Few-Shot Cross-System Anomaly Trace Classification for Microservice-based systems

1st Yuqing Wang* Department of Computer Science University of Helsinki Helsinki, Finland

4th Mutlu Beyazit AnSyMo Universiteit Antwerpen & Flanders Make Antwerp, Belgium 2nd Mika V. Mäntylä Department of Computer Science University of Helsinki Helsinki, Finland

5th Joanna Kisaakye AnSyMo Universiteit Antwerpen & Flanders Make Antwerp, Belgium

3rd Serge Demeyer AnSyMo Universiteit Antwerpen & Flanders Make Antwerp, Belgium

> 6th Jesse Nyyssölä Department of Computer Science University of Helsinki Helsinki, Finland

failures in various fault categories due to their complex and dynamic nature. To effectively handle failures, AIOps tools utilize trace-based anomaly detection and root cause analysis. In this paper, we propose a novel framework for few-shot abnormal trace classification for MSS. Our framework comprises two main components: (1) Multi-Head Attention Autoencoder for constructing system-specific trace representations, which enables (2) Transformer Encoder-based Model-Agnostic Meta-Learning to perform effective and efficient few-shot learning for abnormal trace classification. The proposed framework is evaluated on two representative MSS, Trainticket and OnlineBoutique, with open datasets. The results show that our framework can adapt the learned knowledge to classify new, unseen abnormal traces of novel fault categories both within the same system it was initially trained on and even in the different MSS. Within the same MSS, our framework achieves an average accuracy of 93.26% and 85.2% across 50 meta-testing tasks for Trainticket and OnlineBoutique, respectively, when provided with 10 instances for each task. In a cross-system context, our framework gets an average accuracy of 92.19% and 84.77% for the same metatesting tasks of the respective system, also with 10 instances provided for each task. Our work demonstrates the applicability of achieving few-shot abnormal trace classification for MSS and shows how it can enable cross-system adaptability. This opens an avenue for building more generalized AIOps tools that require less system-specific data labeling for anomaly detection and root cause analysis.

Abstract—Microservice-based systems (MSS) may experience

Index Terms—Microservice, anomaly classification, root cause analysis, tracing, NLP, meta learning, cross-system anomaly detection

I. INTRODUCTION

Microservice architecture is a software design approach where software systems are developed as a collection of small, independent services [1]. Traces are fundamental to understanding and monitoring Microservice-based systems (MSS) [2]. A trace maps the path of a user request and it is composed of interconnected spans [3]. Each span is an individual operation performed by a particular service. Logs record the behaviors of each service instance in a span. Log content varies based on what the developer has decided to log.

The operational workflow for handling failures within MSS typically involves anomaly detection (AD) and root cause analysis (RCA) [4], [5]. AD identifies deviations from normal behavior, signaling potential anomalies. Upon detection of anomalies, RCA is performed to conduct a detailed examination of the anomaly. This examination aims to pinpoint the exact service(s) causing the anomaly and to understand the underlying reasons for its occurrence. Manual AD and RCA can be time-consuming and error-prone for MSS [4] not only due to the numerous services and their intricate interactions but also the complexity and multimodality of monitoring data [4].

Artificial intelligence for IT operations (AIOps) tools place particular emphasis on trace-based AD and RCA due to their importance. Several studies focus on trace-based AD [6]– [8] represent each trace as a sequence of spans and design models to detect abnormal traces. Regarding RCA, substantial efforts [9]–[12] have been devoted to locating the service responsible for abnormal traces for MSS, e.g., using graph neural networks. However, tooling to classify abnormal traces for MSS is still in its infancy, though the demand for such tools is increasing as they provide advantages that can significantly improve RCA. This is primarily due to the following challenges:

- High dimensional and multi-modal trace-related data: In each trace, the length of spans and logs is varied, and they feature attributes in diverse formats, including semantic, numeric, and identical. This makes it impractical to rely solely on raw data for constructing trace representations.
- Imbalanced abnormal trace distribution in MSS. This imbalance arises when some fault categories are more common than others. It poses a challenge for normal machine learning methods as they may become biased towards the more frequent fault categories [13], resulting in poor performance in identifying less frequent.
- Heterogeneity of MSS: Different MSS can have differ-

^{*}Corresponding author: yuqing.wang@helsinki.fi

ent architectures, components, and behaviors, making it difficult to develop a universal solution for both trace representation construction and classification.

Few-shot learning [14] has the potential to address the challenges of imbalanced abnormal trace distribution in MSS for abnormal trace classification, as it can recognize abnormal traces from both frequent and rare fault categories by learning from a minimal number of examples. Autoencoder (AE), an unsupervised algorithm, has been utilized in trace-based AD studies [10], [15], [16] to construct the system-specific low-dimensional trace representations (also known as latent trace representations) by fusing high-dimensional, multi-modal trace-related data, addressing the challenges driven by trace-related data and the heterogeneity of MSS. However, no studies have explored the application of few-shot learning for abnormal trace representations generated by autoencoders for abnormal trace classification has not been explored.

To address the challenges and fill the research gap, we propose a framework for few-shot abnormal trace classification for MSS. Our framework is designed around two central concepts: **C1.** constructing system-specific low-dimensional trace representations by fusing high-dimensional multi-modal trace-related data, which can enable **C2.** the effective and efficient few-shot learning for abnormal trace classification. Figure 1 shows the overview of our framework. Our framework comprises two key components: the Multi-Head Attention Autoencoder (MultiHAttenAE) for **C1**, and, the Transformer Encoder-based Model-Agnostic Meta-Learning (TE-MAML) for **C2**.



Fig. 1. Overview of our framework

Given a MSS, MultiHAttenAE is trained using normal traces to construct low-dimensional trace representations tailored for this MSS. MultiHAttenAE is based on AE [17], and it is in the encoder-decoder structure: an encoder that learns to fuse multi-modal high-dimensional trace-related data into latent trace representations; The decoder reconstructs the original input data from constructed trace representations. The optimized encoder can be used independently to construct low-

dimensional trace representations for new, unseen traces within the MSS it is trained on.

TE-MAML uses the optimized encoder from MultiHAttenAE to construct trace representations for a given MSS. TE-MAML is built upon the meta-learning algorithm MAML [18], which allows to learn a meta-leaner with an optimal set of initial model parameters that can adapt to new tasks using only a few instances and a small number of gradient updates. These new tasks can be either from the same domain or in different domains. In our case, domains are different MSS. By leveraging MAML, TE-MAML learns from meta-training tasks from an MSS and adapts the learned knowledge (i.e., optimized initial parameters) to meta-testing tasks both within the same MSS it was initially trained on and in a different MSS. Here, each of the meta-training and meta-testing tasks is a distinct abnormal trace classification task. TE-MAML uses TE as the meta-learner, which uses a self-mechanism to weigh different parts of trace representations to identify the most relevant features for doing classification.

We evaluate the effectiveness and efficiency of our framework through several experiments on representative benchmark MSS, Trainticket and OnlineBoutique, with open datasets. We evaluated the effectiveness of each component in our framework. We define these research questions:

- **RQ1: Within-system adaptability.** How effective and efficient is our framework, once trained on abnormal trace classification tasks within a MSS, can adapt to new abnormal trace classification tasks within the same MSS?
- **RQ2:** Cross-system adaptability. How effective and efficient is our framework, once trained on abnormal trace classification tasks within a MSS, can adapt to new abnormal trace classification tasks in a different MSS?
- **RQ3: Component Impact.** How does each component of our framework contribute to its overall performance (effectiveness & efficiency)?

