

Article

The Diagnosis-Effective Sampling of Application Traces

Arnak Poghosyan ^{1,2,*} , Ashot Harutyunyan ^{3,*} , Edgar Davtyan ⁴ , Karen Petrosyan ² 
and Nelson Baloian ⁵ 

¹ Institute of Mathematics NAS Armenia, Yerevan 0019, Armenia

² College of Science and Engineering, American University of Armenia, Yerevan 0019, Armenia; karen_petrosyan2@edu.aua.am

³ ML Laboratory, Yerevan State University, Yerevan 0025, Armenia

⁴ Picsart, Miami, FL 33009, USA; edgar.davtyan@picsart.com

⁵ Department of Computer Science, University of Chile, Santiago 8330111, Chile; nbaloian@dcc.uchile.cl

* Correspondence: arnak@instmath.sci.am (A.P.); harutyunyan.ashot@ysu.am (A.H.)

Abstract: Distributed tracing is cutting-edge technology used for monitoring, managing, and troubleshooting native cloud applications. It offers a more comprehensive and continuous observability, surpassing traditional logging methods, and is indispensable for navigating modern complex software architectures. However, the sheer volume of generated traces is staggering in distributed applications, and the direct storage and utilization of every trace is impractical due to associated operational costs. This entails a sampling strategy to select which traces warrant storage and analysis. Historically, sampling methods have included a rate-based approach, often relying heavily on a manual configuration. There is a need for a more intelligent approach, and we propose a hierarchical sampling methodology to address multiple requirements concurrently. Initial rate-based sampling mitigates the overwhelming volume of traces, as no further analysis can be performed on this level. In the next stage, more nuanced analysis is facilitated based on the previous foundation, incorporating information regarding trace properties and ensuring the preservation of vital process details even under extreme conditions. This comprehensive approach not only aids in the visualization and conceptualization of applications but also enables more targeted analysis in later stages. As we delve deeper into the sampling hierarchy, the technique becomes tailored to specific purposes, such as the simplification of application troubleshooting. In this context, the sampling strategy prioritizes the retention of erroneous traces from dominant processes, thus facilitating the identification and resolution of underlying issues. The focus of this paper is to reveal the impact of sampling on troubleshooting efficiency. Leveraging intelligent and explainable artificial intelligence solutions enables the detection of malfunctioning microservices and provides transparent insights into root causes. We advocate for using rule-induction systems, which offer explainability and efficacy in decision-making processes. By integrating advanced sampling techniques with machine-learning-driven intelligence, we empower organizations to navigate the complexities of large-scale distributed cloud environments effectively.

Keywords: native cloud applications; application monitoring; distributed tracing; trace sampling; application troubleshooting; root cause analysis; explainable artificial intelligence; rule-induction systems



Citation: Poghosyan, A.; Harutyunyan, A.; Davtyan, E.; Petrosyan, K.; Baloian, N. The Diagnosis-Effective Sampling of Application Traces. *Appl. Sci.* **2024**, *14*, 5779. <https://doi.org/10.3390/app14135779>

Academic Editors: Milos Manic, Sergiu Dan Stan, Milan Banić, Bogdan Mocan and Aysegül Uçar

Received: 27 April 2024

Revised: 24 June 2024

Accepted: 26 June 2024

Published: 2 July 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Distributed tracing (DT) has emerged as a crucial tool for effectively monitoring and troubleshooting native cloud applications [1–7]. It plays a pivotal role in understanding the behavior and performance of applications across distributed environments. Native cloud applications often adopt microservices architecture, decomposing the application into smaller, loosely coupled services. Each service performs a specific function and communicates with others via Application Programming Interfaces (APIs). DT enables the

inspection of the flow of requests across the microservices, providing insights into how requests are processed and identifying any bottlenecks or issues within the system [8].

Native cloud applications are designed to be highly scalable and dynamically allocate resources based on demand. DT accommodates this dynamic nature by providing visibility regarding the performance of services as they scale up or down in response to changes in workload. This ensures that performance issues are detected and addressed in real time, maintaining the application's reliability and responsiveness. They often utilize containerization and container orchestration via Docker and Kubernetes, respectively. Those technologies enable applications to be deployed and managed efficiently in a distributed environment. DT is integrated with container platforms for tracking requests as they traverse across containers and pods, providing insights into resource utilization and communication patterns.

DT is typically part of a broader observability stack, including metrics and logging [2,9]. Integration with metrics allows for correlation between trace data and performance metrics, enabling a more profound analysis of application behavior. Similarly, integration with logging platforms provides context around specific events or errors captured in logs, enhancing the troubleshooting process [10–18].

Several critical points are made regarding the complexities and challenges associated with DT. One is the volume of traces in which many requests are routine and unremarkable, often called normal traces. Sampling in DT involves capturing a subset of traces for analysis instead of storing every trace to manage costs. This subset includes “interesting” traces encompassing various events in a distributed architecture. Interest can be linked to a trace latency when traces surpassing a latency threshold are selectively sampled. This allows for the identification of performance bottlenecks and areas needing improvement. Interest can also be linked to errors. Erroneous traces or exceptions are sampled to investigate and address system reliability and stability issues. Requests or services can be prioritized to help determine which traces to sample. High-priority components, such as critical services or specific user interactions, are given preference for sampling. Companies typically adopt one of two strategies for capturing these events.

The common sampling strategy is head-based sampling [2], in which traces are randomly sampled based on a predefined rate or probability, such as 0.1–1% of all traces. This approach relies on the principle that a sufficiently large dataset will capture the most interesting traces. However, it is worth noting that sampling meaningfully diminishes the value of DT. While sampling is necessary to manage costs, it can limit the effectiveness of DT for troubleshooting and debugging purposes. Developers may encounter situations where important traces are missed due to sampling, leading to a loss of trust in the tracing data. Consequently, developers may revert to traditional debugging methods, such as logs, undermining the value of DT [8]. A more intelligent strategy is called tail-based sampling, where the system waits until all spans within a request have been completed before determining whether or not to retain the trace based on its entirety. What truly matters here are the around 5% of traces that carry anomalies: errors, exceptions, instances of high latency, or other forms of soft errors. However, the ideal scenario would involve analyzing the entire set of traces, identifying anomalies, and retaining them for thorough examination. In this paper, we advocate for a multi-purpose and hierarchical strategy combining head-based and tail-based scenarios, focusing on the efficiency of application troubleshooting and root cause analysis (RCA). We are considering a multi-layered approach to efficiently capture and preserve different types of information within a system. Overall, this hierarchical sampling approach efficiently manages trace data by prioritizing information based on its relevance and importance for system analysis and troubleshooting. It balances the need to reduce data volume with the requirement to retain critical details for comprehensive understanding and problem resolution.

One of the main benefits of DT is improved application performance, resulting in reduced mean time to detect (MTTD) and mean time to repair (MTTR) IT issues. DT enables penetration to the bottom of application performance issues faster, often before users notice

anything wrong. Upon discovering an issue, tracing can rapidly identify the root cause and address it. It also provides early warning capabilities when microservices are in poor health and highlights performance bottlenecks anywhere. However, this task is only feasible for a few system administrators or site reliability engineers (SREs) due to the large scale of modern distributed environments. The solution is an artificial intelligence (AI) for IT Operations (AIOps) strategy that leverages AI-powered software intelligence to automate developments, service deliveries, and troubleshooting. AIOps is a real game-changer for managing complex IT systems.

