

FT-MoE: Sustainable-learning Mixture of Experts for Fault-Tolerant Computing

Wenjing Xiao^{1,2}, Wenhao Song^{1,2}, Miaojiang Chen^{1,2,*}, Min Chen^{3,4}

¹Guangxi Key Laboratory of Multimedia Communications and Network Technology, Nanning, 530004, China

²School of Computer, Electronics and Information, Guangxi University, Nanning, 530004, China

³School of Computer Science and Engineering, South China University of Technology, Guangzhou 510006, China

⁴Pazhou Laboratory, Guangzhou 510330, China

wenjingx@gxu.edu.cn, wenhaos@st.gxu.edu.cn, mjchen_cs@gxu.edu.cn, minchen@ieee.org

Abstract

Intelligent fault-tolerant (FT) computing has recently demonstrated significant advantages in predicting and diagnosing faults proactively, thereby ensuring reliable service delivery. However, due to the heterogeneity of fault knowledge, dynamic workloads, and limited data support, existing deep learning-based FT algorithms face challenges in fault detection quality and training efficiency. This is primarily because their homogenization of fault knowledge perception difficulties to fully capture diverse and complex fault patterns. To address these challenges, we propose FT-MoE, a sustainable-learning fault-tolerant computing framework based on a dual-path architecture for high-accuracy fault detection and classification. This model employs a mixture-of-experts (MoE) architecture, enabling different parameters to learn distinct fault knowledge. Additionally, we adopt a two-stage learning scheme that combines comprehensive offline training with continual online tuning, allowing the model to adaptively optimize its parameters in response to evolving real-time workloads. To facilitate realistic evaluation, we construct a new fault detection and classification dataset for edge networks, comprising 10,000 intervals with fine-grained resource features, surpassing existing datasets in both scale and granularity. Finally, we conduct extensive experiments on the FT benchmark to verify the effectiveness of FT-MoE. Results demonstrate that our model outperforms state-of-the-art methods.

Code — <https://github.com/1291632523/FT-MoE>

Extended version —

<https://arxiv.org/submit/7074274/view>

Introduction

With the development of Internet of Things (IoT) technology, the number of computation tasks at network edges has increased dramatically (Tuli et al. 2023). When edge devices simultaneously process multiple computation tasks, the limited resource of edge network leads to high device failures and low quality of experience (Xiao et al. 2025a). Traditional restoring methods of replicating or resubmitting tasks is to take action after failure has been observed, and cannot prevent expensive repair costs caused by system failures (Shaikh and Jammal 2024). Thus, to prevent losses

of event failure and system downtime, proactive methods that predict failures before they occur have received increasing attention (Ahmed, Terefe, and Hailemariam 2024). Recently, model-based fault tolerant (FT) computing adopts specialized algorithms for intelligent task scheduling to address event failure problems and receives increasing attention (Xu et al. 2024) (Cai, Qin, and Huang 2025). Among them, deep learning-based FT schemes that do not require any assumptions or specialized algorithms have become the most popular solution (Assiri and Sheneamer 2025) (Amin, Sajid Iqbal, and Hamza Shahbaz 2024). Under the guidance of large amounts of labeled data, deep neural networks capture the feature of abnormal signals well for fault prediction and judgment. However, due to limited computing and communication capacity in edge network (Xiao et al. 2025b), deep learning-based schemes still face difficulties in guaranteeing fault diagnose quality of service systems (Hazra et al. 2024), facing several challenges: (1) **High Knowledge Heterogeneity**: As a multi-task learning, FT task features both fault detection and fault classification. Their sub-tasks have completely different feature learning paths, and their domain knowledge varies from fault to fault (Ledmi, Bendjenna, and Hemam 2018). However, existing approaches (Na et al. 2025) (Qian et al. 2025) (Negi et al. 2021) typically use same parameters and ignore the heterogeneity of fault types and learning tasks, resulting in performance inefficiency. (2) **Highly Dynamic Environment**: Since tasks in edge network are complex and variable, fault detection must adapt to the dynamic environment. This requires the FT model to be not only robust, but also flexible enough to adapt dynamic environments in real time. (3) **Lack of Realistic FT Datasets**: Most existing datasets used in related studies have coarse-grained resource features. These datasets do not reflect the real operating conditions of multitask scheduling and resource-constrained scenarios. As a result, models trained on them struggle to generalize to real-world deployments, which limits the research of edge fault detection algorithms.

In response to the above issues, we propose FT-MoE, a dynamic fault prediction model for fault-tolerant scheduling, which aims to improve the fault detection performance. Unlike existing DL-based approaches, we propose a sustainable learning fault-tolerant model based on a dual-path architecture to achieve efficient FT computing of multiple

tasks. Specifically, to learn heterogeneous knowledge from detection data, we propose a mixture-of-experts(MoE) network with efficient gating modules as the first path, enabling dynamic tuning by adjusting the number of experts based on the activation differences (Guo et al. 2025). To efficiently capture the impact of task migration on device status, we adopt efficient dynamic graph attention network as the second path. Then, we use a cross multi-head attention to model the relationship of heterogeneous knowledge and task scheduling effects extracted from the two paths. Further, to adapt dynamic service environments, FT-MoE employs continual learning. We use a two-stage model optimization scheme, consisting of offline training and online tuning. The Offline training provides enough time to adjust and optimize parameters. The online tuning supports continual learning of FT-MoE, allowing it to dynamically add or remove parameters to adapt to dynamic environments. In addition, to address the lack of realistic fault tolerance datasets, we propose a dataset with fine-grained host resource features and scheduling decisions that better capture the dynamics of real-world environments. Overall, our main contributions are summarized as follows:

- We propose FT-MoE, a dual-path fault-tolerant model that integrates task-adaptive expert selection MoE and dynamic graph attention. The first path uses a task-adaptive MoE to handle heterogeneous knowledge and the second path employs a dynamic graph attention network to capture the effects of task migration. The two paths are connected through a cross multi-head attention mechanism.
- Combining with continual learning, we design a two-stage training scheme for FT-MoE, consisting of offline training and online tuning. Offline training aims to optimize parameters. Online tuning aims to dynamically add or remove parameters to adapt to dynamic environments.
- We construct a new dataset for fault detection and classification. Our dataset includes more fine-grained host resource features and scheduling decisions, making it better reflect the dynamic environment than existing ones.

