A Survey on Automated Log Analysis for Reliability Engineering

SHILIN HE, Department of Computer Science and Engineering, The Chinese University of Hong Kong

PINJIA HE, Department of Computer Science, ETH Zurich

ZHUANGBIN CHEN, TIANYI YANG, YUXIN SU, and MICHAEL R. LYU, Department of Computer Science and Engineering, The Chinese University of Hong Kong

Logs are semi-structured text generated by logging statements in software source code. In recent decades, software logs have become imperative in the reliability assurance mechanism of many software systems because they are often the only data available that record software runtime information. As modern software is evolving into a large scale, the volume of logs has increased rapidly. To enable effective and efficient usage of modern software logs in reliability engineering, a number of studies have been conducted on automated log analysis. This survey presents a detailed overview of automated log analysis research, including how to automate and assist the writing of logging statements, how to compress logs, how to parse logs into structured event templates, and how to employ logs to detect anomalies, predict failures, and facilitate diagnosis. Additionally, we survey work that releases open-source toolkits and datasets. Based on the discussion of the recent advances, we present several promising future directions toward real-world and next-generation automated log analysis.

1 INTRODUCTION

In the recent decades, modern software has been increasingly integrated into our daily lives and becomes indispensable, such as search engine, instant messaging app., and cloud systems. Most of these software systems are expected to be available on a 24×7 basis. Any non-trivial downtime can lead to significant revenue loss, especially for large-scale distributed systems [45, 46, 136]. For example, in 2017, a downtime in Amazon led to a loss of 150+ million US dollars [137]. Thus, the reliability of modern software is of paramount importance.

Software logs have been widely employed in a variety of reliability assurance tasks, because they are often the only data available that record software runtime information. Additionally, logs also play an indispensable role in data-driven decision making in industry [119]. In general, logs are semi-structured text printed by logging statements (*e.g.*, printf(), logger.info()) in the source code. For example, in Fig. 1, the two log messages are printed by the two logging statements

```
1  public void setTemperature(Integer temperature) {
2    // ...
3    logger.debug("Temperature set to {}. Old temperature was {}.", t, oldT);
4    if (temperature.intValue() > 50) {
5      logger.info("Temperature has risen above 50 degrees.");
6    }
7  }
```

$$\Downarrow$$

```
1  0 [setTemperature] DEBUG Wombat - Temperature set to 61. Old temperature was 42.
2  0 [setTemperature] INFO Wombat - Temperature has risen above 50 degrees.
```
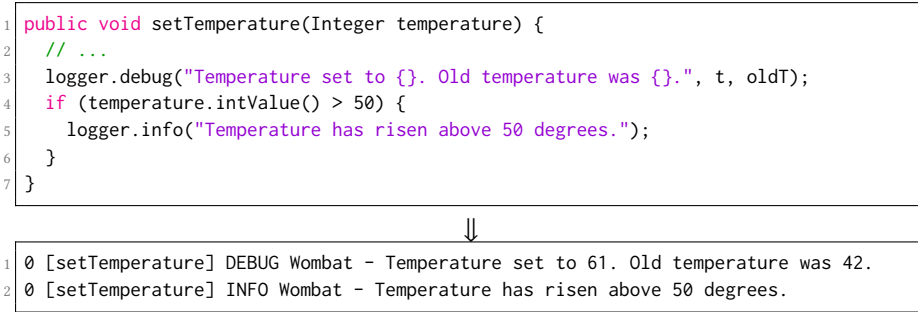
Fig. 1. An example of logging statements by SLF4J and the generated logs.

in the source code. The first few words (*e.g.*, "Wombat") of the log messages are decided by the corresponding logging framework (*e.g.*, SLF4J) and they are in structured form. On the contrary, the remaining words (*e.g.*, "50 degrees") are unstructured because they are written by developers to describe specific system runtime events.

A typical log analysis management framework is illustrated by the upper part of Fig. 2. Particularly, traditional logging practice (*e.g.*, which variables to print) mainly relies on developers' domain knowledge. During system runtime, software logs are collected and compressed as normal files using file compression toolkits [90] (*e.g.*, WinRAR). Additionally, developers leverage the collected logs in various reliability assurance tasks (*i.e.*, log mining), such as anomaly detection. Decades ago, these processes are based on specific rules specified by the developers. For example, to extract specific information related to a task (*e.g.*, thread ID), developers need to design regex (*i.e.*, regular expression) rules for automated log parsing [63]. The traditional anomaly detection process also relies on manually constructed rules [54]. These log analysis techniques were effective at the beginning because most of the widely-used software systems were small and simple.

However, as modern software has become much larger in scale and more complex in structure, traditional log analysis that is mainly based on ad-hoc domain knowledge or manually constructed and maintained rules becomes inefficient and ineffective [117]. This brings four major challenges to modern log analysis. (1) In many practices, while a number of senior developers in the same group share some best logging practices, most of the new developers or developers from different projects write logging statements based on domain knowledge and ad-hoc designs. As a result, the quality of the runtime logs varies to a large extent. (2) The volume of software logs has increased rapidly (*e.g.*, 50 GB/h [109]). Consequently, it is much more difficult to manually dig out the rules (*e.g.*, log event templates or anomalous patterns). (3) With the prevalence of Web service and source code sharing platforms (*e.g.*, Github), software could be written by hundreds of global developers. Developers who should maintain the rules often have no idea of the original logging purpose, which further increases the difficulty in manual maintenance of their rules. (4) Due to the wide adoption of the agile software development concept, a new software version often comes in a short-term manner. Thus, corresponding logging statements update frequently as well (*e.g.*, hundreds of new logging statements per month [142]). However, it is hard for developers to manually update the rules.

To address these challenges, a great amount of work has been accomplished by both researchers and practitioners in recent decades. In particular, starting from 2003, a line of research efforts have been contributed to automated rule construction and critical information extraction from software logs, including the first pieces of work of log parsing [138], anomaly detection [138], and failure prediction [125]. In addition, in the same year, Hätönen *et al.* [57] proposed the first log specific compression technique. Later, many empirical studies were conducted as the first steps towards
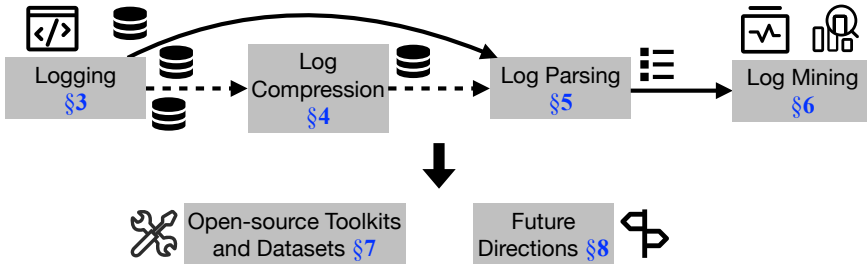
Fig. 2. An overall framework for automated log analysis.

some difficult problems in automated log analysis, including the first study on failure diagnosis by Jiang *et al.* [69] in 2009, the first exploration of logging practice by Yuan *et al.* [152] in 2012, and the first industrial study on logging practice by Fu *et al.* [50] in 2014. Recently, machine learning algorithms, especially deep neural networks, have been widely adopted by the state-of-the-art (SOTA) papers, such as the deep learning-based "what-to-log" approach [82] in logging practice and Deeplog [42] in anomaly detection. Besides machine learning, parallelization has been employed in various recent papers, such as Logzip [90] in log compression and POP [59] in log parsing.

These extensive studies on automated log analysis across multiple core directions have largely boosted the effectiveness and efficiency of systematic usage of software logs. However, the diversity and richness of both the research directions and recent papers could inevitably hinder the non-experts who intend to understand the SOTA and propose further improvements. To address this problem, this paper surveys 139 papers in the last 23 years across a variety of topics in log analysis. The papers under exploration are mainly from top venues in three related fields: software engineering (*e.g.*, ICSE), system (*e.g.*, SOSP), and networking (*e.g.*, NSDI). Thus, the readers can obtain a deep understanding of the advantages and limitations of the SOTA approaches, as well as taking a glance at the existing open-source toolkits and datasets. In addition, the insights and challenges summarized in the paper can help practitioners understand the potential usage of the automated log analysis techniques in practice and realize the gap between academy and industry in this field. We present crucial research efforts on automated log analysis from the following six perspectives as illustrated in Fig. 2:

- Logging: Section 3 introduces approaches that automate or improve logging practices, including *where-to-log*, *what-to-log*, and *how-to-log*.
- Log compression: Section 4 presents approaches to compress software logs in runtime.
- Log parsing: Section 5 discusses how to automatically extract event templates and key parameters from software logs.
- Log mining: Section 6 introduces what automated log mining techniques can do to enhance system reliability. We focus on three main tasks: *anomaly detection*, *failure prediction*, and *failure diagnosis*.
- Open-source toolkits and datasets: Papers providing open-source toolkits and datasets, which facilitate log analysis research, are presented in Section 7.
- Future directions: Section 8 discusses open challenges and promising future directions that can push this field forward beyond current practices.

2 SURVEY METHODOLOGY

To systematically collect the publications for conducting this survey, we maintained a repository for the automated log analysis literature. We first searched relevant papers in online digital libraries

Fig. 3. Paper distribution on each research topic and the associated evolution trend from 1997 to 2020.

and extended the repository by manually inspecting all references of these papers. The repository is now available online.[1]

To begin with, we searched several popular online digital libraries (*e.g.*, IEEE Xplore, ACM Digital Library, Springer Online, Elsevier Online, Wiley Online, and ScienceDirect) with the following keywords: "log", "logging", "log parsing", "log compression", "log + anomaly detection", "log + failure prediction", "log + failure diagnosis". The "log" term is broadly quoted in various domains, such as the query log in database, web search log in recommendation, and the logarithm function in mathematics. Hence, to precisely collect the set of papers of our interest, we mainly focused on regular papers published in top venues (*i.e.*, conferences and journals) of relevant domains, including ICSE, FSE, ASE, TSE, TOSEM, SOSP, OSDI, ATC, NSDI, TDSC, DSN, ASPLOS, and TPDS. Then, we manually inspected each reference of these papers to collect additional publications that are related to the survey topics. We focused on publications studying the reliability issues in software systems, while other issues such as security and privacy are beyond the scope of this survey.

In total, we collected 139 publications in automated log analysis area spanning from year 1997 to 2020. Fig. 3 (on the right part) shows the histogram of annual publications during the period. We can find that automated log analysis has been continuously and actively investigated in the past two decades. Moreover, we can observe a dramatic rise on the number of publications since 2014, indicating that the log analysis field has attracted significant interests since then. One possible reason is that the prosperity of cloud computing leads to the high demand of automated log analysis.

Furthermore, we classify these publications into five categories by research focus: logging, log compression, log parsing, log mining, and empirical study. The distribution is presented on the left part of Fig. 3. Note that some papers may address more than one research focuses, in which case we categorized the paper based on its major contribution. As expected, Fig. 3 shows that a large portion (around a half) of research efforts were devoted to log mining. The reasons are two-fold: First, log mining is a comprehensive task that consists of many sub-tasks (*e.g.*, anomaly detection, failure diagnosis). Each sub-task has a huge space for research exploration and thereby attracts many research studies on it; Second, research areas other than log mining are relatively mature. Besides, logging is also extensively studied in the past decades since it is more closely related to industrial practices. In addition, many log research tasks are empirical studies, and we discussed them in different sections respectively.

Furthermore, we were aware of several existing surveys on logs analysis. Oliner *et al.* [117] briefly reviewed sixteen log analysis papers and pointed out some challenges in 2012. Since then, the log

---

[1]https://github.com/logpai/log-survey

analysis area has been actively studied, which impels a more comprehensive and systematic log survey paper. Other existing surveys mainly studied the log and their use in the security area. Khan *et al.* [75] focused on logs in the cloud computing environment and the logs were leveraged for identifying and preventing suspicious attacks. Zeng *et al.* [155] studied the logging mechanisms and security issues in three computer operating systems. A recent publication [79] reviewed clustering based approaches in cyber security applications. Different from these studies, our survey mainly targets at the automated log analysis for tackling the general reliability issues.

3 LOGGING

Logging is the act to collect numerical and textual data of software behavior, such as low-memory conditions or attempts to access a file. Logging is valuable for both academia [110, 119, 126, 144] and industry [9, 18, 50, 119, 152, 156] across a variety of application domains because logging statements is a fundamental step for all the subsequent log mining tasks.

3.1 Logging Mechanism and Libraries

3.1.1 Logging Mechanism. Logging mechanism is the set of logging statements and their activation code implemented by developers or a given software platform [119]. Fig. 1 shows two example logging statements and the collected logs during the execution of the program. The logging statement at line 3 is executed every time the `setTemperature` method is called, with no specific activation code. The logging statement at line 5 is controlled by the activation code `if (temperature.intValue() > 50)` at line 4. According to the data collected by [119], the most adopted coding pattern that is used to activate the logging statements is `if (condition) then log error()`.

3.1.2 Logging Libraries. To improve flexibility, industrial logging practice often utilizes logging libraries [19, 119], which are software components that facilitate logging and provide advanced features (*e.g.*, thread-safety, log archive configuration, and API separation). Toward this end, a lot of open-source logging libraries have been developed (*e.g.*, Log4j[2], SLF4J[3], AspectJ[4], spdlog[5]).

3.2 Challenges for Logging

The logging mechanism of a software system is usually decided by an empirical process in the development phase [119]. In general, logging practice is scarcely documented or regulated by strict standard, such as the logging mechanism and APIs [9]. Thus, logging relies heavily on human expertise [50, 58, 117, 119, 127]. In the following, we summarize three main challenges for automated logging, which also aligns with the taxonomy mentioned by Chen and Jiang [17]: *where-to-log*, *what-to-log*, and *how-to-log*. Under this categorization, every problem exhibits aspects that represent the primary concerns of the actual practice of logging. Accordingly, we summarized three major aspects, *i.e.*, *diagnosability*, *maintenance*, and *performance*.

3.2.1 where-to-log. Where-to-log is about determining the appropriate location of logging statements. Although logging statements provide rich information and can be inserted almost everywhere, excessive logging results in performance degradation [17] and incurs additional maintenance overhead. In addition, it is challenging to diagnose problems by analyzing a large volume of logs as most logs are unrelated to the problematic scenarios [71]. On the other hand, insufficient logging will also impede the logs' diagnosability. For example, an incomplete sequence of logs may hinder

---

[2]http://logging.apache.org/log4j/
[3]http://www.slf4j.org/
[4]https://www.eclipse.org/aspectj/
[5]https://github.com/gabime/spdlog

the reproduction of precise execution paths [161]. Therefore, developers need to be circumspect in their choices of where-to-log.

3.2.2 what-to-log. What-to-log is about providing sufficient and concise information within the three major components of a logging statement, *i.e.*, verbosity level, static text, and dynamic content. Mis-configured verbosity level has similar consequences with inappropriate logging points. As developers typically filter logs according to the verbosity levels, under-valued verbosity levels may result in missing or ignored log messages while over-valued verbosity levels lead to overwhelming log messages [71]. When composing a snippet of logging code, the static text should be concise and the dynamic content should be coherent and up-to-date. Poorly written static text and inconsistent dynamic content could affect the subsequent diagnosis and maintenance activities [72, 91, 107, 158].

3.2.3 how-to-log. How-to-log is the "design pattern" and maintenance of logging statements systematically. Most software testing techniques focus on verifying the quality of feature code, but a few papers [17, 18, 56, 129, 152] pay attention to the quality and anti-patterns in the logging code. Most industrial and open source systems choose to scatter logging statements across the entire code base, intermixing with feature code [17], which also hardens the maintenance of logging code.

3.3 Logging Approaches

Numerous solutions have been proposed to address the challenges mentioned in Section 3.2. Table 1 summarizes the existing work along with the corresponding problems and aspects. Each row represents a challenge. The approaches fall into three categories: *static code analysis*, *machine learning*, and *empirical study*. A bunch of early work utilized static code analysis to analyze logging in a program without actually executing it. In most cases, it is performed on the source code. Machine learning-based approaches focus on learning from data. By paying attention to the inherent statistical properties of existing logging statements, learning-based approaches automatically give suggestions on improving the logging statements. In the remainder of this section, we discuss solutions by their aspects.

