

LogLM: From Task-based to Instruction-based Automated Log Analysis

Yilun Liu¹, Yuhe Ji¹, Shimin Tao¹, Minggui He¹, Weibin Meng¹, Shenglin Zhang²
Yongqian Sun², Yuming Xie¹, Boxing Chen³, Hao Yang¹

¹Huawei, China

²Nankai University, China

³Huawei Canada, Canada

{liuyilun3, jiyuhe1, taoshimin, heminggui, mengweibin3, yuming.xie, boxing.chen, yanghao30}@huawei.com, {zhangsl, sunyongqian}@nankai.edu.cn

Abstract—Automatic log analysis is essential for the efficient Operation and Maintenance (O&M) of software systems, providing critical insights into system behaviors. However, existing approaches mostly treat log analysis as training a model to perform an isolated task (e.g., anomaly detection, log parsing, etc.) using task-specific log-label pairs. These task-based approaches are inflexible in generalizing to complex scenarios, depend on task-specific training data, and cost significantly when deploying multiple models. In this paper, we propose an instruction-based training approach that transforms log-label pairs from multiple tasks and domains into a unified format of instruction-response pairs. Our trained model, LogLM, can follow complex user instructions and generalize better across different tasks, thereby increasing flexibility and reducing the dependence on task-specific training data. By integrating major log analysis tasks into a single model, our approach also relieves model deployment burden. Experimentally, LogLM outperforms existing approaches across five log analysis capabilities, and exhibits strong generalization abilities on complex instructions and unseen tasks.

Index Terms—log analysis, instruction tuning, large language model, instruction following, multi-task learning

I. INTRODUCTION

In the ever-evolving landscape of software systems, log analysis has become a critical component of system Operation and Maintenance (O&M). Logs serve as an invaluable source of information, capturing events, errors, and other significant activities that can shed light on system performance and potential failures. However, as systems grow in complexity, manually analyzing logs becomes inefficient and requires advanced expertise. Thus, automated log analysis has emerged as a solution, using algorithms and machine learning models to extract meaningful insights from log data. This automation enhances efficiency, enabling timely detection of anomalies, identifying potential patterns, and aiding O&M engineers in taking actions, thereby helping maintain system reliability [1].

Traditionally, the research field of log analysis has been divided into discrete tasks, each with its own methodologies and benchmarking datasets. Two of the most widely studied tasks are log parsing and log-based anomaly detection. Log parsing involves transforming raw, unstructured log data into a structured format, with event templates and key variables extracted from log messages [2]. Anomaly detection, on the other hand, focuses on identifying deviations from normal

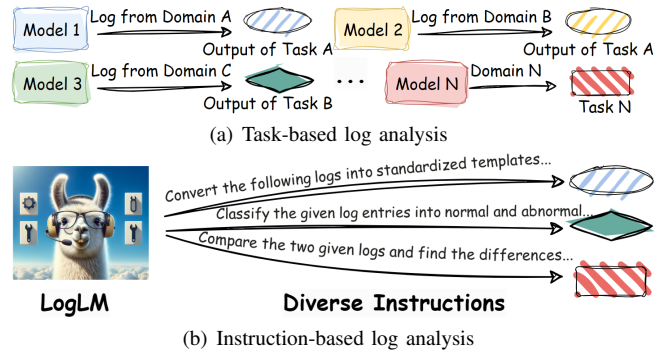


Fig. 1. Illustrated comparison between (a) existing task-based log analysis approaches and (b) our LogLM, an instruction-based log analysis model.

log behavior, which may signal underlying issues within the system [3]. Beyond these, other studied tasks include log interpretation [4], which aims to make logs more understandable to human, root cause analysis [5], which seeks to reveal the causes of system failures from logs, and solution recommendation [6], which suggests corrective actions for the identified problems in logs.

Recently, with the rapid advancements in AI technologies, particularly in large language model (LLM), a variety of LLM-based approaches for log analysis have been introduced [7]–[9]. LLMs such as GPT-4 [10] and Claude-3.5 [11] are capable of following complex instructions and performing a wide range of tasks. However, as shown in Fig. 1(a), most of these approaches continue to treat log analysis as isolated tasks, training LLMs to handle standalone tasks such as log parsing or anomaly detection without considering the interrelations between these tasks. We observed that this task-based training approach fails to fully capitalize on the capabilities of LLMs (e.g., a performance drop of 8.70% in log parsing and 69.47% in anomaly detection in Fig. 5). For LLMs nowadays, the training paradigm is instruction-based (called instruction tuning [12]), involving training on natural instructions across multiple tasks and domains, thereby offering greater flexibility and broader application. Furthermore, current task-based methods encounter several limitations when applied to real-world industrial scenarios:

(1) Inflexibility in generalizing to unseen tasks and complex user instructions. Task-based log analysis models are typically trained to output predictions on a fixed input format, which leads to an inflexible human-machine interactivity style in responding to real-world user instructions, particularly for dynamic O&M environments where new types of tasks often emerge. For example, in real-world cases, instead of parsing a log into a template, an O&M engineer may want the model to compare two logs and find the differences [13]. In addition to reporting anomaly in logs, the engineer may instruct the model to further output a justification along with its conclusion [4]. Existing methods are designed to handle a fixed task (*e.g.*, input a log and output “abnormal” or “normal”) and struggle to adapt to such flexible instructions or unfamiliar tasks, leading to a limited user experience.

(2) Dependency on task-specific and domain-specific training data caused by ignoring cross-task (or domain) connections. As observed by recent studies [14]–[16], existing approaches are dependent on task-specific training data. Their performances can drop significantly in online situation, where available task-specific logs in the specific domain (*e.g.*, modules, versions and devices) is scarce due to frequent software update and third-party module import [14]. The dependency on task-specific or domain-specific training data is due to the ignorance of inherent connections between different analysis tasks and log domains. For example, accurately parsing a log’s structure can significantly improve anomaly detection [17]; a correct diagnosis of the root cause of a failure is crucial for generating a meaningful solution [6]. Also, logs from different domains are heterogeneous in format and structure, but may share some common patterns [15]. Just as human experts draw from prior knowledge across different domains, LLMs may also apply insights from solving one log analysis task to another and enriches analysis abilities by integrating log patterns from diverse domains, thereby requiring less in-domain and task-specific training data.

(3) Significant burden in deployment. In large-scale software systems, the sheer volume of O&M tasks and diversity of logs from different domains demand that engineers deploy multiple specialized models, each fine-tuned for a particular task or domain. This leads to increased complexity in managing these models, as hundreds of specialized models may need to be maintained (*e.g.*, five tasks and nine domains leads to 45 fine-tuned models to deploy), which is also noted as the “pipeline challenge” [18]. Such an approach not only consumes significant computational resources but also burdens engineers with managing and updating a vast array of models.

In this paper, we propose LogLM, a novel instruction-based approach to log analysis that addresses these limitations. (Strictly, this is moving from “task-based” models to an “instruction-tuned unified-task” model.) As shown in Fig. 1(b), our approach transforms task-specific log-label pairs into instruction-response pairs, allowing the model to learn to properly respond to diverse user instructions. Instead of focusing on performing isolated tasks, our goal is to build a model that has comprehensive log analysis abilities. Thus, LogLM

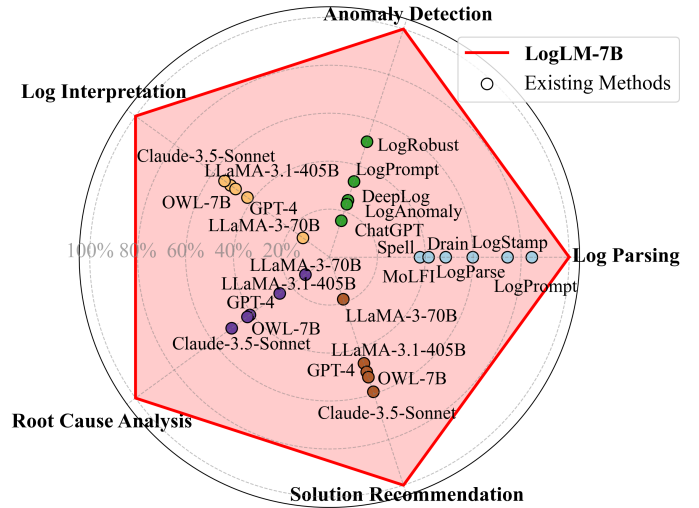


Fig. 2. Comparison of average performance between LogLM and existing methods across five log analysis capabilities. The dots indicate the relative percentage of baselines’ average performances in comparison to LogLM’s. LogLM-7B is fine-tuned from an open-source LLM with 7B parameters.