Related Work

Fault Tolerant Computing Methods

Most contemporary state-of-the-art FT computing methods employ some form of specialized algorithms or machine learning models. Specialized algorithms methods usually rely on predefined models to achieve FT tasks. DFTM (Sivagami and Easwarakumar 2019) detects faults by monitoring the load and network traffic of virtual machines (VM). The faults are classified as single VM faults and multi-node faults. PCFT (Liu et al. 2018) predicts in advance which physical machines are likely to fail based on CPU temperature. ECLB (Sharif, Nickray, and Shahidinejad 2020) technique uses Bayesian methods and neural networks to classify host machines into different categories. CMODLB (Negi et al. 2021) uses k-means and deep learning to cluster nodes and identify overloaded hosts. In a recent work, PreGAN (Tuli, Casale, and Jennings 2022a)

uses gated recurrent units and graph attention networks to construct a deep learning model for fault prediction. PreGAN+ (Tuli, Casale, and Jennings 2024) combines the transformer structure, further improving the accuracy of fault prediction. However, existing methods cannot handle heterogeneous fault knowledge, using the same model for all fault types and ignoring knowledge diversity. Most are offline trained, cannot adapt after deployment and perform worse with changing system conditions. In our experiments, we compare FT-MoE against the state-of-the-art baselines DFTM, ECLB, PCFT, CMODLB, PreGAN and PreGAN+.

Mixture of Experts Model

Mixture of experts (Dai et al. 2024) is a machine learning method based on multiple expert models. Its core idea is to select suitable models for different inputs. Compared to traditional single models, MoE uses multiple experts to handle different types of faults, which helps to improve overall detection accuracy. Due to the flexibility and efficiency, many recent studies have attempted to apply MoE to handle time-series data. MoLE (Ni et al. 2024) trains multiple linear-centric models and a router model that weighs and mixes their outputs to predict time series. MEO (He et al. 2023) proposes combining multiple experts into a single shared expert to significantly improve the computational efficiency of the MoE while maintaining performance. Time-MoE (Shi et al. 2025) combines a series of decoder-only transformer models with MoE, supporting flexible prediction ranges with different input context lengths. However, most existing MoE methods rely on statically trained architecture, which limits their adaptability to changing input distributions and dynamic environments.

Dynamic Fault Detection and Recognizing Algorithm

We propose FT-MoE, a sustainable-learning mixture-of-experts framework for fault-tolerant computing, aiming to improve fault prediction performance. As shown in Fig. 1, FT-MoE consists of three components: a schedule-aware graph encoder to capture task propagation across hosts, a fault-adaptive MoE layer to handle fault heterogeneity and a cross multi-head attention layer to fuse heterogeneous features with scheduling information. In addition, we adopt a two-stage training strategy, including offline optimization for diverse fault scenarios and online tuning for continual adaptation to dynamic environments.

FT-MoE System Model and Problem Statement

In this work, we consider a standard heterogeneous environment, where all computational nodes are interconnected through a Local Area Network (LAN) ¹.

Definition 1. (Host Set and Feature Matrix) We assume that there are M hosts in the environment, denoted as $H = \{h_1, \dots, h_m\}$, where $m \in M$. Each host has N resource utilization features. The feature vector of the host at interval t is represented as $X_t \in \mathbf{R}^{M \times N}$.

¹We provide more details for system model in appendix 1

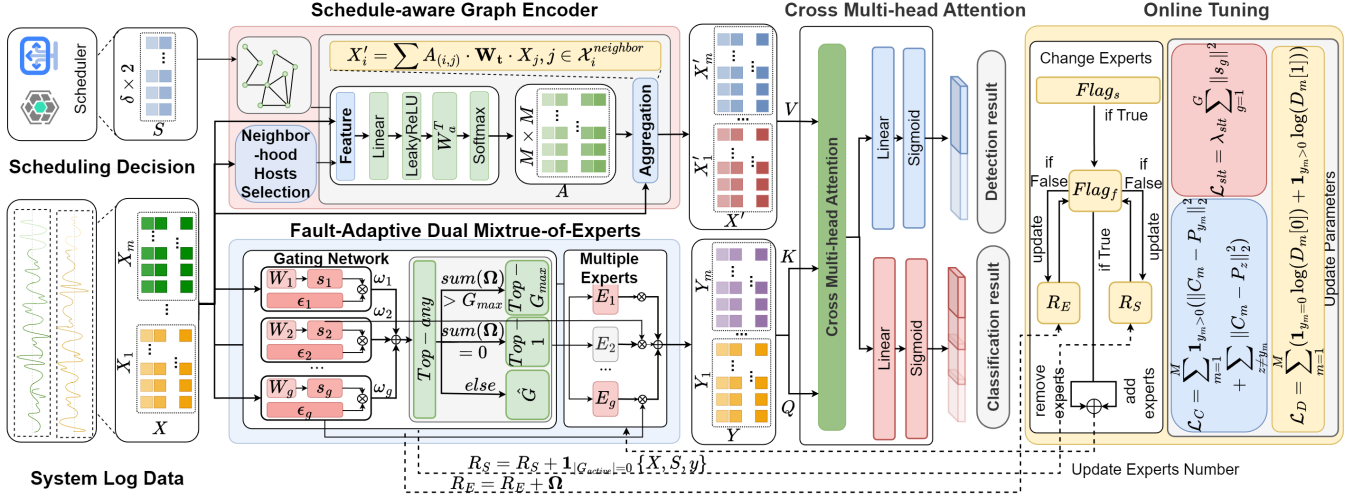


Figure 1: FT-MoE, including the schedule-aware graph encoder, the fault-adaptive mixture-of-experts layer, the cross multi-head attention layer, the offline training module and the online tuning module.