In AIOps, where machine learning (ML) algorithms analyze vast amounts of data to automate IT operations and decision-making processes, transparency and interpretability are paramount. Trust in AIOps hinges on implementing explainable AI (XAI) solutions [19]. Those techniques enable stakeholders to understand and interpret the decisions made by AI models, build confidence in their reliability, and facilitate human–AI collaboration. By providing insights into how AI algorithms arrive at their conclusions, XAI fosters trust among users, reduces the risk of biased or erroneous outcomes, and enhances the adoption and effectiveness of AIOps solutions. We focus on rule-learning systems and show their efficiency in combination with trace sampling. Rule induction systems, such as RIPPER (Repeated Incremental Pruning to Produce Error Reduction) [20] and C5.0 rules [21] are common examples of XAI solutions.

The paper is organized as follows. Section 2 describes the main trends of trace sampling and rule-induction. Section 3 describes the main idea of our approach. Section 4 discusses the main motivation and contribution of this research. Section 5 explains the methods of trace sampling based on properties when the goal is to preserve all available types, latencies, and errors uniformly. Section 6 presents an application troubleshooting strategy and the impact of noise reduction on a rule induction process. Section 7 summarizes the results and discusses future opportunities. Section 8 provides the list of patents related to this research.

2. Related Work

2.1. Distributed Tracing Vendors

Many companies/vendors are developing DT technologies with some sampling capabilities built in. We will mention a few of them.

- Jaeger is an open-source DT system developed by Uber Technologies. It is widely used for monitoring and troubleshooting microservices-based applications [22].
- Zipkin is another open-source DT system originally developed by Twitter. It helps developers gather data for various components of their applications and troubleshoot latency issues [23].
- LightStep is a commercial company that offers tracing solutions for monitoring and troubleshooting distributed systems [24].
- Datadog is a monitoring and analytics platform that offers features for collecting and analyzing traces, logs, and metrics from applications and infrastructure [25].
- Dynatrace is a software intelligence platform that provides features for monitoring and analyzing the performance of applications and infrastructure, including DT capabilities [26].
- New Relic offers a monitoring and observability platform that includes features for collecting and analyzing traces, logs, and metrics from applications and infrastructure [27].
- AppDynamics is an application-performance-monitoring solution that provides features for monitoring and troubleshooting applications, including DT capabilities [28].
- Honeycomb is an observability platform that offers features for collecting and analyzing high-cardinality data, including traces, to debug and optimize applications [29].
- Grafana is an open-source analytics and visualization platform that can visualize traces, logs, and metrics collected from various sources, including DT systems [30].

- Elastic is the company behind Elasticsearch, Kibana, and other products in the Elastic Stack. Their solutions offer features for collecting, analyzing, and visualizing traces, logs, and metrics [31].

2.2. Industrial Sampling Examples

Various companies and researchers are developing end-to-end DT technologies in which sampling is one of the core components as the prevailing approach to reducing tracing overheads. Instead of tracing every request, the sampling only captures and persists traces for a subset of requests to the system. To ensure that the captured data are useful, sampling decisions are coherent per request—a trace is either sampled in its entirety, capturing the full end-to-end execution, or not at all. Sampling effectively reduces computational overheads; these overheads are only paid if a trace is sampled, so they can be easily reduced by reducing the sampling probability ([32–35] with references therein).

Large technology companies such as Google, Microsoft, Amazon, and Facebook often develop methods to monitor and analyze the performance of their own systems and applications. Early tracing systems such as Google’s Dapper [36] and later Facebook’s Canopy [37] make sampling decisions immediately when a request enters the system. This approach, known as head-based sampling, avoids the runtime costs of generating trace data as they occur, uniformly at random, and the resulting data are simply a random subset of requests. In practice, sampling rates can be as low as 0.1% [32].

Tail-based sampling is an alternative to head-based sampling. It captures traces for all requests and only decides whether to keep a trace after it has been generated. While OpenTelemetry [38] offers a tail-based sampling collector, its implementation presents challenges as storing all spans of a trace until the request concludes demands a sophisticated data architecture and solutions. One such end-to-end tracing method to enhance distributed system dependability by dynamically verifying and diagnosing correctness and performance issues is proposed in paper [39]. The idea is based on clustering execution graphs to bias sampling towards diverse and representative traces, even when anomalies are rare.

Google has developed DT technologies, such as Google Cloud Trace and OpenCensus, which allow users to collect and analyze trace data from applications running on Google Cloud Platform and other environments. Each Google Cloud service makes its own sampling decisions. When sampling is supported, that service typically implements a default sample rate or a mechanism to use the parent’s sampling decision as a hint as to whether to sample the span or maximum sampling rate [40]. OpenCensus provides “Always”, “Never”, “Probabilistic”, and “RateLimiting” samplers. The last one tries to sample with a rate per time window, which by default is 0.1 traces/s [41].

Azure Application Insights sampling [42] aims to reduce telemetry traffic, data, and storage costs while preserving a statistically correct analysis of application data. It enables three different types of sampling: adaptive sampling, fixed-rate sampling, and ingestion sampling. Adaptive filtering automatically adjusts the sampling to stay within the given rate limit. If the application generates low telemetry, like during debugging or low usage, it does not drop items as long as the volume stays under the limits. The sampling rate is adjusted to hit the target volume as the telemetry volume rises. Fixed-rate sampling reduces the traffic sent from web servers and web browsers. Unlike adaptive sampling, it reduces telemetry at a fixed rate that a user decides. Ingestion sampling operates where web servers’, browsers’, and devices’ telemetry reaches the Application Insights service endpoint. Although it does not reduce the telemetry traffic sent from an application, it does reduce the amount processed and retained (and charged for) by Application Insights. Microsoft researchers introduced an observability-preserving trace sampling method, denoted as STEAM, based on Graph Neural Networks, which aims to retain as much information as possible in the sampled traces [43].

In 2023, AWS announced the general availability of the tail sampling processor and the group-by-trace processor in the AWS Distro for OpenTelemetry collector [44]. Advanced

sampling refers to a strategy in which the Group By Trace processor and Tail Sampling processor operate together to make sampling decisions based on set policies regarding trace spans. The Group By Trace processor gathers all of the spans of a trace and waits for a pre-defined time before moving them to the next processor. This component is usually used before the tail sampling processor to guarantee that all the spans belonging to the same trace are processed together. Then, the Tail Sampling processor samples traces based on user-defined policies.

2.3. XAI for Application Troubleshooting

One of the main purposes of DT is to ensure application availability and fast troubleshooting in case of performance degradation. However, system administrators can no longer perform real-time decision-making due to the growth of large-scale distributed cloud environments with complicated, invisible underlying processes. Those systems require more advanced and ML/AI-empowered intelligent RCA with explainable and actionable recommendations ([18,45–51] with references therein).

XAI [19] builds user and AI trust, increases solutions' satisfaction, and leads to more actionable and robust prediction and RCA models. Many users think it is risky to trust and follow AI recommendations and predictions blindly, and they need to understand the foundation of those insights. Many ML approaches, like decision trees and rule-induction systems, have sufficient explainability capabilities for RCA. They can detect and predict performance degradations and identify the most critical features (processes) potentially responsible for malfunctioning.

In many applications, explainable outcomes can be more valuable than conclusions based on more powerful approaches that act like black boxes. Rule learners are the best if outcomes' simplicity and human interpretability are superior to the predictive power. The list of known rule learners consists of many exciting approaches. We refer to [52] for a more detailed description of available algorithms, their comparisons, and historical analysis. It contains relatively rich references and describes several applications [20,21,53–56]. Rule learning algorithms have a long history in industrial applications [7,13,57–64]. Many such applications utilize classical classifiers like C5.0Rules [21] and JRip. The last one is the Weka implementation of RIPPER.