is trained on diverse instruction pairs, covering five log analysis capabilities and spanning nine log domains. Compared with existing methods, LogLM offers several key advantages: **(1) Instruction-following capability enables flexibility.** By training on natural instructions, LogLM can handle complex user instructions (Fig. 4), providing greater flexibility. **(2) Deep understanding of cross-task (or domain) connections relieves dependency.** Our training set breaks the barriers among different analysis tasks and log domains, enabling LogLM possessing deeper understanding of connections between tasks and domains. When in-domain (or task) training data is scarce, LogLM learns from out-of-domain (or task) data, as indicated by the increased performance in Fig. 5 after introducing instructions from other tasks and domains in the training set. **(3) Single-model deployment reduces burden.** LogLM’s multi-task ability allows for deploying only a single model for complex systems. Our contributions are:

- We present a new approach to log analysis that uses instruction tuning to combine cross-task and cross-domain log data into a standard instruction-response format, allowing for more effective log analysis capabilities. In Fig. 2, our model aces across all five analysis abilities.
- We develop a highly generalized instruction-following model capable of handling complex user instructions and performing previously unseen tasks (achieving best in 10 out 12 test terms on unseen tasks as shown in Table V).
- We open-source the full training dataset and codes of LogLM, facilitating the development of future models.¹

II. RELATED WORK

A. LLMs & Instruction Tuning

Recent LLMs, such as GPT-4 [10], have demonstrated the capability to execute complex tasks and generate appropriate

¹Dataset and code available at <https://github.com/lunyliu/LogLM>

responses to human instructions [19]–[21]. The development of these capabilities follows a two-stage process. The first stage, known as pre-training, involves training a foundation model to predict subsequent words in large-scale corpora [22]. Despite foundation models like LLaMA [23] being capable of sentence completion, they exhibit limitations in effectively responding to human instructions. To address this deficiency, LLMs undergo a second phase termed human alignment. A common method in this phase is instruction tuning [24], where the foundation model is fine-tuned using a diverse set of human instructions involving various tasks, paired with corresponding desired responses. This process aims to align the model’s outputs with human expectations and enhance its ability to generalize to unseen instructions [24]–[26]. It leverages instruction datasets composed of structured instruction pairs. Each pair consists of an INSTRUCTION, representing a human input, and a RESPONSE, representing the ideal output that resolves the instruction’s task. Through training on these pairs, the model learns to apply its pre-trained knowledge to generate appropriate responses to varied instructions and can generalize to new instructions [12].

Our work, LogLM, can be seen as a pioneering exploration on leveraging the technique of instruction tuning to solve industrial challenges in the field of log analysis. We also noticed a similar work using LLMs to analyze other semi-structured data similar to log data [27], where diverse instructions are synthesized for multiple table-based tasks.

B. Log Analysis

1) *Task of Log Parsing*: Log parsing is a widely studied task designed to reduce the volume of extensive log data by extracting meaningful templates from raw logs, facilitating further analysis tasks such as anomaly detection. Early approaches to log parsing primarily focused on coarse-level parsing, which involves aggregating lexically similar logs into templates. A log template only retains static parts in a log and replaces dynamic variables with a symbol of $\langle * \rangle$. Examples of these methods include cluster-based [2], [28], heuristic [29], [30], and tree-based methods [31]–[33].

Beyond coarse-level methods, there is a growing body of research dedicated to fine-level log parsing, which emphasizes the significance of variables within logs [16], [34]–[36].

Additionally, recent methods have begun to directly leverage LLMs for log parsing, such as LogPPT [37] and LogPrompt [14], while other methods utilize LLMs as an enhancement to existing models [7], [9].

2) *Task of Log-based Anomaly Detection*: Anomaly detection involves distinguishing anomalies within input logs. Following Le *et al.* [3], anomaly is defined as a pattern that violates normal sessions of systems. Techniques for anomaly detection can be categorized into two main classes: session-level methods and template-level methods. Both approaches typically begin by parsing raw logs into templates to facilitate analysis. Session-level methods subsequently aggregate log templates into sessions using a windowing strategy, and treat these log sessions as the minimum unit for anomaly detection.

In contrast, template-level methods directly identify anomalies based on the semantics of individual log templates. If any template within a log session is predicted to be abnormal, template-level methods classify the entire session as anomaly.

Within session-level methods, further distinctions can be made between forecast-based and classification-based approaches. Classification-based methods need both normal and abnormal training logs for training [38]–[40], whereas forecast-based methods require only normal historical logs [41], [42].

Template-level methods include techniques such as LogPrompt [14], which designs chain-of-thought prompts to detect anomalies in log entries using ChatGPT, and RAGLog [8], which employs retrieval-augmented generation to assist anomaly detection using LLMs. Additionally, Cui *et al.* [43] proposed a benchmark for template-level anomaly detection and evaluates various LLMs on this task.

3) *Other Analysis Tasks*: In addition to the two widely investigated tasks, several other tasks in log analysis are also studied. Meng *et al.* [44] proposed the task of log summary, which involves summarizing key phrases from log sequences. Liu *et al.* [4] proposed log interpretation, a task that aims to explain key elements in logs and describe the logs in natural language. Chen *et al.* [5] implemented an on-call system empowered by LLMs that predicts cloud incidents’ root cause category. Ahmed *et al.* [6] utilized fine-tuned LLMs to handle cloud incidents and recommend mitigation steps for engineers.

4) *LogLM v.s. Existing Approaches*: At earlier times there were endeavors which attempted to unify multiple log analysis tasks, such as Unilog [45] and Biglog [15], where a universal log encoder was trained for performing multiple downstream tasks. However, they still require fine-tuning on task-specific and domain-specific log-label pairs to update the parameters in classification layers of the model, leading to multiple different models for varied downstream tasks in the deployment. Also, they can’t follow complex user instructions or perform unseen tasks. LogPrompt [14] is another method supporting both log parsing and anomaly detection. But it requires meticulous prompt engineering for each task. And the underlying LLM is API-based ChatGPT, which limits the application in industry.

Compared with existing approaches, LogLM unifies multiple log tasks and is trained from open-source LLMs, thereby distinguishing by its potentials of single-model deployment for multi-task analysis in industrial applications. Moreover, LogLM can follow complex user instructions and perform analysis tasks that are unseen from training set, increasing the flexibility and intelligence of automated software O&M.

III. METHODOLOGY

Fig. 3 provides an overview of our methodology to build LogLM. We first define five closely interconnected log analysis capabilities for LogLM, which facilitate a profound comprehension of log data and a robust application of log analysis. To construct a high-quality instruction dataset, we identified publicly available data sources comprising real-world logs and user posts from technical communities. We designed

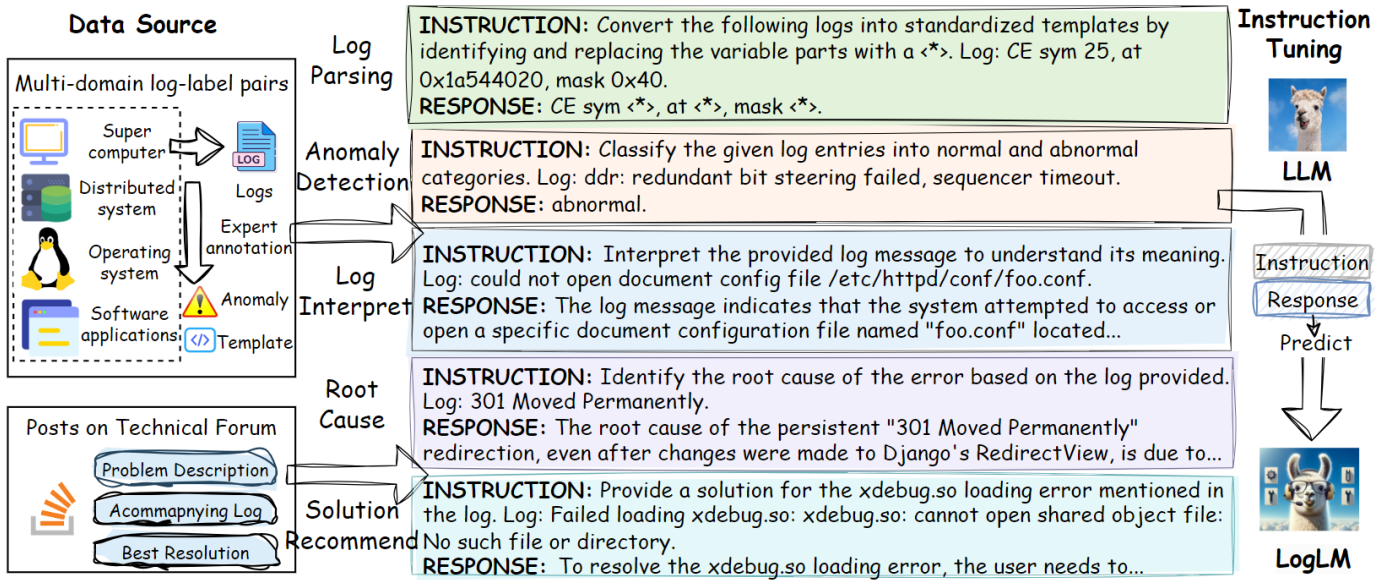


Fig. 3. Illustration on the capabilities composition, training dataset construction and training of LogLM.

an LLM-assisted workflow to transform log-label pairs and log-problem-resolution triplets from these data sources into the format of (INSTRUCTION, RESPONSE), covering all five analysis capabilities. The instruction pairs are then utilized to fine-tune LogLM from foundation LLMs.