Significance. RO1 evaluation is critical because withinsystem adaptability allows the framework to maintain its effectiveness in classifying abnormal traces in the changing context. Since MSS are dynamic with frequent service updates, additions, or removals, new abnormal traces from novel fault categories may appear [1]. A framework that can adapt to new abnormal trace classification tasks without extensive retraining would significantly reduce the costs and rework effort while increasing practical utility in practices. RQ2 evaluates our framework's capability to transfer learned knowledge from one MSS to another MSS with different architectures, components, and behaviors. Prior studies [19]-[21] have investigated crosssystem adaptability in software system AD tasks to enhance the generalization capability of AIOps tools across diverse software systems. Cross-system adaptability is valuable for organizations that run several MSS, as it allows them to use the same framework for abnormal trace classification across different MSS without extensive training on each MSS. Moreover, cross-system adaptability extends the applicability of academic benchmark MSS, enabling to train on the benchmark MSS and adapt the learned knowledge to classify abnormal traces in industry-specific MSS. RQ3 identifies which parts of the framework are most critical for achieving good performance in terms of both effectiveness and efficiency. This information is valuable for others to decide on using the framework in their applications, and also they can focus on optimizing the most critical components for them.

Our main **contributions** are highlighted as follows:

- An unsupervised method to construct system-specific low-dimensional trace representations using highdimensional multi-modal trace-related data. Such trace representations can enable further effective and efficient trace analysis like few-shot abnormal trace classification in our study.
- The evaluated AIOps framework with within-system and cross-system adaptability for few-shot abnormal trace classification for MSS.
- The practical implementation of the meta-learning algorithm MAML for a new NLP classification task: few-shot abnormal trace classification for MSS. Meta-learning has been observed actively in the domain of computer vision, while its exploration within NLP remains relatively limited [22]

II. RELATED WORK AND EXISTING APPROACHES

Existing approaches for abnormal trace classification are very rare for MSS or similar cloud-based systems. The study [23] is the only related work. It uses a convolutional neural network (CNN) to classify abnormal traces from time series-based fault categories (e.g., incremental, mean shift, gradual increase, cylinder). It characterizes the trace as a sequence of spans and uses the time-series data on span attribute "URL" to do the classification. It evaluated the model using a large dataset from a real-world production cloud. The evaluation results showed that the model can accurately classify the tested anomaly patterns. Besides, prior tracebased AD studies [6]-[8] also characterized a trace as a sequence of spans and used span attributes to construct trace representations. However, these representations might not be effective for abnormal trace classification tasks, because each span only offers insight into local behaviors for a service instance. Our study complements the prior studies and makes the novelty from these aspects:

- For a given MSS, our study constructs trace representations by fusing multi-modal trace-related data, spans and logs, which have attributes in diverse formats, including semantic, numeric, and identical.
- Our study aims to achieve effective abnormal trace classification across a wider range of fault categories with fewer instances, making our framework more robust when facing imbalanced abnormal trace distribution in MSS.
- Our study observes the cross-system adaptability of AIOps tools in this field.

III. METHODOLOGY

We first introduce the trace structure and then describe our framework's components, MultiHAttenAE and TE-MAML, in the following sub-sections.

A. Trace structure

Based on OpenTelemetry [3], a trace is structured as a hierarchical tree of spans. The root span is the starting point of the trace, akin to the tree's base, from which all other spans branch out. Each span has a parent span, except for the root span. The span that is currently being executed (i.e, active span) may contain nested sub-spans, which represent smaller units of work that are part of the larger operation encompassed by the active span to which they belong. Figure 2 shows an example tree structure for a trace, where Span A is the root span, and it triggers a sequence of calls to other spans. Figure 3 shows the same spans depicting how a request flows through the execution of each span in sequence and pinpoints the time where relevant logs are generated.



Fig. 2. An example trace structure (Zhang et. al. [9])

	Span A						Span D			
			Span	В			Spar	i E		
m.			S	pan C				Spa	n F	
	iv_ log	_¥_ log	_₩ log	_∀ log	_	_ ¥_¥_ log log	- ¥_	¥ log	- 🛓	_

Fig. 3. Spans and logs in the timeline (modified from Zhang et. al. [9])

B. MultiHAttenAE for constructing trace representation

Figure 4 shows the basic structure of MultiHAttenAE. For a given MSS, we denote a set of traces as $Tr = \{Tr_1, Tr_2, ..., Tr_n\}$, where each Tr_i represents an individual trace consisting of a sequence of spans and logs. Thus, Tr = (Span, Log) represents the combination of spans and logs across all traces.

1) Span prepossessing and vector generation: For each span, we extract all available common attributes within it: start time, end time, duration, span ID, service name (i.e., service component) and URL (i.e., operation). We normalize time related attributes within each span's context, considering the unique characteristics and scale of each span. We denote vector representations of normalized span attributes as $V_{\text{start_time}}$, $V_{\text{end_time}}$ and V_{duration} for Span of Tr.

For MSS, Span IDs are designed with a hierarchical structure reflecting the relationship between the active spans and their nested sub-spans [3]. Take Figure 3 as an instance, assuming that Span A is the active span, both Span B and Span C are nested in Span A. When Span A initiates, it is assigned a Span ID "a480f2.0", while its nested spans, Span B and Span C, are assigned with the derivative Span IDs



Fig. 4. MultiHAttenAE basic architecture

"a480f2.1" and "a480f2.2" respectively. Spans like Span D and Span E, which are not nested sub-spans of any active spans, receive distinct Span IDs, like "a343mc.0" and "a987gq.0", to reflect their separate execution pathways. In our study, to construct a vector representation, we abstract away the shared common prefix in Span IDs and only retain hierarchical-level digits. For example, we reassign Span A, Span B, Span C, Span D, Span E with Span ID "1.0", "1.1", "1.2", "2.0", "3.0". We normalize Span IDs within the trace context. As such, we construct a vector representation $V_{\text{span_id}}$ for *Span* of *Tr*.

We concatenate textual attributes "service name" and "URL" to form a singular attribute termed "service operation". Template-based representation of textual attributes may not be ideal given the diverse range of service operations often present in MSS [24], [25]. Thereby, we use a neural representation method from the study [26]. To construct a neural representation for service operations, we undertake three steps. Step 1. Prepossessing. We convert all uppercase letters to lowercase, substitute specific variables with standard identifiers (e.g., replace instances like "Prod1234" with "ProductID"), and remove any non-alphabetic characters. Step 2. Tokenization. We use WordPiece tokenization [27] to tokenize service operations into subwords. It starts with a vocabulary of individual characters and then incrementally combines them into more extensive sub-word units. We choose WordPiece over other methods because it gives the flexibility to handle any words (including new, unfamiliar and out-of-vocabulary ones) with a relatively small vocabulary of sub-word units, as every word is decomposable into subword units. Step 3. Neural representation. We feed the subwords, into the BERT base model [28] to generate word embeddings for each sub-word. We use the word embeddings generated by the last encoding layer of the model and calculate the sentence embedding of each service operation as the average of its word embeddings. This process yields a vector representation $V_{\text{operation}}$ for service operations of Span of Tr.