Definition 2. (Task Scheduling Matrix) At the beginning of each interval t , a set of new tasks arrives at the system, and the scheduler generates a scheduling decision simultaneously. We assume that there are δ tasks running simultaneously in the system. The scheduling decision at the interval t , denoted as $S_t \in \mathbf{R}^{\delta \times 2}$, shows the task migration from the source host to the target host. Each migration creates an edge between two different hosts.

Schedule-aware Graph Encoder

We design the schedule-aware graph encoder to model the propagation dynamics of task scheduling behaviors across hosts. Specifically, the feature vector of m -th host is denoted as $X_m \in \mathbf{R}^N$. The hosts that are connected to the m -th host in the scheduling decision S are called the neighborhood hosts of the m -th host. The set of neighborhood hosts' indices for the m -th host is denoted as $\mathcal{H}_m^{neighbor}$. Following previous study (Brody, Alon, and Yahav 2022), we calculate attention weights A^2 . $A_{(m,m')}$ represents the attention weight of the edge (m, m') . Then, we aggregate the attention weights of all neighborhood hosts to get the updated feature vector of the i -th host X'_i . The aggregation process can be represented as follows:

$$X'_i = \sum_{j \in \mathcal{H}_i^{neighbor}} A_{(i,j)} \cdot \mathbf{W}_t \cdot X_j, \quad (1)$$

where \mathbf{W}_t is trainable parameter. X' denotes the set of updated feature vector, $X'_i \in X'$.

Fault-Adaptive Mixture-of-Experts

To handle the heterogeneity of FT tasks, we design a fault-adaptive mixture-of-experts layer. It consists of multiple experts that capture diverse fault patterns and a gating network that selects the most relevant experts based on the input.

²We provide more details for calculating the attention weights A in appendix 2

Experts-Adaptation Gating Traditional MoE uses a fixed number of active experts, but this often leads to inefficiency in practice (Muqeth, Liu, and Raffel 2024). To solve this, we introduce Experts-Adaptation Gating (EAGate), which automatically chooses and adjusts the number of experts based on input features for fault prediction. Next, we describe the specific process of the EAGate. EAGate first calculates the cosine similarity between the input and each expert. These similarity scores are then used to decide which experts should be activated. Let E_g denote the g -th expert of the MoE. Let W_g denote the g -th expert's representation matrix of the MoE, which is trainable in EAGate. $g \in G$ and G is the number of experts in an MoE model. Hence, the similarity s_g between X and expert E_g is represented as:

$$s_g = \frac{\langle X, W_g \rangle}{\|X\| \|W_g\|} \quad (2)$$

where $\langle x_1, x_2 \rangle$ function denotes the dot product between x_1 and x_2 and $\|\cdot\|$ is the L2 norms function. By normalizing X and W_g , we obtain all standardized similarity scores which are denoted as $\mathbf{S} = \{s_g\}$.

Then, we apply the sigmoid function to map the similarity score to $[0,1]$ to represent the activation probability of experts, that is $\hat{\mathbf{S}} = \{\hat{s}_g\}$. We define the activation threshold vector as $\epsilon \in \mathbf{R}^g$ and they are trainable parameters. ϵ_g represents the g -th expert's activation threshold. We apply the sigmoid function to map the activation threshold to $[0,1]$, that is $\hat{\epsilon}$. By comparing \hat{s}_g with the activation threshold $\hat{\epsilon}_g$, we can determine whether g -th expert E_g should be activated during a particular computation round. Let ω_g denote the activation state of expert E_g and $\Omega = \{\omega_g\}$ denote the activation state of all experts, its formula is represented as:

$$\omega_g = \text{sign}(\hat{s}_g - \hat{\epsilon}_g), \quad (3)$$

where $\text{sign}(x)$ function is used to convert x to 0 or 1. If $x > 0$, then $\text{sign}(x) = 1$; else, $\text{sign}(x) = 0$. Thus, if $\omega_g = 1$,

it indicates that expert E_g can be activated; otherwise, it can not be activated. We set the activatable expert indices as \hat{G} and the activatable expert as $\hat{E}_{\hat{g}}$, where $\hat{g} \in \hat{G}$. We set the activatable experts as $\hat{E} = \{\hat{E}_{\hat{g}}\}$.

Then, to enhance the model's adaptability and flexibility, we present a Top-any function to select a set of experts to activate, which enables each input to autonomously determine the number of experts to activate. The output of the Top-any function is a set of experts' indices. We set the maximum number of selectable active experts for one input as G_{max} . Specifically, if $\text{sum}(\Omega) > G_{max}$, we select the G_{max} experts with the highest similarity scores to be activated, experts' indices are denoted as $\text{Top-}G_{max}(\hat{s})$; if $\text{sum}(\Omega) = 0$, we select an expert with the highest similarity scores to be activated, expert's index is denoted as $\text{Top-1}(\hat{s})$; else, we select the set of activatable experts \hat{E} , experts' indices are \hat{G} . Its formula is represented as:

$$\text{Top-any} = \begin{cases} \text{Top-}G_{max}(\hat{s}) & \text{if } \text{sum}(\Omega) > G_{max}, \\ \text{Top-1}(\hat{s}) & \text{if } \text{sum}(\Omega) = 0, \\ \hat{G} & \text{else.} \end{cases} \quad (4)$$

Based on Top-any, we can obtain experts selected to be activated, whose indices are denoted as G_{active} .