A. Analysis Capabilities Composition

To facilitate a comprehensive log-related analytical capacity for LogLM, it is imperative to meticulously select the types of instructions utilized during the instruction tuning process. Based on existing studies in the field of log analysis, five most representative log analysis capabilities are considered in composition of the instruction dataset of LogLM. These five capabilities are not discrete entities; rather, they are closely interconnected, thereby fostering a deeper understanding of log-related analytical tasks by the model. Specifically, these capabilities can be categorized into two distinct groups.

(1) **Log Understanding Capabilities**, which consists of structural understanding (*Log Parsing*) and semantic understanding (*Log Interpretation*). *Log Parsing* involves decomposing a log’s structure into its template and variable components, which aids in comprehending the overall structure of logs. Conversely, *Log Interpretation* expands the semantic content of logs into coherent natural language sentences, thereby enhancing the understanding of log semantics. A precise decomposition of log structure and an accurate interpretation of log semantics is the foundation of subsequent applications.

(2) **Log Application Capabilities**, including three log-related application tasks: *Anomaly Detection*, *Root Cause Analysis* and *Solution Recommendation*. These three capabilities follow a logical sequence that reflects the operational procedures typically employed by O&M engineers during log analysis. When system anomaly is detected from logs, engineers need to find the cause of the breakdown and take

measures to solve the problem. Thus, the integration of these three application capabilities, along with the two understanding capabilities, addresses the practical requirements encountered in real-world scenarios, thereby enhancing the overall performance of log analysis.

In real-world practice, the five capabilities are indispensable for an engineer to perform log analysis. For example, Tom, a web O&M engineer, spotted a server log: “2024-12-25 08:32:00, ERROR 500 /checkout - Database Timeout”. By filtering out irrelevant variables (through *Log Parsing*), Tom focused on “ERROR 500” and “Database Timeout”. Based on his expertise (through *Log Interpretation*), this log was interpreted by Tom as a failed database connection, which suggested a potential system anomaly (*Anomaly Detection*). By examining server resources, Tom found the root cause of this incident: the recent database update had introduced a change in query execution plans, which overloaded the server during peak traffic (*Root Cause Analysis*). Tom then recommended several solutions in a group meeting, including rolling back the update, scaling up the database server and setting up proactive alerts (*Solution Recommendation*). **By including these capabilities into the training set, the LLM can comprehensively learn essential skills for log analysis and maximally aligns with real-world human engineers.**

B. Instruction Dataset Construction

Upon determining the capabilities composition, we seek to construct pairs of (INSTRUCTION, RESPONSE) within these capabilities. Log parsing and anomaly detection are two well-established tasks, thereby the instruction pairs of the two capabilities were constructed from log-label pairs in representative open-source datasets. For the other three capabilities, it is challenging to find off-the-shelf open-source datasets. Thus, we designed an LLM-assisted workflow to decompose log-

TABLE I
STATISTICS OF THE INSTRUCTION DATASET FOR TRAINING LOGLM

Capabilities	Data Source	# Instruction Pairs	Domain
Log Parsing	Loghub2k [46]	200	HDFS
		200	Hadoop
		200	Zookeeper
		200	BGL
		200	HPC
		200	Linux
		200	Proxifier
Anomaly Detection	LogPrompt [14]	194	BGL
		138	Spirit
Log Interpretation	LogExpert [47]	300	Apache
Root Cause Analysis		300	
Solution Recommendation		300	
All	-	2632	-

related high-quality community posts on technical forums, into instruction pairs of the three capabilities. The statistics of the constructed instruction dataset for training is shown in Table I, and the specific construction process is discussed below.

1) *Log Parsing*: As shown in Table I, the source data utilized for constructing instruction pairs related to *Log Parsing* is the Loghub2k dataset [46], which comprises real-world logs derived from a variety of domains. Zhu *et al.* [2] manually extracted log templates from a subset of 2000 logs across multiple domains, thereby establishing a benchmark that is considered most representative for the log parsing task. Consistent with prior research [15], [48], [49], log-template pairs from seven prominent domains—including supercomputers, distributed systems, operating systems, and software applications—were selected to formulate instruction pairs.

As discussed in Section I, our objective is to mitigate reliance on in-domain and task-specific log data. Thus, we incorporated only a limited selection of log-template pairs into the instruction dataset, reserving the majority for testing purposes. This approach effectively simulates real-world online scenarios where in-domain log data for training is often scarce [16]. Specifically, we adhered to the setting in LogPrompt [14], wherein the chronologically first 10% of logs were selected from each domain to construct the training set and the remaining 90% were for evaluation. This yielded a total of 1400 log-template pairs across the seven domains.

In constructing instruction pairs for *Log Parsing*, as shown in Fig. 3, we employed a straightforward prompt proposed by Liu *et al.* [14] as INSTRUCTION. This prompt has demonstrated consistent parsing performance when applied to ChatGPT, thereby qualifying it as a high-quality INSTRUCTION. For the RESPONSE, the manually extracted templates from the dataset provide a human-preferred answer to the INSTRUCTION, thereby enhancing human alignment in subsequent instruction tuning processes.

2) *Anomaly Detection*: Instruction pairs related to *Anomaly Detection* were derived from the BGL and Spirit benchmark datasets, which were meticulously curated by Oliner *et al.*

[50]. These datasets comprise system logs collected from supercomputers, with annotations performed in collaboration with domain experts to identify anomalous events for the anomaly detection task. Liu *et al.* [14] extracted a total of 1766 log templates from the BGL dataset and 1297 from the Spirit dataset, and released these pairs of log templates and anomaly labels. From this collection, we randomly selected a minor subset—approximately 10%—from each dataset to construct the instruction pairs, while reserving the remaining data for evaluation purposes. To ensure a balanced representation of abnormal and normal samples, we ensured that each subset contained around 10% abnormal samples, thereby reflecting the original distribution of abnormal and normal log messages in the source data. For the construction of the instruction pairs, as shown in Fig. 3, we employed a straightforward prompt for the anomaly detection task proposed by Liu *et al.* [14], with a single log as an input and the normal/abnormal label serving directly as the RESPONSE. For new logs, multiple entries can be concatenated and input as a single log.

3) *Interpretation, Root Cause & Solution (IRS)*: For IRS capabilities, the availability of off-the-shelf datasets that encompass logs and their corresponding task outputs is limited. However, community contributions on technical forums may offer valuable data sources useful for the development of IRS capabilities for LogLM. Specifically, when users post technical issues accompanied by error logs generated by software, their descriptions often contain the direct causes of the errors present in the logs, as well as the contextual circumstances surrounding those logs, which help in interpreting the logs. Furthermore, the highest-voted recommended resolutions typically include interpretations of the logs, identifications of the root causes of the errors, and user-preferred solutions.

To achieve this, we utilize a dataset released by Wang *et al.* [47], which consists of log-related user posts from Stack Overflow tagged with "apache". This dataset was meticulously curated to include only those user posts that contained at least one accompanying log. Additionally, the authors manually verified the feasibility of the highest-voted resolutions for each post, resulting in a high-quality dataset comprising 384 log-problem-resolution triplets. Subsequently, we composed a prompt to instruct GPT-4 to decompose three instruction pairs, regarding to *Log Interpretation*, *Root Cause Analysis*, and *Solution Recommendation*, from each log-problem-resolution triplet within the dataset. The specific prompt utilized is:

Assume you are a dedicated IT specialist focusing on log analysis for system O&M. The following real-world case includes a log, a title, a description of related user posted problem and a community solution. Your task is to decompose three INSTRUCTION-INPUT-RESPONSE pairs based on the real-world case. Requirement for INSTRUCTION: The topic of three instructions are interpretation of the log, root cause of the log and solution of the log, respectively. Organize each instruction to be a concise user query on the log. Requirement for INPUT:

The three inputs are all the given log in the real-world case. If the original log is quite long, retain only necessary part. Requirement for RESPONSE: The three responses must properly meet the instructions. Organize your language to be precise and professional. Format your answer like this: INSTRUCTION 1: xxx\n INPUT 1: xxx\n RESPONSE 1: xxx\n INSTRUCTION 2: xxx\n INPUT 2: xxx\n RESPONSE 2: xxx\n INSTRUCTION 3: xxx\n INPUT 3: xxx\n RESPONSE 3: xxx. The case begins: **{Input}**

By applying this prompt to GPT-4 (model version: *gpt-4-turbo-preview*), we generated 384 instruction pairs corresponding to each capability, with INSTRUCTION reflecting the specific requests merged with the input log, and RESPONSE being derived from human-authored posts. To mitigate the risk of generating inaccurate or irrelevant content, we implemented a human calibration process to ensure that the generated instruction pairs were consistent with the original user posts and expert resolutions. Approximately 2% of the generated content was excluded following this inspection, resulting in 376 calibrated instruction pairs for each of the three capabilities. These instruction pairs were subsequently divided into training and testing sets in an 8:2 ratio (*i.e.*, 300:76).