For Tr, we concatenate the individual vector representations acquired from the preceding phase, thereby establishing a composite vector $V_{\text{span}} \in \mathbb{R}^{d_{\text{span}}}$ for Span, where d_{span} represents the dimensionality of the vector space for V_{span} :

$$V_{\text{span}} = \text{Concat}(v_{\text{start_time}}, v_{\text{end_time}}, v_{\text{duration}}, v_{\text{span_id}}, v_{\text{operation}})$$
(1)

2) Log prepossessing and vector generation: To capture the semantic contents of each log, we extract all textual attributes: severity level (such as INFO, WARN, ERROR), component (a

part of the system that generated the message), log message content (written by developers reflecting the system state). We concatenate these attributes to form a singular attribute termed "log event". We build a neural representation for log events that omits the usual step of log parsing. This approach leverages the capability of neural representations to more effectively comprehend the semantic meanings of log events [29]. Some past works indicate that log parsing might provide limited benefits [26], [30]. Our method for building a neural representation of log events follows the same three-step process as the one that is used for service operations in spans in Section III-B1. This process yields a vector representation $V_{log} \in \mathbb{R}^{d_{log}}$, which comprises the sentence embeddings of log events, for Log of Tr. Here, d_{log} represents the dimensionality of the vector space for V_{log} .

3) Trace representation construction: For a given MSS, we construct trace representations for its traces Tr utilizing our MultiHAttenAE, which consists of two main components: encoder and decoder. The encoder first projects the input vectors V_{span} and V_{log} into a common feature space $\mathbb{R}^{d'}$:

$$V'_{\text{span}} = g(W_{\text{span}}V_{\text{span}} + b_{\text{span}})$$

$$V'_{\text{log}} = g(W_{log}V_{\text{log}} + b_{log})$$
(2)

where, g denotes the activation function, W_{span} and W_{log} are the respective weight matrices, and b_{span} and b_{log} are the bias vectors.

The encoder incorporates the multi-head attention mechanism [31], as shown in Equation 3. First, the mechanism calculates the attention each sequence element in Q should put on all other elements in K and V, where Q, K and V denote the matrices for queries, keys, and values, respectively. It computes initial attention scores by taking the dot product of Q and K, regulates these scores by $\sqrt{d_k}$ for numerical stability, and then processes the scores through a softmax function, yielding the final attention distribution. This final attention distribution assigns weights to the elements in V. Second, the mechanism characterizes each attention head (head_i), which is computed by separate learned linear projections of Q, K, and V using the matrices W^Q , W^K , W^V as learnable weights. This process transforms the original matrices to produce distinct queries, keys, and values for each head, enabling a unique attention mechanism per head. Third, the mechanism details the formation of the multi-head attention, termed as MultiHead. It concatenates the outputs of all the individual attention heads into a single vector and linearly transforms this single vector using the weight matrix W^O to produce the final representation for multihead attention representation.

Attention
$$(Q, K, V) = \operatorname{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V,$$

head_i = Attention $(QW^Q, KW^K, VW^V),$
MultiHead (Q, K, V) = Concatenate $(\operatorname{head}_1, \dots, \operatorname{head}_h)W^O$
(3)

Our encoder takes V'_{span} and V'_{\log} as input into the above multi-head attention mechanism to fuse them into a lowdimensional trace representation Z for the set of traces T_r . This approach effectively captures the essential patterns and dependencies present within the multi-modal data V'_{span} and V'_{\log} . The fusion process is delivered as:

$$\mathbf{Z} = MultiHead(V'_{\text{span}}, V'_{\text{log}}, V'_{\text{log}})$$
(4)

where, we set V'_{span} as Q, and V'_{log} as both K and V. This setup aligns with the roles of spans in reflecting the trace structure and service communications, while logs provide detailed contextual event information.

For the set of traces $Tr = \{Tr_1, Tr_2, ..., Tr_n\}$, we generate the corresponding trace representations $Z = \{Z_1, Z_2, ..., Z_n\}$, where Z_i corresponds to Tr_i . The decoder reconstructs trace representations Z into the original span and log vectors, effectively inverting the encoder's process:

$$\hat{V}_{\text{span}} = g(W'_{\text{span}}Z + b'_{\text{span}})$$

$$\hat{V}_{\text{log}} = g(W'_{\text{log}}Z + b'_{\text{log}})$$
(5)

where, $\hat{V}_{\text{span}} \in \mathbb{R}^{d_{\text{span}}}$, and $\hat{V}_{\log} \in \mathbb{R}^{d_{\log}}$, g is the activation function, and W'_{span} and W'_{\log} are the respective weight matrices, and b'_{span} and b'_{\log} are the bias vectors.

Training MultiHAttenAE includes optimizing its parameters Ψ to minimize the overall loss \mathcal{L} which between original and reconstructed vectors:

$$\min_{\Psi} \mathcal{L} = \|\hat{V}_{\text{span}} - V_{\text{span}}\|^2 + \|\hat{V}_{\log} - V_{\log}\|^2$$
(6)

here, the mean squared error (MSE) loss is utilized to quantify the euclidean distance between the original $(V_{\text{span}}, V_{\text{log}})$ and their respective reconstructed $(\hat{V}_{\text{span}}, \hat{V}_{\text{log}})$.

C. TE-MAML for few-shot abnormal trace classification

TE-MAML constructs a classifier to do few-shot abnormal trace classification. Following the principle of few-shot learning, we define our classification problem as the N-way K-shot setup. This setup involves N distinct fault categories, with each category having K labeled support examples.

TE-MAML is built upon the meta-learning algorithm MAML, which aims to learn a meta-leaner with an optimal set of initial model parameters that can quickly adapt to new tasks using only a few instances and a small number of gradient updates [18]. TE-MAML uses TE [31] as the meta-learner. The core objective of TE-MAML is to learn the initial model parameters of this meta-learner, such that when provided with the support set S of an N-way K-shot task, it can be quickly and robustly updated to perform well on the query set Q (i.e.,

classify the query set Q well). Figure 5 illustrates TE-MAML's basic architecture. TE-MAML progresses through two phases: meta-training, where the meta-leaner is trained, and meta-testing, where the meta-leaner is adapted to the new tasks. We explain the meta-leaner, and meta-training and meta-testing phases in detail below.

1) Meta learner: In TE-MAML, the meta-learner TE is denoted as f, which takes trace representations Z of a given MSS as input. We choose TE as the meta-leaner due to its self-attention mechanism, which can weigh the most relevant elements in each trace representation Z_i and capture the dependencies among them for doing classification. This selfattention mechanism follows the multi-head attention mechanism presented in Equation 3. It is termed 'self-attention' because it uses the same input Z as Q, K, and V in the attention process, as illustrated below:

$$output = MultiHead(Z, Z, Z)$$
(7)

For classifying abnormal traces, the output from the above self-attention mechanism is further passed through a pooling layer to highlight key features, a dropout layer to prevent overfitting, and a fully connected layer that transforms the refined output into a suitable form for classification. Finally, the softmax classifier takes the output from the fully connected layer and computes the probabilities for each fault category.

2) Meta-training: During this phase, TE-MAML is trained on a set of meta-training tasks (i.e., few-shot abnormal trace classification tasks), denoted as T = (S, Q), which can be sampled from a single MSS or multiple MSS. This phase aims at optimizing the meta-learner's initial model parameters to enable rapid and effective adaptation to new tasks. Each of these meta-training tasks is unique and it is denoted by $T_i =$ (S_i, Q_i) , where S_i is a support set and Q_i is a query set for the *i*-th task. The support set $S_i = \{(z_{ij}^{spt}, y_{ij}^{spt})\}_{i=1}^{N \times K}$ is indexed by j from 1 to $N \times K$. Here, $N \times K$ follows our N-way K-shot setup, indicating that there are N distinct fault categories and each has K labeled trace instances. Each $(z_{ij}^{spt}, y_{ij}^{spt})$ is a pair of a trace representation and its corresponding fault category label, where $y_{ij}^{spt} \in \{1, 2, \dots, N\} = [N]$. Meta-training tasks form the foundation for the meta-learner's adaptive learning process. Similarly, the query set $Q_i = \{(z_{ig}^{qry}, y_{ig}^{qry})\}_{g=1}^{N \times M}$, indexed by g from 1 to $N \times M$. Here, $N \times M$ indicates there are N distinct fault categories and each has M labeled trace instances. M is typically greater than K (i.e., $|Q_i| > |S_i|$) to ensure robust optimization across all meta-training tasks. Trace representations z_{ij}^{spt} and z_{ig}^{qry} for S_i and Q_i respectively are generated using the optimized encoder of MultiHAttenAE for a given MSS.