3.3.1 Diagnosability. Logs are valuable for investigating and diagnosing failures. However, a logging statement is only as helpful as the information it provides. The research on this aspect aims at (1) understanding the helpfulness of logs for failure diagnosis and (2) making logs informative for diagnosis.

(1) *Understanding the helpfulness of logs for failure diagnosis.* This is of great importance because logs are widely adopted for failure diagnosis. According to a survey [50] involving 54 experienced developers in Microsoft, almost all the participants agreed that "logging statements are important in system development and maintenance" and "logs are a primary source for problem diagnosis". Fu *et al.* [50] also studied the types of logging statements in industrial software systems by source code analysis, and summarized five types of logging snippets, *i.e.*, assertion-check logging, return-value-check logging, exception logging, logic-branch logging, and observing-point logging. Besides, Fu *et al.* [50] further demonstrated the potential feasibility of predicting where to log. Shang *et al.* [130] conducted the first empirical study to provide a taxonomy for user inquiries of logs. They identified five types of information that were often sought from log lines by practitioners, *i.e.*, meaning, cause, context, impact, and solution. Shang *et al.* [130] were also the first to associate the development knowledge at present in various development repositories (*e.g.*, code commits and issues reports) with the log lines and to assist practitioners in resolving real-life log inquiries.

(2) *Making logs informative for failure diagnosis.* As printed logs are often the only run-time information source for debugging and analysis, the quality of log data is critically important. *LogEnhancer* [153] made the first attempt to systematically and automatically augment existing

Table 1. Summary of logging approaches.

Problems	Aspects	Objectives
where-to-log	Diagnosability	Suggest appropriate placement of logging statements into source code [29, 30, 82, 151, 160, 161, 165] ; Study logging practices in industry [50].
	Performance	Minimize or reduce performance overhead [39, 151].
what-to-log	Diagnosability	Enhance existing logging code to aid debugging [153]; Suggest proper variables and text description in log [58, 91].
	Maintenance	Determine whether a logging statement is likely to change in the future [73]; Characterize and detect duplicate logging code [83].
	Performance	Study the performance overhead and energy impact of logging in mobile app [24, 156]; Automatically change the log level of a system in case of anomaly [112].
how-to-log	Diagnosability	Characterize the anti-patterns in the logging code [130]; Optimize the implementation of logging mechanism to facilitate failure diagnosis [101].
	Maintenance	Characterize and detect the anti-patterns in the logging code [17, 18, 56, 152]; Characterize and prioritize the maintenance of logging statements [72]; Study the relationship between logging characteristics and the code quality [129]; Propose new abstraction or programming paradigm of logging [76, 92].
	Performance	Optimize the compilation and execution of logging code [146].

logging statements in order to reduce the number of possible code paths and execution states for developers to pinpoint the root cause of a failure. Zhao *et al.* [160, 161] followed the idea of *LogEnhancer* and proposed an algorithm capable of completely disambiguating the call path of HDFS requests. Yuan *et al.* [151] found that the majority of unreported failures were manifested via a generic set of error patterns (*e.g.*, system call return errors) and proposed the tool *Errlog* to proactively add pattern-specific logging statements by static code analysis.

As modern software becomes more complex, where-to-log has become an important but difficult decision, largely limited to the developer's domain knowledge. Around 60% of failures due to software faults do not leave any trace in logs, and 70% of the logging pattern aims to detect errors via a checking code placed at the end of a block of instructions [29, 30]. Cinque *et al.* [30] concluded that the traditional logging mechanism has limited capacity due to the lack of a systematic error model. They further formalized the placement of the logging instruction and proposed to use system design artifacts to manually define *rule-based logging* which utilizes error models about what causes errors to fail. Zhu *et al.* [165] made an important first step towards the goal of "learning to log". They proposed a logging recommendation tool, *LogAdvisor*, that learns the common logging rules on where-to-log from existing code via training a classifier and further leverages it for informative and viable recommendations to developers. Li *et al.* [82] proposed a deep learning framework to suggest where-to-log at the block level. Li et al. also concluded that there might be similar rules

regarding the implementation of logging mechanism across different systems and development teams, which agreed with the industrial survey by Pecchia *et al.* [119].

The lack of strict logging guidance and domain-specific knowledge makes it difficult for developers to decide what-to-log. To address this need, He *et al.* [58] conducted the first empirical study on the usage of natural language in logging statements. They showed the global (*i.e.*, in a project) and local (*i.e.*, in a file) repeatability of text descriptions. Furthermore, they demonstrate the potential of automated description text generation for logging statements. Liu *et al.* [91] proposed a deep learning-based approach to recommend variables in logging statements by learning embeddings of program tokens. In order to troubleshoot transiently-recurring problems in cloud-based production systems, Luo *et al.* [101] put forward a new logging mechanism that assigns a blame rank to methods based on their likelihood of being relevant to the root cause of the problem. With the blame rank, logs generated by a method over a period of time are proportional to how often it is blamed for various misbehavior, thus facilitating diagnosis.

3.3.2 Maintenance. The maintenance of logging code has also attracted researchers' interest. The research on this aspect aims at (1) characterizing the maintenance and detecting anti-patterns of logging statements and (2) proposing new abstractions of logging.

(1) Characterizing the maintenance and detecting anti-patterns of logging statements. Anti-patterns in logging statements are bad coding patterns that undermine the quality and effectiveness of logging statements and increases the maintenance effort of projects. Many papers [18, 56, 129] performed empirical studies to reveal the link between logs and defects. These papers observed a positive correlation between logging characteristics and post-release defects. Therefore, practitioners should allocate more effort to source code files with more logs.

Yuan *et al.* [152] made the first attempt to conduct a quantitative characteristic study of how developers log within four pieces of large open-source software. They described common anti-patterns and provided insights into where developers spend most of their efforts in modifying the log messages and how to improve logging practice. They further implemented a prototype checker to verify the feasibility of detecting unknown problematic statements using historical commit data. Chen and Jiang [17] and Hassani *et al.* [56] both studied the problem of how-to-log by characterizing and detecting the anti-patterns in the logging code. The analysis [17] of well-maintained open-source systems revealed six anti-patterns that are endorsed by developers. Chen and Jiang [17] then encoded these anti-patterns into a static code analysis tool to automatically detect anti-patterns in the source code. Li *et al.* [83] developed an automated static analysis tool, *DLFinder*, to detect duplicate logging statements that have the same static text description.

Just like feature code, logging code updates with time [73]. Moreover, logging statements are often changed without consideration for other stakeholders, resulting in sudden failures of log analysis tools and increased maintenance costs for such tools. Pecchia *et al.* [119] reviewed the industrial practice in the reengineering of logging code. Kabinna *et al.* [72] empirically studied the migration of logging libraries and the main reasons for the migration. Kabinna *et al.* [73] later examined the important metrics for determining the stability of logging statements and further leveraged learning-based models (random forest classifier and Cox proportional hazards) to determine whether a logging statement is likely to remain unchanged in the future. Their findings were helpful to build robust log analysis tools by ensuring that these tools relied on logs generated by more stable logging statements.

(2) Proposing new abstractions of logging. Maintaining logging code along with feature code has proven to be error-prone [18, 56, 129]. Hence, additional logging approaches [76, 92] have been proposed to resolve this issue. Kiczales *et al.* [76] proposed a new programming paradigm that improves the modularity of the logging code. To tackle the ordering problem of logs in distributed

systems, Lockerman *et al.* [92] introduced *FuzzyLog* that featured strong consistency, durability, and failure atomicity.

3.3.3 Performance. The intermixing nature of the logging code and feature code usually incurs performance overhead, storage cost, and development and maintenance efforts [24, 39, 112, 153, 156]. Tools like *LogEnhancer* [153], *Errlog* [151], *Log2* [39], and *INFO-logging* [160] all took performance into consideration while dealing with diagnosability and maintenance issues. Mizouchi *et al.* [112] proposed a dynamical adjusting verbosity level to record irregular events while reducing performance overhead. Yang *et al.* [146] proposed *NanoLog*, a nanosecond scale logging system that achieved relatively low latency and high throughput by moving the workload of logging from the runtime hot path to the post-compilation and execution phases of the application. Chowdhury *et al.* [24] were the first to explore the energy impact of logging in mobile apps. Zeng *et al.* [156] conducted a case study that characterized the performance impact of logging on Android apps.

4   LOG COMPRESSION

After collecting logs by executing logging statements during runtime, logs are stored for failure diagnosis or sensitive operations auditing. Then, how to store logs efficiently becomes a challenging problem. Since large-scale software systems run on a 7*24 basis, the generated log size has been very huge (*e.g.*, 50 gigabytes per hour [109]). Besides, many logs require the long-term storage for identifying duplicate problems and mining failure patterns from historical logs [3, 38, 154]. Auditing logs that record sensitive user operations are often stored for two years and even more to track system misuse in future. Archiving logs in such a huge volume for a long period brings inconceivably heavy burden to storage space, electrical power, and network bandwidth for transmission. To tackle these problems, a line of research has been focusing on the log compression, which aims to remove the redundancy in a large log file and reduce its storage consumption.

4.1   Challenges for Log Compression

In practice, log compression can be achieved by various approaches. The most straightforward way is to reduce the amount of logging statements in the source code or to set a higher verbosity level of logs in the runtime phase. Although effective, this method leads to the information loss in logs, which impedes the log-based troubleshooting. Besides, it is a common practice to compress log files by some general file compressors, such as gzip. The general file compressor often serves as a core component in existing logging libraries (*e.g.*, log4j). However, the method is not tailored for the semi-structured log format, making it very ineffective in log compression. The structured fields (*i.e.*, constant parts) are often repetitive across many raw logs, which consume a lot of storage space and should be handled specifically. To address these issues, a tailored log compression algorithm is highly in demand.

4.2   Characteristics for Log Compression

An ideal log compressor should achieve a high compression ratio while imposing short compression and decompression time. The compression ratio denotes the ratio between the file sizes before and after the compression. A higher compression ratio indicates that less storage space is consumed and the compression algorithm is more effective. The log compression algorithm takes a certain amount of time to compress and decompress the file respectively, which is supposed to be as short as possible. In this paper, we mainly explored 12 existing log compressing approaches and several empirical studies. As shown in Table 2, to clearly demonstrate the advantages and disadvantages of these approaches, we list several characteristics that concern the log compressors based on the following descriptions in the surveyed papers. *(1) Use of general compressors (GC).* Some log

Table 2. Summary of log compression approaches.

Methods		GC	Target Data	Scalability	Heuristics
Bucket	Balakrishnan *et al.* [7]	Yes	Blue Gene/L logs	High	Yes
	Skibinski *et al.* [132]	Yes	All logs	Low	No
	Hassan *et al.* [55]	No	Telecom logs	High	Yes
	Christensen *et al.* [25]	Yes	All logs	High	Yes
	Feng *et al.* [48]	Yes	All logs	Middle	No
Dictionary	Deorowicz *et al.* [37]	No	Apache web log	Middle	Yes
	Lin *et al.* [86]	No	Structured Logs	High	Yes
	Mell *et al.* [105]	Yes	Structured Logs	High	No
	Racz *et al.* [122]	Yes	Web log	High	Yes
Statistics	Hatonen *et al.* [57]	No	Structured Logs	High	Yes
	Meinig *et al.* [104]	No	All logs	High	Yes
	Liu *et al.* [90]	Yes	Structured Logs	Low	No
Industry	Splunk [133]	Yes	All logs	Low	No
	ElasticSearch [44]	Yes	All logs	Low	No

compressors utilize the general compressor as the backend technique after reformatting the logs. Specifically, these reformattings can transform raw logs to better formats that achieve a higher compression ratio. *(2) General applicability* denotes whether the log compressor can be generally applicable to various log data formats. A general log compressor shows better utility in different scenarios. For example, as shown in the "target data" of Table 2, the log compressor proposed in [37] was specialized for Apache web log while other log compressors [25] can be applied to different log data types. *(3) Scalability.* Since logs are often of a great volume in practice, the efficiency of log compression is crucial to the practical employment. A highly scalable compressor could save both the compression and decompression time, which thereby saves the log query time and supports real-time analysis. *(4) Heuristics.* The deployment would be more efficient and smoother if the log compressor requires little prior knowledge, *e.g.*, heuristic rules on log formats and preprocessing procedures. The less heuristics a log compressor needs, the more generally applicable it would be.

4.3 Log Compression Approaches

According to the techniques employed behind, there are three categories of log compression methods: (1) *bucket-based compression* that divides the log data into different blocks and compresses each block in parallel; (2) *dictionary-based compression* that builds a dictionary for fields in the log and replaces strings by referring to the dictionary; and (3) *statistics-based compressors* that apply complex statistical models to find correlations between logs before compression. Note that some log compression methods may involve two or more techniques, in which case we took the major compression technique as their category.

4.3.1 Bucket-based Compression. Bucket-based compression methods break the log data into different blocks by some log-specific characteristics and then compresses these blocks in parallel. For instance, many logs show the strong temporal locality or high similarity in certain fields, therefore, these logs can be gathered and then compressed to achieve a high compression ratio. To tackle the problem that log data were often heterogeneous with varying patterns over time,

Christensen *et al.* [25] proposed to partition the log data into different buckets by leveraging the temporal locality and then compressed these buckets in parallel. Hassan *et al.* [55] adopted a similar idea for telecom logs; however the objective was not the log message but instead the sequence of log events. The log data is firstly broken into equal sized periods and compressed separately. Different from other studies, the proposed approach uses the compressed log sequences to identify noteworthy usage scenarios for operational profile customization.

Besides, the bucket partition often serves as an initial step towards achieving a high compression ratio. For example, LogPack [132] was a multi-tiered log compression method, in which each tier addressed one notion of redundancy. The first tier handled the local redundancy between neighboring lines. The second tier handled the global repetitiveness of tokens while the third tier handled all the remaining redundancy by employing a general compressor. Similarly, Balakrishnan *et al.* [7] first observed a number of trends in the Blue Gene/L system logs and proposed several solution to compress the logs accordingly. For example, one trend was that most columns in adjacent logs tended to be the same. At last, the generic compression utilities were applied to further compress the log file. Hence, they utilized a variant of delta encoding method which compared the log with preceding logs and encoded the differences only. In addition, Multi-level Log Compression (MLC) [48] divided logs with redundancy into different buckets by calculating the Jaccard similarity among logs. Then, the logs were condensed by a variant of delta encoding, followed by a general compressor to further improve the compression ratio.

4.3.2   Dictionary-based Compression. Dictionary-based compression removes the redundancy in log files by replacing repetitive strings with references to a dictionary. For example, some frequent strings (*e.g.*, IP addresses) can be mapped to a condensed string in the dictionary. Lin *et al.* [86] presented a column-wise independent compression for structured logs, which was also very similar to the work presented in [37, 122] though their focuses were web logs. The general idea of these methods is as follows: At first, it splits a log entry into several columns by its fields. After observing common properties in columns (*e.g.*, common prefixes and suffixes), the method builds a dictionary-based model with Huffman code to compress each column separately. Many general compression techniques can then be applied to achieve a higher compression ratio, including move-to-front coding, phrase sequence substitution, etc. Similarly, Mell *et al.* [105] proposed a multi-step method for log compression. It first separated and sorted logs according to particular properties (*e.g.*, value of the first column), and then it created a dictionary where logs are stored by hashing the field name, followed by serializing the dictionary and compressing logs with a general compressor.

4.3.3   Statistics-based Compression. Statistics-based methods realize the log compression by building a statistical model to identify possible redundancy in the log data. Unlike previous two categories of methods that are mainly heuristic-based, the statistically-based compression is more automated. Hatonen *et al.* [57] demonstrated a Comprehensive Log Compression (CLC) method to dynamically characterize and combine log data entries. Particularly, the method first identified frequently occurring patterns from dense log data byfrequent pattern mining, and then linked patterns to the data as a data directory. Meinig *et al.* [104] proposed an approach adopted from the rough set in the uncertainty theory. It treated the log file as a decision table with removable attributes identified by an one-time analysis of log data. Specifically, attributes that could roughly express each other are then collapsed to one representative attribute. The drawback was that the method might lose some information in the raw log messages. Liu *et al.* [90] proposed the Logzip method based on log parsing introduced in Section 5. Logzip first automatically extracted log templates from raw log messages by a statistical log parsing model. Then Logzip structuralized log templates and other information into intermediate representations, which were finally fed into a general compressor.