C. Training of LogLM

The instruction dataset in Table I serves as a paradigm for properly responding to human instructions in log analysis, which the LLM learns to emulate during instruction tuning. Starting with a pre-trained foundation model, the LLM gradually aligns its pre-existing knowledge with human-preferred analysis outputs by predicting the next tokens in the RESPONSE conditioned on an input INSTRUCTION.

Denote the instruction pair by $x = (\text{INSTRUCTION}, \text{RESPONSE})$, where $x \in C$ represents elements in the constructed instruction dataset C . The parameters of the foundation model are denoted as θ , and after instruction tuning, the parameters are updated to θ_c , which represents the adapted model, LogLM. The goal is to maximize the log likelihood of next tokens (*i.e.*, minimal unit of a tokenized sentence) in the RESPONSE conditioned on the corresponding INSTRUCTION. By denoting the tokens in the RESPONSE of x as $(y_1, y_2, \dots, y_{N_x})$, where N_x is the number of tokens in the RESPONSE, the training objective can be formulated as:

$$\theta_c = \arg \max_{\theta} \sum_{x \in C} \sum_{i=1}^{N_x} \log P(y_i | \text{INSTRUCTION}, y_{1:i-1}; \theta, x). \quad (1)$$

IV. EXPERIMENT

A. Implementation Details

In the experiment, the foundation model θ is LLaMA-2-7B [51], an open-source LLM with 7B parameters. Our main implementation of LogLM, denoted by LogLM-7B, was trained on the full instruction dataset in Table I according to Eq. (1). During instruction tuning, the learning rate is

2×10^{-5} , the batch size is 32 and the number of training epochs is six. Other parameters follow the default settings in LLaMAFactory [52], the framework we utilized for LLM training and inferencing.

B. Research Questions & Key Findings

In this section, we describe the research questions (RQ) on evaluating LogLM and the key findings in experiments.

RQ1: Can the instruction-based approach outperform task-based methods in the five log analysis capabilities?

Key Findings of RQ1: To assess LogLM’s log analysis capabilities, in Section IV-C, we compare its performance with existing task-based methods across five benchmarks representing different log analysis capabilities. As a single model, LogLM-7B consistently outperforms 20 existing approaches across different analysis tasks, highlighting its strong industrial application potential due to its cost-effectiveness.

RQ2: (A) What contributes to LogLM’s superior performance? (B) Can LogLM reduce reliance on in-domain and task-specific log data? (C) Does the model benefit from increased diversity and quantity of training instructions?

Key Findings of RQ2: Section IV-D addresses the three sub-questions through an ablation study on the training data composition of LogLM. Beginning with merely in-domain training data for a specific task, we observe a substantial improvement in LogLM’s average performance as instruction pairs from other domains and tasks are introduced. This demonstrates that LogLM’s advantage stems from learning connections between instruction pairs across diverse domains and tasks (RQ2-A). Furthermore, this supports answering RQ2-B, as the majority of LogLM’s training data come from other domains and tasks, reducing the need for extensive data collection in a private domain or for an emerging task in industrial applications. Regarding RQ2-C, the finding suggests that training on instruction pairs from a diverse set of log domains and analysis tasks enhances LogLM’s performance, highlighting its application potential in complex systems.

RQ3: Can LogLM generalize to unseen tasks and follow complex user instructions?

Key Findings of RQ3: In Section IV-E1, we train LogLM using instruction pairs from only four capabilities, leaving the remaining one as an unseen task. Despite this, LogLM achieved best performance on 10 of 12 test items in these unseen tasks. Additionally, a case study in Section IV-E2 demonstrate LogLM-7B’s ability to answer knowledge-based questions related to log analysis, handle previously unseen instructions, and perform a combination of tasks.

The rest of Section IV is organized as follows. Section IV-C addresses RQ1, evaluating *Log Parsing* in Section IV-C1, *Anomaly Detection* in Section IV-C2 and the other three capabilities in Section IV-C3. Section IV-D and IV-E address RQ2 and RQ3, respectively.

C. RQ1: Benchmarking on Log Analysis Capabilities

1) Log Parsing:

TABLE II
BENCHMARKING ON THE CAPABILITY OF LOG PARSING EVALUATED IN BOTH COARSE-LEVEL AND FINE-LEVEL

Methods	HDFS		Hadoop		Zookeeper		BGL		HPC		Linux		Proxifier		Avg.	
	RI ^a	F1	RI	F1	RI	F1	RI	F1	RI	F1	RI	F1	RI	F1	RI	F1
IPLoM [30]	0.914	0.389	0.636	0.068	0.787	0.225	0.858	0.391	0.228	0.002	0.695	0.225	0.822	0.500	0.706	0.257
LKE [28]	0.861	0.424	0.150	0.198	0.787	0.225	0.848	0.379	0.119	0.381	0.825	0.388	0.379	0.309	0.567	0.329
LogSig [53]	0.872	0.344	0.651	0.050	0.787	0.225	0.806	0.333	0.119	0.002	0.715	0.146	0.559	0.339	0.644	0.206
FT-tree [31]	0.908	0.385	0.668	0.046	0.773	0.186	0.275	0.497	0.119	0.002	0.709	0.211	0.722	0.420	0.596	0.250
Spell [29]	0.871	0.000	0.721	0.058	0.102	0.045	0.503	0.536	0.882	0.000	0.706	0.091	0.621	0.000	0.629	0.104
Drain [32]	0.914	0.389	0.647	0.068	0.787	0.225	0.822	0.397	0.119	0.002	0.695	0.225	0.822	0.500	0.687	0.258
MoLFI [54]	0.871	0.000	0.699	0.095	0.899	0.000	0.792	0.333	0.881	0.000	0.410	0.026	0.621	0.000	0.739	0.065
LogParse [34]	0.907	0.632	0.349	0.502	0.982	0.348	0.992	0.665	0.194	0.330	0.825	0.588	0.490	0.334	0.677	0.486
LogStamp [16]	0.954	0.523	0.927	0.594	0.992	0.275	0.984	0.818	0.949	0.434	0.760	0.658	0.811	0.438	0.911	0.534
LogPrompt [4]	0.890	0.863	0.879	0.763	0.948	0.889	0.964	0.865	0.934	0.759	0.758	0.766	0.567	0.653	0.849	0.794
LogLM-7B	1.000	0.998	1.000	0.973	1.000	0.995	0.999	0.977	0.999	0.935	0.994	0.934	0.879	0.940	0.982	0.965

^a **RI** stands for coarse-level RandIndex. **F1** stands for fine-level F1-score.

a) *Evaluation Setting*: As is discussed in Section III-B1, this benchmark involves testing the performance of log parsing on the chronologically last 90% of the logs from seven domains in the Loghub2k [46] dataset. We compare LogLM with 10 existing approaches specifically designed for the task of log parsing, encompassing cluster-based methods [28], [53], heuristic methods [29], [30], [54], tree-based methods [31], [32], machine learning methods [34] and LLM-based methods [4], [16]. Following the experiment setting in Liu *et al.* [14], each baseline is trained on the first 10% of the logs in each domain, except for LogPrompt [4], which directly utilize ChatGPT for log parsing without training. In other words, for each domain, LogLM and other approaches only see the first 10% of logs in the specific domain during training and the rest 90% of logs are unseen, thereby simulating online situations while ensuring a fair comparison.