The recommendations derived from rule-learning systems can be additionally verified regarding uncertainty based on the Dempster–Shafer theory (DST) of evidence [65]. This is an inference framework with uncertainty modeling, where independent sources of knowledge (expert opinions) can be combined for reasoning or decision-making. In [66], it is leveraged to build a classification method that enables a user-defined rule validation mechanism. These rules “encode” expert hypotheses as conditions on the dataset features that might be associated with certain classes while making predictions for observations. The theory provides a “what-if” analysis framework for comprehending the underlying application, as applied in [67]. This means that in the current study context, we can apply the same framework to estimate the effectiveness of the trace sampling approaches while comparing the rule validation results in the pre- and post-sampling stages. DST is overlooked in terms of learning algorithms that could be utilized in predictive system diagnostics ([45], which emphasizes that gap). Therefore, its application in our use-case of RCA-effective trace sampling is an additional novelty that we introduce.

3. The Main Idea

This section describes our multi-layer strategy with the final goal of application troubleshooting. One common approach to collecting traces is to use a tracer based on OpenTelemetry, Zipkin, Jaeger, etc. The tracer is configured to forward traces, metrics, and logs from an application to a proxy. The proxy securely, quickly, and reliably sends data to a Trace Manager. A cloud-scale application generates a large number of traces. The proxy should be configured to sample data to reduce the volume. A common strategy

is to apply head-based sampling, also known as a rate-based strategy, which randomly samples without going deep into the context (see Figure 1).

In our case, regardless of the sampling strategy, the tracer allows all error spans. Based on the traces, it collects and reports request, error, and duration (RED) metrics to provide full application observability. RED metrics measure requests (the number of requests being served per second), errors (the number of failed requests per second), and durations (per-minute histogram distributions of the amount of time that each request takes). These metrics can be used for application troubleshooting.




Figure 1. Trace sampling multi-layer design.

The Trace Manager collects and samples traces based on their properties, such as type, duration, and error, although other properties can be considered. Then, the Trace Manager visualizes and analyzes the traces. At this stage, the main purpose of the property-based sampling is to preserve all available properties for further analysis. Trace-type sampling involves categorizing traces based on their characteristics or origins within the system. On the other hand, duration-based sampling prioritizes the collection of traces associated with significant delays or performance bottlenecks.

Error-based sampling identifies and retains traces associated with errors, exceptions, or other abnormal conditions within the application. Combining these approaches allows the system to store more diverse traces, ensuring that valuable information across different categories and performance metrics is captured. By isolating and storing these traces, developers and system administrators have the necessary data to diagnose and address issues effectively, improving overall system reliability and performance.

However, the basic purpose of the Trace Manager is application troubleshooting in the case of performance degradation. It has diverse tools for different management tasks. One of them is the Traces Browser, which can reveal the context and details of traces. In the Traces Browser, one can search for traces that include spans for a particular operation or examine the spans that belong to a selected trace. It is also possible to view the corresponding RED metrics for troubleshooting purposes.

Another sophisticated troubleshooting engine is an applications map that overviews how the applications and services are linked. It allows us to focus on a specific service, view RED metrics for each service, and filter out trace traffic for the malfunctioning microservices. Here, noise reduction is performed based on the rate-based sampling of erroneous traces. It preserves all dominant normal and erroneous traces and removes the rare ones. This approach is acceptable when the problem has already been detected, and the goal is to explain it.

The final layer applies XAI on the sampled data for RCA, bottleneck detection, and optimizations. In particular, rule learning systems which applied to the reduced tracing traffic

can generate recommendations in the form of rules which should be inspected by system administrators to accelerate the remediation of IT issues.

4. Motivations and the Main Contributions

A common method for troubleshooting applications involves utilizing Traces Browser or Application Maps, which enable the examination of malfunctioning microservices. These tools allow for the isolation of trace traffic, facilitating the feeding of trace data into the browser for an in-depth analysis of traces and spans. Following this manual investigative process, the metadata associated with the spans can uncover underlying issues and suggest potential remediation procedures. However, those approaches have serious drawbacks. The volume of traces without sampling is so large that it is not feasible to draw the Application Maps within a realistic period. Moreover, manual inspection requires expertise and a thorough understanding of the application's internal flows. Additionally, it does not allow us to properly automate the procedure in case of many applications, traces, and spans. Our primary objective was to automate and optimize this intricate process using ML methodologies, focusing on achieving explanations through XAI solutions such as RIPPER or Dempster–Shafer theory. However, we encountered a significant challenge due to the overwhelming volume of traces, even after the initial stage of head-based sampling. Hence, the critical task was the identification of effective sampling strategies to tackle this issue.

Deliberating between ML approaches and simpler statistical methods, we initially opted for the latter for several reasons, including leveraging expert knowledge, maintaining transparency, and retaining control over the sampling process. This decision led us to develop a dynamic streaming solution that proved highly practical. Our next goal was to assess the impact of this sampling technique on XAI methods.

This paper's contribution is multi-fold. We discuss the trace sampling approach that has been successfully implemented, and demonstrate its capability for diverse applications. We developed application trace sampling technology that leverages the statistical properties of traces, such as type, duration, and errors. It is easy to implement and has serious benefits compared to more sophisticated ML approaches. Statistical methods provide more interpretable results. This transparency can be crucial for understanding the underlying processes or relationships in the data. The corresponding parameters are easy to fine-tune and control. Statistical methods are built on clear assumptions about the data distribution, which can help to guide the analysis and ensure the validity of the results. In contrast, ML models are more ambivalent to data assumptions, which can sometimes lead to unexpected or inaccurate outcomes. Our approach is suitable for limited data availability. ML models, on the other hand, often require larger datasets to achieve reliable performance. Statistical methods often allow for the easier incorporation of domain knowledge or expert insights into the analysis process, enhancing the interpretability and relevance of the results.

Then, we explore the impact of sampling on application troubleshooting. We demonstrated that sampling positively impacted rule-learning systems by reducing time and generating clearer rules. Hence, sample size reduction does not impact the efficiency of the application troubleshooting approach. We also experimented with the Dempster–Shafer classifier, which was successfully implemented and proved applicable to the RCA in different IT issues.

However, the efficiency of ML strategies is an open problem. There is a compelling need to explore ML solutions for sampling that offer enhanced scalability, speed, and accuracy, particularly in application domains where such capabilities are crucial [68–71]. One of the important problems is related to trace type identification. ML methods can categorize traces based on specific characteristics or outcomes as a supervised learning problem. Alternatively, clustering approaches can be used to group similar traces together based on patterns or similarities. Another important problem is anomaly trace detection. We assume that anomalies are known and traces arrive as either normal or abnormal. This is an interesting problem, as anomalies have different sources, and ML approaches can help

to identify unusual or abnormal traces that deviate from the “norm”. The next important problem is detecting the change in the dynamic flow of traces. Appropriate actions or decisions in dynamic trace analysis scenarios may be required. In some scenarios, when modeling and forecasting trace data over time, the capture of temporal dependencies can be an important problem. In general, diversifying our focus to include other ML methods could provide valuable insights into their effectiveness within our sampling strategy.

While our primary emphasis has been on rule induction methods, driven by their interpretability features, there is potential value in exploring alternative ML approaches and assessing how sampling strategies influence their performance. Furthermore, investigating how different sampling techniques impact these methods can offer a deeper understanding of their behavior and potential for optimization. In instances where interpretability may be compromised, supplementary techniques such as SHAP [72] or LIME [73] could be employed to mitigate this limitation and enhance the overall interpretability of the results.

5. Trace Sampling Based on Type, Duration, and Errors

Modern distributed applications trigger an enormous number of traces. The direct storage of hundreds of thousands of traces heavily impacts users’ budgets. A sampling strategy is a procedure that decides which traces to store for further utilization. We describe a procedure that samples traces based on their types, duration, and/or errors. The final goal is to preserve sufficient information for application troubleshooting without possible distortion.