4.3.4   Others. Recently, Yao *et al.* [147] provided an empirical study on comparing the performance of general compressors on compressing log data against natural language data. Although the study was targeted on general compressors instead of specified log compressors, their findings could potentially guide the design of tailored log compressors. Besides, Otten *et al.* [118] presented a systematic review on general compression techniques, based on which they investigated the use of compression for log data reduction and the use of semantic knowledge to improve data compression.

5   LOG PARSING

After log collection, log messages will be input into different downstream log mining tasks (*e.g.*, anomaly detection) for further analysis. However, most of the existing log mining tools [62, 144] requires structured input data (*e.g.*, a list of structured log events or a matrix). Thus, a crucial step of automated log analysis is to parse the semi-structured log messages into structured log events.

081109 204655 556 INFO dfs.DataNode$PacketResponder: Received block blk_3587508140051953248 of size 67108864 from /10.251.42.84	

Structured Log

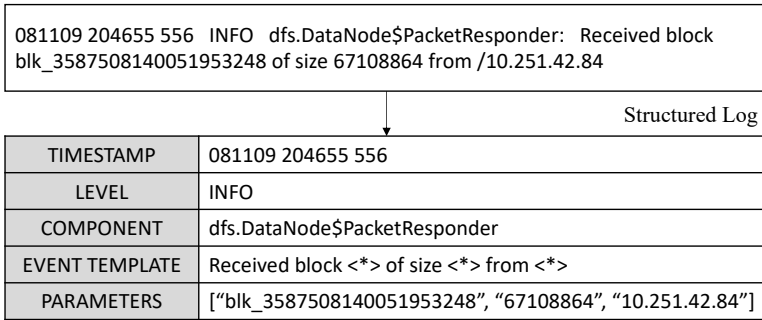| TIMESTAMP	081109 204655 556
LEVEL	INFO
COMPONENT	dfs.DataNode$PacketResponder
EVENT TEMPLATE	Received block <*> of size <*> from <*>
PARAMETERS	["blk_3587508140051953248", "67108864", "10.251.42.84"]

Fig. 4. A log parsing example.

Fig. 4 presents a log parsing example, where the input is a log message collected from Hadoop Distributed File System (HDFS) [144]. A log message is composed of message header and message content. The message header is determined by the logging framework and thus it is relatively easy to extract, such as verbosity levels (*e.g.*, "INFO"). In contrast, it is difficult to extract key information from the message content because it is mainly written by developers in free-form natural language. Typically, the message content contains *constants* and *variables*. Constants are the fixed text written by the developers (*e.g.*, "Received") and describe a system event, while variables are the values of the program variables which carry dynamic runtime information. The goal of log parsing is to distinguish between *constants* and *variables*. All the constants form the *event template*. The output of a log parser is a structured log message, containing an event template and the key parameters.

5.1   Challenges for Log Parsing

As a key component, log parsing has become an appealing selling-point [93, 95, 123] of many industrial log management solutions [166] (*e.g.*, Splunk [133]). However, these industrial solutions only support common log types such as Apache logs [123]. When parsing general log messages, they rely on regexs and ad-hoc scripts provided by developers. These scripts separate log messages into different groups, where log messages in the same group have the same event template. In modern software engineering, even with these existing industrial solutions, log parsing is still a challenging task due to three main reasons: (1) the large volume of logs and thus the great effort on manual regex construction; (2) the complexity of software and thus the diversity of event templates; and (3) the frequency of software updates and thus the frequent update of logging statements.

Table 3. Summary of log parsing approaches.

	Methods	Coverage	Preprocessing	Technique
Offline	SLCT [138]	Partial	No	Frequent pattern mining
	AEL [70]	All	Yes	Heuristics
	LKE [49]	All	Yes	Clustering
	LFA [114]	All	No	Frequent pattern mining
	LogSig [135]	All	No	Clustering
	IPLoM [102, 103]	All	No	Iterative partitioning
	LogCluster [139]	Partial	No	Frequent pattern mining
	LogMine [53]	All	Yes	Clustering
	POP [59]	All	Yes	Iterative partitioning
	MoLFI [108]	All	Yes	Evolutionary algorithms
Online	SHISO [113]	All	No	Clustering
	LenMa [131]	All	No	Clustering
	Spell [41]	All	No	Longest common subsequence
	Drain [60, 61]	All	Yes	Heuristics
	Logram [33]	All	Yes	Frequent pattern mining

5.2 Log Parser Characteristics

We explore 15 existing automated log parsing approaches. To provide a clear glimpse of these parsers, we focus on three main characteristics: mode, coverage, and preprocessing. We follow the definition of these characteristics introduced by Zhu *et al.* [166]. Note different from their work, which was an evaluation study providing benchmarks and open-source implementation of existing log parsers, this survey focuses on the methodological comparison of these parsers. In additional, while parsing rules can be constructed by analyzing source code [144], we focus on parsers that only utilize logs as input.

Mode. Mode is the most important log parsing characteristic to consider. Depending on the usage scenarios of parsing, log parsers can be categorized into two modes: offline and online. Offline log parsers require all the log messages beforehand and parse log messages in a batch manner. To cope with the frequent software update, developers need to periodically re-run the offline parser to obtain the newest event templates. In contrast, online parsers parse the log messages in a streaming manner, which work seamlessly with the downstream log mining tasks.

Coverage. Coverage is the ability to match all the log messages with their corresponding event templates. In Table 3, "Partial" indicates the parser can only parse part of the log messages (*i.e.*, some log messages will be matched with no event templates). This might lead to the neglection of crucial system anomalies.

Preprocessing. Preprocessing refers to the process that removes some *variables* or replaces them by *constants* based on domain knowledge. For example, IP addresses (*e.g.*, 10.251.42.84) are typical variables in cloud systems' log messages. This step requires some manual efforts (*i.e.*, constructing regexs for these variables).

5.3   Offline Log Parsing Approaches

Frequent Pattern Mining. SLCT (Simple Logfile Clustering Tool) [138] is the first research paper on automated log parsing. SLCT conducted two passes in total to obtain associated words. In the first pass, SLCT counted the occurrence of all the tokens and marked down the frequent words. These frequent words were utilized in the second pass to find out associated frequent words. Finally, for a log message, if it contained a pattern of associated frequent words, these words would be regarded as constants and employed to generate event templates. Otherwise, the log message would be placed into an outlier cluster without matched event templates. *LFA* [114] adopted a similar strategy as SLCT. Differently, LFA could cover all the log messages.

Clustering. LogCluster [139] is similar to SLCT [138]. Differently, LogCluster allowed variable length of parameters in between via a clustering algorithm. Thus, compared with SLCT, LogCluster is better at handling log messages of which the parameter length is flexible. For example, "Download Facebook and install" and "Download Whats App and install" have the same event template "Download <*> and install" while the length of the parameter (*i.e.*, an app name) is flexible. *LKE* (Log Key Extraction) [49] was developed by Microsoft. LKE adopted a hierarchical clustering algorithms with a customized weighted edit distance metric. Additionally, the clusters were further partitioned by heuristic rules. *LogSig* [135] was a more recent clustering-based parser than LKE. Instead of directly clustering log messages, LogSig transformed each log message into a set of word pairs and clustered logs based on the corresponding pairs. *LogMine* [53] adopted an agglomerative clustering algorithm. It was implemented in map-reduce framework for better efficiency.

Heuristics. AEL [70] employed a list of specialized heuristic rules. For example, for all the pairs like "word=value," AEL regarded the "value" as a variable and replaced it with a "$v" symbol.

Evolutionary Algorithms. MoLFI [108] formulate log parsing as a multi-objective optimization problem and propose an evolutionary algorithm-based approach. Specifically, MoLFI employs the Non-dominated Sorting Genetic Algorithm II [36] to search for a Pareto optimal set of event templates. Compared with other log parsers, the strength of MoLFI is that it requires little parameter tuning effort because the four parameters required by MoLFI has effective default values. However, because of the adoption of evolutionary algorithm, the MoLFI is slower than most of the parsers.

Iterative Partitioning. IPLoM [102, 103] contained three steps and partitioned log messages into groups in a hierarchical manner. (1) Partition by log message length. (2) Partition by *token position*. The position containing the least number of unique words is "token position". Partitioning was conducted according to the words in the token position. (3) Partition by mapping. Mapping relationships were searched between the set of unique tokens in two token positions, which were selected using a heuristic criterion. *POP* [59] is a parallel log parser that utilizes distributed computing to accelerate the parsing of large-scale software logs. POP can parse 200 million HDFS log messages in 7 mins, while most of the parsers (*e.g.*, LogSig) failed to terminate in reasonable time.

5.4   Online Log Parsing Approaches

Clustering. SHISO [113] is the first online log parsing approach. SHISO used a tree-form structure to guide the parsing process, where each node was correlated with a log group and an event template. The numbers of children nodes in all the layers were the same and were manually configured beforehand. During the parsing process, SHISO traversed the tree to find the most suitable log group by comparing the log message and the event templates in the corresponding log groups. SHISO is sensitive to path explosion and thus its efficiency is often unsatisfactory. *LenMa* [131] is similar to SHISO. LenMa encodes each log message into a *length vector*, where each dimension

records the number of characters of a token. For example, "Receive a file." would be vectorized as [7, 1, 5]. During parsing, LenMa would compare the length vectors of the log messages.

Longest Common Subsequence. Similar to SHISO and LenMa, *Spell* [41] maintained a list of log groups. To accelerate the parsing process, Spell utilized specialized data structures: prefix tree and inverted index. In addition, Spell provided a parallel implementation.

Heuristics. Drain [61] maintained log groups via the leaf nodes in the tree. The internal nodes of the tree embedded different heuristic rules. The extended version of Drain [60] was based on a directed acyclic graph that allowed log group online merging. In addition, it provided the first automated parameter tuning mechanism for log parsing.

Frequent Pattern Mining. Logram [33] is the current state-of-the-art parser. Different from the existing approaches that count frequent tokens, Logram considered frequent n-gram. The core insight of Logram is: frequent n-gram are more likely to be constants. Note that Logram assumed developers had some log messages on hand to construct the dictionary.

6 LOG MINING

Log mining employs statistics, data mining, and machine learning techniques for automatically exploring and analyzing large volume of log data to glean meaningful patterns and informative trends. The extracted patterns and knowledge could guide and facilitate monitoring, administering, and troubleshooting of software systems. In this section, we first elaborate on challenges encountered in log mining (Section 6.1). Then, we describe the general workflow of log mining (Section 6.2). Finally, we introduce three major log mining tasks for reliability engineering, including *anomaly detection* (Section 6.3), *failure prediction* (Section 6.4), and *failure diagnosis* (Section 6.5). Other relevant studies with relatively lesser popularity are laid out in Section 6.6.

6.1 Challenges of Log Mining

Traditionally, engineers perform simple keyword search (such as "error", "exception", and "failed") to mine suspicious logs that might be associated with software problems, *e.g.*, component failures. Some rule-based tools [54, 121, 124] have been developed to detect software problems by comparing logs against a set of manually defined rules which describe normal software behaviors. However, due to the ever-increasing volume, variety, and velocity of logs produced by modern software, such approaches fall short for being labor-intensive and error-prone. Moreover, suspicious logs are often overwhelmed by logs generated during software normal executions. Manually sifting through a massive amount of logs to identify failure-relevant ones is like finding a needle in a haystack.

Additionally, inspecting logs for software troubleshooting often requires engineers to possess descent knowledge about the software. However, modern software systems usually consist of many components developed by different engineers, leading to the generation of heterogeneous logs, which makes troubleshooting beyond the ability of a single engineer. Moreover, due to the high complexity of modern software systems, failures could stem from various sources of software and hardware issues. Examples include software bugs, hardware damage, OS crash, service exception, etc. In addition, promptly pinpointing to the root cause by inspecting logs highly relies on engineers' expertise and experience. However, such knowledge is often not well accumulated, organized, and documented. Therefore, sophisticated ways to conduct automatic log mining are in high demand.

6.2 A General Workflow of Log Mining

The general workflow of log mining is illustrated in Fig. 5, which mainly consists of four steps, *i.e.*, log partition, feature extraction, model training, and online deployment.

Fig. 5. A general workflow of log mining.

6.2.1 Log Partition. Modern software systems often adopt a microservice architecture and comprise a large number of modules operating in a multi-threaded environment. Different microservices or modules often aggregate their execution logs into a single log file, which hinders automated log mining. To tackle the problem, interleaved logs should be partitioned into different groups, each of which represents the execution of individual system tasks. The partitioned log group is the basic unit in extracting features before constructing log mining models. Moreover, studies [63] show that the way to partition logs impacts the log mining performance. In the literature, we observe that two elements often serve as the log partitioner, namely, *timestamp* and *log identifier*.

Timestamp. It records the occurrence time of each log, which is a fundamental feature supported by many logging libraries, such as Log4j. Through regexs, the timestamp with different formats can be easily extracted from raw logs in log parsing phase (Section 5). In general, two strategies are often adopted to conduct timestamp-based log partition, *i.e.*, fixed window and sliding window. Fixed window has a predefined window size, which means the time span or time interval (*e.g.*, 30 minutes) used to split chronologically sorted logs. Extended from the fixed window, sliding window allows the overlapping between two consecutive fixed windows. The sliding window has two attributes, *i.e.*, window size and step size. The step size indicates the forwarding distance of the window alone the time axis to generate log partitions, which is often smaller than the window size.

Log Identifier. It is a token identifying a series of related operations or message exchanges of the system. For instance, HDFS logs adopt *block_id* to record the operations (*e.g.*, allocation, replication, and deletion) on a specific block. Common log identifiers include *user ID*, *task/session/job ID*, *variable/component name*, etc., which can be extracted by log parsing. Compared to timestamp, log identifier is a clearer and more definite signal for partitioning logs. Therefore, many research studies [49] employed the log identifier for its ability in distinguishing logs of different task executions. However, the log identifiers might be non-unique in representing distinct system entities in different logs, such as virtual machines, physical machines, and networks [149]. This is because, for example, identifiers may not be propagated and synchronized across different services, one thread or process may serve more than one request through multiplexing, etc. Therefore, dedicated algorithms have been developed to address this problem and we elaborate on them later.

6.2.2 Feature Extraction. To analyze logs automatically, textual logs in a log partition should be transformed into appropriate formats that could fit machine learning algorithms. Particularly, through reviewing the literature, we identified two categories of log-based features, namely, numerical feature and graphical feature.

Numerical Feature. It represents log's statistical properties including numerical and categorical fields that can be directly extracted from logs. It conveys the information of a log partition into a numerical vector representation. In the following, we summarize the following four types of features, which are employed by most of the existing work.

- *Log event sequence*: A sequence of log events recording system's activities. Particularly, each element can simply be the log event ID or log embedding vector, *e.g.*, learned by word2vec algorithms [111].
- *Log event count vector*: A feature vector recording the log events occurrence in a log partition, where each feature denotes a log event type and the value counts the number of occurrence.
- *Parameter value vector*: A vector recording the value of parameters (*i.e.*, variables extracted by log parsing) that appear in logs.
- *Ad hoc features*: A set of relevant and representative features extracted from logs, which are defined using domain knowledge on the object software system and problem context.

Graphical Feature. To discover the hierarchical and sequential relations (*e.g.*, dependency and co-occurrence) between system components and events with logs, the graphical feature usually produces a directed graph model characterizing system behaviors, *e.g.*, the execution path of a process. The graphical features serve as the foundation for a variety of downstream log mining tasks (such as monitoring [149, 162] and diagnosis [4, 8, 116]). For example, log-based behavioral differencing [52] can identify system executions that are derived from system normal behaviors, which has various applications in system evolution, testing, and security [8]. In this section, we briefly introduce the algorithms used by existing work for graphical feature extraction, which can be roughly categorized into two lines of work.