As shown in Figure 5, at the beginning of the metatraining phase, model parameters of the meta-learner f are randomly initialized. Here, the meta-learner f is considered as a parameterized function f_{θ} with its parameters θ .

The meta-training phase is composed of two integral parts: the inner loop and the outer loop. The inner loop is responsible for task-level learning, wherein the model is adapted to each



Fig. 5. TE-MAML learning process

meta-training task. The outer loop operates at the meta-level, with the objective of adjusting the model's initial parameters in such a way that a few gradient steps will produce maximum performance on new tasks.

In the inner loop, when adapting to each training task T_i , the meta-learner f's parameters θ are transformed into taskspecific parameters θ'_i , corresponding to the modified function $f_{\theta'_i}$. θ'_i is computed using gradient descent updates on the support set S_i of the training task T_i . Each gradient descent update is computed as:

$$\theta_i' = \theta - \alpha \nabla_\theta \mathcal{L}_{T_i}(f_\theta(S_i)) \tag{8}$$

where α is the learning rate for the inner loop update process and \mathcal{L}_{T_i} is the loss function for task T_i .

In the outer loop, the meta-learner f's model parameters are trained by optimizing for the performance of θ'_i with respect to θ across meta-training tasks T. More concretely, the metatraining objective is to minimize the overall loss \mathcal{L} on the query set Q_i of the training task T_i as follows:

$$\min_{\theta} \mathcal{L}(\theta) = \sum_{T_i \in T} \mathcal{L}_{T_i}(f_{\theta'_i}(Q_i))$$
(9)

During the outer loop, the meta-training optimization across training tasks T is performed. The model parameters θ are updated p times, where p is defined as the predetermined number of optimization steps. Each update gets the updated model parameters θ_c . For example, when using one gradient update (i.e., p = 1),

$$\theta \leftarrow \theta - \beta \nabla_{\theta} \sum_{T_i \in T} \mathcal{L}_{T_i}(f_{\theta'_i}(Q_i))$$
(10)

where β is the learning rate for the outer loop update process.

The standard MAML outer loop update involves a computationally intensive process of computing a gradient through a gradient, which requests an additional backward pass through the function f to compute Hessian-vector products [18]. We use a first-order approximation [18] to simplify the update by such second-order derivatives.

As the final result of the outer loop, the optimal initial model parameters θ^* are obtained. The meta-learner f is then initialized with θ^* , transforming it into f_{θ^*} . This optimally initialized f_{θ^*} is then used in the meta-testing phase.

3) Meta-testing: This phase adapts the optimized metalearner f_{θ^*} to the meta-testing task $T_{ts} = (S_{ts}, Q_{ts})$, which has the same structure as the meta-learning task $T_i = (S_i, Q_i)$. In this phase, the support set S_{ts} is fed into the meta-learner f_{θ^*} . To finely adapt to the meta-testing task T_{ts} , a limited number of gradient updates are applied, mirroring gradient updates in the meta-training's inner loop. These adaptation steps are aimed at fine-tuning the meta-learner f's parameters specifically for T_{ts} , resulting in adapted parameters θ'_{ts} . After adaptation, the meta-leaner, now as $f_{\theta'_{ts}}$, proceeds to classify abnormal traces in Q_{ts} into fault categories. Typically, the optimized meta-learner f_{θ^*} will be tested on many metatesting tasks to ensure its generalization capability.

IV. EVALUATION

A. Experimental design

1) Dataset: We use two open-source datasets to establish our own dataset: DeepTraLog [32] and Nezha [33]. (1) Deep-TraLog was collected from TrainTicket, a medium-scale, widely-used benchmark MSS for train ticket booking, comprising 45 services. TrainTicket contains various real-industrial faults in different branches. DeepTraLog was created by Zhang et al. [9], who executed such fault branches of TrainTicket and generated abnormal traces of 14 fault categories. These fault categories pertain to related to asynchronous interaction, multiinstance, configuration, and monolithic dimensions. (2) Nezha was created by Yu et al. [34]. It includes data from two MSS: TrainTicket, and a small-scale e-commerce benchmark MSS OnlineBoutique composed of 12 services. This dataset was constructed by injecting faults into the pods of TrainTicket and OnlineBoutique and collocating abnormal traces. Within each system, faults injected into each pod belong to four categories: CPU-intention, CPU-consumption, service exception, and message return. Each fault case within a pod here is considered as a unique fault category.

2) Training and evaluation: To train MultiHAttenAE for the TrainTicket and OnlinBoutique, we utilized a random selection of 3960 normal traces (3360 training/570 validation) for each system from both DeepTraLog and Nezha datasets. Traces within each system were input into MultiHAttenAE to train and optimize MultiHAttenAE. This process is detailed in the specified section III-B.

To train and evaluate TE-MAML, we establish our fault dataset using DeepTraLog and Nezha. Our fault dataset in-

 TABLE I

 DESCRIPTIVE STATISTICS ON TRACES IN OUR FAULT DATASET.

Trainticket.Base fault categories	Mean	Min	Max
Unique traces per fault category:	1117	26	2309
Spans per trace:	71	1	303
Logs per trace:	44	1	219
Trainticket.Novel fault categories	Meam	Min	Max
Unique traces per fault category:	1275	45	2546
Spans per trace:	87	1	345
Logs per trace:	60	1	340
OnlineBoutique.Base fault categories	Mean	Min	Max
Unique traces per fault category:	565	32	1018
Spans per trace:	53	1	180
Logs per trace :	53	4	184
OnlineBoutique.Novel fault categories	Mean	Min	Max
Unique traces per fault category:	320	34	902
Spans per trace:	52	1	190
Logs per trace:	49	4	164

cludes abnormal traces from 30 fault categories for TrainTicket and 32 fault categories for OnlineBoutique. We divide 30 fault categories of Trainticket into 20 base and 10 novel fault categories, and 32 fault categories of OnlineBoutique into 22 base and 10 novel fault categories. These categories within each system do not overlap. That is, there are 20 and 22 base fault categories in TrainTicket and OnlineBoutique, respectively, along with 10 novel fault categories for each system. To provide a consistent basis to compare experimental results, we standardized the number of novel fault categories for both systems. Considering the variability in the length of spans and logs within traces, and the significant role semantic attributes of these spans and logs play in constructing trace representations via MultiHAttenAE (Section III-B). We design the composition of both base and novel fault categories for each system to include a random mix of fault categories from our dataset. This mix incorporates fault categories that comprise traces with varying lengths of spans and logs. It ensures the representation of fault categories with traces containing both short and long spans and logs within base and novel fault categories for each system. Table I presents summary statistics for abnormal traces in base and novel categories of each system. The full list of base and novel fault categories of each system will be presented in our replication package. Using the established base and novel fault categories on both systems, we perform the following experiments:

- E1 (TrainTicket to TrainTicket). We train our framework on TrainTicket's base fault categories and adapt it to TrainTicket's novel fault categories.
- E2 (OnlineBoutique to OnlineBoutique). We train our model on OnlineBoutique's base fault categories and then adapt it to OnlineBoutique's novel fault categories.
- E3 (OnlineBoutique to TrainTicket). We train our model on OnlineBoutique's base fault categories and then adapt it to TrainTicket's novel fault categories.
- E4 (TrainTicket to OnlineBoutique). We train our model on TrainTicket's base fault categories and then adapt it to OnlineBoutique's novel fault categories.

