spatiotemporal swin-transformer network for short time weather forecasting. In *CIKM Workshops*, 2021.
- [Campbell and Diebold, 2005] Sean D Campbell and Francis X Diebold. Weather forecasting for weather derivatives. *Journal of the American Statistical Association*, 100(469):6–16, 2005.
- [Chattopadhyay *et al.*, 2020] Ashesh Chattopadhyay, Ebrahim Nabizadeh, and Pedram Hassanzadeh. Analog forecasting of extreme-causing weather patterns using deep learning. *Journal of Advances in Modeling Earth Systems*, 12(2):e2019MS001958, 2020.
- [Chen and Lai, 2011] Ling Chen and Xu Lai. Comparison between arima and ann models used in short-term wind speed forecasting. In *2011 Asia-Pacific Power and Energy Engineering Conference*, pages 1–4. IEEE, 2011.
- [Chen *et al.*, 2022a] Fengwen Chen, Guodong Long, Zonghan Wu, Tianyi Zhou, and Jing Jiang. Personalized federated learning with graph. *arXiv preprint arXiv:2203.00829*, 2022.
- [Chen *et al.*, 2022b] Shengchao Chen, Ting Shu, Huan Zhao, Qilin Wan, Jincan Huang, and Cailing Li. Dynamic multiscale fusion generative adversarial network for radar image extrapolation. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–11, 2022.
- [Chen *et al.*, 2023] Shengchao Chen, Ting Shu, Huan Zhao, Guo Zhong, and Xunlai Chen. Tempee: Temporal-spatial parallel transformer for radar echo extrapolation beyond auto-regression. *arXiv preprint arXiv:2304.14131*, 2023.
- [Gao *et al.*, 2022] Liang Gao, Huazhu Fu, Li Li, Yingwen Chen, Ming Xu, and Cheng-Zhong Xu. Feddc: Federated learning with non-iid data via local drift decoupling and correction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10112–10121, 2022.
- [Graves, 2012] Alex Graves. Long short-term memory. *Supervised sequence labelling with recurrent neural networks*, pages 37–45, 2012.
- [Guo *et al.*, 2022] Tao Guo, Song Guo, Junxiao Wang, and Wenchao Xu. Promptfl: Let federated participants cooperatively learn prompts instead of models—federated learning in age of foundation model. *arXiv preprint arXiv:2208.11625*, 2022.
- [Hagemann *et al.*, 2013] Stefan Hagemann, Cui Chen, Douglas B Clark, Sonja Folwell, Simon N Gosling, Ingjerd Haddeland, Naota Hanasaki, Jens Heinke, Fulco Ludwig, Frank Voss, et al. Climate change impact on available water resources obtained using multiple global climate and hydrology models. *Earth System Dynamics*, 4(1):129–144, 2013.
- [Jiang *et al.*, 2020] Jing Jiang, Shaoxiong Ji, and Guodong Long. Decentralized knowledge acquisition for mobile internet applications. *World Wide Web*, 23(5):2653–2669, 2020.
- [Karevan and Suykens, 2020] Zahra Karevan and Johan AK Suykens. Transductive lstm for time-series prediction: An application to weather forecasting. *Neural Networks*, 125:1–9, 2020.
- [Karimireddy *et al.*, 2020] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, pages 5132–5143. PMLR, 2020.
- [Kingma and Ba, 2014] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [Kipf and Welling, 2016] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [Kjellstrom *et al.*, 2016] Tord Kjellstrom, David Briggs, Chris Freyberg, Bruno Lemke, Matthias Otto, and Olivia Hyatt. Heat, human performance, and occupational health: a key issue for the assessment of global climate change impacts. *Annual review of public health*, 37:97–112, 2016.
- [Li and Liang, 2021] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021.
- [Li *et al.*, 2019] Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zhihua Zhang. On the convergence of fedavg on non-iid data. *arXiv preprint arXiv:1907.02189*, 2019.
- [Li *et al.*, 2020] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2:429–450, 2020.
- [Li *et al.*, 2023] Zhiwei Li, Guodong Long, and Tianyi Zhou. Federated recommendation with additive personalization. *arXiv preprint arXiv:2301.09109*, 2023.
- [Liu *et al.*, 2021] Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. Gpt understands, too. *arXiv preprint arXiv:2103.10385*, 2021.
- [Long *et al.*, 2020] Guodong Long, Yue Tan, Jing Jiang, and Chengqi Zhang. Federated learning for open banking. In *Federated Learning: Privacy and Incentive*, pages 240–254. Springer, 2020.