The first line of work leverages the objects identified in logs, *e.g.*, process ID and system component name, to construct graphs for system state monitoring. As objects often demonstrate complicated hierarchical relations, it usually needs a more sophisticated algorithm for log partition. For instance, to represent execution structure and object hierarchy in logs, Zhao *et al.* [162] constructed a stack structure graph by considering the 1:1, 1:n, and n:m mappings among different objects. To tackle the challenge that unique identifiers are often unavailable, Yu *et al.* [149] proposed to group interleaved logs based on a common set of identifiers. Each log contains an identifier set that acts as a state node of the graph, and transitions are added by examining the subset and superset relations of different identifier sets.

The other line of work makes use of the underlying statistical distribution of log events, *e.g.*, the order of log events, the temporal and spatial locality of dependent log events, to build various graphs for system behavior analysis. Particularly, some work aims at recovering the exact behavioral model from log traces. Log partition therefore requires that identifiers can tie together the set of events associated with a program execution. For example, Amar *et al.* [4] presented two invariants of k-Tails [13], namely 2KDiff and nKDiff, to build Finite State Machine (FSM) for log differencing. 2KDiff computes and highlights the differences between two log files containing a set of partitioned logs; while nKDiff is able to conduct the comparison for multiple log files at once. Later Bao *et al.* [8] extended their work by proposing s2KDiff and snKDiff, which take into consideration the frequencies of different behaviors. Busany *et al.* [15] studied the scalability problem of existing behavioral log analysis algorithms by extracting Finite State Automaton (FSA) models or temporal properties from logs. Particularly, they suggested to analyze only a sample of log traces and provided statistical guarantees to support the correctness of the analysis result. We refer readers to this paper for more related studies.

However, as previously discussed, a unique identifier for each execution trace is not always available, especially in distributed applications and systems. Therefore, some work attempts to mine behavioral models using interleaved logs. For example, Lou *et al.* [96] employed statistical inference to learn the temporal dependencies among log events, which is also studied in similar work including [11, 12]. Nandi *et al.* [116] computed the nearest neighbor groups to capture the temporal co-occurrence of log events more accurately. The resulted model is a program Control
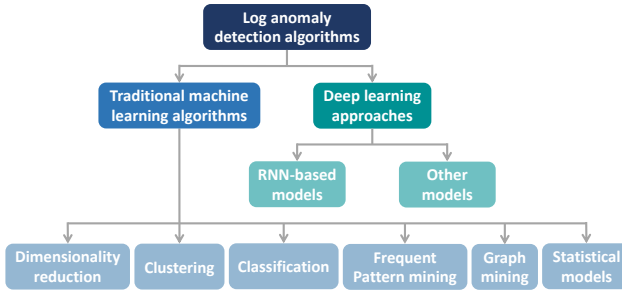
Fig. 6. A taxonomy of log anomaly detection algorithms.

Flow Graph (CFG) spanning distributed components. Particularly, the correlation between two components is calculated using either the Jaccard similarity or the Bayesian conditional probability approach. Du *et al.* [42] described two methods to capture service executions by FSA models. The first method leverages the trained log anomaly detection model, whose predictions encode the underlying path of task execution. The second method builds a matrix where each entry represents the co-occurrence probability of two log keys appearing together within a predefined distance.

6.2.3 Model Training. In this process, appropriate algorithms are selected based on the problem at hand and selected models are trained based on the extracted features. A variety of machine learning algorithms have been proposed, and we leave more details in the following sections. This step is often conducted in an offline manner. Moreover, a fundamental assumption of various log mining tasks is that the majority of logs should exhibit patterns that conform to system's normal behaviors. For example, in anomaly detection, different models are trained to capture various patterns from different perspectives and used to detect anomalies that lack the desireable properties.

6.2.4 Online Deployment. Once the model is trained offline, it could be deployed to the real-world software systems for various log mining tasks. For example, an anomaly detection model can be integrated into software products to detect malicious system behaviors and raise alarms in real time. Since software systems might be upgraded in their product life cycle, a previously trained model may not fit the pattern changes. Several studies [42, 88] supported online update to adapt to unprecedented log patterns.

6.3 Anomaly Detection

Log anomaly detection is the task of identifying system anomalous patterns that do not conform to expected behaviors on log data. In this section, we elaborate on existing research work in this field based on the algorithms adopted therein. As shown in Fig. 6, we categorize the algorithms into two broad classes: traditional machine learning algorithms and deep learning approaches. The surveyed approaches are listed in Table 4 together with some interesting properties. Specifically, we summarize the algorithm/model and feature used by the approach and whether or not each approach is unsupervised and online. Particularly, an unsupervised approach is one that does not require labels for model training; and an online approach is one that can process its input piece-by-piece in a streaming fashion, *i.e.*, prediction can be made without knowing the entire input from the start. In this paper, we mainly focus on general types of anomalies that are natural imperfections of software systems and will directly impact system's reliability. External attacks or intrusions (such as cyber attacks and malicious activities) that are more relevant to system's security go out of the scope for this paper.

Table 4. Summary of log anomaly detection approaches.

	Methods	Algorithm/Model	Feature	Unsupervised	Online
Traditional machine learning	Xu *et al.* [144]	PCA	★ †	Yes	No
	Lin *et al.* [88]	Clustering	∗	Yes	No
	He *et al.* [62]	Clustering	∗ ★	Yes	No
	Liang *et al.* [84]	SVM	‡	No	No
	Kimura *et al.* [77]	SVM	‡	No	No
	Xu *et al.* [143]	Frequent pattern mining	∗ ★	Yes	Yes
	Shang *et al.* [128]	Frequent pattern mining	∗	Yes	No
	Lou *et al.* [97]	Frequent pattern mining	★	Yes	No
	Farshchi *et al.* [47]	Frequent pattern mining	★	Yes	No
	Nandi *et al.* [116]	Graph mining	¶	Yes	No
	Lou *et al.* [96]	Graph mining	¶	Yes	No
	Yamanishi *et al.* [145]	Statistical model	∗	Yes	No
	He *et al.* [63]	Logistic regression	★	No	No
Deep learning	Du *et al.* [42]	LSTM model	∗ †	Yes	Yes
	Zhang *et al.* [158]	LSTM classification model	∗	No	No
	Meng *et al.* [107]	LSTM model	∗ ★	Yes	Yes
	Xia *et al.* [141]	LSTM-based GAN model	∗	Yes	Yes
	Lu *et al.* [100]	CNN model	∗	No	No
	Liu *et al.* [89]	Graph embedding model	¶	Yes	No

∗ Log event sequence, ★ Log event count vector, † Parameter value vector
‡ Ad hoc features, ¶ Graphical feature

6.3.1 Traditional Machine Learning Algorithms. Traditional machine learning algorithms usually perform on top of features that are explicitly provided by practitioners, *e.g.*, log event count vector. Particularly, the anomaly detection task can be formalized into different types and solved with different algorithms such as clustering, classification, regression, etc.

Dimensionality Reduction. It transforms data of high dimension into a low-dimensional space such that some meaningful properties of the original data can be retained in the low-dimensional representation. Principal Components Analysis (PCA) is one of the most popular algorithms of this kind. By calculating the projection of data points to the first k principal components, anomalies can be identified if the distance is larger than a threshold. PCA was first applied by Xu *et al.* [143, 144] for mining system problems from console logs. Particularly, two types of feature vectors, namely, log event count vector and parameter value vector (*i.e.*, denotes the ratio of state variables), are constructed and fed into a PCA model for anomaly detection.

Clustering. Clustering-based log anomaly detection is an unsupervised method that groups a set of log-based feature vectors in such a way that vectors in the same cluster are more similar to each other (and as dissimilar as possible to vectors from other clusters). Clusters that contain very few data instances tend to be anomalous. For example, Lin *et al.* [88] proposed LogCluster, which clusters log sequences and recommends a representative sequence to assist developers quickly identify potential problems. After clustering logs, the centroid of each cluster is selected as the representative log sequence. He *et al.* [62] proposed a novel framework called Log3C, which incorporates system KPIs into the identification of impactful problems in service systems. In

particular, they proposed Cascading Clustering to group numerous log sequences promptly and precisely. Finally, a multivariate linear regression model is used to identify impactful problems that lead to KPI degradation.

Classification. Anomaly detection with classification categorizes the log partition into either normal or anomalous type, where anomalous examples are derived from normal ones in terms of some statistical properties. Support Vector Machine (SVM) is a supervised classification method that is commonly used for log anomaly detection. In [84], Liang et al. vectorized log partitions by identifying six types of features, including events number in a time window, accumulated events number, etc. Based on these features, they applied four models for anomaly detection, *i.e.*, RIPPER (a rule-based classifier), SVM, a traditional nearest neighbor predictor, and a customerized one. In [77], Kimura et al. proposed a log analysis method for proactive failure detection based on logs' characteristics, including frequency, periodicity, burstiness, and correlation with maintenance and failures. SVM with Gaussian kernel is employed for detecting failures.

Frequent Pattern Mining. It aims at discovering the most frequent item sets and sub-sequences in a log dataset that characterize system's normal behaviors. Data instances that do not conform to the frequent patterns will be reported as anomalies. The presence of specific log events and the order of log events can both constitute patterns. For example, Xu *et al.* [143] mined sets of log messages that often co-occur to detect abnormal execution traces in an online setting. Compared to offline approaches, online pattern matching can quickly identify benign system executions, thus striking a balance between accuracy and efficiency. Similarly, Lim *et al.* [85] searched for item sets that were commonly associated with different failure types. Such item sets can assist developers in predicting and characterizing failures. Other approaches mine the sequential patterns by log events to discover anomalies or bugs [98, 128]. Shang *et al.* [128] proposed an approach to assist developers in maintaining big data analytics applications such as Hadoop. They examined the differences between execution log sequences produced in lab and in production environments.

Lou et al. [97] were the first to consider mining invariants among log messages for system anomaly detection. Two kinds of invariants that model the relations between the number of different log messages are derived: (1) invariants in textual logs that describe the equivalence relation; and (2) invariants as a linear equation, which is the linear independence relation. Farshchi *et al.* [47] proposed a similar approach which mines the correlation and causation relations between log events and cloud system metric changes. Specifically, they employed a regression-based approach to learn a set of assertions modeling the linear relations. Anomaly detection is then conducted by monitoring log event streams and checking the compliance of metrics against the assertions.

Graph Mining. The set of approaches mainly employs graphical features, *i.e.*, various graph models (Section 6.2), to identify behavioral changes of complex systems, which can be used for early detection of anomalies and allowing proactive actions for correction. For example, Nandi *et al.* [116] introduced a CFG mining approach to detect anomalous runtime behaviors of distributed applications from execution logs. Two types of anomalies are checked, *i.e.*, sequence anomaly and distribution anomaly. A sequence anomaly is raised when an expected child is missing for a parent node within the given time interval; a distribution anomaly is raised when an edge probability is violated. In [49], Fu et al. modeled the execution behaviors of each system module as a state transition graph by learning an FSA from log sequences. Each transition in the learned FSAs corresponds to a log key. For each state transition, two types of information (*i.e.*, the time consumed and circulation number) are recorded and applied to detect two performance issues: low transition time and low transition loop. By employing the Gaussian distribution to model the state transition in a distributed system, low-performance transitions can be automatically identified by setting a proper threshold.

Statistical Model. Some work detects a system's anomalous symptoms by resorting to statistical models. For example, Yamanishi *et al.* [145] employed a mixture of Hidden Markov Models to monitor syslog behaviors. Particularly, the model is learned with an online discounting learning algorithm by dynamically selecting an optimal number of mixture components. Anomaly scores are assigned using universal test statistics whose threshold can be dynamically optimized. In [63], He et al. employed a logistic regression model to detect anomalies. They adopted event count vectors as the feature and trained the model with a set of labeled data. A testing instance is declared as anomalous when the probability estimated by the logistic function is greater than 0.5.

6.3.2 Deep learning approaches. Deep learning uses an architecture of multiple layers, called neural networks, to progressively extract features from the raw input with different layers addressing different level of feature abstraction. As neural networks have exhibited exceptional ability in modeling complex relationships, they are widely applied in log-based anomaly detection. In literature, we observed a significant portion of papers adopting the Recurrent Neural Network (RNN) model or its variants. Therefore, we classify the deep learning approaches into RNN-based models and other models.

RNN-based Models. Models belonging to the RNN family (such as LSTM model and GRU model) are commonly used to automatically learn the sequential patterns in log data. Anomalies are raised when log patterns deviate from the model's predictions. For example, Du *et al.* [42] proposed DeepLog, which utilizes an LSTM model to learn system's normal execution patterns by predicting the next log event given a sequence of preceding log events. However, some anomalies may manifest themselves as an irregular parameter value instead of a deviation from a normal execution path. Hence, DeepLog also applies an LSTM model for checking the validity of parameter value vectors. Many existing studies have a closed-world hypothesis, which assumes that the log data are stable over time and the set of distinct log events is fixed and known. However, Zhang *et al.* [158] found that log data often contain previously unseen log events or log sequences, demonstrating log instability. To tackle this problem, they proposed LogRobust to extract the semantic information of log events by leveraging off-the-shelf word vectors. Then, a bidirectional LSTM model is used to detect anomalies.

In terms of capturing log's semantics, Meng *et al.* [107] found existing word2vec models did not distinguish well between synonyms and antonyms. Therefore, they trained a word embedding model to explicitly consider the information of synonyms and antonyms. Meng *et al.* [106] further extended their work by proposing a semantic-aware representation framework for online log analysis. The issues of log-specific word embedding and out-of-vocabulary (OOV) are both addressed. Recently, Zuo *et al.* [167] combined the transaction-level topic modeling for learning the embedding of logs, where a transaction is a group of logs sequentially occurring in a time window. To address the problem of insufficient labels, Chen *et al.* [22] applied transfer learning to share anomalous knowledge between two software systems. Specifically, they first trained an LSTM model on the data with sufficient anomaly labels to extract sequential log features, which were then fed into fully connected layers for anomaly classification. Next, the LSTM model was fine-tuned with logs from another system with limited labels, while the fully connected layers were fixed.

Other Deep Learning Models. Besides the RNN-based models, other model architectures also play a role in detecting anomalies with logs. For example Xia *et al.* [141] proposed LogGAN, an LSTM-based Generative Adversarial Network (GAN). Like all GAN-style models, LogGAN comprises a generator and a discriminator. The generator attempts to capture the distribution of real training data and synthesizes plausible examples; while the discriminator tries to distinguish fake instances from real and synthetic data. Effort is also devoted to explore the feasibility of Convolutional Neural Networks (CNN) for anomaly detection. Specifically, Lu *et al.* [100] first

utilized a word embedding technique to encode logs into two-dimension feature matrices, upon which CNN models with different filters were then applied for anomaly detection. Liu *et al.* [89] proposed log2vec, a graph embedding based method for cyber threat detection. Specifically, they first converted logs into a heterogeneous graph using heuristic rules and then learned the embedding of each log entry by a graph representation learning approach. Based on the embedding vectors, logs were grouped into clusters, whose size smaller than a threshold would be reported as malicious.

6.4 Failure Prediction

When software systems deviate from fulfilling a required system function [80], a failure occurs and it often shows human-perceivable symptoms. The failure could lead to unintended results and user dissatisfaction, especially for large-scale software systems. Traditional ways of failure management are mostly passive, which deal with failures after they have happened. In contrast, failure prediction aims to proactively predict the failure before it happens. A common practice is to leverage the valuable logs to proactively predict failures. For example, in [157], switch failures in data center networks are predicted from current status and curated historical hardware failure cases. Usually, input data of the predictive model are system logs, which record the system status, changes in configuration, and operational maintenance, etc.

According to the source of failures, failure prediction could mainly be categorized into two scenarios: a) Prediction of independent failures in homogeneous systems, *e.g.*, high-performance computing (HPC) systems. b) Prediction of outage caused by a collection of heterogeneous devices or components, which is widely observed in large-scale cloud systems.