E1 and E3 are within system experiments while E2 and

E4 are cross-system experiments. Considering prior MAML studies [18], [35] and our fault dataset, we investigate 5-way 5-shot and 10-shot setups across experiments E1-E4. In these experiments, each meta-training and meta-testing task is a unique abnormal trace classification task, where each task involves classifying abnormal traces from a distinct set of 5-way fault categories. In each experiment, we train and evaluate our TE-MAML using the randomly established meta-training and meta-testing tasks on base and novel fault categories of the involved system(s), as described below:

In the meta-training phase of each experiment, four sets of unique meta-training tasks are randomly generated using the involved system's base fault categories. Each task is either a 5-way 5-shot or 10-shot abnormal trace classification task. Specifically, for each task, 5 unique fault categories (with 5shot or 10-shot labeled trace instances) are randomly selected from base fault categories of the involved system. A total of four unique meta-learning tasks are established.

In the meta-testing phase of each experiment, 50 distinct meta-testing tasks are randomly created by selecting 5 distinct fault categories out of 10 novel fault categories of the involved system. The mathematical permutations for selecting 5 out of 10 categories yields a total of 252 possible permutations, calculated as C(10,5). We focus on 50 distinct meta-testing tasks that represent 252 possible permutations to evaluate our framework following the prior studies [18], [35]. This approach allows for a focused comprehensive evaluation, ensuring a broad coverage of various fault categories while keeping the data size within practical limits for detailed examination. Each meta-testing task is unique and is either a 5-way 5-shot or 10-shot abnormal trace classification task.

Upon randomly established meta-training and meta-testing tasks during respective meta-training and meta-testing phases in each experiment, traces from the corresponding fault categories within these tasks are automatically sampled from our dataset for the involved system(s). The given system's optimized encoder of MultiAttenAE is employed to generate representations for sampled traces for this system.

3) Implementation details: E1-E4 are conducted on a Linux server with a 32-core CPU and an NVIDIA Ampere A100 GPU with 40 GB of memory, utilizing Python 3.10.6. We train MultiHAttenAE and TE-MAML using the AdamW optimizer. Further details regarding hyperparameter settings of Multi-HAttenAE and TE-MAML will be provided in our replication package.

4) Baselines: Given the absence of directly comparable research within this domain, our baselines are established via an ablation study by systematically exploring alternative configurations to evaluate the impact of different components within our model. These configurations were designed by replacing key elements of our framework. We carried out experiments E1-E4 on these baselines:

As the MultiHAttenAE alternative, OnlySpan follows the related work [23] to consider each trace as a sequence of spans and construct trace representations using span attributes. As alternatives to the multi-head attention fusion (used in

MultiHAttenAE alternative: Feed **only spans (OnlySpan)** into TE-MAML for classification.

Multihead atten fusion alternatives: Linear-based AE (Linear-AE) fusion +TE-MAML Gated linear unit-based AE (GLU-AE) fusion +TE-MAML

Meta-learner alternatives: MultiHAttenAE+Linear-MAML MultiHAttenAE+RNN-MAML MultiHAttenAE+LSTM-MAML MultiHAttenAE+CNN-MAML

MAML alternatives: MultiHAttenAE+Prototypical network (ProtoNet) MultiHAttenAE+TE-based Matching network (TE-MatchingNet) MultiHAttenAE+Nearest neighbor (NearNeighbor) MultiHAttenAE+Decison tree

our MultiHAttenAE), Linear-AE and GLU-AE use the linear projection and gate mechanism, respectively, to fuse spans and logs for generating trace representations. These approaches provide simpler alternatives to multi-head attention mechanisms. Meta-learner alternatives include the basic linear model, sequence models RNN and LSTM, and CNN. MAML alternatives include metric-based meta-learning algorithms (ProtoNet [36], MatchingNet [37]) and traditional classification models (NearNeighbor and Decision tree). These alternatives are simpler than MAML. Note that, MAML alternative baselines are only evaluated on E1 and E3 since the fundamental algorithms of these do not have cross-system adaptability.

5) Evaluation metrics: In each experiment's meta-testing phase, for our framework and baselines, each is evaluated on 50 meta-testing tasks of the involved system, randomly selected from 252 possible permutations. To ensure a consistent comparison, the same random seed is used for all evaluations, ensuring that both our framework and all baselines are subjected to the same set of task permutations in each experiment.

Effectiveness. To assess the effectiveness of our framework and baselines, we have chosen accuracy as the evaluation metric, a standard measure for multi-class classification tasks. We reported the average accuracy along with the 95% confidence interval computed across 50 meta-testing tasks in each experiment, as well as the range of task accuracies (minimum and maximum). This provides a comprehensive view of the approach's effectiveness and offers insights into its consistency and reliability in various scenarios. Given that both DeepTraLog and Nezha datasets (used to establish our fault dataset) were created using automatic labeling methods, with normal and abnormal traces collected by executing normal and faulty versions of the system interchangeably, there is a chance that some latent normal traces are labeled as abnormal in the respective fault category and vice wise. To mitigate the risk of obtaining an inaccurate evaluation of the approach's performance due to such automatic labeling methods, we calculated each meta-testing task's accuracy by running this task 5 times and taking the highest. Efficiency. To evaluate the efficiency of our framework and baselines, we calculate the time each

takes at different phases of the training and adaptation process. This analysis provides insights into the computational costs of our framework compared with baselines.

B. Evaluation results

Table II and III present effectiveness evaluation results of our framework and all baselines. They report the average accuracy (in %) and the corresponding 95% confidence interval, which are computed across 50 meta-testing tasks in each experiment. The range of task accuracies (minimum and maximum) is also reported in these tables. Regarding efficiency, Table IV presents the average time (in seconds) taken by our framework and other effective baselines to adapt to each meta-testing task in each experiment. Table V presents the average time (in seconds) taken by our framework to construct trace representations for each meta-testing task in each experiment, compared with other effective baselines and the baseline OnlySpan+TE-MAML. For our framework and AE-based baselines, this time is calculated as the time taken by the optimized encoder to construct trace representations (see Figure 1) after training. For the baseline OnlySpan+TE-MAML, the time is calculated on how long it takes to use span vectors to directly construct trace representations. We analyze the above evaluation results to answer our research questions in the following sub-sections.

1) RQ1: Within-system adaptability: Effectiveness. Table II shows that both our framework and multi-head attention alternatives get similar high average accuracy in both 5-shot and 10-shot setups of E1 for Trainticket. Referring to Table III, our framework outperforms all baselines in both 5-shot and 10-shot setups of E2 for OnlineBoutique. Efficiency. Table IV shows that our framework and other baselines (Linear-AE and GLU-AE) that have a pre-training phase with AE can adapt more quickly to new tasks compared to the baseline MultiHAttenAE+NearNeighbor, which does not have a pretraining phase and thus requires instant training on new tasks. Our framework and these approaches also require time to construct trace representations before adaptation, as shown in Table V. Answer to RQ1. Overall, the evaluation results confirm our framework's effective and efficient within-system adaptability.