- [Long *et al.*, 2021] Guodong Long, Tao Shen, Yue Tan, Leah Gerrard, Allison Clarke, and Jing Jiang. Federated learn-

- ing for privacy-preserving open innovation future on digital health. In *Humanity Driven AI: Productivity, Well-being, Sustainability and Partnership*, pages 113–133. Springer, 2021.
- [Long *et al.*, 2023] Guodong Long, Ming Xie, Tao Shen, Tianyi Zhou, Xianzhi Wang, and Jing Jiang. Multi-center federated learning: clients clustering for better personalization. *World Wide Web*, 26(1):481–500, 2023.
- [Ma *et al.*, 2022] Jie Ma, Guodong Long, Tianyi Zhou, Jing Jiang, and Chengqi Zhang. On the convergence of clustered federated learning. *arXiv preprint arXiv:2202.06187*, 2022.
- [McMahan *et al.*, 2017] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agueria y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017.
- [Radford *et al.*, 2021] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021.
- [Sahu *et al.*, 2018] Anit Kumar Sahu, Tian Li, Maziar Sanjabi, Manzil Zaheer, Ameet Talwalkar, and Virginia Smith. On the convergence of federated optimization in heterogeneous networks. *arXiv preprint arXiv:1812.06127*, 3:3, 2018.
- [Sapankevych and Sankar, 2009] Nicholas I Sapankevych and Ravi Sankar. Time series prediction using support vector machines: a survey. *IEEE computational intelligence magazine*, 4(2):24–38, 2009.
- [Shi *et al.*, 2015] Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. *Advances in neural information processing systems*, 28, 2015.
- [Tan *et al.*, 2021] Yue Tan, Guodong Long, Lu Liu, Tianyi Zhou, Qinghua Lu, Jing Jiang, and Chengqi Zhang. Fedproto: Federated prototype learning over heterogeneous devices. *arXiv preprint arXiv:2105.00243*, 3, 2021.
- [Tan *et al.*, 2022a] Yue Tan, Yixin Liu, Guodong Long, Jing Jiang, Qinghua Lu, and Chengqi Zhang. Federated learning on non-iid graphs via structural knowledge sharing. *arXiv preprint arXiv:2211.13009*, 2022.
- [Tan *et al.*, 2022b] Yue Tan, Guodong Long, Jie Ma, Lu Liu, Tianyi Zhou, and Jing Jiang. Federated learning from pre-trained models: A contrastive learning approach. *arXiv preprint arXiv:2209.10083*, 2022.
- [Veličković *et al.*, 2017] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- [Voyant *et al.*, 2012] Cyril Voyant, Marc Muselli, Christophe Paoli, and Marie-Laure Nivet. Numerical weather prediction (nwp) and hybrid arma/ann model to predict global radiation. *Energy*, 39(1):341–355, 2012.
- [Wang *et al.*, 2022] Zhuowei Wang, Tianyi Zhou, Guodong Long, Bo Han, and Jing Jiang. Fednoil: a simple two-level sampling method for federated learning with noisy labels. *arXiv preprint arXiv:2205.10110*, 2022.
- [Wu *et al.*, 2021] Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in Neural Information Processing Systems*, 34:22419–22430, 2021.
- [Xu *et al.*, 2022] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. Vitpose+: Vision transformer foundation model for generic body pose estimation. *arXiv preprint arXiv:2212.04246*, 2022.
- [Yates *et al.*, 2021] Andrew Yates, Rodrigo Nogueira, and Jimmy Lin. Pretrained transformers for text ranking: Bert and beyond. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, pages 1154–1156, 2021.
- [Yu *et al.*, 2017] Bing Yu, Haoteng Yin, and Zhanxing Zhu. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. *arXiv preprint arXiv:1709.04875*, 2017.
- [Zerveas *et al.*, 2021] George Zerveas, Srideepika Jayaraman, Dhaval Patel, Anuradha Bhamidipaty, and Carsten Eickhoff. A transformer-based framework for multivariate time series representation learning. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 2114–2124, 2021.
- [Zhang *et al.*, 2023] Chunxu Zhang, Guodong Long, Tianyi Zhou, Peng Yan, Zijian Zhang, Chengqi Zhang, and Bo Yang. Dual personalization on federated recommendation. *arXiv preprint arXiv:2301.08143*, 2023.
- [Zhou *et al.*, 2021] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 11106–11115, 2021.
- [Zhou *et al.*, 2022a] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022.
- [Zhou *et al.*, 2022b] Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. *arXiv preprint arXiv:2201.12740*, 2022.