A trace type can be approximately identified by the root span or the span that arrived the earliest. More strict identification is based on the analysis of its structure. However, the latest method is rather resourceful and will require additional grouping/clustering approaches. The durations of traces (in milliseconds) are available in the corresponding metadata. It is possible to extract the durations for different trace types and estimate the average/typical durations of the corresponding processes. The information on errors is also available in trace metadata. Traces are composed of spans. We have detailed information regarding the spans in the corresponding tags that describe the micro-processes. One of the tags indicates the health of a span. We can assume that the entire trace is erroneous if at least one of its spans has a “True” error label. We suggest a parametric approach for all sampling scenarios, allowing us to modify the required percentage of compression rates accordingly.

5.1. Sampling by Trace Type

Assume that we know how to determine a trace type. Let M be the number of trace types for a specific period and $N_k, k = 1, \dots, M$ be the number of traces in each type. This information should be collected/updated for a recent period, as the number of trace types and the trace-flow velocity are subject to rapid changes in modern applications. Let N be the number of all traces before sampling:

$$N = \sum_{k=1}^M N_k, \quad (1)$$

and $p_k(\text{type}) = N_k/N, k = 1, \dots, M$ be the probability of the occurrence of the k -th type. Let $r_k(\text{type})$ and $N_k^*, k = 1, \dots, M$ be the sampling rate (the percentage/fraction of stored traces) and the number of traces after the sampling of the k -th type, respectively:

$$N_k^* = r_k(\text{type})N_k, \quad k = 1, \dots, M, \quad (2)$$

where N_k^* should be rounded to the closest integer. Also, we denote by N^* the total number of all traces after the sampling

$$N^* = \sum_{k=1}^M N_k^*. \quad (3)$$

The ratio $r = N^*/N$ is the final sampling rate. We propose to compute the sampling rate r_k as inversely proportional to the probability $p_k(type)$:

$$r_k(type) = 1 - (p_k(type))^\alpha, \quad \alpha > 0, \quad (4)$$

where α is a parameter that a user should determine to satisfy the needed final sampling rate r . Generally, the more frequent a trace type, the lower the corresponding sampling rate r_k .

Let us see how the parameter α should be determined if the final user requirement is known:

$$N^* = \sum_{k=1}^M N_k^* = \sum_{k=1}^M r_k N_k = \sum_{k=1}^M (1 - (p_k)^\alpha) N_k = N \sum_{k=1}^M p_k (1 - (p_k)^\alpha) = N G_{type}(\alpha), \quad (5)$$

where

$$r = N^*/N = G_{type}(\alpha) = \sum_{k=1}^M p_k(type)(1 - (p_k(type))^\alpha). \quad (6)$$

This can be considered an analog of the Gini index, showing the total sampling rate. For a specific dataset of traces, we can compute the values of $G(\alpha)$ across different α and determine its value to satisfy a user's requirement. We can try to estimate the proper value of α based on the recent collection of traces (say, for the last 2 h) and dynamically update those values to meet the requirement.

Let us illustrate how this procedure works for a specific application. Figure 2 shows the distribution of traces across different trace types.




Figure 2. The distribution of traces across different types for a specific application.

There are types with almost 4000 traces and very rare ones, like ones with 3 representatives. Overall, we detected 20,425 traces distributed across 28 different trace types. Figure 3 and Table 1 reveal how different rates compress traces across types. We experimented with $\alpha = 1, 0.5, 0.25$, and 0.1 . The sampling rate is always smaller for common trace types and larger for rare ones (see Table 1).

Smaller values of α correspond to smaller final sampling rates. When $\alpha = 1$, we store 17,903 traces across all types: 87.7% of all traces. When $\alpha = 0.5$, we store 13,629 traces, 66.7% of all traces. For $\alpha = 0.25$, we store 8882 traces, 43.5%. When $\alpha = 0.1$, we store 4273 traces, 20.9% of all traces.

Let us inspect some of the sampling rates for specific types. The most frequent type contains 3996 traces with the probability $p = 0.2$. The corresponding sampling rate (for $\alpha = 1$) is 80.5%, and 3215 traces from this class were stored. One of the rarest types contains

3 traces with almost zero probability. The sampling rate is almost 100%, and we store two or three traces for different α .




Figure 3. The sampling of traces of Figure 2 for different values of the parameter α .

Table 1. The number of traces before and after the sampling.

Type Index	Original Numbers	$\alpha = 1$	$\alpha = 0.5$	$\alpha = 0.25$	$\alpha = 0.1$
1	788	758	634	439	219
2	3996	3215	2229	1339	602
3	3778	3080	2154	1301	587
4	2082	1870	1418	906	426
5	249	246	222	167	89
6	2043	1839	1397	895	421
7	142	142	131	101	56
8	104	104	97	77	43
9	64	64	61	49	29
10	621	603	513	362	184
11	216	214	194	147	79
12	62	62	59	48	28
13	241	239	215	162	87
14	93	93	87	69	39
15	49	49	47	39	23
16	44	44	42	35	21
17	98	98	92	73	41
18	41	41	40	33	19
19	3	3	3	3	2
20	63	63	60	49	28
21	3	3	3	3	2
22	3028	2580	1863	1150	527
23	1413	1316	1042	689	332
24	23	23	23	19	12
25	546	532	457	326	166
26	58	58	55	45	26
27	51	51	49	40	23
28	526	513	442	316	162
Total	20,425	17,903	13,629	8882	4273
Rate	-	87.7%	66.7%	43.5%	20.9%

The correct value of α that will satisfy a user requirement could be estimated by computing the $G(\alpha)$ values. Figure 4 explains the procedure. Let us assume that the required sampling rate is 10%, meaning that we want to store around that percentage of traces. $\alpha = 0.044$ provides a sampling rate $r = 0.099$.

This algorithm should be applied if trace types carry important information for an application, and both rare and frequent types should be preserved for further analysis. Mathematically, it means no more than, say, 50 trace types with diverse frequencies. Diversity can be inspected via the Gini index or entropy. If those measures are close to zero, the sampling will try to maximize them and store equal portions from all types.




Figure 4. The values of $G_{type}(\alpha)$ show the final sampling rates for different α . The red-cross corresponds to $\alpha = 0.044$ with the final sampling rate $r = 0.099$ (around 10%).

5.2. Sampling by a Trace Duration

In this subsection, we consider the sampling of traces based on their durations. This setup is reasonable if we have a few trace types where traces have almost the same average durations. We would like to keep all traces with some average durations and those with extraordinarily short or long ones.

Assume N traces with durations $\{d_k\}_{k=1}^N$. Let $D(m)$ be the corresponding histogram of durations with m number of bins:

$$D(m) = \{n_1, \dots, n_m\}, \quad (7)$$

where n_s , $s = 1, \dots, m-1$ is the number of traces with durations within intervals $[t_{s-1}, t_s)$, $s = 1, \dots, m-1$ and n_m with durations within $[t_{m-1}, t_m]$. Let $p_s(dur)$, $s = 1, \dots, m$ be the probability that a trace has a duration from the s -th bin:

$$p_s(dur) = \frac{n_s}{N}, s = 1, \dots, m. \quad (8)$$

We determine the sampling rate $r_s(dur)$ of a trace with a duration within the s -th bin to be inversely proportional to the corresponding probability $p_s(dur)$:

$$r_s(dur) = 1 - (p_s(dur))^\beta, s = 1, \dots, m, \beta > 0, \quad (9)$$

where β is the parameter to meet a user's requirement.

Let n_s^* be the number of traces after sampling in the s -th bin:

$$n_s^* = r_s(dur)n_s. \quad (10)$$

Let N^* be the total number of traces after the sampling:

$$N^* = \sum_{s=1}^m n_s^* = \sum_{s=1}^m r_s(dur)n_s = N \sum_{s=1}^m (1 - (p_s)^\beta) p_s = N G_{dur}(\beta), \quad (11)$$

where

$$r = N^*/N = G_{dur}(\beta) = \sum_{k=1}^M p_k(dur) \left(1 - (p_k(dur))^\beta\right) \quad (12)$$

is the total sampling rate for different values of parameter β .