Exploratory analysis. Besides, for OnlineBoutique, through analysis, we found that both our framework and effective baselines (MultiHAttenAE+CNN-MAML, MultiHAttenAE+NearNeighbor) are less effective in differentiating abnormal traces between those caused by CPU contention and network delay. This issue arises because both fault cases may result in service latency, as reflected by anomalies in the span start time, end time, and duration. Such specific fault cases, which are part of tasks that get lower accuracy among all meta-testing tasks in both 5-shot and 10-shot setups of E2, contribute to a reduction in the overall average accuracy. Similarly, for Trainticket, the meta-testing task that receives the lowest average accuracy (62.00-78.00%) on both our framework and the highest baselines, including Multihead attention fusion alternatives

 TABLE II

 Comparison of our framework and baselines on Trainticket.

Model	E1. TrainTicket	t to TrainTicket	E3. OnlineBoutique to TrainTicket		
	5-shot	10-shot	5-shot	10-shot	
Our MultiHAttenAE+TE-MAML	92.91±2.10 (74.67-100.0)	93.26±1.40 (76.00-100.0)	86.35±2.00 (70.67-100.0)	92.19±1.99 (74.67-100.0)	
MultiHAttenAE alternatives:					
OnlySpan+TE-MAML	80.64±2.84 (57.33-97.33)	78.77±2.80 (60.00-97.33)	79.25±2.89 (57.33-97.33)	80.19±3.10 (56.00-97.33)	
Multihead atten fusion alternatives:					
Linear-AE+TE-MAML	89.15±2.29 (73.33-100.0)	90.59±2.43 (70.67-100.0)	83.09±2.55 (62.67-97.33)	90.61±2.01 (72.00-100.0)	
GLU-AE +TE-MAML	92.21±1.73 (77.33-100.0)	93.07±1.64 (77.33-100.0)	85.07±2.38 (66.67-100.0)	94.40±2.19 (72.00-100.0)	
Transformer encoder alternatives:					
MultiHAttenAE+Linear-MAML	45.84±2.21 (25.33-61.33)	45.36±2.16 (30.67-60.00)	43.81±1.99 (28.00-58.67)	43.87±1.93 (29.33-60.00)	
MultiHAttenAE+RNN-MAML	49.65±2.09 (37.33-64.00)	42.88±1.93 (24.00-58.67)	48.45±1.75 (38.67-65.33)	47.07±1.88 (34.67-58.67)	
MultiHAttenAE+LSTM-MAML	41.39±2.20 (21.33-56.00)	42.67±1.91 (29.33-56.00)	40.32±1.68 (22.67-52.00)	42.29±1.79 (25.33-56.00)	
MultiHAttenAE+CNN-MAML	57.06±2.85 (41.33-81.00)	69.20±2.19 (56.00-88.00)	49.04±1.80 (38.67-64.00)	69.47±2.65 (48.00-89.33)	
MAML alternatives:					
MultiHAttenAE+TE-MatchingNet	76.56±2.80 (49.33-93.33)	76.05±2.37 (50.67-94.67)	_	_	
MultiHAttenAE+ProtoNet	57.25±0.03 (40.00-74.67)	59.68±0.03 (44.00-76.00)	_	_	
MultiHAttenAE+NearNeighbor	88.19±0.02 (74.67-98.67)	92.56±0.02 (78.00-100.0)			
MultiHAttenAE+DecisonTree	66.80±0.03 (46.67-88.00)	77.09±0.03 (54.67-96.00)	_	_	

TABLE III

COMPARISON OF OUR FRAMEWORK AND BASELINES ON ONLINEBOUTIQUE

Model	E2. OnlineBoutique	to OnlineBoutique	E4. TrainTicket to OnlineBoutique		
	5-shot	10-shot	5-shot	10-shot	
Our MultiHAttenAE+TE-MAML	82.50±2.35 (65.33-98.67)	85.20±2.33 (66.67-98.67)	82.37±2.07 (64.00-97.33)	84.77±2.28 (68.00-98.67)	
MultiHAttenAE alternatives:					
OnlySpan+TE-MAML	72.83±2.40 (57.33-88.00)	73.15±2.81 (46.67-92.00)	71.81±2.25 (56.00-85.33)	73.60±2.21 (57.33-85.33)	
Multihead atten fusion alternatives:					
Linear-AE+TE-MAML	76.15±2.59 (60.00-95.00)	78.21±2.50 (64.00-96.00)	75.81±2.45 (52.00-89.33)	74.32±2.44 (53.33-88.00)	
GLU-AE +TE-MAML	80.61±2.96 (58.67-98.67)	77.49±2.67 (48.00-94.67)	74.96±2.76 (54.67-94.67)	77.57±2.7 (56.00-94.67)	
Transformer encoder alternatives:					
MultiHAttenAE+Linear-MAML	42.59±3.63 (20.00-77.33)	40.75±3.53 (21.33-68.00)	47.01±3.59 (25.33-74.67)	44.35±4.10 (20.00-89.33)	
MultiHAttenAE+RNN-MAML	72.59±2.50 (54.67-94.67)	64.75±2.53 (46.67-80.00)	72.58±2.58 (54.70-94.70)	71.01±2.80 (56.00-89.33)	
MultiHAttenAE+LSTM-MAML	54.80±2.11 (40.00-70.67)	55.97±2.25 (41.33-69.33)	56.19±1.92 (38.67-72.00)	59.71±2.05 (42.67-77.33)	
MultiHAttenAE+CNN-MAML	80.10±2.16 (60.00-94.67)	83.07±3.29 (68.00-97.33)	79.01±2.63 (56.00-97.33)	84.08±2.76 (65.33-100.0)	
MAML alternatives:					
MultiHAttenAE+TE-MatchingNet	76.29±3.00 (54.67-96.00)	73.11±2.94 (50.67-94.67)	_		
MultiHAttenAE+ProtoNet	74.51±0.03 (53.33-92.00)	76.59±0.04 (58.66-94.67)	_		
MultiHAttenAE+NearNeighbor	80.96±0.03 (64.00-98.67)	84.75±0.03 (62.80-98.67)	_		
MultiHAttenAE+DecisonTree	66.99±0.02 (54.67-80.00)	73.79±0.03 (58.67-82.67)	_	—	

 TABLE IV

 Adaptation efficiency (task average in seconds)

	E	E1	E	E3
	5-shot	10-shot	5-shot	10-shot
Our framework	0.0460	0.0680	0.0651	0.0953
GLU-AE +TE-MAML	0.0640	0.0942	0.0651	0.0954
Linear-AE +TE-MAML	0.0748	0.0946	0.0640	0.0944
MultiHAttenAE+NearNeighbor	0.1860	0.1890	—	—
	E	E2	E	24
	5-shot	10-shot	5-shot	10-shot
Our framework	0.0848	0.0977	0.0875	0.0974
GLU-AE +TE-MAML	0.0801	0.0978	0.0731	0.0969
CNN-MAML	0.0470	0.0472	0.0417	0.0490
MultiHAttenAE+NearNeighbor	0.5140	0.5307	—	_

 TABLE V

 TRACE CONSTRUCTION EFFICIENCY (TASK AVERAGE IN SECONDS)

	Trainticket		OnlineBoutique	
	5-shot	10-shot	5-shot	10-shot
Our MultiHAttenAE	1.8436	2.1524	4.1426	5.0189
GLU-AE+TE-MAML	1.8395	2.1541	4.3121	5.1921
Linear-AE+TE-MAML	1.8286	2.1486	4.1223	5.2066
OnlySpan+TE-MAML	1.1671	1.1818	1.9065	2.1822

certain service latency.