Following Liu *et al.* [14], two most popular evaluation metrics for the task of log parsing among existing studies [15], [16], [34] are utilized, encompassing both coarse-level (RandIndex [55]) and fine-level (F1-score). RandIndex assesses the accuracy of log clustering (*i.e.*, whether two logs with the same template are accurately clustered together), regardless of the correctness of variables in the extracted templates. In contrast, F1-score measures the precise identification of variable parts in logs, thereby serving as a fine-level metric. To calculate the F1-score, the predicted log template is tokenized into a list of tokens. Then, for each token, count the terms TP , TN , FP , and FN . If the token is truly a variable and is correctly predicted as such (or not), increment the TP (or FP) by one. If the token is not a variable and is predicted not to be a variable (or be), increment the TN (or FN) by one. The F1-score is computed as the harmonic average of Recall ($Recall = \frac{TP}{TP+FN}$) and Precision ($Precision = \frac{TP}{TP+FP}$).

b) *Result*: As shown in Table II, our model, LogLM-7B, exhibited remarkable performances on this benchmark, outperforming existing approaches significantly in both coarse-level evaluation and fine-level evaluation. Averagely, LogLM-7B outperforms best baselines by 7.79% in RandIndex and by 10.03% in F1-score. This advantages indicate that LogLM

TABLE III
BENCHMARKING ON THE CAPABILITY OF ANOMALY DETECTION

Methods	BGL		Spirit	
	S-F1 ^a	T-F1	S-F1	T-F1
DeepLog [41]	0.194	-	0.092	-
LogAnomaly [42]	0.129	-	0.138	-
LogRobust [38]	0.536	-	0.045	-
ChatGPT [14]	0.129	0.067	0.122	0.050
LogPrompt [14]	0.314	0.233	0.144	0.071
LogLM-7B	0.811	0.625	0.584	0.278

^a **S-F1/T-F1** means F1-Score in session/template-level.

can accurately recognize variable parts in logs while extracting correct coarse-level templates, even when the in-domain training logs are limited in quantity.

2) *Anomaly Detection*:

a) *Evaluation Setting*: This evaluation involves comparing LogLM with both template-level methods [14] and session-level methods [38], [41], [42], as discussed in Section II-B2. Thus, the evaluation settings are split into template-level and session-level. For template-level, the test set is the split template-label pairs as discussed in Section III-B2, containing around 90% of the templates extracted by Liu *et al.* [14] from the dataset of BGL and Spirit. For session-level, following the setting in LogPrompt, log sessions were built using fixed-window grouping with a length of 100 chronologically adjacent logs in BGL and Spirit. The first 4000 logs in each dataset are used for training baselines and the rest logs are for testing. To avoid data leakage, logs in the training set of LogLM were excluded from the session-level test set, leading to a final test set of 40521 and 7515 sessions for BGL and Spirit.

The evaluation metric for both the template-level and session-level is the F1-score of anomaly. F1-score in template-level (or session-level) considers accurate recall of abnormal logs (or sessions) from test cases and the precise prediction of anomaly in template-level (or session-level). Similar to Section IV-C1a, TP represents the successful identification of an anomaly (similarly for TN , FP , and FN) and the F1-score

TABLE IV
BENCHMARKING WITH EXISTING LLMs ON LOG INTERPRETATION, ROOT CAUSE ANALYSIS, AND SOLUTION GENERATION

Methods	Log Interpretation				Root Cause Analysis				Solution Recommendation			
	BLEU	R-1 ^a	R-2	R-L	BLEU	R-1	R-2	R-L	BLEU	R-1	R-2	R-L
LLaMA-3-70B [56]	0.507	7.984	2.121	5.864	0.172	6.876	1.273	4.382	0.529	7.640	1.557	5.368
LLaMA-3.1-405B [56]	4.466	28.563	11.106	17.071	1.416	12.984	3.694	8.162	2.018	18.768	5.287	12.045
Claude-3.5-Sonnet [11]	4.324	30.512	11.786	18.307	2.563	25.824	7.909	15.221	2.698	23.618	6.936	15.109
GPT-4 [10]	2.831	25.205	9.147	13.683	2.057	21.810	6.544	11.518	1.502	21.964	5.532	12.274
OWL-7B [57]	2.566	28.211	8.289	19.123	1.947	20.893	5.671	14.718	0.953	21.620	5.006	15.574
LogLM-7B	15.584	46.488	23.087	34.769	12.398	40.602	19.042	30.227	8.241	34.415	13.911	25.431

^a **R-1** stands for ROUGE-1. **R-2** stands for ROUGE-2. **R-L** stands for ROUGE-L.

is calculated as the harmonic average of Recall and Precision.

b) Result: The result is shown in Table III. As is observed in the study of Liu *et al.* [14], detecting anomaly in online situations with limited number of training logs can be challenging for most methods, including even ChatGPT. Normally, without fitting on enough historical data, it is hard to precisely model the anomaly patterns in domain-specific logs. However, by training on diverse log tasks, LogLM possesses more knowledge in log analysis, thereby can model in-domain anomaly patterns more efficiently. As a result, LogLM-7B achieves a strong result in Table III both for BGL and Spirit, while utilizing only 100+ in-domain logs in training.

3) Interpretation, Root Cause & Solution (IRS) Tasks:

a) Evaluation Setting: The evaluation sets for IRS tasks encompass 76 Q&A pairs for each capability, constructed from user posts and highest-voted resolutions in real-world technical community as described in Section III-B3. For baselines, general-purpose LLMs are involved given their strong abilities of solving complex problems [19], [21], including LLaMA-3-70B [56], LLaMA-3.1-405B [56], Claude-3.5-Sonnet [11] and GPT-4 [10]. In addition, OWL-7B [57], a domain-purpose LLM designed for IT operations, is also included in the evaluation. All baselines are directly prompted with instructions in the evaluation sets (*i.e.*, zero-shot). For fairness, we examine the zero-shot (*i.e.*, no in-domain training data) performance of LogLM in Section IV-E1.

The evaluation employs four commonly-used metrics in Q&A tasks for measuring the accuracy of generated answers compared with reference answers. ROUGE-1 and ROUGE-2 [58] assess the overlap of unigrams and bigrams (*i.e.*, a sequence of one or two consecutive words), respectively, between the generated and reference texts, focusing on successful recalls of information in reference answers. ROUGE-L [58] measures the longest common subsequence (LCS) between the two texts, capturing sentence-level structural similarity. BLEU [59] evaluates precision of generated answers by counting percentage of n-gram (*i.e.*, n consecutive words) overlap in the generated texts, penalizing brevity and rewarding fluency.

Each of these metrics ranges from 0 to 100. A score of 100 indicates a perfect match between the generated and reference text, while a score of 0 indicates no overlap. ROUGE-1 and ROUGE-2 are computed as the ratio of overlapping n-grams in the generated text to number of n-grams in the reference text.

ROUGE-L is calculated as the length of LCS to the length of the reference text. BLEU, however, is calculated as the geometric mean of n-gram precision (*i.e.*, overlapping n-grams to number of n-grams in the predicted text) from $n = 1$ to $n = 4$, with a brevity penalty to penalize short outputs. These four metrics capture both surface-level lexical overlap and deeper structural similarities, which are essential for Q&A evaluation.

b) Result: As shown in Table IV, LogLM-7B outperforms existing LLMs in all three capabilities. The advantages in ROUGE-1 and ROUGE-2 indicate that LogLM’s responses recalled more key information in the reference responses, which is essential for providing correct root causes of logged errors and feasible solutions. Also, the higher scores in ROUGE-L indicate the sentence structure of LogLM’s responses are more similar to human-calibrated reference responses derived from community expert resolutions. This suggests that LogLM’s answers are more natural in style and layout, which is especially advantageous for generating human-preferred interpretations of logs. Moreover, a significant advantage of LogLM-7B in BLEU score is observed across the three capabilities, indicating that LogLM generated more precise and readable contents with less hallucination.