6.4.1 Homogeneous Systems. A typical homogeneous system is large-scale supercomputers which may encounter faults (e.g., component failures) every day. Exascale systems are expected to combat with higher fault rates due to enormous number of system components and higher workload [35]. However, triggering resilience-mitigating mechanisms is difficult since there are no obvious, well-defined failure indicators, which rely on a deeper understanding of the faults caused by hardware and software components. Therefore, logs become the most reliable information source to monitor the system health. Furthermore, to catch early warning signs of these failures, there are additional applications that scan through the monitoring data from the system logs and proactively generate alerts. Some of these alerts are generated from the manually defined thresholds on the collected system performance data (*e.g.*, CPU utilization rate) or error logs (*e.g.*, certain types of errors appearing too frequently within some time range). The objective of such alerts is to provide early warnings so that protective actions (*e.g.*, virtual machine migration or software rejuvenation) can be performed to mitigate or minimize the impact of such failures. In this context, efficient failure prediction via system log mining can enable proactive recovery mechanisms to increase reliability.

In homogeneous systems, component failures or node soft lock-ups typically lead to crashes of user jobs scheduled on the affected nodes, and they may cause undesired downtime. One approach to mitigate such problems is to predict node failures with a sufficient lead time in order to take proactive measures. Sahoo *et al.* [125] first addressed this problem by collecting system logs representing the components' health status. Then, several time series models are employed to predict the system health of each node through indication metrics, such as the percentage of system utilization, usage of network IO, and system idle time. Das *et al.* [35] proposed Desh to identify failure indicators in each node with enhanced training and classification for failure event log chain. Desh is a deep learning-based method that contains three phases: (1) It first recognizes chains of log events leading to a failure. (2) Then it re-trains chain recognition of events augmented with expected lead times to failure; (3) Finally Desh predicts lead times during testing/inference deployment to predict which specific node fails in how many minutes. To speedup the recognition

of the failure chain from log events, Das *et al.* [34] proposed a new node failure predicting method called Aarohi to extend their previous work to online learning setting. Aarohi first trains an offline deep learning model with log parsing, then utilizes grammar-based rules to provide online testing. Another approach from feature engineering tries to extract specific patterns about a certain time frame before a critical event. This anomaly pattern can serve as a proactive alerts. Klinkenberg *et al.* [78] adopted descriptive statistics and supervised machine learning techniques to create such proactive alert from the system logs. They trained a binary classification model to detect the potential node failure based on a given time sequence of monitoring data collected from each node. Berrocal *et al.* [10] used environmental logs to extract numerical indicators and conducted a Void Search (VS) algorithm on these numerical values for the failure prediction task.

6.4.2   Heterogeneous Systems. Different from homogeneous systems which focus on predicting failures based on raw signals sealed in the event logs or trace logs, the prediction of critical failures (also called outage) in heterogeneous systems relies on the failure signals from log information collected from diversified system components. In current real-time heterogeneous systems, such as the cloud system, outage prediction is an important and challenging task to perform. On one hand, outages are critical system failures that could lead to severe consequences. On the other hand, outages occur without a significant alerting signals pattern, which makes it hard to predict. Moreover, the scope of impacting signals is a complex process to define. Therefore, compared to homogeneous system settings, it is more valuable to detect node outages and distinguish regular internal failures from those caused by external factors, such as maintenance, human errors, and others.

From the perspective of methodology, the advancement of failure prediction in heterogeneous systems relies on effective detection of the relationship between early alerting signals from system logs and node/component outage. Chen *et al.* [23] proposed AirAlert framework based on Bayesian network to find the conditional dependence between the alerting signal extracted from system logs and the outage. Then, given several alert signals, AirAlert employed gradient boosting tree method to conduct outage prediction. Lin *et al.* [87] designed MING framework to find the relationship between complex failure-indicating alerts and the outage from temporal and spatial features. They carefully designed a ranking model that combines the intermediate results from LSTM model capturing temporal signals and a Random Forest model incorporating spatial data. Zhou *et al.* [164] proposed MEPFL to narrow down the scope of failure prediction into microservice level. Based on a set of manual selected features defined on the system trace logs, MEPFL trains prediction models at both the trace level and the microservice level to predict three common types of microservice application faults: multi-instance faults, configuration faults, and asynchronous interaction faults.

6.5   Failure Diagnosis

Although anomaly detection and failure prediction can pinpoint whether a problem occurs or will occur, there is a huge gap between the detection and removal of a problem or a failure. To completely resolve the problems, failure diagnosis is a crucial step, but the diagnosis process is notoriously expensive and inefficient. It is reported that failure diagnosis takes over 100 billion dollars, and developers spend more than half of their time on debugging [159]. Following the concept of error, fault, and failure, as defined in [80], failure diagnosis aims to identify the fault that has led to the user-perceived impairment in a software system. Log-based failure diagnosis is now a standard practice for software developers. However, failure diagnosis is very challenging. The complexity of modern software systems grows rapidly, where different services, software, and hardware are tightly coupled. It is too complex to correctly and efficiently disentangle the relations among the fault, failure, and human-observed symptoms. Moreover, as software systems become

more mature, failures are becoming more and more hard to detect and diagnose, from perceptible software functionality issues to imperceptible problems [74], *e.g.*, performance issues [1, 109, 148].

To tackle these challenges, in recent decades, techniques to automate the diagnosis process have been widely developed, for example, using the informative logs [27, 68, 164]. Jiang *et al.* [69] provided one of the first characteristic studies of log-based problem troubleshooting on real-world cases. They concluded that problem troubleshooting is a time-consuming and challenging task, which can be significantly facilitated by logs. In addition, they appealed to engineers to employ automated method to speedup the problem resolution time, which is also the focus of this survey. Likewise, Zhou *et al.* [163] empirically investigated the faults and debugging practices in microservice systems. The results show that proper tracing and visualization techniques can improve the diagnosis, which also suggest the strong need for more intelligent log analysis. Other empirical studies include understanding failures with logs in high-performance computing (HPC) systems [43] and cloud management systems [32], and analyzing the failure reporting capability of logs in an industrial air traffic control system [28].

In this section, we review recent work on automated log-based failure diagnosis. Particularly, these research studies mainly focus on the large-scale software systems where failure diagnosis is burdensome, including the software applications [148, 150], general distributed systems [16, 51, 68, 98, 115], storage systems [69], microservice systems [163, 164], big data systems such as Hadoop [128] and Spark [99], HPC systems [26, 27, 43], and cloud computing systems [65, 109, 120, 134, 154], such as Openstack. As the characteristic and functionality may vary, these systems demonstrate different failure behaviors. However, the overall methodology for failure diagnosis still shares similar techniques, which can be categorized into four types, *i.e.*, execution replay, model-based, statistics-based, and retrieval-based methods. There are also studies on diagnosing the hardware fault by traces (*e.g.*, [81]), which are beyond the score of this survey.

6.5.1 Execution Replay Methods. Traditional rule-based methods heavily rely on a set of predefined rules (*e.g.*, in the format of "if-then") from the expert knowledge to diagnose failures. However, the methods cannot be well generalized to unseen failures that are not included in the rules. On the contrary, execution replay methods aim to automatically infer the execution flow from logs and trace back to the software system failure, as previously introduced in Section 6.2. The set of methods is human-interpretable and automated.

Yuan *et al.* [150] proposed the SherLog, which analyzes source code by run-time logs to represent the detailed execution process of a failure. Specifically, SherLog infers both control and data value information of a failed execution to help developers understand the failure. Similarly, LogMap [16] first retrieves log messages from bug reports and then applies static analysis technique to identify corresponding logging lines in source code. At last, it traverses through the logging lines to derive the potential code paths. By following the method, we can reconstruct the execution path and assist the debugging process.

6.5.2 Model-based Methods. Model-based methods utilize logs to build the reference model (*e.g.*, execution path) for a software system and then check which log events violate the reference model. In short, it sets up standards for normal software system executions and diagnoses failures by detecting the possible inconsistency.

The majority of existing model-based studies leverage the log information to reconstruct the system execution flow as a graph representation [6, 66, 67]. To achieve so, Jia et al. proposed to mine a time-weighted control flow graphs (TCFG) [67] and the service topology [66] from interleaved logs during the offline phase. In the online phase, a failure can be easily diagnosed by observing the deviation between execution log sequences and the mined graph models. After observing that log sequences generated by a normal execution are consistent across multiple runs, Tak *et*

al. [134] proposed the LOGAN to capture normal log patterns by log grouping, log parsing, and log alignment. When a failure occurs, LOGAN highlights the divergence of current log sequence from a reference model and suggests the possible root cause. Different from these research tasks, to detect concurrency bugs in distributed systems, Lu *et al.* [98] proposed to mine logs from historical executions to uncover feasible but undertested log message orders. The log message orders represent possible execution flow that are likely to expose errors.

6.5.3   *Statistics-based Methods.* Since software systems generate logs to record normal and abnormal executions, it is intuitive to employ some statistical techniques (*e.g.*, statistical distribution, correlation analysis, clustering) to capture the relationship between the log information and the consequent failures.

Chuah *et al.* [27] developed a diagnostics tool, FDiag, which parses the log messages and employs statistical correlation analysis to attribute the observed failure to a possible root cause. However, the proposed method limits its diagnostic capability to known failures. To overcome the difficulty, Chuah *et al.* [26] further extended their approach to an advanced FDiagV3 method, in which a PCA and ICA-based correlation approach is newly added. Unlike the previous FDiag method, FDiagV3 can identify those unknown failures by automatically extracting the outlier issues. In [51], a three-step approach is designed to establish the causal dependency graph through log events, which help identify the process that a failure occurs. The general idea resembles execution replay methods, but the causal dependency graph is statistically mined by grouping similar log events and utilizing the temporal order information. Similarly, Yu *et al.* [148] designed a two-phase approach to comprehend performance problem of device drivers in Windows by narrowing down the diagnosis scope and pattern mining. To achieve so, it first measures the impact of suspicious components on performance by extracting the graph from execution traces. Then, by representing the behavior pattern with the signature set tuple, thresholds are set to identify highly suspicious and high-impact components that are likely to cause performance problems.

Lu *et al.* [99] proposed to detect and diagnose anomalies in the Spark system after extracting information from logs. In complementary to previous studies, the proposed method builds statistical models from the data distribution of several task-related features. It then diagnoses the anomaly by setting threshold and analyzes the root causes with weighted factors. Simlar to the method in anomaly detection [144], SCMiner [3] firstly utilizes a PCA method to detect abnormal system call sequences. Then, it maps the abnormal sequence to specific application functions, which are mined by frequent pattern mining based on system call traces using binary instrumentation. At last, SCMiner identifies and ranks buggy functions by matching with the function call traces. Likewise, CloudDiag [109] employs a statistical technique (*i.e.*, based on the data distribution) to identify the category of a fine-grained performance problem. Then, a fast matrix recovery algorithm, RPCA, is adopted to identify the root cause (*i.e.*, method invocations) for the performance failure.

6.5.4   *Retrieval-based Methods.* In practice, failures that previously occurred are valuable since they can aid developers in better diagnosing newly-occurred failures. Retrieval-based methods, as indicated by the name, retrieve similar failures in a knowledge base composed of failures in history or populated by injected faults in the test environment.

Shang *et al.* [128] focused on diagnosing big data analytics applications in Hadoop system by injecting failures manually and analyzing the logs. Particularly, they proposed to uncover differences between pseudo and cloud deployments by log parsing, execution sequence recovery, and sequence comparison. Pham *et al.* [120] adopted a similar idea but they targeted on the general distributed system. Under the assumption that similar faults generate similar failures, we can locate the root cause of a reported failure by inspecting matched failures in the knowledge base. The method first reconstructs execution flows between system components, computes the similarity of the

reconstructed flows, and performs precise fault injection. Nagaraj *et al.* [115] proposed to diagnose the performance failures by comparing the logs of system behaviors in good and bad performance. Based on the two collected sets of logs, they applied a number of machine learning techniques to automatically compare and infer the strongest associations between performance and system components, where failures were likely to happen.

CAM [68] applies the same idea to the cause analysis for test alarms in system and integration testing. In detail, the failure matching is achieved by using the K nearest neighbors (KNN) algorithm to find similar attribute vectors, which are built on test log terms extracted by term frequencyâĂŞinverse document frequency (TF-IDF). Furthermore, Amar et al [3] extended the CAM by keeping failing logs and removing logs that passed the test. Then, the most relevant logs in historical failures were extracted by a modified TF-IDF and then vectorized. The vectors were utilized to train an exclusive version of KNN to identify possible log lines that led to the failure. Similarly, Yuan *et al.* [154] facilitated the failure diagnosis by leveraging historical failures. They built four classifiers by vectorizing historical failures logs using a natural language processing tool word2vec. When a new failure occurs, the corresponding classifier was employed to identify the root cause.

Moreover, some failures are indistinguishable because they generate very similar logs. To tackle the problem, Ikeuchi et al. [65] proposed a two-phrase framework with user actions participated. In the learning phase, a log-based failure database is constructed by performing various user actions. In the inferring phase, when a user action is performed, operators match the failures with the database and obtain additional informative log data.

6.6 Others

There are other valuable topics to broaden the log mining literature. In this section, we mainly introduce two reliability engineering tasks: duplicate issue mining and log-based software testing.

Duplicate Issue Mining. Developers may submit the same issue or bug report to the repository. These duplicate issue reports should be detected when new report arrives. Otherwise, issue maintainer will encounter heavy but unnecessary workload to process these reports. Wang *et al.* [140] uses classification techniques to detect duplicate issues with higher accuracy from informative execution traces as a reliable bug information source. Ding *et al.* [38] transferred the duplicate-issue detection problem to similar context mining from the history records. They generated a signature for each issue from its corresponding transaction logs and then retrieved historical issues with similar signatures from a historical issue repository.

Log-based software testing. Log-based software testing aims to stimulate the overall software workflow under testing by the generated log data. Many existing work solves this problem by proposing a variety of state-machine-based formalism, similar to the techniques introduced in Section 6.2. A pioneer work [5] discussed the availability of using formal log analyzer to enhance software reliability. The majority of log-based testing approaches utilize cues from log files to improve testing coverage or completeness, Chen *et al.* [20] designed a framework called LogCoCo to estimate code coverage measures using the readily available execution logs. Chen *et al.* [21] derived the state machine framework to conduct workload testing by extracting representative user behaviors from large-scale systems. Cohen et al. [31] proposed a concept called log confidence to evaluate the completeness of any log-based dynamic specification mining task.

7 OPEN-SOURCE TOOLKITS AND DATASETS

In recent years, there are many research studies focusing on automated log analysis. At the same time, a number of log management toolkits have been developed to facilitate the real-world

Table 5. A list of open-source log analysis tools and log datasets.

	Category	Description	Link
Tools	GrayLog	A centralized log collection and real-time analysis tool.	link
	GoAccess	A real-time terminal-based log analyzer for web logs.	link
	Fluentd	A data collector that unifies multiple logging infrastructures.	link
	Logstash	A server-side processor that ingests and transforms log data.	link
	Logalyze	A centralized log management and network monitoring software.	link
	Prometheus	A monitoring solution for metrics collection and alerting.	link
	Syslog-ng	A scalable log collector and processor across infrastructures.	link
	LogPAI	A comprehensive automated log analysis toolkit	link
Datasets	Loghub	A large collection of log datasets from various systems.	link
	Failure dataset	Error logs produced by OpenStack cloud system.	link
	CFDR	Computer failure data repository of different systems.	link
	SecRepo	A list of security log data such as malware and system logs.	link
	EDGAR	Apache log files that record and store user access statistic.	link

deployment in industry. In this section, we collected a set of open-source toolkits and datasets for automated log analysis. Besides, we did not consider log analysis tools that are mainly for enterprise use, such as Splunk [133] and Loggly [94]. After carefully studying the related tools, we curated a list of open-source log analysis toolkits, as shown in Table 5. These log analysis tools are GrayLog, GoAccess, Fluentd, Logstash, Logalyze, Prometheus, Syslog-ng, and LogPAI. Among them, LogPAI toolkit is an open-source project built upon pure research outputs while other tools are packed as enterprise products. By observing these toolkits, we listed some common and crucial features that they support in industrial deployment, including log collection, searching, routing, parsing, visualization, alerting and automated analysis. Among all these log analysis tools, most of them focus primarily on phases before the log mining, and the resulting outcomes (*e.g.*, visualization) are presented to developers for inspecting and diagnosis. Particularly, LogPAI provides a set of automated log analysis tools such as log parsing, anomaly detection and problem identification.