and MultiHAttenAE+NearNeighbor, is less effective in differentiating abnormal traces caused by network delay and

2) RQ2: Cross-system adaptability: Effectiveness. Table III shows that our framework achieves similar average accuracy in 5-shot/10-shot setup of E2 and E4. Likewise, Table II illustrates that E1 and E3 get similar levels of average accuracy in their respective 10-shot setup. These results evidenced our framework's cross-system adaptability. However, E3 gets lower average accuracy than E1 in the 5-short setup. A possible reason is: abnormal traces from OnlineBoutique are less semantically rich compared to the ones from TrainTicket. This can be attributed to two factors. First, fault categories in OnlineBoutique are less diverse than those in TrainTicket, as described in Section IV-A. Second, OnlineBoutique's traces are simpler than TrainTicket's, with the shorter average length of spans and logs, as shown in Table I. In such a context, when adapting the knowledge learned from OnlineBoutique tasks to new tasks from Trainticket, more instances may be needed to provide to TE-MAML during the meta-testing phase to ensure high effectiveness. Efficiency. Our framework takes similar adaptation times across meta-testing tasks in E1 and E3, as well as the E2 and E4, on both 5-shot and 10-shot setups, as shown in Table IV. This indicates that our framework requires a similar amount of time for withinsystem adaptation and cross-system adaptation. Answer to RQ2. Overall, the evaluation results confirm our framework's effective and efficient cross-system adaptability.

3) RQ3: Component impact: As shown in Table II and III, each part of our framework contributes to its overall effectiveness. Among them, MultiHAttenAE is the most significant contributor to the framework's effectiveness. This can be seen from these tables, our framework outperforms the baseline OnlySpan+TE-MAML by about 10% in most setups of E1-E4. However, MultiHAttenAE is also the most computationally expensive part of our framework.

We conduct several measurements to provide an overview of the computational costs of the main components of MultiHAttenAE, aiming to help others make decisions about its use in their applications. Table VI shows the time taken (in seconds) by our MultiHAttenAE to construct neural representations for 3960 normal traces used to train our MultiHAttenAE and other AE-based baselines. We compare this time with the time taken to construct neural representations for semantic attributes of only spans. Neural representation is the most computationally expensive part of our framework, as it utilizes a large language model, BERT.

Table VII compares the average time taken by our MultiHAttenAE and other AE-based baselines to train through the AE encoder-decoder structure (as shown in Figure 1) for each epoch, using a batch size of 32 and a total of 3960 normal traces for the respective system. We measure this average time by running 100 epochs, but in our experiments, MultiHAttenAE requires only 20-30 epochs to train effectively with appropriate hyperparameters for Trainticket and OnlineBoutique. Also, others can refer to Table V to get an idea of how much time the optimized MultiHAttenAE takes to construct trace representations for a given MSS.

 TABLE VI

 NEURAL REPRESENTATION CONSTRUCTION EFFICIENCY (IN SECONDS)

For span and log semantic attributes	Trainticket 323.4698	OnlineBoutique 120.3414
For only span semantic attribute	77.0094	57.9034

 TABLE VII

 TRAINING EFFICIENCY (PER EPOCH IN SECONDS)

Our MultiHAttenAE	Trainticket 89.6450	OnlineBoutique 62.5008
GLU-AE Linear-AE	86.3644 88.0473	62.1378 62.2194

V. THREATS TO VALIDITY

The main internal threats may lay in the correctness of baselines' development. Given that there is only one related work [23], we designed most of the baselines ourselves based on our understanding via an ablation study. To mitigate this threat, we carefully reviewed the core principles of the baseline methodologies, conducted extensive testing, and consulted domain experts for their insights.

The external threats may arise from the limitation of our fault dataset sourced from open datasets DeepTraLog and Nezha. First, we intended to include trace-related metrics in constructing the latent trace representations for MSS. Including trace-related metrics could potentially help in distinguishing resource-related anomalies within a trace. This may address the issue referred to RQ1 results, where our framework and baselines are less effective in differentiating abnormal traces caused by CPU contention and network delay in OnlineBoutique, as well as those caused by network delay and service latency in Trainticket. Nezha is a dataset that contains sufficient modality data and enables the extraction of spans, logs, and metrics within a trace to construct its representation for the observed MSS. However, we did not find other similar MSS open datasets that meet such needs. Second, we aimed to expand the evaluation of our framework by including additional datasets that contain abnormal traces from a broader range of fault categories, as well as datasets from various MSS. Despite our efforts, we were unable to find such datasets. Thus, our study used abnormal traces from 30 and 32 fault categories in Trainticket and OnlineBoutique, respectively. This limitation led us to adopt a 5-way setup in order to ensure an adequate distribution of meta-training and meta-testing tasks in each experiment. However, prior studies have demonstrated that MAML is highly effective for complex text classification tasks (e.g., 20-way and 50way classifications) across a variety of contexts [38]-[40]. The above two external threats pose the possibility that our application of our framework for classifying abnormal traces for MSS might not fully utilize its generalization capabilities. Addressing these threats may improve our framework, thereby enhancing the precision and robustness of trace-level RCA in MSS. To address these threats, we are in the process of deploying MSS and generating new open datasets.

VI. CONCLUSION

This paper proposes a framework for few-shot abnormal trace classification for MSS with two main components: (1) MultiHAttenAE for constructing system-specific trace representations, which enables (2) TE-MAML to perform effective

and efficient few-shot learning for abnormal trace classification. The proposed framework is evaluated on representative benchmark MSS with open datasets. The evaluation results shows our framework's effective and efficient within-system and cross-system adaptability. Overall, our framework provides a solution for a few-shot abnormal trace classification for MSS. Future work will focus on further improving the generalizability, scalability, and interpretability of the proposed framework. Also, we plan to evaluate its performance on real industrial MSS.

VII. REPLICATION PACKAGE

To ensure the transparency and reproducibility of our research, we will publish a replication package upon acceptance of the paper. The replication package will include all the necessary artifacts to replicate our experiments, such as the source code, data, and implementation details. We will provide detailed instructions on how to use the replication package and will make it available on a GitHub public repository.

VIII. ACKNOWLEDGMENT

The authors would like to thank Nauman Ali for reading the paper and providing valuable feedback.