Let us illustrate how this procedure works for the same dataset of traces presented earlier. We will skip the trace types and will look into the trace durations.

Figure 5 shows the distribution of traces with durations in different time intervals. It shows three dominant time intervals in which the most traces are concentrated. We can aggressively sample those frequent traces and moderately sample the rare ones outside the dense areas. Table 2 shows how the sampling procedure works for different β parameter values (we use seven bins).




Figure 5. The distribution of traces with different durations (in milliseconds).

Table 2. The number of traces before and after the sampling.

Time Intervals	Original Numbers	$\beta = 1$	$\beta = 0.5$	$\beta = 0.25$	$\beta = 0.1$
[9, 1828)	6685	4498	2861	1629	707
[1828, 3647)	5061	3808	2542	1491	660
[3647, 5466)	6328	4368	2806	1607	700
[5466, 7284)	81	81	76	61	35
[7284, 9103)	235	233	210	159	85
[9103, 10922)	1111	1051	852	575	281
[10922, 12741]	925	884	729	499	247
Total	20,425	14,923	10,076	6021	2715
Rate	-	73.1%	49.3%	29.5%	13.3%

The first column of Table 2, “Time Intervals”, shows the seven intervals of duration. The intervals have equal lengths, although the bins could be nonuniform. The second column shows the initial number of traces with the durations within the corresponding time intervals. The remaining columns show the numbers after sampling with the corresponding β . The last row of Table 2 reveals the final sampling rates. The bigger the value of β , the more severe the reduction in traces. The value $\beta = 0.1$ provides a sampling rate of around 13%. Figure 6 visualizes the results of Table 2.

If the requirement is to sample exactly 10% of traces, then Figure 7 shows the parameter β estimation based on the values of $G_{dur}(\beta)$. The red cross indicates that $\beta = 0.083$ provides a sampling rate of 0.0996.




Figure 6. The sampling of traces of Figure 5 for different values of the parameter β .




Figure 7. The values of $G_{dur}(\beta)$ show the sampling rates for the different parameter values β . The red cross corresponds to $\beta = 0.083$ with a total sampling rate of 0.0996 (around 10%).

Proper histogram construction is a crucial milestone for frequency-based sampling. In general, the problem is the impact of outliers on the bin construction process. Standard procedures use an equidistant split of the data range, and data outliers can enlarge it, resulting in big bins with resolution distortion. A more accurate procedure should involve outlier detection to preserve the data in separate bins with further application of the classical procedure for the main part of the data. We consider outlier detection via a modified MAD (median absolute deviation) algorithm [74] with a slight modification. We define upper and lower baselines as 0.9 and 0.1 quantiles of data:

$$M(up) = q_{0.9}(data), \quad M(low) = q_{0.1}(data), \quad (13)$$

where $q_s(data)$ is the s -th quantile of the data with $0 \leq s \leq 1$. We calculate the upper and lower distances:

$$dist(up) = |data(up) - M(up)|, \quad dist(low) = |data(low) - M(low)|, \quad (14)$$

where $data(up)$ are data points bigger than or equal to $M(up)$ and $data(low)$ are data points smaller than or equal to $M(low)$. We set upper and lower MADs as

$$MAD(up) = q_{0.8}(dist(up)), \quad MAD(low) = q_{0.8}(dist(low)). \quad (15)$$

Based on the MADs, upper and lower thresholds are defined as

$$upper = \min(M(up) + 2.5MAD(up), \max(data)), \quad (16)$$

and

$$lower = \max(M(low) - 2.5MAD(low), \min(data)). \quad (17)$$

All data points lying above or below than the corresponding thresholds are assumed to be upper and lower outliers, respectively. We cut outliers from the data and construct the classical histogram for the remaining data with some predefined numbers of bins (say bins = 5). Then, we calculate the number of lower and upper outliers and append them to the main histogram from left and right, respectively.

Hence, if data have small-value outliers, then the first bin of the histogram contains the number of data points within the interval $[data(min), upper]$.

If data have big-value outliers, the last bin contains the number of data points within the interval $(upper, \max(data)]$. The main part of the histogram relates to data points within the interval $[lower, upper]$. This procedure allows the proper construction of the corresponding histograms with a small number of bins. It is worth noting that the main part of the histogram consists of uniform intervals, while the widths of the first and last bins could be different from those bins. Finally, if data have outliers from both sides, the corresponding histogram has $bins + 2$ final bins. If data have outliers from only one side, then the final number of bins is $bins + 1$. No other bins are added to the classical histogram if data come without outliers.

5.3. Hybrid Approach Based on Types and Durations

The hybrid approach takes into account both trace type and duration information. This algorithm performs accurate sampling if both durations and types are important. Assume that traces are already grouped into trace types. We generate the histogram of durations for each group and apply the procedure described in the previous subsection. We also consider the probability of a trace type. Common types will be sampled more aggressively.

Let

$$H^{(k)} = \{n_1^{(k)}, \dots, n_m^{(k)}\}, \quad k = 1, \dots, M, \quad (18)$$

be the histogram of durations of the k -th trace type and $n_s^{(k)}$ be the number of traces in the s -th bin of the k -th type. As above, M is the number of different types, and m is the number of bins in the histogram. Let N_k be the number of traces in the k -th trace type:

$$N_k = \sum_{s=1}^m n_s^{(k)}. \quad (19)$$

Let N be the total number of traces:

$$N = \sum_{k=1}^M N_k. \quad (20)$$

Let

$$P^{(k)} = \{p_1^{(k)}, \dots, p_m^{(k)}\}, \quad p_s^{(k)} = \frac{n_s^{(k)}}{N_k}, \quad k = 1, \dots, M, \quad (21)$$

and

$$P = \{p_1, \dots, p_M\}, \quad p_k = \frac{N_k}{N}. \quad (22)$$

We denote the sampling rate of a trace from the s -th bin of the k -th trace type by $r_s^{(k)}$. It shows the fraction of traces that should be stored. We compute it by the following formula as inversely proportional to the corresponding probabilities:

$$r_s^{(k)} = 1 - (p_k)^\alpha \left(p_s^{(k)}\right)^\beta, \quad (23)$$

where $\alpha, \beta \geq 0$ are some unknown parameters to be tuned to meet the requirement for the final sampling rate r .

Let us show how it could be carried out. Let N_k^* be the number of traces in the k -th trace type after sampling. Let $n_s^{(k)*}$ be the number of traces in the s -th bin and the k -th trace type after sampling:

$$N_k^* = \sum_{s=1}^m n_s^{(k)*} = \sum_{s=1}^m n_s^{(k)} r_s^{(k)} = N_k \sum_{s=1}^m p_s^{(k)} \left(1 - (p_k)^\alpha (p_s^{(k)})^\beta\right). \quad (24)$$

Let N^* be the total number of traces after sampling across all types and durations:

$$N^* = \sum_{k=1}^M N_k^* = \sum_{k=1}^M N_k \sum_{s=1}^m p_s^{(k)} \left(1 - (p_k)^\alpha (p_s^{(k)})^\beta\right) = NG(\alpha, \beta), \quad (25)$$

where

$$r = \frac{N^*}{N} = G(\alpha, \beta) = \sum_{k=1}^M \sum_{s=1}^m p_k p_s^{(k)} \left(1 - (p_k)^\alpha (p_s^{(k)})^\beta\right) \quad (26)$$

is the required final sampling rate that can be accomplished by appropriately selecting parameters α and β .