Besides the open-source toolkits, a number of log datasets collected from various software systems are available for research studies in academia. Table 5 shows the list of public log datasets and their descriptions. The Loghub [64] maintains a large log collection that is freely accessible for research purposes. Some log datasets are production data released from previous studies, while others are collected in the lab environment. Another research study [32] presented error logs produced by injecting failures in the OpenStack cloud management system. The computer failure data repository (CFDR) focuses mainly on component failure logs in a variety of large production systems. Besides these log datasets, there are also logs collected for other purposes, such as security and web search. SecRepo is a curated list of security data such as malware and system logs. The EDGAR dataset contains Apache logs that record user access statistic through internet search.

8 FUTURE DIRECTIONS

In this section, we would like to discuss promising future directions. Particularly, we start with concluding future directions for each topic reviewed in this survey. Then we share our thoughts on the next-generation log analysis framework, which might shed light on future studies.

8.1 Logging

As the foundations of automated log analysis, current logging practices mostly focus on characterization and recommendation of logging (*i.e.*, where, what, and how to log). We could take a step further to analysis-oriented and automated logging.

Analysis Oriented Logging. Section 6 has demonstrated that the value of logs in automatically mining patterns and diagnosing failures. A recent industrial study [9] pointed out that the inconsistent presentation format with-in and among teams poses significant challenges in developing automated log analysis tools. Hence, an interesting research direction is to design advanced logging mechanisms that can coordinate with the subsequent log analysis steps during the development.

Automated Generation of Logging Statements. The automated generation of logging statements can alleviate developers' burden. Recent research at the intersection of machine learning, programming languages, and software engineering has taken an important step toward proposing learnable code-generating models that utilize rich patterns in source code [2]. Automatically generating logging statements is approachable by incorporating code-generating models with existing approaches that suggest where, what, and how to log.

8.2 Log Compression

Although log compression has been studied for decades, its performance is still unsatisfactory in varying storage and practical use demands. We summarize the future directions on automated log compression and query-efficient log compression.

Automated Log Compression. Existing log compressors often require many heuristic rules and cannot be generalized well to other log data [7, 37, 55], which makes a general log compressor inefficiently designed. To tackle the problem, a fully automated log compression algorithm is highly in demand. It should be automatically adapted for different log formats and require few heuristic rules. One possible solution is to use machine learning techniques to uncover the redundant log patterns and then compress the redundancy respectively.

Query-efficient Log Compression. Different logs may play different roles and have different storage requirements. For example, failure logs could be frequently queried by developers and the corresponding log compressor should promptly respond to the query request. Therefore, how to incorporate the query demands into the log compression becomes vital to log service providers. Particularly, several sub-problems might be: how to predict the user query request, how to store logs in a systematically modularized way for different demands, etc.

8.3 Log Parsing

Although there have been much progress in log parsing in recent years, we find it still have a lot of room for further improvement. In this section, we present our insights and suggestions on heterogeneous log parsing and distributed online log parsing.

Heterogeneous Log Parsing. Most existing log parsers assume that log messages of a system should have the same headers (*e.g.*, timestamp and component name), which are often extracted before hand by ad-hoc scripts. However, modern software is very complex and often adopts different logging frameworks, leading to a mixture of log messages with different headers, consequently degrading the performance of existing approaches. Therefore, a log parsing method that can automatically extract the header of log messages is called for.

Distributed Online Log Parsing. Existing approaches [41, 59] provide offline parallel mechanisms to accelerate the log parsing process. However, there is a lack of distributed online parsing solution in large-scale distributed systems. The desired log parser should have multiple running threads in either the master or the working node and parse the log messages in a streaming manner. With a distributed online parser, we could generate structured log messages on the fly, which facilities both log compression and log mining tasks.

8.4 Log Mining

Although being extensively studied, log mining is still very challenging due to the complex problem definition and poor data quality in different scenarios. Unlike other research topics that are relatively mature, log mining still has a lot of space to explore. We summarize the following insights encountered during the log data mining and propose some inspiring future directions.

Concept Drift-adapted Log Mining. Most of the existing approaches train models with historical log data offline. However, modern software systems continuously undergo feature upgrades and system renewal; hence, the patterns of logs may drift accordingly [158]. Common strategies of online updating, *e.g.*, periodically retraining the model, suffer from high false positive rates, and strategies like zero-positive learning [40] have limited generalization ability. Online learning and incremental learning that fit the log mining scenario are appealing research directions.

Log Mining with Semantic Understanding. Logs, as a structured language [58], inherit abundant semantic information. Existing studies often focus on the event representation while ignoring the semantic information embedded in logs. Recently, efforts [106, 107, 158] have been devoted towards the automatic understanding of logs, which would be effective ways to eliminate the problems brought by log concept drift. Moreover, it can also play an important role in the automated generation of logging statements.

Interpretable Log Mining. Whether a log analysis algorithm can provide interpretable results is vital for administrator and analysts to trust and act on the automated analysis. Some traditional machine learning algorithms possess the merit of interpretability [97, 116, 128, 143]. However, deep learning models, though having achieved impressive performance, are criticized as black-box oracles. Therefore, the community of log analysis should continue pursuing more explainable and trustworthy algorithms, for example, by utilizing attention mechanisms [14].

8.5 Next-generation log analysis framework

Log analysis, as a core component in Artificial Intelligence for IT Operations (AIOps), is crucial to the reliability engineering of software systems. Particularly, the demand for automated log analysis draws dramatic attention as the increasing needs of cloud computing. We believe that the next-generation log analysis would be the "autopilot" in the cloud environment. We list as below two major directions that the next-generation log analysis framework might contain.

Intelligent Log Analysis. Current automated log analysis contains several consecutive steps, but together they are not highly intelligent. For example, they do not form an end-to-end framework, indicating the possible loss of some critical information in processing logs. The next generation log analysis should be an intelligent and end-to-end framework. In this framework, logging statements are automatically suggested by considering the failure diagnosis demand. Then, the framework intelligently stores raw logs following practical requirements. At last, logs are processed without information loss and correlated with human-observable failures for diagnosis using advanced techniques. Moreover, the framework would automatically generate failure reports by aggregating

essential information including all analyzed system logs together. In addition, it should auto-repair the detected failures by observing experiences from human experts.

Human-computer Interactive Log Analysis. Current automated log analysis research facilitates the failure diagnosis process by providing insightful decisions and clues for root cause analysis. In essence, developers are still at the core of failure diagnosis. For example, the diagnosis results should be confirmed and then debugged by developers. However, there is a lack of interactions between advanced analysis techniques and developers-in-charge. Besides, the diagnosis process and results are seldom reused in the future log analysis. Therefore, it is interesting and promising to involve developers into the loop of automated log analysis. We could build an interactive dialogue system with automated analysis tools. The automated log analysis tools first provide suggestions to developers for decision making. Then, the developers can request for additional information of failures, and the diagnosis process with human involvement can be recorded and reused to improve the automated log analysis tools.

9 CONCLUSION

Recent years have witnessed the blossom of log analysis. Given the importance of software logs in reliability engineering, extensive efforts have been contributed to efficient and effective log analysis. This survey mainly explores the four main steps in automated log analysis framework: logging, log compression, log parsing, and log mining. Additionally, we introduce the available open-source toolkits and datasets. Our article enables the outsiders to step in this promising and practical field, and allows the experts to fill in the gaps of their knowledge background. Based on the investigation of these recent advances, we proposed new insights and discussed several future directions, including how to make automated log analysis more feasible under the modern agile and distributed development style, and the concept of a next-generation log analysis framework.

A ELECTRONIC APPENDIX

REFERENCES

[1] Marcos K Aguilera, Jeffrey C Mogul, Janet L Wiener, Patrick Reynolds, and Athicha Muthitacharoen. 2003. Performance debugging for distributed systems of black boxes. *ACM SIGOPS Operating Systems Review* 37, 5 (2003), 74–89.

[2] Miltiadis Allamanis, Earl T. Barr, Premkumar T. Devanbu, and Charles A. Sutton. 2018. A survey of machine learning for big code and naturalness. *ACM Computing Survey* (2018), 81:1–81:37.

[3] Anunay Amar and Peter C Rigby. 2019. Mining historical test logs to predict bugs and localize faults in the test logs. In *Proc. of the IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. 140–151.

[4] Hen Amar, Lingfeng Bao, Nimrod Busany, David Lo, and Shahar Maoz. 2018. Using finite-state models for log differencing. In *Proc. of the 26th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. 49–59.

[5] James H. Andrews. 1998. Testing using log file analysis: tools, methods, and issues. In *Proc. of the thirteenth IEEE conference on automated software engineering ASE*. 157.

[6] Anton Babenko, Leonardo Mariani, and Fabrizio Pastore. 2009. AVA: automated interpretation of dynamically detected anomalies. In *Proc. of the eighteenth international symposium on Software testing and analysis (ISSTA)*. 237–248.

[7] Raju Balakrishnan and Ramendra K Sahoo. 2006. Lossless compression for large scale cluster logs. In *Proc. 20th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*. 7–pp.

[8] Lingfeng Bao, Nimrod Busany, David Lo, and Shahar Maoz. 2019. Statistical log differencing. In *Proc. of the 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 851–862.

[9] Titus Barik, Robert DeLine, Steven Mark Drucker, and Danyel Fisher. 2016. The bones of the system: a case study of logging and telemetry at Microsoft. In *Proc. of the 38th International Conference on Software Engineering (ICSE)*. 92–101.

[10] Eduardo Berrocal, Li Yu, Sean Wallace, Michael E. Papka, and Zhiling Lan. 2014. Exploring void search for fault detection on extreme scale systems. In *Proc. of the IEEE international conference on cluster computing (CLUSTER)*. 1–9.

[11] Ivan Beschastnikh, Yuriy Brun, Michael D Ernst, and Arvind Krishnamurthy. 2014. Inferring models of concurrent systems from logs of their behavior with CSight. In *Proc. of the IEEE/ACM 36th International Conference on Software

Engineering (ICSE). 468–479.

[12] Ivan Beschastnikh, Yuriy Brun, Sigurd Schneider, Michael Sloan, and Michael D Ernst. 2011. Leveraging existing instrumentation to automatically infer invariant-constrained models. In *Proc. of the 19th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. 267–277.

[13] Alan W Biermann and Jerome A Feldman. 1972. On the synthesis of finite-state machines from samples of their behavior. *IEEE transactions on Computers* (1972), 592–597.

[14] Andy Brown, Aaron Tuor, Brian Hutchinson, and Nicole Nichols. 2018. Recurrent neural network attention mechanisms for interpretable system log anomaly detection. In *Proc. of the First Workshop on Machine Learning for Computing Systems*. 1–8.

[15] Nimrod Busany and Shahar Maoz. 2016. Behavioral log analysis with statistical guarantees. In *Proc. of the IEEE/ACM 38th International Conference on Software Engineering (ICSE)*. 877–887.

[16] An Ran Chen. 2019. An empirical study on leveraging logs for debugging production failures. In *Proc. of the IEEE/ACM 41st International Conference on Software Engineering: Companion (ICSE-Companion)*. 126–128.

[17] Boyuan Chen and Zhen Ming (Jack) Jiang. 2017. Characterizing and detecting anti-patterns in the logging code. In *Proceedings of the 39th International Conference on Software Engineering (ICSE)*. 71–81.

[18] Boyuan Chen and Zhen Ming (Jack) Jiang. 2017. Characterizing logging practices in Java-based open source software projects - a replication study in Apache Software Foundation. *Empirical Software Engineering* (2017), 330–374.

[19] Boyuan Chen and Zhen Ming (Jack) Jiang. 2020. Studying the use of java logging utilities in the wild. In *Proc. of the 42nd International Conference on Software Engineering (ICSE)*.

[20] Boyuan Chen, Jian Song, Peng Xu, Xing Hu, and Zhen Ming Jiang. 2018. An automated approach to estimating code coverage measures via execution logs. In *Proc. of the 33rd ACM/IEEE International Conference on Automated Software Engineering (ASE)*. 305–316.

[21] Jinfu Chen, Weiyi Shang, Ahmed E. Hassan, Yong Wang, and Jiangbin Lin. 2019. An experience report of generating load tests using log-recovered workloads at varying granularities of user behaviour. In *Proc. of the 34th IEEE/ACM international conference on automated software engineering (ASE)*. 669–681.

[22] Rui Chen, Shenglin Zhang, Dongwen Li, Yuzhe Zhang, Fangrui Guo, Weibin Meng, Dan Pei, Yuzhi Zhang, Xu Chen, and Yuqing Liu. 2020. Cross-system log anomaly detection for software systems. In *Proc. of the 31th International Symposium on Software Reliability Engineering (ISSRE)*.

[23] Yujun Chen, Yong Xu, Hao Li, Yu Kang, Xian Yang, Qingwei Lin, Hongyu Zhang, Feng Gao, Zhangwei Xu, Yingnong Dang, Dongmei Zhang, and Hang Dong. 2019. Outage prediction and diagnosis for cloud service systems. In *Proc. of the World Wide Web Conference (WWW)*. 2659–2665.

[24] Shaiful Alam Chowdhury, Silvia Di Nardo, Abram Hindle, and Zhen Ming (Jack) Jiang. 2018. An exploratory study on assessing the energy impact of logging on Android applications. *Empirical Software Engineering* (2018), 1422–1456.

[25] Robert Christensen and Feifei Li. 2013. Adaptive log compression for massive log data. In *Proc. of the SIGMOD Conference*. 1283–1284.

[26] Edward Chuah, Arshad Jhumka, James C Browne, Bill Barth, and Sai Narasimhamurthy. 2015. Insights into the diagnosis of system failures from cluster message logs. In *Proc. of European Dependable Computing Conference (EDCC)*. 225–232.

[27] Edward Chuah, Shyh-hao Kuo, Paul Hiew, William-Chandra Tjhi, Gary Lee, John Hammond, Marek T Michalewicz, Terence Hung, and James C Browne. 2010. Diagnosing the root-causes of failures from cluster log files. In *Proc. of the International Conference on High Performance Computing (HPC)*. 1–10.

[28] Marcello Cinque, Domenico Cotroneo, Raffaele Della Corte, and Antonio Pecchia. 2014. What logs should you look at when an application fails? insights from an industrial case study. In *Proc. of the 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 690–695.

[29] Marcello Cinque, Domenico Cotroneo, Roberto Natella, and Antonio Pecchia. 2010. Assessing and improving the effectiveness of logs for the analysis of software faults. In *Proc. of the 2010 IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 457–466.

[30] Marcello Cinque, Domenico Cotroneo, and Antonio Pecchia. 2013. Event logs for the analysis of software failures: a rule-based approach. *IEEE Transactions on Software Engineering* (2013), 806–821.

[31] Hila Cohen and Shahar Maoz. 2015. Have we seen enough traces?. In *Proc. of the 30th IEEE/ACM international conference on automated software engineering (ASE)*. 93–103.

[32] Domenico Cotroneo, Luigi De Simone, Pietro Liguori, Roberto Natella, and Nematollah Bidokhti. 2019. How bad can a bug get? an empirical analysis of software failures in the OpenStack cloud computing platform. In *Proc. of the 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. 200–211.

[33] Hetong Dai, Heng Li, Weiyi Shang, Tse-Hsun Chen, and Che-Shao Chen. 2020. Logram: efficient log parsing using n-gram dictionaries. *IEEE Transactions on Software Engineering* (2020).

[34] Anwesha Das, Frank Mueller, and Barry Rountree. 2020. Aarohi: making real-time node failure prediction feasible. In *Proc. pf the IEEE international parallel and distributed processing symposium (IPDPS)*. 1092–1101.

[35] Anwesha Das, Frank Mueller, Charles Siegel, and Abhinav Vishnu. 2018. Desh: deep learning for system health prediction of lead times to failure in HPC. In *Proc. of the 27th International Symposium on High-Performance Parallel and Distributed Computing (HPDC)*. 40–51.