D. RQ2: Ablation Study on Training Data Composition

1) Evaluation Setting: In this section, we trained models from LLaMA-2-7B with the same training settings, except using four distinct groups of instruction pairs, and evaluated their performances on *Log Parsing* and *Anomaly Detection*. The four groups are as follows: (1) LogLM: This group utilizes the complete instruction dataset in Table I, resulting in the exact LogLM-7B model. (2) W/o IRS instructions: In this group, instructions related to *Log Interpretation*, *Root Cause Analysis*, and *Solution Recommendation* were excluded from the training process. (3) W/o IRS+Parsing instructions (or w/o IRS+Anomaly instructions): In addition to the IRS-related instructions, this group further excludes instructions involving *Log Parsing* (or *Anomaly Detection*) to assess the model’s ability to perform *Anomaly Detection* (or *Log Parsing*). In this case, only task-relevant instruction pairs remain in the training set. (4) W/o IRS+Parsing+Domain instructions (or w/o IRS+Anomaly+Domain instructions): This group involves the further exclusion of instruction pairs containing logs from

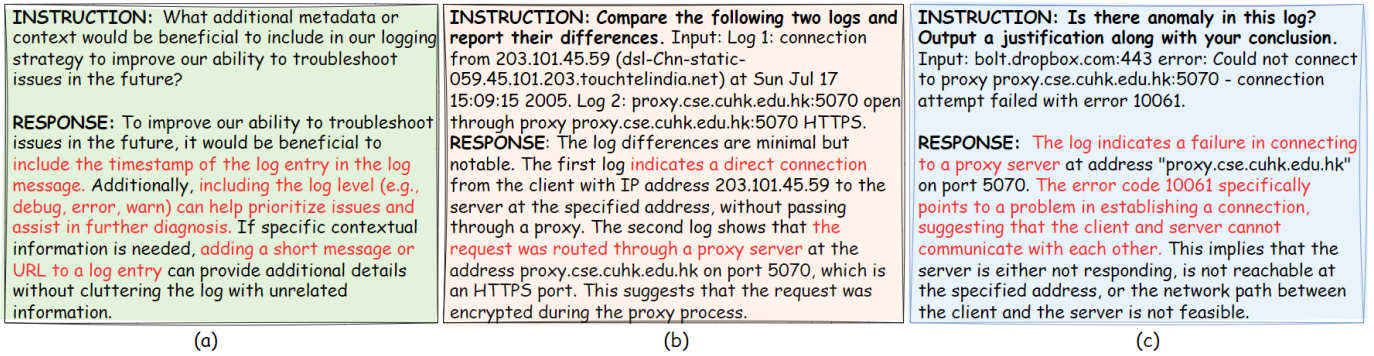
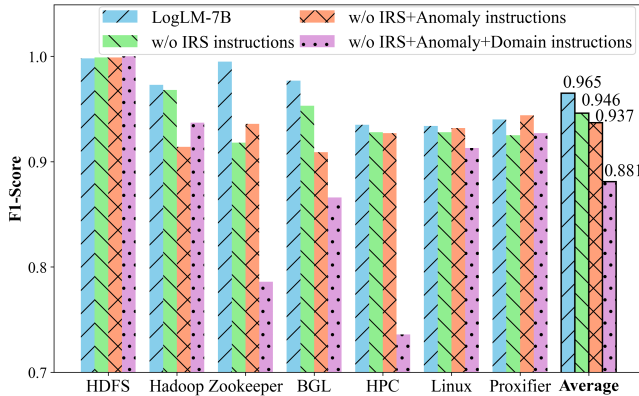
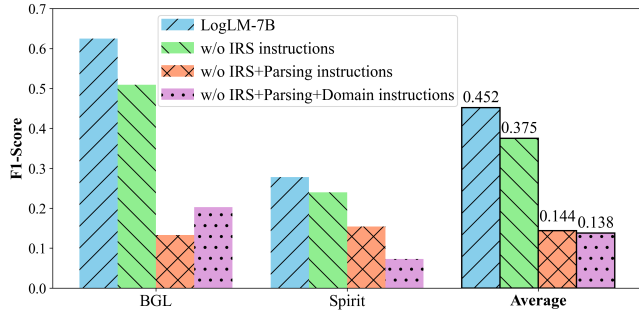


Fig. 4. Three cases of LogLM-7B responding to complex user instructions: (a) log-related Q&A, (b) unseen new task, and (c) combination of tasks.



(a) Log Parsing



(b) Anomaly Detection

Fig. 5. Ablation study on the training data of LogLM, evaluated on (a) Log Parsing and (b) Anomaly Detection. See additional results (other three tasks, and an upsampling group to control quantity) in our GitHub Page.

other domains, resulting in a training set composed solely of in-domain logs relevant to the specific task.

2) *Result:* As shown in Fig. 5, the average performance for both tasks consistently declines when instruction pairs from other tasks and domains are excluded. Since the aim of instruction tuning is to train the model to generate human-preferred responses to user instructions, incorporating diverse instruction pairs from various tasks and domains enhances the model’s alignment with human preferences. Rather than merely fitting on task labels from a specific domain, the model is trained to leverage its pre-existing knowledge and inference capabilities to address a wide range of log-related problems,

TABLE V
EVALUATION OF LOGLM ON TASKS UNSEEN FROM TRAINING

Unseen Task ^a	BLEU	ROUGE-1	ROUGE-2	ROUGE-L
Log	6.451	36.841	13.212	25.364
Interpretation	(v.s. 4.466) ^b	(v.s. 30.512)	(v.s. 11.786)	(v.s. 19.123)
Root Cause	9.253	36.089	14.747	26.890
Analysis	(v.s. 2.563)	(v.s. 25.824)	(v.s. 7.909)	(v.s. 15.221)
Solution	1.233	24.221	6.033	17.475
Recommendation	(v.s. 2.698)	(v.s. 23.618)	(v.s. 6.936)	(v.s. 15.574)

^a Unseen task X : testing on X after excluding X from training.

^b Compare LogLM with best baselines of the task from Table IV.

thereby developing a robust and comprehensive log analysis capability while reducing dependency on in-domain task data. This suggests LogLM’s potential to benefit from the continual emergence of new logs and tasks in real-world systems.

E. RQ3: Generalization on Unseen & Complex Instructions

1) *Performing Unseen Analysis Tasks:* Table V displays the performance of LogLM on unseen tasks, which were excluded from training set in advance (*i.e.*, four capabilities for training and the remaining one capability for testing). A universal decline on performance of corresponding capabilities can be observed when being treated as the unseen task. The capability with the sharpest performance drop is *Solution Recommendation*, possibly due to its challenging nature compared with other capabilities. However, despite not being trained on the excluded task, LogLM still outperforms existing LLMs (highest scores from Table IV achieved by baselines are displayed in the parentheses) in 10 out of the 12 test terms. This result not only indicates a strong generalization ability of LogLM on unseen instructions, but also suggests that LogLM acquired a comprehensive problem-solving ability related to log analysis through capturing cross-task connections.

2) *Following Complex User Instructions:* Fig. 4 shows three cases in which LogLM-7B responds to complex user instructions. In case (a), an open-ended question related to log analysis is posed, seeking advice on improving logging strategies. LogLM’s response includes several recommendations, such as incorporating timestamps, establishing log levels, and

adding concise messages, demonstrating that LogLM possesses a solid understanding of log analysis and can effectively apply log-related knowledge to generate appropriate responses. In case (b), LogLM is tasked with a novel problem: comparing two logs and identifying the differences. The model successfully identifies key distinctions, such as the use of proxies and encryption methods, indicating that LogLM’s analytical capabilities are not confined to predefined tasks or domains but can extend to emerging problems with logical accuracy. Case (c) involves a combination of two known capabilities from training: *Anomaly Detection* and *Log Interpretation*, requiring the model to detect anomalies while providing justification for its conclusions. This type of instruction is common in real-world scenarios where O&M engineers need more than just predicted values from a model—they require detailed explanations to inform further actions [4]. In response, LogLM effectively follows the instruction, delivering both a conclusion and a well-reasoned justification for the “10061 error”.

The flexibility exhibited by LogLM may be attributed to the systematic composition of instruction pairs in Table I, which span two log-understanding and three log-application capabilities. Consequently, through training on these diverse pairs of user instructions and human-preferred responses, LogLM has developed not only a deep comprehension of log semantics and structures but also a robust capacity to address practical log-related problems and generate well-formatted responses. This instruction-following ability highlights the model’s potential in dynamic and complex industrial environments.

V. DISCUSSION

A. LogLM in Practice

We further examine the feasibility of our approach in industrial applications, by deploying LogLM as the log analysis component within a software and network O&M platform of Huawei. The deployment process involved instruction tuning based on Huawei’s proprietary LLM, Pangu [60], using the proposed instruction dataset in Table I. To support retrieving real-time logs from local devices via private API, instruction pairs involving function calling [61] were also added to the instruction dataset as a new capability of LogLM. The fine-tuned model is denoted by LogLM-Huawei.

In practice, the platform operates by allowing users to issue instructions, which are then responded by LogLM-Huawei. The model first retrieves relevant logs from the specified devices by calling system APIs. Upon receiving the logs, LogLM-Huawei performs a thorough analysis and responds to users. The instructions from users span a variety of operational needs, such as resource checks and fault diagnosis.

For instance, a typical query might involve a client asking for an investigation into whether a specific user experienced issues over the past three days. LogLM-Huawei firstly requests an API calling to retrieve logs within past three days from the user’s associated devices. The platform will parse the request, collect the logs retrieved from the API and send them to the model. LogLM-Huawei then evaluates the devices’ statuses through examining on multiple logs, and provides analysis

results on possible issues the user experienced, along with recommendations on actions. Other frequent user instructions include inquiries about specific device statuses and system-generated logs, and related knowledge-based questions.