First, we will illustrate this procedure when $\alpha = \beta$. Figure 8 reveals the procedure for a specific trace type containing 3996 traces. The left figure shows the scatter plot of durations. The right figure shows the histogram of durations with 7 bins before and after the samplings. The value $\beta = 0.1$ corresponds to the sampling that stores 1104 traces. The next value $\beta = 0.05$ stores 599 traces. The value $\beta = 0.025$ corresponds to the sampling with 315 preserved traces. Finally, the value $\beta = 0.01$ stores 131 traces. The last one corresponds to the 3.3% sampling rate. In addition, the probability of this trace type is 0.2, which also impacts the sampling rates with $\alpha = \beta$ setting.

Figure 9 shows the hybrid procedure for another trace type with the probability 0.0048. It contains only 98 traces and is a rare type compared to the previous example. It also corresponds to the setting $\alpha = \beta$. The right figure shows the sampling results for different values of β . For example, the value $\beta = 0.01$ corresponds to the sampling that stores only 7 traces of this specific type. It corresponds to the 7.1% sampling rate. The value $\beta = 0.025$ corresponds to the 18.4% sampling rate. The value $\beta = 0.05$ stores 31 traces and corresponds to the 31.6% sampling rate. Finally, the value $\beta = 0.1$ corresponds to the sampling rate 52%.




Figure 8. The hybrid approach for a specific trace type (N2 in Figure 10). The left figure shows the plot of durations. The right figure shows the counts of traces in different bins before and after sampling corresponding to different values of $\alpha = \beta$.




Figure 9. The hybrid approach for a specific trace type (N17 in Figure 10). The left figure shows the plot of durations. The right figure shows the counts of traces for $\alpha = \beta$.

Figure 10 shows the result of the hybrid approach across different trace types when $\alpha = \beta$. The total sampling rate is 4% for $\alpha = \beta = 0.01$. The values $\alpha = \beta = 0.1$ correspond to the sampling rate 31%. This means that the values of α and β can control the sampling rates in a wide range of values. If we need to accomplish a strict requirement, we can turn to the $G(\alpha, \beta)$ values corresponding to the sampling rates. Figure 11 shows the values of $G(\beta, \beta)$, where the red cross corresponds to the sampling rate of 10% (the exact value of G is 0.099) with $\beta = 0.029$. We can also tune the values of α and β independently. Figure 12 shows the surface of sampling rates corresponding to different values. For example, the total sampling rate of 10% can be accomplished by $\alpha = 0.03, \beta = 0.022$, or $\alpha = 0.026, \beta = 0.029$, or $\alpha = 0.011, \beta = 0.056$, or $\alpha = 0.023, \beta = 0.035$, or $\alpha = 0.01, \beta = 0.058$, etc.




Figure 10. The hybrid sampling approach for $\alpha = \beta$.




Figure 11. The sampling rates corresponding to different values $\alpha = \beta$. The red cross corresponds to 10% with $\beta = 0.03$.




Figure 12. The sampling rates that correspond to different values of α and β . The colors correspond to the different ranges of sampling rates for a better visualization.

Parameter optimization can be performed without the utilization of long historical data. We can take an initially random value for the parameters and, after each hour, verify the actual compression ratio. Then, by increasing or decreasing the values, we achieve the required sampling. This will especially work for dynamic applications when relying on available historical information is impossible.

How, in practice, can we sample a trace? When a specified trace (with known type and duration) is detected with a known sampling rate r , a random variable from the $Bernoulli(r)$ distribution should be generated. This random variable will have a Boolean outcome with the value 1 that has probability r and value 0 with the probability $1 - r$. If the outcome is 1, we will store the specified trace; otherwise, we will ignore it. This will allow us to store the traces with the required sampling rates in the long run.

Returning to the problem of histogram construction, we refer to a powerful approach known as the t-digest algorithm [75], which addresses several problems in histogram construction. The first concern corresponds to the storing of time series data with trace durations. For each trace type, we need to store the corresponding time series of durations for further histogram construction, and we need those time series with sufficient statistics. For 28 trace types, as in our experiments, we need to store the 28 time series. The second concern corresponds to the process of histogram construction. Each time we need those histograms, sorting should be applied to the corresponding duration data, and the procedure must be repeated each time in case of some new arrivals. The concern is the impact of outliers. An efficient solution to the mentioned problems is the t-digest algorithm. Instead of storing the entire time series data, it stores only the result of data cluster centroids and data counts in each cluster. The efficient merging approach allows for the combination of different t-digest histograms, streamlining the entire process. The t-digest is very precise while estimating extreme data quantiles (close to 0 and 1), making the procedure robust to outliers.

5.4. The Sampling of Erroneous Traces

In specific frameworks, sampling should preserve other important properties besides duration and type. One such property is the normality/abnormality of a trace, which

characterizes microservices' performance and could be used for troubleshooting and root-cause analysis. In the described approach, we cannot guarantee a sufficient number of important erroneous traces in the sampled dataset, as there was no requirement to preserve them. Now, we will try to address that requirement by discussing various approaches.

The first approach is the natural modification of the previous one, ensuring the existence of enough erroneous traces in each trace type after the sampling. More precisely, we verify the existence of erroneous traces in each trace type. If they exist, we divide the corresponding type into two types containing only normal or erroneous traces. Then, we can apply the approach that has already been considered (sampling by type, duration, or hybrid) to the same dataset but with renewed trace types.

Dividing erroneous traces into additional types according to the corresponding error codes is also possible. Assume that trace type "A" has erroneous traces containing error codes "400-Bad Request" and "401-unauthorized". Then, we can divide it into three new trace types: "A-we" (without errors), "A-400" (with 400 code), and "A-401" (with 401 code). It is also possible to divide erroneous traces into types by their error codes independent of the types. For example, collect all traces with "400" error codes together in a single type, independent of their original types. Let us show how this approach works.

Assume two trace types; for example, 300 traces of type "A" and 130 of type "B", and we must store only 30% of the traces. The sampling, based only on the trace types, provides the value $\alpha = 0.619$. It means storing 60 traces of type "A" and 69 of type "B". Figure 13 illustrates those numbers.

Now, let us also consider erroneous traces. Assume type "A" has 200 normal and 100 erroneous traces. Thus, types "A" and "A*" contain 200 and 100 traces, respectively. Then, assume type "B" has 60 normal and 70 erroneous traces. Thus, "B" and "B*" contain 60 and 70 traces, respectively. Now, instead of two types, there are four. The type-based sampling provides the value $\alpha = 0.283$, which leads to sampling rates 0.2, 0.43, 0.34, and 0.4, respectively.




Figure 13. The sampling of two trace types without counting the errors.

Figure 14 shows the distributions before and after sampling. As a result, from 260 normal and 170 erroneous traces, the approach sampled 66 and 63, respectively (see the right figure of Figure 14). Now, we can guarantee that the final sampled set also contains erroneous traces across different types. It is possible to apply the hybrid approach that will also consider the duration of the erroneous traces.




Figure 14. The sampling of two trace types also counts the errors.

The second approach tries to control the percentage of erroneous traces better in the sampled dataset. Let $0 < h < 1$ be the final required sampling rate. Assume that h_e and h_n are the sampling rates of erroneous and normal traces, respectively:

$$N_n^* = h_n N_n, \quad (27)$$

and

$$N_e^* = h_e N_e, \quad (28)$$

where N_n and N_e are the number of normal and erroneous traces before sampling, respectively, and N_n^* and N_e^* after sampling.

Can we also put some requirements on h_e ? We have

$$h_n \frac{N_n^*}{N_n} + h_e \frac{N_e^*}{N_e} = h, \quad (29)$$

and

$$h_n = \frac{N}{N_n} \left(h - h_e \frac{N_e}{N} \right). \quad (30)$$

If the requirements on h and h_e lead to the value $0 < h_n < 1$, then the requirements can be accomplished. Otherwise, if h_n turns out to be a negative number, that means it is impossible, and we must ask to change the requirements.