REFERENCES

- X. Zhou, X. Peng, T. Xie, J. Sun, C. Ji, W. Li, and D. Ding, "Fault analysis and debugging of microservice systems: Industrial survey, benchmark system, and empirical study," *IEEE Transactions on Software Engineering*, vol. 47, no. 2, pp. 243–260, 2018.
- [2] B. Li, X. Peng, Q. Xiang, H. Wang, T. Xie, J. Sun, and X. Liu, "Enjoy your observability: an industrial survey of microservice tracing and analysis," *Empirical Software Engineering*, vol. 27, pp. 1–28, 2022.
- [3] OpenTelemetry, "OpenTelemetry," 2024, accessed: 2024-03-14.[Online]. Available: https://opentelemetry.io/
- [4] J. Soldani and A. Brogi, "Anomaly detection and failure root cause analysis in (micro) service-based cloud applications: A survey," ACM Computing Surveys (CSUR), vol. 55, no. 3, pp. 1–39, 2022.
- [5] Á. Brandón, M. Solé, A. Huélamo, D. Solans, M. S. Pérez, and V. Muntés-Mulero, "Graph-based root cause analysis for service-oriented and microservice architectures," *Journal of Systems and Software*, vol. 159, p. 110432, 2020.
- [6] S. Nedelkoski, J. Cardoso, and O. Kao, "Anomaly detection from system tracing data using multimodal deep learning," in 2019 IEEE 12th International Conference on Cloud Computing (CLOUD). IEEE, 2019, pp. 179–186.
- [7] J. Kaldor, J. Mace, M. Bejda, E. Gao, W. Kuropatwa, J. O'Neill, K. W. Ong, B. Schaller, P. Shan, B. Viscomi *et al.*, "Canopy: An end-to-end peng xinrformance tracing and analysis system," in *Proceedings of the 26th symposium on operating systems principles*, 2017, pp. 34–50.
- [8] H. Chen, K. Wei, A. Li, T. Wang, and W. Zhang, "Trace-based intelligent fault diagnosis for microservices with deep learning," in 2021 IEEE 45th Annual Computers, Software, and Applications Conference (COMPSAC). IEEE, 2021, pp. 884–893.
- [9] C. Zhang, X. Peng, C. Sha, K. Zhang, Z. Fu, X. Wu, Q. Lin, and D. Zhang, "Deeptralog: Trace-log combined microservice anomaly detection through graph-based deep learning," in *Proceedings of the 44th International Conference on Software Engineering*, 2022, pp. 623–634.
- [10] J. Chen, F. Liu, J. Jiang, G. Zhong, D. Xu, Z. Tan, and S. Shi, "Tracegra: A trace-based anomaly detection for microservice using graph deep learning," *Computer Communications*, vol. 204, pp. 109–117, 2023.
- [11] K. Zhang, C. Zhang, X. Peng xinng, and C. Sha, "Putracead: Trace anomaly detection with partial labels based on gnn and pu learning," in 2022 IEEE 33rd International Symposium on Software Reliability Engineering (ISSRE). IEEE, 2022, pp. 239–250.

- [12] R. Chen, J. Ren, L. Wang, Y. Pu, K. Yang, and W. Wu, "Microegrcl: An edge-attention-based graph neural network approach for root cause localization in microservice systems," in *International Conference on Service-Oriented Computing*. Springer, 2022, pp. 264–272.
- [13] C. Padurariu and M. E. Breaban, "Dealing with data imbalance in text classification," *Proceedia Computer Science*, vol. 159, pp. 736–745, 2019.
- [14] Y. Wang, Q. Yao, J. T. Kwok, and L. M. Ni, "Generalizing from a few examples: A survey on few-shot learning," ACM computing surveys (csur), vol. 53, no. 3, pp. 1–34, 2020.
- [15] P. Liu, H. Xu, Q. Ouyang, R. Jiao, Z. Chen, S. Zhang, J. Yang, L. Mo, J. Zeng, W. Xue *et al.*, "Unsupervised detection of microservice trace anomalies through service-level deep bayesian networks," in 2020 *IEEE 31st International Symposium on Software Reliability Engineering* (ISSRE). IEEE, 2020, pp. 48–58.
- [16] Z. Xie, H. Xu, W. Chen, W. Li, H. Jiang, L. Su, H. Wang, and D. Pei, "Unsupervised anomaly detection on microservice traces through graph vae," in *Proceedings of the ACM Web Conference 2023*, 2023, pp. 2874– 2884.
- [17] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *science*, vol. 313, no. 5786, pp. 504–507, 2006.
- [18] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *International conference on machine learning*. PMLR, 2017, pp. 1126–1135.
- [19] R. Chen, S. Zhang, D. Li, Y. Zhang, F. Guo, W. Meng, D. Pei, Y. Zhang, X. Chen, and Y. Liu, "Logtransfer: Cross-system log anomaly detection for software systems with transfer learning," in 2020 IEEE 31st International Symposium on Software Reliability Engineering (ISSRE). IEEE, 2020, pp. 37–47.
- [20] C. Zhang, T. Jia, G. x. Shen, P. Zhu, and Y. Li, "Metalog: Generalizable cross-system anomaly detection from logs with meta-learning," in 2024 IEEE/ACM 46th International Conference on Software Engineering (ICSE). IEEE Computer Society, 2024, pp. 938–938.
- [21] X. Han and S. Yuan, "Unsupervised cross-system log anomaly detection via domain adaptation," in *Proceedings of the 30th ACM international conference on information & knowledge management*, 2021, pp. 3068– 3072.
- [22] N. Holla, P. Mishra, H. Yannakoudakis, and E. Shutova, "Learning to learn to disambiguate: Meta-learning for few-shot word sense disambiguation," arXiv preprint arXiv:2004.14355, 2020.
- [23] S. Nedelkoski, J. Cardoso, and O. Kao, "Anomaly detection and classification using distributed tracing and deep learning," in 2019 19th IEEE/ACM international symposium on cluster, cloud and grid computing (CCGRID). IEEE, 2019, pp. 241–250.
- [24] Y. Fu, M. Yan, J. Xu, J. Li, Z. Liu, X. Zhang, and D. Yang, "Investigating and improving log parsing in practice," in *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on* the Foundations of Software Engineering, 2022, pp. 1566–1577.
- [25] N. Dragoni, S. Giallorenzo, A. L. Lafuente, M. Mazzara, F. Montesi, R. Mustafin, and L. Safina, "Microservices: yesterday, today, and tomorrow," *Present and ulterior software engineering*, pp. 195–216, 2017.
- [26] V.-H. Le and H. Zhang, "Log-based anomaly detection without log parsing," in 2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE). IEEE, 2021, pp. 492–504.
- [27] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey *et al.*, "Google's neural machine translation system: Bridging the gap between human and machine translation," *arXiv preprint arXiv:1609.08144*, 2016.
- [28] Google Research, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," https://github.com/google-research/ bert, 2018, accessed: 2024-03-14.
- [29] M. Mäntylä, Y. Wang, and J. Nyyssölä, "Loglead-fast and integrated log loader, enhancer, and anomaly detector," *arXiv preprint* arXiv:2311.11809, 2023.
- [30] S. Hashemi and M. Mäntylä, "Onelog: Towards end-to-end training in software log anomaly detection," arXiv preprint arXiv:2104.07324, 2021.
- [31] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [32] FudanSELab, "Deeptralog," https://github.com/FudanSELab/ DeepTraLog, 2024.
- [33] IntelligentDDS, "Nezha," https://github.com/IntelligentDDS/Nezha, 2024.

- [34] G. Yu, P. Chen, Y. Li, H. Chen, X. Li, and Z. Zheng, "Nezha: Interpretable fine-grained root causes analysis for microservices on multi-modal observability data," in *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2023, pp. 553–565.
- [35] H.-J. Ye and W.-L. Chao, "How to train your maml to excel in few-shot classification," in *International Conference on Learning Representations*, 2021.
- [36] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," Advances in neural information processing systems, vol. 30, 2017.
- [37] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra *et al.*, "Matching networks for one shot learning," 2016.
- [38] D. Jung, D. Kang, S. Kwak, and M. Cho, "Few-shot metric learning: Online adaptation of embedding for retrieval," in *Proceedings of the Asian Conference on Computer Vision*, 2022, pp. 1875–1891.
- [39] M. A. Jamal and G.-J. Qi, "Task agnostic meta-learning for few-shot learning," in *Proceedings of the IEEE/CVF conference on computer* vision and pattern recognition, 2019, pp. 11719–11727.
- [40] Z. Yu, L. Chen, Z. Cheng, and J. Luo, "Transmatch: A transfer-learning scheme for semi-supervised few-shot learning," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 12 856–12 864.