[36] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE transactions on evolutionary computation* 6, 2 (2002), 182–197.

[37] Sebastian Deorowicz and Szymon Grabowski. 2008. Sub-atomic field processing for improved web log compression. In *Proc. of the IEEE International Conference on "Modern Problems of Radio Engineering, Telecommunications and Computer Science" (TCSET)*. 551–556.

[38] Rui Ding, Qiang Fu, Jian-Guang Lou, Qingwei Lin, Dongmei Zhang, and Tao Xie. 2014. Mining historical issue repositories to heal large-scale online service systems. In *Proc. of the 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 311–322.

[39] Rui Ding, Hucheng Zhou, Jian-Guang Lou, Hongyu Zhang, Qingwei Lin, Qiang Fu, Dongmei Zhang, and Tao Xie. 2015. Log2: A cost-aware logging mechanism for performance diagnosis. In *Proc. of the 2015 USENIX Annual Technical Conference (ATC)*. 139–150.

[40] Min Du, Zhi Chen, Chang Liu, Rajvardhan Oak, and Dawn Song. 2019. Lifelong anomaly detection through unlearning. In *Proc. of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*. 1283–1297.

[41] Min Du and Feifei Li. 2018. Spell: online streaming parsing of large unstructured system logs (tkde). *IEEE Transactions on Knowledge and Data Engineering* (2018).

[42] Min Du, Feifei Li, Guineng Zheng, and Vivek Srikumar. 2017. Deeplog: anomaly detection and diagnosis from system logs through deep learning. In *Proc. of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*. 1285–1298.

[43] Nosayba El-Sayed and Bianca Schroeder. 2013. Reading between the lines of failure logs: Understanding how HPC systems fail. In *Proc. of the 43rd annual IEEE/IFIP international conference on dependable systems and networks (DSN)*. 1–12.

[44] ELK. 2012. ElasticSearch. https://www.elastic.co/elk-stack Online; accessed September 1, 2020.

[45] Facebook. 2019. Downtime, outages and failures - understanding their true costs. http://www.evolven.com/blog/downtime-outages-and-failures-understanding-their-true-costs.html

[46] Facebook. 2019. Facebook loses $24,420 a minute during outages. https://www.theatlantic.com/technology/archive/2014/10/facebook-is-losing-24420-per-minute/382054/

[47] Mostafa Farshchi, Jean-Guy Schneider, Ingo Weber, and John Grundy. 2015. Experience report: Anomaly detection of cloud application operations using log and cloud metric correlation analysis. In *Proc. of the 26th International Symposium on Software Reliability Engineering (ISSRE)*. 24–34.

[48] Bo Feng, Chentao Wu, and Jie Li. 2016. MLC: an efficient multi-level log compression method for cloud backup systems. In *Proc. of the 2016 IEEE Trustcom/BigDataSE/ISPA*. 1358–1365.

[49] Qiang Fu, Jian-Guang Lou, Yi Wang, and Jiang Li. 2009. Execution Anomaly Detection in Distributed Systems through Unstructured Log Analysis. In *Proc. of International Conference on Data Mining (ICDM)*.

[50] Qiang Fu, Jieming Zhu, Wenlu Hu, Jian-Guang Lou, Rui Ding, Qingwei Lin, Dongmei Zhang, and Tao Xie. 2014. Where do developers log? an empirical study on logging practices in industry. In *Proc. of the 36th International Conference on Software Engineering (ICSE)*. 24–33.

[51] Xiaoyu Fu, Rui Ren, Sally A McKee, Jianfeng Zhan, and Ninghui Sun. 2014. Digging deeper into cluster system logs for failure prediction and root cause diagnosis. In *2014 IEEE International Conference on Cluster Computing (CLUSTER)*. IEEE, 103–112.

[52] Maayan Goldstein, Danny Raz, and Itai Segall. 2017. Experience report: Log-based behavioral differencing. In *Proc. of the 28th IEEE International Symposium on Software Reliability Engineering (ISSRE)*. 282–293.

[53] Hossein Hamooni, Biplob Debnath, Jianwu Xu, Hui Zhang, Guofei Jiang, and Abdullah Mueen. 2016. LogMine: fast pattern recognition for log analytics. In *Proc. of the 25th ACM international conference on Information and knowledge management (CIKM)*. 1573–1582.

[54] Stephen E Hansen and E Todd Atkins. 1993. Automated system monitoring and notification with swatch.. In *Proc. of the 7th USENIX Large Installation System Administration Conference (LISA)*, Vol. 93. 145–152.

[55] Ahmed E Hassan, Daryl J Martin, Parminder Flora, Paul Mansfield, and Dave Dietz. 2008. An industrial case study of customizing operational profiles using log compression. In *Proc. of the 30th international conference on software engineering (ICSE)*. 713–723.

[56] Mehran Hassani, Weiyi Shang, Emad Shihab, and Nikolaos Tsantalis. 2018. Studying and detecting log-related issues. *Empirical Software Engineering* (2018), 3248–3280.

[57] Kimmo Hätönen, Jean François Boulicaut, Mika Klemettinen, Markus Miettinen, and Cyrille Masson. 2003. Comprehensive log compression with frequent patterns. In *Proc. of the International Conference on Data Warehousing and Knowledge Discovery (DaWaK)*. Springer, 360–370.

[58] Pinjia He, Zhuangbin Chen, Shilin He, and Michael R. Lyu. 2018. Characterizing the natural language descriptions in software logging statements. In *Proc. of the 33rd ACM/IEEE International Conference on Automated Software Engineering (ASE)*. 178–189.

[59] Pinjia He, Jieming Zhu, Shilin He, Jian Li, and Michael R Lyu. 2017. Towards automated log parsing for large-scale log data analysis. *IEEE Transactions on Dependable and Secure Computing* (2017).

[60] Pinjia He, Jieming Zhu, Pengcheng Xu, Zibin Zheng, and Michael R Lyu. 2018. A directed acyclic graph approach to online log parsing. *arXiv preprint arXiv:1806.04356* (2018).

[61] Pinjia. He, Jieming. Zhu, Zibin. Zheng, and Michael R. Lyu. 2017. Drain: an online log parsing approach with fixed depth tree. In *Proc. of the 24th International Conference on Web Services (ICWS)*.

[62] Shilin He, Qingwei Lin, Jian-Guang Lou, Hongyu Zhang, Michael R Lyu, and Dongmei Zhang. 2018. Identifying impactful service system problems via log analysis. In *Proc. of the 26th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. 60–70.

[63] Shilin He, Jieming Zhu, Pinjia He, and Michael R Lyu. 2016. Experience report: System log analysis for anomaly detection. In *Proc. of the 27th IEEE International Symposium on Software Reliability Engineering (ISSRE)*.

[64] Shilin He, Jieming Zhu, Pinjia He, and Michael R Lyu. 2020. Loghub: A Large Collection of System Log Datasets towards Automated Log Analytics. *arXiv preprint arXiv:2008.06448* (2020).

[65] Hiroki Ikeuchi, Akio Watanabe, Takehiro Kawata, and Ryoichi Kawahara. 2018. Root-cause diagnosis using logs generated by user actions. In *Proc. of the IEEE Global Communications Conference (GLOBECOM)*. 1–7.

[66] Tong Jia, Pengfei Chen, Lin Yang, Ying Li, Fanjing Meng, and Jingmin Xu. 2017. An approach for anomaly diagnosis based on hybrid graph model with logs for distributed services. In *Proc. of the IEEE International Conference on Web Services (ICWS)*. 25–32.

[67] Tong Jia, Lin Yang, Pengfei Chen, Ying Li, Fanjing Meng, and Jingmin Xu. 2017. Logsed: anomaly diagnosis through mining time-weighted control flow graph in logs. In *Proc. of the IEEE 10th International Conference on Cloud Computing (CLOUD)*. 447–455.

[68] He Jiang, Xiaochen Li, Zijiang Yang, and Jifeng Xuan. 2017. What causes my test alarm? Automatic cause analysis for test alarms in system and integration testing. In *Proc. of the IEEE/ACM 39th International Conference on Software Engineering (ICSE)*. 712–723.

[69] Weihang Jiang, Chongfeng Hu, Shankar Pasupathy, Arkady Kanevsky, Zhenmin Li, and Yuanyuan Zhou. 2009. Understanding customer problem troubleshooting from storage system logs. In *Proc. of the 7th conference on File and storage technologies (FAST)*. 43–56.

[70] Zhen Ming Jiang, Ahmed E Hassan, Parminder Flora, and Gilbert Hamann. 2008. Abstracting execution logs to execution events for enterprise applications. In *Proc. of the Eighth International Conference on Quality Software (QSIC)*. 181–186.

[71] Zhen Ming Jiang, Ahmed E. Hassan, Gilbert Hamann, and Parminder Flora. 2008. Automatic identification of load testing problems. In *Proc. of the 24th IEEE International Conference on Software Maintenance (ICSM)*. 307–316.

[72] Suhas Kabinna, Cor-Paul Bezemer, Weiyi Shang, and Ahmed E. Hassan. 2016. Logging library migrations: a case study for the apache software foundation projects. In *Proc. of the 13th International Conference on Mining Software Repositories (MSR)*. 154–164.

[73] Suhas Kabinna, Cor-Paul Bezemer, Weiyi Shang, Mark D. Syer, and Ahmed E. Hassan. 2018. Examining the stability of logging statements. *Empirical Software Engineering* (2018), 290–333.

[74] Soila P Kavulya, Kaustubh Joshi, Felicita Di Giandomenico, and Priya Narasimhan. 2012. Failure diagnosis of complex systems. In *Resilience assessment and evaluation of computing systems*. Springer, 239–261.

[75] Suleman Khan, Abdullah Gani, Ainuddin Wahid Abdul Wahab, Mustapha Aminu Bagiwa, Muhammad Shiraz, Samee U Khan, Rajkumar Buyya, and Albert Y Zomaya. 2016. Cloud log forensics: Foundations, state of the art, and future directions. *ACM Computing Surveys (CSUR)* (2016), 178–184.

[76] Gregor Kiczales, John Lamping, Anurag Mendhekar, Chris Maeda, Cristina Lopes, Jean-Marc Loingtier, and John Irwin. 1997. Aspect-oriented programming. In *European conference on object-oriented programming*. Springer.

[77] Tatsuaki Kimura, Akio Watanabe, Tsuyoshi Toyono, and Keisuke Ishibashi. 2018. Proactive failure detection learning generation patterns of large-scale network logs. *IEICE Transactions on Communications* (2018).

[78] Jannis Klinkenberg, Christian Terboven, Stefan Lankes, and Matthias S. MÃijller. 2017. Data mining-based analysis of hpc center operations. In *Proc. of the IEEE International Conference on Cluster Computing (CLUSTER)*. 766–773.

[79] Max Landauer, Florian Skopik, Markus Wurzenberger, and Andreas Rauber. 2020. System log clustering approaches for cyber security applications: A survey. *Computers & Security* 92 (2020), 101739.

[80] Jean-Claude Laprie. 1995. Dependable computing: Concepts, limits, challenges. In *Special issue of the 25th international symposium on fault-tolerant computing (FTCS)*. 42–54.

[81] Man-Lap Li, Pradeep Ramachandran, Swarup K Sahoo, Sarita V Adve, Vikram S Adve, and Yuanyuan Zhou. 2008. Trace-based microarchitecture-level diagnosis of permanent hardware faults. In *Proc. of the IEEE International Conference on Dependable Systems and Networks (DSN)*. 22–31.

[82] Zhenhao Li, Tse-Hsun Chen, and Weiyi Shang. 2020. Where shall we log? studying and suggesting logging locations in code blocks. In *Proc. of the 35rd IEEE/ACM International Conference on Automated Software Engineering (ASE)*.

[83] Zhenhao Li, Tse-Hsun (Peter) Chen, Jinqiu Yang, and Weiyi Shang. 2019. Dlfinder: characterizing and detecting duplicate logging code smells. In *Proc. of the 41st International Conference on Software Engineering (ICSE)*. 152–163.

[84] Yinglung Liang, Yanyong Zhang, Hui Xiong, and Ramendra Sahoo. 2007. Failure prediction in ibm bluegene/l event logs. In *Proc. of the 7th IEEE International Conference on Data Mining (ICDM)*. 583–588.

[85] Chinghway Lim, Navjot Singh, and Shalini Yajnik. 2008. A log mining approach to failure analysis of enterprise telephony systems. In *Proc. of the 38th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 398–403.

[86] Hao Lin, Jingyu Zhou, Bin Yao, Minyi Guo, and Jie Li. 2015. Cowic: A column-wise independent compression for log stream analysis. In *Proc. of the 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID)*. 21–30.

[87] Qingwei Lin, Randolph Yao, Murali Chintalapati, Dongmei Zhang, Ken Hsieh, Yingnong Dang, Hongyu Zhang, Kaixin Sui, Yong Xu, Jian-Guang Lou, Chenggang Li, and Youjiang Wu. 2018. Predicting node failure in cloud service systems. In *Proc. of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. ACM Press, 480–490.

[88] Qingwei Lin, Hongyu Zhang, Jian-Guang Lou, Yu Zhang, and Xuewei Chen. 2016. Log clustering based problem identification for online service systems. In *Proc. of the 38th International Conference on Software Engineering Companion (ICSE-Companion)*. 102–111.

[89] Fucheng Liu, Yu Wen, Dongxue Zhang, Xihe Jiang, Xinyu Xing, and Dan Meng. 2019. Log2vec: a heterogeneous graph embedding based approach for detecting cyber threats within enterprise. In *Proc. of the 26th ACM SIGSAC Conference on Computer and Communications Security (CCS)*. 1777–1794.

[90] Jinyang Liu, Jieming Zhu, Shilin He, Pinjia He, Zibin Zheng, and Michael R Lyu. 2019. Logzip: extracting hidden structures via iterative clustering for log compression. In *Proc. of the 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 863–873.

[91] Zhongxin Liu, Xin Xia, David Lo, Zhenchang Xing, Ahmed E Hassan, and Shanping Li. 2019. Which variables should i log? *IEEE Transactions on Software Engineering* (2019).

[92] Joshua Lockerman, Jose M. Faleiro, Juno Kim, Soham Sankaran, Daniel J. Abadi, James Aspnes, Siddhartha Sen, and Mahesh Balakrishnan. 2018. The fuzzylog: a partially ordered shared log. In *Proc. of the 13th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 357–372.

[93] Loggly. 2009. Automated parsing log types. https://www.loggly.com/docs/automated-parsing Online; accessed September 1, 2020.

[94] Loggly. 2009. Loggly - log management by loggly. https://www.loggly.com Online; accessed September 1, 2020.

[95] logz. 2014. Log parsing - automated, easy to use, and efficient. https://logz.io/product/log-parsing Online; accessed September 1, 2020.

[96] Jian-Guang Lou, Qiang Fu, Shengqi Yang, Jiang Li, and Bin Wu. 2010. Mining program workflow from interleaved traces. In *Proc. of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining (SIGKDD)*. 613–622.

[97] Jian-Guang Lou, Qiang Fu, Shengqi Yang, Ye Xu, and Jiang Li. 2010. Mining invariants from console logs for system problem detection.. In *Proc. of the 2010 USENIX Annual Technical Conference (ATC)*. 1–14.

[98] Jie Lu, Feng Li, Lian Li, and Xiaobing Feng. 2018. Cloudraid: hunting concurrency bugs in the cloud via log-mining. In *Proc. of the 26th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. 3–14.

[99] Siyang Lu, BingBing Rao, Xiang Wei, Byungchul Tak, Long Wang, and Liqiang Wang. 2017. Log-based abnormal task detection and root cause analysis for spark. In *Proc. of the IEEE International Conference on Web Services (ICWS)*. 389–396.

[100] Siyang Lu, Xiang Wei, Yandong Li, and Liqiang Wang. 2018. Detecting anomaly in big data system logs using convolutional neural network. In *Proc. of the 16th IEEE International Conference on Dependable, Autonomic and Secure Computing (DASC)*. 151–158.