Let us return to the previous example with two trace types. Can we sample 30% of traces while preserving 60% of erroneous ones? We input $h = 0.3$ and $h_e = 0.6$ and tried to find appropriate values for h_n . Our calculations show that $h_n = 0.104$ will work. We will preserve 60% of erroneous traces and 10% of normal traces by sampling 30% of all traces. Figure 15 illustrates the choices.




Figure 15. The sampling of two trace types with stricter requirements on the percentage of erroneous traces. Now, we preserve 10% of normal traces and 60% of erroneous ones. The final sampling rate is 30%.

6. Troubleshooting of Applications

The ultimate goal of DT is to monitor application performance, detect malfunctioning microservices, and explain the root causes of problems so that they can be resolved quickly. This is feasible by inspecting the tracing traffic passing through specific microservices, detecting erroneous traces/spans, and trying to explain their origin. Thus, the explainability of ML models is the most crucial property. We consider two highly explainable approaches.

One is RIPPER [20], which is state-of-the-art in inductive rule learning [53]. It has important technical characteristics, like supporting missing values and numerical and categorical variables, and multiple classes. We experimented with Weka's RIPPER implementation, or JRip [76]. The application of RIPPER for the inspection of erroneous traces is discussed in [7]. We are not going into such details, but will show the impact of noise reduction on the rules.

Next is the Dempster–Shafer classifier considered in [66]. It is a far more complex approach but has interesting implications. One important characteristic is the ability to measure the uncertainty of specific rules. It should be interesting to compare the uncertainties of the same rules before and after the sampling.

The analysis of tracing traffic passing through a specific microservice starts with data preprocessing. We transform the trace traffic into tabular data to apply ML algorithms. A single trace shows an individual request passage through the microservices. It contains a series of tagged time intervals known as spans. A span contains metadata known as tags and application tags for better process resolution. The traffic can contain hundreds or thousands of traces with different numbers of spans and tags, which are application-specific.

In our examples [7], traces contain the names as trace-IDs, and then after each trace, the list of spans with process names, and after each span, the corresponding tags. During the preprocessing stage, we remove some of the fields containing redundant information, like “traceID”, “spanID”, “startMs”, and many others. It should be very straightforward to denoise trace traffic based on expert knowledge or user feedback. Firstly, we preprocess span names. We make a list of all distinct span names existing in trace traffic and put them as the names of the columns of a dataframe. Then, for each trace (as a row), we verify the names of its spans, and in the corresponding columns, enter value = 1. The other columns contain missing values, as that specific trace does not contain those spans. The entire dataframe consists of ones and missing values. Fortunately, the RIPPER and Dempster–Shafer classifiers ignore those missing values and explain the output based on the existing spans in a trace. Eventually, the number of rows in the dataframe coincides with the number of traces in trace traffic, and the number of columns coincides with the number of distinct spans (processes).

We label traces via the tag “error”, where the value “true” indicates that the corresponding trace is erroneous. A missing tag “error” indicates that the trace is normal. It can also be useful to construct a column in a dataframe that will indicate the type of trace. Secondly, we construct columns revealing metadata information. We put the metadata value in the corresponding cell, where the row corresponds to the trace and the column corresponds to the distinct metadata name.

Tracing traffic passing through a malfunctioning microservice can still contain many traces, even after a series of samplings. In this final stage, we apply the sampling of erroneous traces for reducing the volume and simplifying the application of rule induction methods. We aim to understand how the sampling impacts the rule-generation process regarding precision, recall, or uncertainty. It is worth noting that in this stage, the goal of the sampling is not to preserve all possible types and durations of erroneous traces. Exactly the opposite: the goal should be to remove rare erroneous traces and preserve dominant/common ones, as they probably explain the problems of a microservice. We separate normal and erroneous traces by trace types. Then, in all those groups, we sample at the same rate. As a result, some of the groups (rare ones) can vanish. Error-based sampling may downplay the influence of rare events by sampling them less frequently or

excluding them altogether. We show how this reduction in outliers helps to stabilize the analyses and make recommendations more robust to extreme cases.

Then, we apply rule-induction classifiers to explain the origin of erroneous traces. These learning systems generate understandable rules based on input data. Humans can easily comprehend these rules, which provide clear insights into the decision-making process of the AI model. Unlike black-box ML models such as deep neural networks, where it is often challenging to understand how the model arrives at its predictions, rule induction systems offer transparency. Users can inspect the rules to understand the underlying logic and reasoning behind the rules. This enhances the trustworthiness of the AI system.

By interpreting the induced rules, developers and system administrators can gain insights regarding the possible remediation of issues. The rules provide transparent and interpretable explanations of certain errors occurring under specific conditions or scenarios. Armed with insights from the RCA, appropriate measures can be taken to optimize the system, fix bugs, or implement preventive measures to reduce the likelihood of similar errors occurring in the future.

Let us apply this procedure to a real customer cloud environment. The visualization of tracing traffic is known as an application map. It shows which microservices malfunction and which traces must be collected for troubleshooting. We collected 5899 traces with some 4917 features that should explain the origin of errors. We had 3145 normal and 2754 erroneous traces. JRip (RIPPER) ran for 11.09 s, outputting 13 rules (see Figure 16).

The first 12 rules describe the erroneous traces. They are mostly connected with the type. Only one of the rules is connected with the tag “-annotations-matched”. The fraction at the end of the rule shows how many traces were fired by the rule (numerator) and how many misfired (denominator). The first two rules have large coverage and are 100% precise. The others should be connected with the noise in the dataset. We hope that the sampling will remove the noise. However, the accuracy of this classifier is 99.4%. The precision and recall of both classes are also bigger than 99%.

Figure 17 shows the distribution of normal traces across different types before (the top figure) and after (the bottom figure) the sampling. There are many types containing just a single representative. The sampling preserves only the common groups and removes all rare types. Here, we applied the 10% sampling rate. Similarly, Figure 18 shows the distribution of erroneous traces. In both figures, the labels on the horizontal axes show the names of trace types. We identified the trace type by the root span.

Figure 19 reveals the JRip rules after the sampling, which preserved the first two important rules and removed the others connected with the noise. The sampled dataset contained 300 normal and 268 erroneous traces. The classifier applied to the sampling dataset showed 99.1% accuracy. It misclassified 5 traces. The execution time was 0.1 s. As a result, the sampling reduced the execution time, especially for noisy datasets. Moreover, it shortened the list of top recommendations so that users can focus on the most important ones.

As we mentioned before, the sampling degraded the statistical evidence. The classifiers could not describe it. Quite the opposite; the accuracy of classifiers increased or remained almost the same. We can refer to the Dempster–Shafer theory and compare the uncertainties of the rules. The theory of belief functions, also referred to as evidence theory or Dempster–Shafer theory, is a general framework for reasoning with uncertainty. It offers an alternative to traditional probabilistic theory for the mathematical representation of uncertainty. Figure 20 shows how the sampling affects the uncertainties. It shows the uncertainties of the first two rules before and after the sampling.

JRIP rules before the sampling:

(type = Zipkin.ad-selector-canary.user-db-with-retry.getuserinfostruct) → **output = 1** (1317/0)

(type = Zipkin.ad-selector-canary.user-db.getuserinfostruct) → **output = 1** (1321/0)

(type = Zipkin.ad-selector-canary.infer-scores) → **output** = 1 (21/3)

(type = Zipkin.ad-selector-canary.get-native-ads) → **output** = 1 (22/4)

(type = Zipkin.prod-user-db-data-0.impression-scan) → **output** = **1** (14/0)

(Zipkin.ad-selector-canary.get-native-ads_annotations__matched = false) → **output = 1** (12/0)

(type = Zipkin.user-db-api-prod.getuserinfostruct) → **output** = **1** (10/0)

(type = Zipkin.prod-user-db-data-1.impression-scan) → **output** = 1 (8/0)

(type = Zipkin.prod-user-db-data-2.impression-scan) → **output** = 1 (6/0)

(type = Zipkin.ad-selector-canary.run-filters) → **output** = 1 (2/0)

(type = Zipkin.prod-user-db-data-3.impression-scan) → **output** = 1 (2/0)

(type = Zipkin.user-db-api-prod.user-db.getuserinfostruct) → **output** = 1 (2/0)

=> output = 0 (3162/24)

Figure 16. JRip rules before the sampling.