Multiple Experts Module The multiple experts module of an MoE consists of multiple sparse activation experts. To avoid model learning bias caused by unusually strong or weak outputs from some experts, we adopt a weighted sum approach to integrate the outputs of all activation experts (Sakib, Mustajab, and Alam 2024). Let G_{active} denote the set of experts' indices selected to be activated. The output of the multiple experts module is denoted as Y , which is calculated as follows:

$$Y = \frac{1}{|G_{active}|} \sum_{g \in G_{active}} E_g s_g, \quad (5)$$

where $|G_{active}|$ denotes the number of activation experts' indices in G_{active} .

Cross Multi-head Attention Module

The cross multi-head attention (CMHA) module is used to model the relationship between the task scheduling behavior and the heterogeneous resource knowledge. Specifically, the attention is represented as:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax} \left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}} \right) \mathbf{V}, \quad (6)$$

where \mathbf{Q} is the query, \mathbf{K} is the key and \mathbf{V} is the value. d_k is the dimension of \mathbf{K} . CMHA computes attention in parallel using multiple sets of \mathbf{Q} , \mathbf{K} and \mathbf{V} . Each set has its own weight matrix, is called a *head* and captures different input aspects. β is the number of heads. $head_a$ is the result of the computation for a -th *head* and $a \in A$, represented as:

$$head_a = \text{Attention}(\mathbf{Q}W_{\beta}^Q, \mathbf{K}W_{\beta}^K, \mathbf{V}W_{\beta}^V), \quad (7)$$

$$\text{MultiHead}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = [head_1; \dots; head_{\beta}]W_O, \quad (8)$$

where W_{β}^Q , W_{β}^K , W_{β}^V and W_O are all trainable parameters. We set $Q = Y$ and $K = V = X'$. Let each host's features actively search the graph to find which scheduling relationships affect them most, and use this to make better fault predictions. O represents the context vector obtained from CMHA, which is represented as:

$$O = \text{MultiHead}(Y, X', X'), \quad (9)$$

To adjust the feature dimensions and scale the output, we use two separate feed-forward networks: one for fault detection and one for fault classification. For fault detection, we apply a linear layer Linear^D followed by a softmax to produce probabilities. For fault classification, we use another linear layer Linear^C followed by a sigmoid. These two FFNs are referred to as FFN^D and FFN^C , respectively. FFN^D and FFN^C are represented as:

$$\text{FFN}^D = \text{Softmax}(\text{Linear}^D(O)). \quad (10)$$

$$\text{FFN}^C = \text{Sigmoid}(\text{Linear}^C(O)). \quad (11)$$

Finally, FFN^D and FFN^C output the fault detection result $D \in \mathbf{R}^{M \times 2}$ and fault classification result $C \in \mathbf{R}^{M \times Q}$, respectively. Here, D gives the fault prediction for all M hosts, where D_m represents the prediction for the m -th host. A fault is predicted in host m if $D_m[1] > D_m[0]$. The classification output C gives each host's fault feature vector, where Q is the feature dimension and C_m is the vector for host m . We assume there are Z fault types, with each type $z \in Z$ represented by a prototype vector P_z in a matrix $P \in \mathbf{R}^{Z \times Q}$. These prototype vectors are trainable and initialized randomly in the range $[0, 1]$. To classify a fault, we compute the Euclidean distance between C_m and each P_z , and select the fault type whose prototype is closest to C_m .

Offline Training and Online Tuning

In this section, we describe the offline training and the online fine-tuning. The offline training describes how to train the model with datasets acquired in a real environment. The online fine-tuning describes how to fine-tune the model based on real-time data when running the simulation platform.

Offline Training In this section, we describe the offline training process for the FT-MoE to detect and classify faults. Datasets of FT computing consist of time-series data X , scheduling decisions S and fault type labels y . y_m represents m -th host's fault class, where $y_m = 0$ denotes no fault. For the fault detection, we use cross-entropy loss as the loss function and sum the loss over all hosts to get the final loss, its formula is represented as:

$$\mathcal{L}_D = \sum_{m=1}^M (\mathbf{1}_{y_m=0} \log(D_m[0]) + \mathbf{1}_{y_m>0} \log(D_m[1])), \quad (12)$$

where $\mathbf{1}_{\text{cond}}$ is an indicator function, represented as:

$$\mathbf{1}_{\text{cond}} = \begin{cases} 1, & \text{if cond is true,} \\ 0, & \text{if cond is false,} \end{cases} \quad (13)$$

where *cond* denotes the condition.

For the fault classification, we use the triplet loss as the loss function (Dong and Shen 2018). We calculate the euclidean distance between the model-generated fault type prediction for m -th host C_m and each prototype vector in P . Then the euclidean distance between C_m and the real category prototype vector P_{y_m} is minimized while increasing the euclidean distance between C_m and the incorrect z -th class prototype vector P_z . Therefore, the fault classification loss \mathcal{L}_C for all hosts with system faults is represented as:

$$\mathcal{L}_C = \sum_{m=1}^M \mathbf{1}_{y_m > 0} (\|C_m - P_{y_m}\|_2^2 + \sum_{z \neq y_m} \|C_m - P_z\|_2^2). \quad (14)$$

Besides, we also introduce the expert selection loss, ensuring balanced activation frequencies for each expert (Yang et al. 2024). Specifically, the similarity scores of experts are regularized using the L2 norms function. Let λ_{slt} be the regularization weight and \mathcal{L}_{slt} denote the expert selection loss. its formula is expressed as:

$$\mathcal{L}_{slt} = \lambda_{slt} \sum_{g=1}^G \|s_g\|^2 \quad (15)$$

Thus, the final loss function \mathcal{L}_{final} combines \mathcal{L}_D , \mathcal{L}_C and \mathcal{L}_{slt} , which is represented as:

$$\mathcal{L}_{final} = \alpha_1 \mathcal{L}_D + \alpha_2 \mathcal{L}_C + \alpha_3 \mathcal{L}_{slt} \quad (16)$$

where α_1 , α_2 and α_3 are the weight parameters.