Since its deployment in the platform managing arrays of devices, LogLM has processed over 30k queries over six months, averaging 200+ queries per day. LogLM’s robust log analysis and adaptive instruction-following abilities make it the central intelligence of the O&M platform, driving real-world applications and delivering valuable insights to customers.

B. Threats to Validity

Our study has several limitations:

(1) *Hallucination in LLMs*: Despite the competitive performance of LogLM-7B in log analysis benchmarks, occasional hallucinations observed in generated outputs may affect accuracy. While hallucination is a known issue in LLMs and has garnered attention from AI researchers [62], mitigation strategies exist for industrial deployment, such as incorporating post-processing modules, optimizing prompts, or utilizing larger foundation models.

(2) *Fairness of Experimental Comparisons*: In Section IV-C, the training data for LogLM differs from that of baseline methods, raising concerns about the fairness of comparisons. However, due to the proprietary and pre-trained nature of many baselines (e.g., relying on ChatGPT), identical data usage is unfeasible. To ensure fairness, as discussed in Sections IV-C1a and IV-C2a, LogLM uses the same or fewer in-domain and task-specific log-label pairs than existing methods for *Log Parsing* and *Anomaly Detection*. Furthermore, for the three IRS capabilities, Table V shows that LogLM outperforms most baselines, even without task-specific training data.

(3) *Biases in Studied Logs*: Most studied logs in this paper are system logs, from middlewares, frameworks to operation systems. In practice, logs from software applications are also analyzed, which may leads to biases in LogLM. However, the availability of open-source datasets for application logs is limited. To address this, logs from the Apache server application are included in the study, mitigating this issue.

(4) *Limited Practice Verification*: In practice, LogLM has been verified only on Huawei’s O&M platform. Other organizations have different environments and requirements, which may introduce potential external threats to LogLM. Expanding its deployment across various organizations would enhance its generalizability and help mitigate external threats

VI. CONCLUSION

In this paper, we present LogLM, an instruction-based log analysis approach. Through our empirical study, three key advantages of LogLM over existing task-based methods are highlighted: (1) greater flexibility, (2) multi-task and generalization capability and (3) enabling single-model deployment. The released instruction dataset of LogLM can facilitate future endeavors in building LLMs for software O&M. Future work include adding more instruction sets, testing with larger foundation models and further deploying LogLM across various organizations to validate its generalizability.

REFERENCES

- [1] A. Oliner, A. Ganapathi, and W. Xu, "Advances and challenges in log analysis," *Communications of the ACM*, vol. 55, no. 2, pp. 55–61, 2012.
- [2] J. Zhu, S. He, J. Liu, P. He, Q. Xie, Z. Zheng, and M. R. Lyu, "Tools and benchmarks for automated log parsing," in *2019 IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*. IEEE, 2019, pp. 121–130.
- [3] V.-H. Le and H. Zhang, "Log-based anomaly detection with deep learning: How far are we?" in *2022 IEEE/ACM International Conference on Software Engineering (ICSE)*. IEEE, 2022, pp. 1356–1367.
- [4] Y. Liu, S. Tao, W. Meng, F. Yao, X. Zhao, and H. Yang, "Logprompt: Prompt engineering towards zero-shot and interpretable log analysis," in *Proceedings of the 2024 IEEE/ACM 46th International Conference on Software Engineering: Companion Proceedings*, 2024, pp. 364–365.
- [5] Y. Chen, H. Xie, M. Ma, Y. Kang, X. Gao, L. Shi, Y. Cao, X. Gao, H. Fan, M. Wen *et al.*, "Automatic root cause analysis via large language models for cloud incidents," in *Proceedings of the Nineteenth European Conference on Computer Systems*, 2024, pp. 674–688.
- [6] T. Ahmed, S. Ghosh, C. Bansal, T. Zimmermann, X. Zhang, and S. Rajmohan, "Recommending root-cause and mitigation steps for cloud incidents using large language models," in *ICSE 2023*, May 2023.
- [7] Z. Jiang, J. Liu, Z. Chen, Y. Li, J. Huang, Y. Huo, P. He, J. Gu, and M. R. Lyu, "Lilac: Log parsing using llms with adaptive parsing cache," *Proceedings of the ACM on Software Engineering*, vol. 1, no. FSE, pp. 137–160, 2024.
- [8] J. Pan, W. S. Liang, and Y. Yidi, "Raglog: Log anomaly detection using retrieval augmented generation," in *2024 IEEE World Forum on Public Safety Technology (WFPST)*. IEEE, 2024, pp. 169–174.
- [9] A. Zhong, D. Mo, G. Liu, J. Liu, Q. Lu, Q. Zhou, J. Wu, Q. Li, and Q. Wen, "Logparser-llm: Advancing efficient log parsing with large language models," in *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024, pp. 4559–4570.
- [10] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, "Gpt-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.
- [11] A. Anthropic, "The claude 3 model family: Opus, sonnet, haiku," *Claude-3 Model Card*, 2024.
- [12] S. Zhang, L. Dong, X. Li, S. Zhang, X. Sun, S. Wang, J. Li, R. Hu, T. Zhang, F. Wu *et al.*, "Instruction tuning for large language models: A survey," *arXiv preprint arXiv:2308.10792*, 2023.
- [13] M. Hartley, "How to use chatgpt to compare documents," <https://www.evolution.ai/post/use-chatgpt-to-compare-documents>, 2024.
- [14] Y. Liu, S. Tao, W. Meng, J. Wang, W. Ma, Y. Chen, Y. Zhao, H. Yang, and Y. Jiang, "Interpretable online log analysis using large language models with prompt strategies," in *Proceedings of the 32nd IEEE/ACM International Conference on Program Comprehension*, 2024, pp. 35–46.
- [15] S. Tao, Y. Liu, W. Meng, Z. Ren, H. Yang, X. Chen, L. Zhang, Y. Xie, C. Su, X. Oiao *et al.*, "Biglog: Unsupervised large-scale pre-training for a unified log representation," in *2023 IEEE/ACM 31st International Symposium on Quality of Service (IWQoS)*. IEEE, 2023, pp. 1–11.
- [16] S. Tao, W. Meng, Y. Cheng, Y. Zhu, Y. Liu, C. Du, T. Han, Y. Zhao, X. Wang, and H. Yang, "Logstamp: Automatic online log parsing based on sequence labelling," *ACM SIGMETRICS Performance Evaluation Review*, vol. 49, no. 4, pp. 93–98, 2022.
- [17] Y. Fu, M. Yan, Z. Xu, X. Xia, X. Zhang, and D. Yang, "An empirical study of the impact of log parsers on the performance of log-based anomaly detection," *Empirical Software Engineering*, vol. 28, no. 1, p. 6, 2023.
- [18] A. S. van Linschoten, "The state of llm operations or llmops: Why everything is hard," <https://www.zenml.io/blog/state-of-llmops-why-everything-is-hard>, 2024.
- [19] K. S. Kalyan, "A survey of gpt-3 family large language models including chatgpt and gpt-4," *Natural Language Processing Journal*, p. 100048, 2023.
- [20] Y. Liu, S. Tao, X. Zhao, M. Zhu, W. Ma, J. Zhu, C. Su, Y. Hou, M. Zhang, M. Zhang *et al.*, "Coachlm: Automatic instruction revisions improve the data quality in llm instruction tuning," in *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. IEEE, 2024, pp. 5184–5197.
- [21] T. H. Kung, M. Cheatham, A. Medenilla, C. Sillos, L. De Leon, C. Elepaño, M. Madiaga, R. Aggabao, G. Diaz-Candido, J. Maningo *et al.*, "Performance of chatgpt on usmle: Potential for ai-assisted medical education using large language models," *PLoS digital health*, vol. 2, no. 2, p. e0000198, 2023.
- [22] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [23] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar *et al.*, "Llama: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.
- [24] J. Wei, M. Bosma, V. Zhao, K. Guu, A. W. Yu, B. Lester, N. Du, A. M. Dai, and Q. V. Le, "Finetuned language models are zero-shot learners," in *International Conference on Learning Representations*, 2021.
- [25] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," *The Journal of Machine Learning Research*, vol. 21, no. 1, pp. 5485–5551, 2020.
- [26] S. Mishra, D. Khashabi, C. Baral, and H. Hajishirzi, "Cross-task generalization via natural language crowdsourcing instructions," in *60th Annual Meeting of the Association for Computational Linguistics, ACL 2022*. Association for Computational Linguistics (ACL), 2022, pp. 3470–3487.
- [27] P. Li, Y. He, D. Yashar, W. Cui, S. Ge, H. Zhang, D. Rifinski Fainman, D. Zhang, and S. Chaudhuri, "Table-gpt: Table fine-tuned gpt for diverse table tasks," *Proc. ACM Manag. Data*, vol. 2, no. 3, May 2024. [Online]. Available: <https://doi.org/10.1145/3654979>
- [28] Q. Fu, J.-G. Lou, Y. Wang, and J. Li, "Execution anomaly detection in distributed systems through unstructured log analysis," in *2009 ninth IEEE international conference on data mining*, 2009, pp. 149–158.
- [29] M. Du and F. Li, "Spell: Streaming parsing of system event logs," in *2016 IEEE 16th International Conference on Data Mining (ICDM)*. IEEE, 2016, pp. 859–864.
- [30] A. A. Makanju, A. N. Zincir-Heywood, and E. E. Milios, "Clustering event logs using iterative partitioning," in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2009, pp. 1255–1264.
- [31] S. Zhang, W. Meng *et al.*, "Syslog processing for switch failure diagnosis and prediction in datacenter networks," in *IEEE/ACM 25th International Symposium on Quality of Service (IWQoS'17)*, 2007, pp. 1–10.
- [32] P. He, J. Zhu, Z. Zheng, and M. R. Lyu, "Drain: An online log parsing approach with fixed depth tree," in *2017 IEEE international conference on web services (ICWS)*. IEEE, 2017, pp. 33–40.
- [33] G. Chu, J. Wang, Q. Qi, H. Sun, S. Tao, and J. Liao, "Prefix-graph: A versatile log parsing approach merging prefix tree with probabilistic graph," in *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. IEEE, 2021, pp. 2411–2422.
- [34] W. Meng, Y. Liu, F. Zaiter *et al.*, "Logparse: Making log parsing adaptive through word classification," in *2020 29th International Conference on Computer Communications and Networks (ICCCN)*, 2020, pp. 1–9.
- [35] Y. Huo, Y. Su, C. Lee, and M. R. Lyu, "Semparser: A semantic parser for log analytics," in *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 2023, pp. 881–893.
- [36] Z. Li, C. Luo, T.-H. P. Chen, W. Shang, S. He, Q. Lin, and D. Zhang, "Did we miss something important? studying and exploring variable-aware log abstraction," in *ICSE 2023*, May 2023.
- [37] V.-H. Le and H. Zhang, "Log parsing with prompt-based few-shot learning," in *2023 IEEE/ACM International Conference on Software Engineering (ICSE)*. IEEE, 2023.
- [38] X. Zhang, Y. Xu, Q. Lin, B. Qiao, H. Zhang, Y. Dang, C. Xie, X. Yang, Q. Cheng, Z. Li *et al.*, "Robust log-based anomaly detection on unstable log data," in *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2019, pp. 807–817.
- [39] S. Lu, X. Wei, Y. Li, and L. Wang, "Detecting anomaly in big data system logs using convolutional neural network," in *2018 IEEE 16th Intl Conf on Dependable (DASC/PiCom/DataCom/CyberSciTech)*. IEEE, 2018, pp. 151–158.
- [40] 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.
- [41] M. Du, F. Li, G. Zheng, and V. Srikumar, "Deeplog: Anomaly detection and diagnosis from system logs through deep learning," in *Proceedings of the 2017 ACM SIGSAC conference on computer and communications security*, 2017, pp. 1285–1298.