[101] Liang Luo, Suman Nath, Lenin Ravindranath Sivalingam, Madan Musuvathi, and Luis Ceze. 2018. Troubleshooting transiently-recurring errors in production systems with blame-proportional logging. In *Proc. of the 2018 USENIX Annual Technical Conference (ATC)*. 321–334.

[102] Adetokunbo Makanju, A Nur Zincir-Heywood, and Evangelos E Milios. 2012. A lightweight algorithm for message type extraction in system application logs. *IEEE Transactions on Knowledge and Data Engineering (TKDE)* (2012).

[103] Adetokunbo AO Makanju, A Nur Zincir-Heywood, and Evangelos E Milios. 2009. Clustering event logs using iterative partitioning. In *Proc. of International Conference on Knowledge Discovery and Data Mining (KDD)*.

[104] Michael Meinig, Peter Tröger, and Christoph Meinel. 2019. Rough logs: a data reduction approach for log files.. In *Proc. of the International Conference on Enterprise Information Systems (ICEIS)*. 295–302.

[105] Peter Mell and Richard E Harang. 2014. Lightweight packing of log files for improved compression in mobile tactical networks. In *Proc. of the IEEE Military Communications Conference (MILCOM)*. 192–197.

[106] Weibin Meng, Ying Liu, Yuheng Huang, Shenglin Zhang, Federico Zaiter, Bingjin Chen, and Dan Pei. 2020. A semantic-aware representation framework for online log analysis. (2020).

[107] Weibin Meng, Ying Liu, Yichen Zhu, Shenglin Zhang, Dan Pei, Yuqing Liu, Yihao Chen, Ruizhi Zhang, Shimin Tao, Pei Sun, et al. 2019. LogAnomaly: unsupervised detection of sequential and quantitative anomalies in unstructured logs.. In *Proc. of the 2019 International Joint Conferences on Artificial Intelligence (IJCAI)*. 4739–4745.

[108] Salma Messaoudi, Annibale Panichella, Domenico Bianculli, Lionel Briand, and Raimondas Sasnauskas. 2018. A search-based approach for accurate identification of log message formats. In *Proc. of the IEEE/ACM 26th International Conference on Program Comprehension (ICPC)*.

[109] Haibo Mi, Huaimin Wang, Yangfan Zhou, Michael Rung-Tsong Lyu, and Hua Cai. 2013. Toward fine-grained, unsupervised, scalable performance diagnosis for production cloud computing systems. *Transactions on Parallel and Distributed Systems* 24, 6 (2013), 1245–1255.

[110] Microsoft. 2018. Event logging. https://docs.microsoft.com/en-us/windows/win32/eventlog/event-logging Online; accessed September 1, 2020.

[111] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Proc. of the 27th Conference on Neural Information Processing Systems (NIPS)*. 3111–3119.

[112] Tsuyoshi Mizouchi, Kazumasa Shimari, Takashi Ishio, and Katsuro Inoue. 2019. PADLA: a dynamic log level adapter using online phase detection. In *Proc. of the 27th International Conference on Program Comprehension (ICPC)*. 135–138.

[113] Masayoshi Mizutani. 2013. Incremental mining of system log format. In *Proc. of the IEEE International Conference on Services Computing (SCC)*. 595–602.

[114] Meiyappan Nagappan and Mladen A Vouk. 2010. Abstracting log lines to log event types for mining software system logs. In *Proc. of the 7th IEEE Working Conference on Mining Software Repositories (MSR)*. 114–117.

[115] Karthik Nagaraj, Charles Killian, and Jennifer Neville. 2012. Structured comparative analysis of systems logs to diagnose performance problems. In *Proc. of the 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. 353–366.

[116] Animesh Nandi, Atri Mandal, Shubham Atreja, Gargi B Dasgupta, and Subhrajit Bhattacharya. 2016. Anomaly detection using program control flow graph mining from execution logs. In *Proc. of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (SIGKDD)*. 215–224.

[117] Adam J. Oliner, Archana Ganapathi, and Wei Xu. 2012. Advances and challenges in log analysis. *ACM Communication* (2012), 55–61.

[118] Frederick John Otten et al. 2008. *Using semantic knowledge to improve compression on log files*. Ph.D. Dissertation. Rhodes University.

[119] Antonio Pecchia, Marcello Cinque, Gabriella Carrozza, and Domenico Cotroneo. 2015. Industry practices and event logging: assessment of a critical software development process. In *Proc. of the 37th IEEE/ACM International Conference on Software Engineering (ICSE)*. 169–178.

[120] Cuong Pham, Long Wang, Byung Chul Tak, Salman Baset, Chunqiang Tang, Zbigniew Kalbarczyk, and Ravishankar K Iyer. 2016. Failure diagnosis for distributed systems using targeted fault injection. *Transactions on Parallel and Distributed Systems* 28, 2 (2016), 503–516.

[121] James E Prewett. 2003. Analyzing cluster log files using logsurfer. In *Proc. of the 4th Annual Conference on Linux Clusters*. Citeseer.

[122] Balázs Rácz and András Lukács. 2004. High density compression of log files. In *Proc. of the Data Compression Conference (DCC)*. IEEE, 557.

[123] Rapid7. 2000. New automated log parsing. https://blog.rapid7.com/2016/03/03/new-automated-log-parsing Online; accessed September 1, 2020.

[124] John P Rouillard. 2004. Real-time log file analysis using the simple event correlator (sec).. In *Proc. of the 18th USENIX Large Installation System Administration Conference (LISA)*, Vol. 4. 133–150.

[125] Ramendra K. Sahoo, Adam J. Oliner, Irina Rish, Manish Gupta, José E. Moreira, Sheng Ma, Ricardo Vilalta, and Anand Sivasubramaniam. 2003. Critical event prediction for proactive management in large-scale computer clusters. In *Proc. of the ninth ACM SIGKDD international conference on knowledge discovery and data mining (SIGKDD)*. 426–435.

[126] Bianca Schroeder and Garth A. Gibson. 2006. A large-scale study of failures in high-performance computing systems. In *Proc. of the 2006 International Conference on Dependable Systems and Networks (DSN)*. 249–258.

[127] Weiyi Shang. 2012. Bridging the divide between software developers and operators using logs. In *Proc. of the 34th International Conference on Software Engineering (ICSE)*. 1583–1586.

[128] Weiyi Shang, Zhen Ming Jiang, Hadi Hemmati, Brain Adams, Ahmed E Hassan, and Patrick Martin. 2013. Assisting developers of big data analytics applications when deploying on hadoop clouds. In *Proc. of the 35th International Conference on Software Engineering (ICSE)*. 402–411.

[129] Weiyi Shang, Meiyappan Nagappan, and Ahmed E. Hassan. 2015. Studying the relationship between logging characteristics and the code quality of platform software. *Empirical Software Engineering* (2015), 1–27.

[130] Weiyi Shang, Meiyappan Nagappan, Ahmed E. Hassan, and Zhen Ming Jiang. 2014. Understanding log lines using development knowledge. In *Proc. of the 30th International Conference on Software Maintenance and Evolution (ICSME)*. 21–30.

[131] Keiichi Shima. 2016. Length matters: clustering system log messages using length of words. *arXiv:1611.03213* (2016).

[132] Przemysław Skibiński and Jakub Swacha. 2007. Fast and efficient log file compression. In *Proc. of 11th east-European conference on advances in databases and information systems (ADBIS)*. 56–69.

[133] Splunk. 2005. Splunk platform. http://www.splunk.com Online; accessed September 1, 2020.

[134] Byung Chul Tak, Shu Tao, Lin Yang, Chao Zhu, and Yaoping Ruan. 2016. Logan: problem diagnosis in the cloud using log-based reference models. In *Proc. of the IEEE International Conference on Cloud Engineering (IC2E)*. 62–67.

[135] Liang Tang, Tao Li, and Chang-Shing Perng. 2011. LogSig: Generating system events from raw textual logs. In *Proc. of the 20th ACM international conference on Information and knowledge management (CIKM)*. 785–794.

[136] StatusCake Team. 2020. The Most Expensive Website Downtime Periods in History. https://www.statuscake.com/the-most-expensive-website-downtime-periods-in-history/

[137] UpGuard. 2019. The cost of downtime at the world's biggest online retailer. https://www.upguard.com/blog/the-cost-of-downtime-at-the-worlds-biggest-online-retailer Online; accessed September 1, 2020.

[138] Risto Vaarandi. 2003. A data clustering algorithm for mining patterns from event logs. In *Proc. of the 3rd IEEE Workshop on IP Operations & Management (IPOM)*. 119–126.

[139] Risto Vaarandi and Mauno Pihelgas. 2015. LogCluster - A data clustering and pattern mining algorithm for event logs. In *Proc. of the 11th International conference on network and service management (CNSM)*. 1–7.

[140] Xiaoyin Wang, Lu Zhang, Tao Xie, John Anvik, and Jiasu Sun. 2008. An approach to detecting duplicate bug reports using natural language and execution information. In *Proc. of the 30th international conference on software engineering (ICSE)*. 461–470.

[141] Bin Xia, Yuxuan Bai, Junjie Yin, Yun Li, and Jian Xu. 2020. LogGAN: a log-level generative adversarial network for anomaly detection using permutation event modeling. *Information Systems Frontiers* (2020), 1–14.

[142] Wei Xu. 2010. *System problem detection by mining console logs*. Ph.D. Dissertation. University of California, Berkeley.

[143] Wei Xu, Ling Huang, Armando Fox, David Patterson, and Michael Jordan. 2009. Online system problem detection by mining patterns of console logs. In *Proc. of the Ninth IEEE International Conference on Data Mining (ICDM)*. 588–597.

[144] Wei Xu, Ling Huang, Armando Fox, David Patterson, and Michael I Jordan. 2009. Detecting large-scale system problems by mining console logs. In *Proc. of the 22nd ACM Symposium on Operating Systems Principles (SOSP)*. 117–132.

[145] Kenji Yamanishi and Yuko Maruyama. 2005. Dynamic syslog mining for network failure monitoring. In *Proc. of the 11nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (SIGKDD)*. 499–508.

[146] Stephen Yang, Seo Jin Park, and John K. Ousterhout. 2018. NanoLog: A nanosecond scale logging system. In *Proc. of the 2018 USENIX Annual Technical Conference (ATC)*. 335–350.

[147] Kundi Yao, Heng Li, Weiyi Shang, and Ahmed E Hassan. 2019. A study of the performance of general compressors on log files. *Empirical Software Engineering* (2019).

[148] Xiao Yu, Shi Han, Dongmei Zhang, and Tao Xie. 2014. Comprehending performance from real-world execution traces: A device-driver case. In *Proc. of the 19th international conference on Architectural support for programming languages and operating systems (ASPLOS)*. 193–206.

[149] Xiao Yu, Pallavi Joshi, Jianwu Xu, Guoliang Jin, Hui Zhang, and Guofei Jiang. 2016. Cloudseer: workflow monitoring of cloud infrastructures via interleaved logs. In *Proc. of the 21st International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. 489–502.

[150] Ding Yuan, Haohui Mai, Weiwei Xiong, Lin Tan, Yuanyuan Zhou, and Shankar Pasupathy. 2010. SherLog: error diagnosis by connecting clues from run-time logs. In *Proc. of the fifteenth International Conference on Architectural support for programming languages and operating systems (ASPLOS)*. 143–154.

[151] Ding Yuan, Soyeon Park, Peng Huang, Yang Liu, Michael Mihn-Jong Lee, Xiaoming Tang, Yuanyuan Zhou, and Stefan Savage. 2012. Be conservative: enhancing failure diagnosis with proactive logging. In *Proc. of the 10th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 293–306.

[152] Ding Yuan, Soyeon Park, and Yuanyuan Zhou. 2012. Characterizing logging practices in open-source software. In *Proc. of 34th International Conference on Software Engineering (ICSE)*. 102–112.

[153] Ding Yuan, Jing Zheng, Soyeon Park, Yuanyuan Zhou, and Stefan Savage. 2011. Improving software diagnosability via log enhancement. In *Proc. of the 16th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. 3–14.

[154] Yue Yuan, Wenchang Shi, Bin Liang, and Bo Qin. 2019. An approach to cloud execution failure diagnosis based on exception logs in openstack. In *Proc. of the IEEE 12th International Conference on Cloud Computing (CLOUD)*. 124–131.

[155] Lei Zeng, Yang Xiao, Hui Chen, Bo Sun, and Wenlin Han. 2016. Computer operating system logging and security issues: a survey. *Security and Communication Networks* 9, 17 (2016), 4804–4821.

[156] Yi Zeng, Jinfu Chen, Weiyi Shang, and Tse-Hsun (Peter) Chen. 2019. Studying the characteristics of logging practices in mobile apps: a case study on F-Droid. *Empirical Software Engineering* (2019), 3394–3434.

[157] Shenglin Zhang, Ying Liu, Weibin Meng, Zhiling Luo, Jiahao Bu, Sen Yang, Peixian Liang, Dan Pei, Jun Xu, Yuzhi Zhang, Yu Chen, Hui Dong, Xianping Qu, and Lei Song. 2018. PreFix: switch failure prediction in datacenter networks. *Proc. ACM Meas. Anal. Comput. Syst.* 2 (2018), 2:1–2:29.

[158] Xu Zhang, Yong Xu, Qingwei Lin, Bo Qiao, Hongyu Zhang, Yingnong Dang, Chunyu Xie, Xinsheng Yang, Qian Cheng, Ze Li, et al. 2019. Robust log-based anomaly detection on unstable log data. In *Proc. of the 27th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. 807–817.

[159] Yongle Zhang, Kirk Rodrigues, Yu Luo, Michael Stumm, and Ding Yuan. 2019. The inflection point hypothesis: a principled debugging approach for locating the root cause of a failure. In *Proc. of the 27th ACM Symposium on Operating Systems Principles (SOSP)*. 131–146.

[160] Xu Zhao, Kirk Rodrigues, Yu Luo, Michael Stumm, Ding Yuan, and Yuanyuan Zhou. 2017. The game of twenty questions: do you know where to log?. In *Proc. of the 16th Workshop on Hot Topics in Operating Systems (HotOS)*. 125–131.

[161] Xu Zhao, Kirk Rodrigues, Yu Luo, Michael Stumm, Ding Yuan, and Yuanyuan Zhou. 2017. Log20: fully automated optimal placement of log printing statements under specified overhead threshold. In *Proc. of the 26th Symposium on Operating Systems Principles (SOSP)*. 565–581.

[162] Xu Zhao, Kirk Rodrigues, Yu Luo, Ding Yuan, and Michael Stumm. 2016. Non-intrusive performance profiling for entire software stacks based on the flow reconstruction principle. In *Proc. of the 12th Symposium on Operating Systems Design and Implementation (OSDI)*. 603–618.

[163] Xiang Zhou, Xin Peng, Tao Xie, Jun Sun, Chao Ji, Wenhai Li, and Dan Ding. 2018. Fault analysis and debugging of microservice systems: Industrial survey, benchmark system, and empirical study. *IEEE Transactions on Software Engineering* (2018).

[164] Xiang Zhou, Xin Peng, Tao Xie, Jun Sun, Chao Ji, Dewei Liu, Qilin Xiang, and Chuan He. 2019. Latent error prediction and fault localization for microservice applications by learning from system trace logs. In *Proc. of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*. 683–694.

[165] Jieming Zhu, Pinjia He, Qiang Fu, Hongyu Zhang, Michael R. Lyu, and Dongmei Zhang. 2015. Learning to log: helping developers make informed logging decisions. In *Proc. of the 37th IEEE/ACM International Conference on Software Engineering (ICSE)*. 415–425.

[166] Jieming Zhu, Shilin He, Jinyang Liu, Pinjia He, Qi Xie, Zibin Zheng, and Michael R. Lyu. 2019. Tools and benchmarks for automated log parsing. In *Proc. of the 41st International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*. 121–130.

[167] Yuan Zuo, Yulei Wu, Geyong Min, Chengqiang Huang, and Ke Pei. 2020. An intelligent anomaly detection scheme for micro-services architectures with temporal and spatial data analysis. *IEEE Transactions on Cognitive Communications and Networking* (2020).