Performance-aware Online Tuning To adapt to dynamic service environments, we design a performance-aware online tuning strategy for FT-MoE. It fine-tunes the fault-adaptive MoE using the host features of the previous interval X_{t-1} , scheduling decisions of the previous interval S_{t-1} and fault labels of the current interval y_t . For simplicity, we denote the input as $\{X, S, y\}$. We still use \mathcal{L}_{final} to update the parameters. During online tuning, all components of FT-MoE are frozen except for the MoE module. Its key idea is to remove experts with low contribution and add new ones to handle previously unseen data, enabling efficient continual learning under changing conditions.

In detail, we record the active state of all experts to keep track of which experts are frequently activated, denoted as $R_E \in \mathbf{R}^G$. $R_{E,g}$ represents the times the g -th expert is activated during this inference. Its formula is represented as:

$$R_E = R_E + \Omega \quad (17)$$

Furthermore, we can also record the activated experts number of each input, denoted as R_S . So we can find inputs that did not activate any experts. Its formula is represented as:

$$R_S = R_S + \mathbf{1}_{|G_{active}|=0} \{X, S, y\} \quad (18)$$

We use $flag_s$ and $flag_f$ to determine when to start and stop the routing recording. We initialize $flag_s = 1$ and $flag_f = 0$ and set an epoch threshold I . If the current epoch number is a multiple of I , we set $flag_f = 1$. Based on these history states, we can fine-tune fault-adaptive MoE according to the following strategy:

(1) If $R_{E,g} = 0$, it means that the g -th expert has not been activated by any input during the training process. The g -th expert will be removed from the network to reduce the computational resources and improve the training efficiency.

(2) If $R_S = NULL$, it means that there are inputs that do not have suitable experts and the system will automatically add a new expert. In addition, the new expert threshold $\hat{\epsilon}_{new}$ is initially set to 0 to ensure that the inactive input can activate the newly added expert.

After adding experts and deleting experts, we set $flag_f = 0$ and restart the routing recording.

Our Dataset

We build a new dataset to enable accurate and adaptable fault prediction in dynamic environments. It supports both supervised learning and continual adaptation. Compared to datasets like Google Cluster Trace, Alibaba Cluster Trace v2021, and Bitbrains GWA-T-12 VM Trace, ours uniquely includes full resource metrics, detailed scheduling logs, and host-level fault labels³.

To simulate realistic and dynamic workloads, we design a task generator using the Bitbrains GWA-T-12 dataset. We select the 'rnd' subset with 500 VMs and filter those whose CPU usage at the 10th interval is between 500 and 3000 MIPS to build a candidate VM pool. In each simulation step, we randomly generate containers based on a Gaussian distribution. Each container uses a VM trace from the pool, and its full resource usage over time is used to simulate container behavior. We use the COSCO framework (Tuli et al. 2022) to simulate our fault-tolerant system with a 16 node distributed cluster. Eight nodes have 4 GB of RAM, and the other eight have 8 GB. The default scheduler, Gradient Based Optimization task scheduling strategy (GOBI) (Tuli et al. 2022), assigns containers based on host load and task demands. During execution, we log scheduling decisions and resource usage. Based on this simulator, we obtain host resource status data in various service environments as the training dataset. Anomaly Detection Engine for Linux Logs (ADE) tool (Agyemang 2024) is used to generate ground-truth labels for our dataset. Fault types consist of CPU fault, RAM fault and disc fault⁴. Our dataset contains 115,010 no fault samples, 8,001 CPU faults, 22,771 RAM faults and 14,218 Disk faults.

Experiment

Experimental Setting

For FT-MoE, we set the dimensionality of the prototype vector Q at 8. For gradual updating of prototype vectors P , we adopt a hybrid learning prototype vector updating scheme⁵. The number of experts in an MoE model G is numbered 12 and the maximum number of selectable active experts for

³We provide more details for public datasets and comparison between them and ours in appendix 3 and 4

⁴We provide more details on how to determine the type of fault in appendix 5

⁵We provide more details for hybrid-learning prototype vectors updating scheme in appendix 6

Method	Detection				Classification	
	Accuracy	Precision	Recall	F1 Score	HR	NDCG
DFTM	0.8210 ± 0.0200	0.7100 ± 0.0700	0.8100 ± 0.0200	0.7560 ± 0.0350	0.4800 ± 0.0200	0.4400 ± 0.0020
ECLB	0.8800 ± 0.0180	0.7350 ± 0.0650	0.8500 ± 0.0200	0.7880 ± 0.0400	0.5050 ± 0.0100	0.4700 ± 0.0020
PCFT	0.8500 ± 0.0120	0.7600 ± 0.0500	0.8700 ± 0.0180	0.8110 ± 0.0300	0.5650 ± 0.0090	0.5200 ± 0.0020
CMODLB	0.8700 ± 0.0100	0.7750 ± 0.0300	0.8750 ± 0.0150	0.8225 ± 0.0250	0.5950 ± 0.0030	0.5300 ± 0.0030
PreGAN	0.9000 ± 0.0080	0.8200 ± 0.0220	0.8950 ± 0.0120	0.8555 ± 0.0200	0.6050 ± 0.0060	0.5650 ± 0.0050
PreGAN+	0.9060 ± 0.0070	0.8290 ± 0.0060	0.9180 ± 0.0080	0.8710 ± 0.0180	0.6250 ± 0.0030	0.5800 ± 0.0040
FT-MoE	0.9053 ± 0.0028	0.8273 ± 0.0079	0.9322 ± 0.0094	0.8766 ± 0.0061	0.6496 ± 0.0055	0.6021 ± 0.0027

Table 1: Overall comparison of methods across the fault detection and the fault classification.

one input G_{max} is numbered 8. For the offline training FT-MoE model, we use the AdamW optimizer to train FT-MoE offline (Saleh and Mashaly 2019), with an initial learning rate of 0.001 and decays by 0.1 every 10 epochs. We set the weights of the loss function L_{fin} : α_1 , α_2 and α_3 to 0.35, 0.5 and 0.15, respectively. The update weight η is set to 0.9. Note that we run all experiments for 100 scheduling intervals and average over 5 runs as final performance results. We use the COSCO framework as the service simulator for our experimental.