- [42] W. Meng, Y. Liu, Y. Zhu *et al.*, “Loganomaly: Unsupervised detection of sequential and quantitative anomalies in unstructured logs.” in *IJCAI*, vol. 19, no. 7, 2019, pp. 4739–4745.
- [43] T. Cui, S. Ma, Z. Chen, T. Xiao, S. Tao, Y. Liu, S. Zhang, D. Lin, C. Liu, Y. Cai *et al.*, “Logeval: A comprehensive benchmark suite for large language models in log analysis,” *arXiv preprint arXiv:2407.01896*, 2024.
- [44] W. Meng, F. Zaiter, Y. Zhang, Y. Liu, S. Zhang, S. Tao, Y. Zhu, T. Han, Y. Zhao, E. Wang *et al.*, “Logsummary: Unstructured log summarization for software systems,” *IEEE Transactions on Network and Service Management*, vol. 20, no. 3, pp. 3803–3815, 2023.
- [45] Y. Zhu, W. Meng, Y. Liu, S. Zhang, T. Han, S. Tao, and D. Pei, “Unilog: Deploy one model and specialize it for all log analysis tasks,” *arXiv preprint arXiv:2112.03159*, 2021.
- [46] S. He, J. Zhu, P. He, and M. R. Lyu, “Loghub: a large collection of system log datasets towards automated log analytics,” *arXiv preprint arXiv:2008.06448*, 2020.
- [47] J. Wang, G. Chu, J. Wang, H. Sun, Q. Qi, Y. Wang, J. Qi, and J. Liao, “Logexpert: Log-based recommended resolutions generation using large language model,” in *Proceedings of the 2024 ACM/IEEE 44th International Conference on Software Engineering: New Ideas and Emerging Results*, 2024, pp. 42–46.
- [48] S. Tao, Y. Liu, W. Meng, J. Wang, Y. Zhao, C. Su, W. Tian, M. Zhang, H. Yang, and X. Chen, “Da-parser: A pre-trained domain-aware parsing framework for heterogeneous log analysis,” in *2023 IEEE 47th Annual Computers, Software, and Applications Conference (COMPSAC)*. IEEE, 2023, pp. 322–327.
- [49] Y. Liu, S. Tao, W. Meng, J. Wang, H. Yang, and Y. Jiang, “Multi-source log parsing with pre-trained domain classifier,” *IEEE Transactions on Network and Service Management*, 2023.
- [50] A. Oliner and J. Stearley, “What supercomputers say: A study of five system logs,” in *37th annual IEEE/IFIP international conference on dependable systems and networks (DSN’07)*. IEEE, 2007, pp. 575–584.
- [51] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale *et al.*, “Llama 2: Open foundation and fine-tuned chat models,” *arXiv preprint arXiv:2307.09288*, 2023.
- [52] Y. Zheng, R. Zhang, J. Zhang, Y. YeYanhan, and Z. Luo, “LlamaFactory: Unified efficient fine-tuning of 100+ language models,” in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, Y. Cao, Y. Feng, and D. Xiong, Eds. Bangkok, Thailand: Association for Computational Linguistics, Aug. 2024, pp. 400–410.
- [53] L. Tang, T. Li, and C.-S. Perng, “Logsig: Generating system events from raw textual logs,” in *Proceedings of the 20th ACM international conference on Information and knowledge management*, 2011, pp. 785–794.
- [54] S. Messaoudi, A. Panichella, D. Bianculli, L. Briand, and R. Sasnauskas, “A search-based approach for accurate identification of log message formats,” in *2018 IEEE/ACM 26th International Conference on Program Comprehension (ICPC)*. IEEE, 2018, pp. 167–16710.
- [55] W. M. Rand, “Objective criteria for the evaluation of clustering methods,” *Journal of the American Statistical association*, vol. 66, no. 336, pp. 846–850, 1971.
- [56] A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Yang, A. Fan *et al.*, “The llama 3 herd of models,” *arXiv preprint arXiv:2407.21783*, 2024.
- [57] H. Guo, J. Yang, J. Liu, L. Yang, L. Chai, J. Bai, J. Peng, X. Hu, C. Chen, D. Zhang *et al.*, “Owl: A large language model for it operations,” in *The Twelfth International Conference on Learning Representations*, 2024.
- [58] C.-Y. Lin, “Rouge: A package for automatic evaluation of summaries,” in *Text summarization branches out*, 2004, pp. 74–81.
- [59] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, “Bleu: a method for automatic evaluation of machine translation,” in *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 2002, pp. 311–318.
- [60] X. Ren, P. Zhou, X. Meng, X. Huang, Y. Wang, W. Wang, P. Li, X. Zhang, A. Podolskiy, G. Arshinov *et al.*, “Pangu- $\{\Sigma\}$: Towards trillion parameter language model with sparse heterogeneous computing,” *arXiv preprint arXiv:2303.10845*, 2023.
- [61] T. Schick, J. Dwivedi-Yu, R. Dessì, R. Raileanu, M. Lomeli, E. Hambro, L. Zettlemoyer, N. Cancedda, and T. Scialom, “Toolformer: Language models can teach themselves to use tools,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [62] Y. Zhang, Y. Li, L. Cui, D. Cai, L. Liu, T. Fu, X. Huang, E. Zhao, Y. Zhang, Y. Chen *et al.*, “Siren’s song in the ai ocean: a survey on hallucination in large language models,” *arXiv preprint arXiv:2309.01219*, 2023.