Evaluation Metrics

Fault Prediction. Our framework includes two tasks of the fault detection and the fault classification. They have different evaluation metrics: (1) For the fault detection, we use precision, recall and F1 score as prior work (Anwar et al. 2025). (2) For the fault classification, we use two metrics, that is HitRate (HR) (Tuli, Casale, and Jennings 2022b) and Normalized Discounted Cumulative Gain (NDCG) (Järvelin and Kekäläinen 2002)⁶.

Fault-Tolerant Scheduling. Quality of Service (QoS) scores are some metrics for evaluating the quality of user service. Here, we adopt the CPU utilization, the RAM utilization, the energy consumption and the response time metrics to measure the system performance⁷.

Experimental Results

To evaluate the performance of FT-MOE, we conducted comparative experiments with the latest baselines: DFTM (Sivagami and Easwarakumar 2019), ECLB (Sharif, Nickray, and Shahidinejad 2020), PCFT (Liu et al. 2018), CMODLB (Negi et al. 2021), PreGAN (Tuli, Casale, and Jennings 2022a) and PreGAN+ (Tuli, Casale, and Jennings 2024). First, we evaluate the performance of the fault detection and the fault classification, which have different evaluation metrics. Then, to verify the effect of FT-MoE on avoiding system overload, we further evaluate the task migration performance after a successful fault detection.

Following (Tuli, Casale, and Jennings 2024), we use the Generative Adversarial Networks (GAN) to decide tasks scheduling decisions based on predicted faults, which contains a generator and a discriminator network. We analyze

⁶We provide details for experimental metrics in appendix 7

⁷We provide additional QoS results in appendix 8

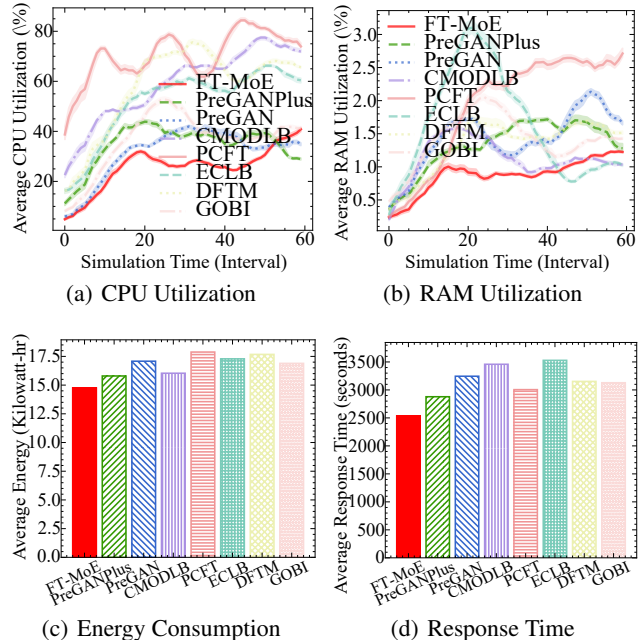


Figure 2: QoS performance of FT-MoE with baselines.

the QoS scores after task scheduling decisions based on the prediction results of our model. **Comparison of Fault Prediction Methods.** The experimental results are shown in Table 1. (1) In the fault detection, FT-MoE outperforms other baseline methods in recall and F1 score respectively. FT-MoE achieves 0.9041 in accuracy, which outperforms most other baseline methods. This is due to FT-MoE capturing high-dimension features of input X_t and S_t better than other methods. (2) In the fault classification, FT-MoE’s HR and NDCG showed an improvement of 3.936% and 3.81% respectively, over the best baseline method. This is due to FT-MoE considering knowledge heterogeneity and using different sets of parameters for inputs with different types of feature. While baseline models use identical parameters.

System Scheduling QoS Breakdown. Fig. 2 compares some QoS scores of all models, including the original GOBI without any fault-tolerant method. For all the fault prediction models, we use the same scheduling decision generator (Tuli, Casale, and Jennings 2024) to generate scheduling

decisions based on the fault prediction results. (1) Fig. 2(a) and Fig. 2(b) indicates that FT-MoE has lower average CPU and RAM utilization compared to the baselines. This is due to the higher fault detection and fault classification accuracy of our model, which helps the generator generate better scheduling decisions. (2) Due to the lower average CPU and RAM utilization, FT-MoE also has the lowest average energy consumption of 14.1641 KW-hr, followed by PreGAN+ with 16.4324 KW-hr, as Fig. 2(c) shows. (3) Fig. 2(d) shows a comparison of FT-MoE and the baselines with the response time. It shows that FT-MoE has the minimum response time. This is due to FT-MoE, which has a faster inference speed compared to the baseline methods.

Ablation Study

In this section, we conduct ablation studies to analyze the effects of different components in FT-MoE⁸.




Figure 3: Training curves for the fault detection and classification.

Analysis of EAGate In this section, we evaluate the impact of the proposed EAGate. Fig. 3 presents the accuracy trends for both the fault detection and the fault classification during training, comparing FT-MoE models with the EAGate (W/EAGate) and without it (Wo/EAGate). In the Wo/EAGate configuration, the EAGate is replaced by a static gating. We could observe that: (1) The fault classification accuracy of the Wo/EAGate fluctuates drastically and the optimal accuracy only reaches 87.956%. This is because static gating treats all diverse fault type inputs uniformly, ignoring their differences. (2) In contrast, the W/EAGate configuration exhibits a steady improvement in classification accuracy, ultimately reaching 90.479%. Because the EAGate can adaptively select experts based on input features, enabling more fine-grained training and capturing heterogeneous knowledge of different types of fault. In general, these results show that the EAGate plays a key role in stabilizing training and improving the model’s ability to extract meaningful features for accurate fault classification.

Analysis of Online Tuning In this section, our aim is to demonstrate the effects of online tuning. Fig. 4 shows the changes in FT tasks accuracy during running in the edge environment of the FT-MoE with online tuning (W/OnlineTuning) and without it (Wo/OnlineTuning). Fig. 4(a) shows that after a period of running, the performance of Wo/OnlineTuning decreases greatly and finally reduces to 52.624% for

⁸We provide more experiment results in appendix 9




Figure 4: Fine-tuning curves for the fault detection and classification.

the fault detection and 57.223% for the fault classification. This is because Wo/MoE uses unchanged parameters, so its ability to generalize decreases as new data comes in. This phenomenon is common in continual learning and is called plasticity loss (Zang et al. 2025). Fig. 4(b) shows the running results for W/OnlineTuning. Although there is also a performance loss for a while after the start of the fine-tuning, it then stabilizes. After that, the performance is maintained at a high level. This is because W/OnlineTuning has a stronger generalization ability compared to Wo/OnlineTuning.




Figure 5: T-SNE plot for the fault classification using Wo/CMHA and W/CMHA.

Analysis of CMHA In this section, our aim is to demonstrate the necessity of the CMHA module. Fig. 5 includes t-SNE visualizations of the predicted class prototypes for the fault classification. We observe that the embeddings produced by the model with CMHA (W/CMHA) are more compact and well separated by fault type. This indicates stronger intra-class consistency and higher inter-class discrimination. In contrast, the embeddings of the model without CMHA (Wo/CMHA) are more dispersed, leading to poorer classification performance. Additionally, W/CMHA also achieves significantly better alignment with the ground truth than Wo/CMHA, particularly in accurately localizing faults across different hosts and time periods. These results show that CMHA improves feature interaction and fault prediction by better associating task behavior with host status.

Conclusion

In this paper, we propose FT-MoE, a multitask learning framework for fault awareness and adaptive fault-tolerant learning. It leverages a MoE network with efficient gating to

capture diverse fault patterns, and adopts a two-stage training strategy, including offline optimization and online continual tuning, to adapt to dynamic environments. Experiments on a simulated edge setting show that FT-MoE outperforms existing methods. Ablation studies confirm the effectiveness of each module in supporting multitask fault tolerance. In future work, our goal is to validate FT-MoE in real-world deployments and incorporate federated learning for privacy to assess its robustness and scalability.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (Nos. 62502101, 62462002) and partially supported by the Natural Science Foundation of Guangxi, China (Nos. 2025GXNSFAA069958, 2025GXNSFBA069394)

References

- Agyemang, E. F. 2024. Anomaly detection using unsupervised machine learning algorithms: A simulation study. *Scientific African*, 26: e02386.
- Ahmed, S.; Terefe, T.; and Hailemariam, D. 2024. Machine learning for base transceiver stations power failure prediction: A multivariate approach. *e-Prime - Advances in Electrical Engineering, Electronics and Energy*, 10: 100814.
- Amin, A. A.; Sajid Iqbal, M.; and Hamza Shahbaz, M. 2024. Development of Intelligent Fault-Tolerant Control Systems with Machine Learning, Deep Learning, and Transfer Learning Algorithms: A Review. *Expert Systems with Applications*, 238: 121956.
- Anwar, T.; Mu, C.; Yousaf, M. Z.; Khan, W.; Khalid, S. M.; Hourani, A. O.; and Zaitsev, I. 2025. Robust fault detection and classification in power transmission lines via ensemble machine learning models. *Scientific Reports*, 15.
- Assiri, B.; and Sheneamer, A. 2025. Fault tolerance in distributed systems using deep learning approaches. *PLOS ONE*, 20.
- Brody, S.; Alon, U.; and Yahav, E. 2022. How Attentive are Graph Attention Networks? In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Cai, W.; Qin, L.; and Huang, J. 2025. MoC-System: Efficient Fault Tolerance for Sparse Mixture-of-Experts Model Training. In *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2, ASPLOS '25*, 655–671. New York, NY, USA: Association for Computing Machinery. ISBN 9798400710797.
- Dai, D.; Deng, C.; Zhao, C.; Xu, R.; Gao, H.; Chen, D.; Li, J.; Zeng, W.; Yu, X.; Wu, Y.; Xie, Z.; Li, Y.; Huang, P.; Luo, F.; Ruan, C.; Sui, Z.; and Liang, W. 2024. DeepSeek-MoE: Towards Ultimate Expert Specialization in Mixture-of-Experts Language Models. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1280–1297. Bangkok, Thailand: Association for Computational Linguistics.
- Dong, X.; and Shen, J. 2018. Triplet Loss in Siamese Network for Object Tracking. In Ferrari, V.; Hebert, M.; Sminchisescu, C.; and Weiss, Y., eds., *Computer Vision – ECCV 2018*, 472–488. Cham: Springer International Publishing. ISBN 978-3-030-01261-8.
- Guo, B.; Zhou, C.; He, S.; Chen, J.; and Shen, X. 2025. Mixture-of-Experts as Continual Knowledge Adapter for Mobile Vision Understanding. *IEEE Transactions on Mobile Computing*, 1–15.
- Hazra, A.; Tummala, V. M. R.; Mazumdar, N.; Sah, D. K.; and Adhikari, M. 2024. Deep reinforcement learning in edge networks: Challenges and future directions. *Physical Communication*, 66: 102460.
- He, S.; Fan, R.-Z.; Ding, L.; Shen, L.; Zhou, T.; and Tao, D. 2023. Merging Experts into One: Improving Computational Efficiency of Mixture of Experts. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 14685–14691. Singapore: Association for Computational Linguistics.
- Järvelin, K.; and Kekäläinen, J. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Trans. Inf. Syst.*, 20(4): 422–446.
- Ledmi, A.; Bendjenna, H.; and Hemam, S. M. 2018. Fault Tolerance in Distributed Systems: A Survey. In *2018 3rd International Conference on Pattern Analysis and Intelligent Systems (PAIS)*, 1–5.
- Liu, J.; Wang, S.; Zhou, A.; Kumar, S. A. P.; Yang, F.; and Buyya, R. 2018. Using Proactive Fault-Tolerance Approach to Enhance Cloud Service Reliability. *IEEE Transactions on Cloud Computing*, 6(4): 1191–1202.
- Muqeeth, M.; Liu, H.; and Raffel, C. 2024. Soft Merging of Experts with Adaptive Routing. *Trans. Mach. Learn. Res.*, 2024.
- Na, L.; Cai, B.; Zhang, C.; Liu, J.; and Li, Z. 2025. A heterogeneous transfer learning method for fault prediction of railway track circuit. *Engineering Applications of Artificial Intelligence*, 140: 109740.
- Negi, S.; Rauthan, M. S.; Vaisla, K. S.; and Panwar, N. 2021. CMODLB: an efficient load balancing approach in cloud computing environment. *The Journal of Supercomputing*, 77: 8787 – 8839.
- Ni, R.; Lin, Z.; Wang, S.; and Fantì, G. 2024. Mixture-of-Linear-Experts for Long-term Time Series Forecasting. In Dasgupta, S.; Mandt, S.; and Li, Y., eds., *International Conference on Artificial Intelligence and Statistics, 2-4 May 2024, Palau de Congressos, Valencia, Spain*, volume 238 of *Proceedings of Machine Learning Research*, 4672–4680. PMLR.
- Qian, Q.; Zhang, B.; Li, C.; Mao, Y.; and Qin, Y. 2025. Federated transfer learning for machinery fault diagnosis: A comprehensive review of technique and application. *Mechanical Systems and Signal Processing*, 223: 111837.

- Sakib, M.; Mustajab, S.; and Alam, M. 2024. Ensemble deep learning techniques for time series analysis: a comprehensive review, applications, open issues, challenges, and future directions. *Clust. Comput.*, 28: 73.
- Saleh, N.; and Mashaly, M. 2019. A Dynamic Simulation Environment for Container-based Cloud Data Centers using ContainerCloudSim. In *2019 Ninth International Conference on Intelligent Computing and Information Systems (ICICIS)*, 332–336.
- Shaikh, S.; and Jammal, M. 2024. Survey of fault management techniques for edge-enabled distributed metaverse applications. *Computer Networks*, 254: 110803.
- Sharif, A.; Nickray, M.; and Shahidinejad, A. 2020. Fault-tolerant with load balancing scheduling in a Fog-based IoT application. *IET Communications*, 14.
- Shi, X.; Wang, S.; Nie, Y.; Li, D.; Ye, Z.; Wen, Q.; and Jin, M. 2025. Time-MoE: Billion-Scale Time Series Foundation Models with Mixture of Experts. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Sivagami, V.; and Easwarakumar, K. 2019. An Improved Dynamic Fault Tolerant Management Algorithm during VM migration in Cloud Data Center. *Future Generation Computer Systems*, 98: 35–43.
- Tuli, S.; Casale, G.; and Jennings, N. R. 2022a. PreGAN: Preemptive Migration Prediction Network for Proactive Fault-Tolerant Edge Computing. In *IEEE INFOCOM 2022 - IEEE Conference on Computer Communications*, 670–679.
- Tuli, S.; Casale, G.; and Jennings, N. R. 2022b. TranAD: deep transformer networks for anomaly detection in multivariate time series data. *Proc. VLDB Endow.*, 15(6): 1201–1214.
- Tuli, S.; Casale, G.; and Jennings, N. R. 2024. PreGAN+: Semi-Supervised Fault Prediction and Preemptive Migration in Dynamic Mobile Edge Environments. *IEEE Transactions on Mobile Computing*, 23(6): 6881–6895.
- Tuli, S.; Mirhakimi, F.; Pallewatta, S.; Zawad, S.; Casale, G.; Javadi, B.; Yan, F.; Buyya, R.; and Jennings, N. R. 2023. AI augmented Edge and Fog computing: Trends and challenges. *Journal of Network and Computer Applications*, 216: 103648.
- Tuli, S.; Poojara, S. R.; Srirama, S. N.; Casale, G.; and Jennings, N. R. 2022. COSCO: Container Orchestration Using Co-Simulation and Gradient Based Optimization for Fog Computing Environments. *IEEE Transactions on Parallel and Distributed Systems*, 33(1): 101–116.
- Xiao, W.; Li, X.; Hu, L.; Hao, Y.; and Chen, M. 2025a. DT-SNet: Dynamic Transformer Slimming for Efficient Vision Recognition. *IEEE Transactions on Multimedia*, 1–12.
- Xiao, W.; Shi, C.; Chen, M.; Liu, Z.; Chen, M.; and Song, H. H. 2025b. GraphEdge: Dynamic graph partition and task scheduling for GNNs computing in edge network. *Information Fusion*, 103329.
- Xu, H.; Liao, L.; Liu, X.; Chen, S.; Chen, J.; Liang, Z.; and Yu, Y. 2024. Fault-tolerant deep learning inference on CPU-GPU integrated edge devices with TEEs. *Future Generation Computer Systems*, 161: 404–414.
- Yang, Y.; Qi, S.; Gu, W.; Wang, C.; Gao, C.; and Xu, Z. 2024. XMoE: Sparse Models with Fine-grained and Adaptive Expert Selection. In Ku, L.-W.; Martins, A.; and Sriku-mar, V., eds., *Findings of the Association for Computational Linguistics: ACL 2024*, 11664–11674. Bangkok, Thailand: Association for Computational Linguistics.
- Zang, Z.; Sun, C.; Liu, L.; Sun, F.; and Zheng, C. 2025. Loss of Plasticity: A New Perspective on Solving Multi-Agent Exploration for Sparse Reward Tasks. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '25*, 2299–2308. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9798400714269.