Figure 17. The distribution of normal traces across the types before and after the sampling.




Figure 18. The distribution of erroneous traces across the types before and after the sampling.

JRIP rules after the sampling:

(type = Zipkin.ad-selector-canary.user-db-with-retry.getuserinfostruct) → **output = 1 (131/0)**

(type = Zipkin.ad-selector-canary.user-db.getuserinfostruct) → **output = 1 (132/0)**

=> **output = 0 (305/5)**

Figure 19. JRip rules after the sampling.

Uncertainties before the sampling:

(type = Zipkin.ad-selector-canary.user-db-with-retry.getuserinfostruct) →

uncertainty = 0.00022

(type = Zipkin.ad-selector-canary.user-db.getuserinfostruct) →

uncertainty = 0.00041

Uncertainties after the sampling:

(type = Zipkin.ad-selector-canary.user-db-with-retry.getuserinfostruct) →

uncertainty = 0.0048

(type = Zipkin.ad-selector-canary.user-db.getuserinfostruct) →

uncertainty = 0.061

Figure 20. The uncertainties of rules before and after the sampling.

7. Conclusions and Future Work

DT is essential for increasing the visibility of complex interactions and dependencies within native cloud applications. By tracing the requests across microservices, containers, and dynamic environments, users can effectively monitor, manage, and troubleshoot their native cloud applications to ensure optimal performance and reliability. Despite the efforts, the adoption of this technology encounters several challenges, and one of the most crucial is the volume of data and the corresponding resource consumption needed to handle it.

Sampling is a technique that alleviates the overhead in collecting, storing, and processing vast amounts of trace data. It reduces the volume of data by selectively capturing only a fraction of traces. This approach offers benefits such as decreased latency, reduced resource consumption, and improved scalability. However, sampling introduces challenges, particularly in maintaining representative samples and preserving the accuracy of analysis results. Striking a balance between sampling rate and data fidelity is crucial to ensure effective troubleshooting and performance analysis in distributed systems. Overall, DT sampling plays a vital role in managing the complexity of distributed environments while optimizing resource utilization and maintaining analytical efficacy.

We explored several approaches for sampling that preserved traces with some specific properties. One such property was the type of trace that described the transaction similarity. Traces of the same type should normally have almost the same structure. The goal of such sampling was to preserve traces across all available types. Another property was the duration of the trace. This approach could or might not be combined with the information regarding the type. Trace durations were important as they described the transaction duration. Similar transactions had some typical/ average durations. Atypical durations indicated a transaction/microservice malfunction. The next important trace characteristic was its normality. Erroneous traces carried important information regarding problems. Hence, it was natural to try to keep all representatives while monitoring the performance of microservices. The flow of erroneous traces would show which microservices had degraded performance. Further, those traces were critical sources of information for troubleshooting application issues.

We sampled only dominant/common errors at the troubleshooting stage, trying to remove the rare ones. This removed the noise and made explanations more confident. RCA could be performed by tracing traffic passing through a malfunctioning microservice. Rule-learning ML methods could help generate explicit rules that explain the problems and clarify the remediation process. We showed how rule-induction systems like RIPPER solved this problem and provided recommendations that system administrators could follow to accelerate the resolution process. We also showed that sampling could dramatically decrease the program execution time and provide more clear recommendations. However, as mentioned before, sampling degraded the statistics, and sometimes rare but important evidence could escape the analysis.

Several limitations of our approach are crucial to consider. Acknowledging and addressing these limitations is vital for enhancing the robustness and applicability of the sampling strategy. Firstly, a significant limitation is the dependence of trace types on the root span. This approach overlooks the fact that a substantial volume of traces lack a root span, necessitating the utilization of the first span instead. Overall, this introduces potential inaccuracies in type definitions, as many traces may exhibit distinct structures not adequately captured by this method. Secondly, the procedure for calculating trace durations may be prone to misinterpretation due to the challenges associated with defining trace types, as previously discussed. Thirdly, reliance on the system-defined definition of erroneous traces may not always align directly with the underlying issues that rule induction systems aim to address. This mismatch could hinder the effectiveness of the rule induction process by focusing on inaccurately identified erroneous traces. Fourthly, the current sampling strategy is primarily based on trace types and durations, overlooking potentially crucial properties for diverse applications. Lastly, the approach's scalability is

constrained by its lack of full streaming capability, which could pose challenges in efficiently handling large volumes of trace data.

In future work, we plan to explore the scalability of our strategy to handle larger datasets or more complex applications. We must investigate how our strategy performs as the volume of data increases, and consider further adjustments to maintain efficiency. Our sampling strategy is based on the statistical properties of application traces. We must delve deeper into the process of refining this by exploring different ML solutions to see how they influence the efficiency of the process. This can involve experimenting with various algorithms and parameters to optimize the sampling approach. It is important to develop specific evaluation metrics to quantify the impact of the sampling strategy on rule induction methods. These could include accuracy, efficiency, the interpretability of the rules generated, or any other relevant performance indicators. Based on those indicators, we should conduct comparative studies to compare the performance of our sampling strategy with existing sampling methods or other rule induction techniques. This can highlight the peculiarities of our approach. It is also important to examine related research areas, such as predictive modeling, anomaly detection, or decision support systems, to identify the impact of the sampling strategies. This will help to broaden the scope of our research and uncover new avenues for exploration.

8. Patents

Poghosyan A., Harutyunyan A., Grigoryan N., Pang C., Oganessian G., Baghdasaryan D., Automated methods and systems that facilitate root-cause analysis of distributed-application operational problems and failures by generating noise-subtracted call-trace-classification rules. Filed: 1 October 2021. Application No.: US 17/492,099. Patent No.: US 11880272 B2. Granted: 23 January 2024.

Poghosyan A., Harutyunyan A., Grigoryan N., Pang C., Oganessian G., Baghdasaryan D., Automated methods and systems that facilitate root cause analysis of distributed-application operational problems and failures. Filed: 1 October 2021. Application No.: US 17/491,967. Patent No.: US 11880271 B2. Granted: 23 January 2024.

Poghosyan, A., Harutyunyan, A., Grigoryan, N., Pang, C., Oganessian, G., and Avagyan, K., Methods and systems for intelligent sampling of application traces. Application filed by VMware LLC in 2021. Application No.: US 17/367,490. Patent No.: US 11940895 B2. Granted: 26 March 2024.

Poghosyan, A.V., Harutyunyan, A.N., Grigoryan, N.M., Pang, C., Oganessian, G., and Avagyan, K., Methods and systems for intelligent sampling of normal and erroneous application traces. Application filed by VMware LLC in 2021. Application No.: US 17/374,682. Publication of US 20220291982 A1 in 2022.

Poghosyan, A.V., Harutyunyan, A.N., Grigoryan, N.M., Pang, C., Oganessian, G., and Baghdasaryan, D., Automated methods and systems that facilitate root cause analysis of distributed-application operational problems and failures. Application filed by VMware LLC in 2021. Application No.: US 17/491,967 and 17/492,099.

Grigoryan N.M., Poghosyan A., Harutyunyan A.N., Pang C., Nag D.A., Methods and Systems that Identify Dimensions Related to Anomalies in System Components of Distributed Computer Systems using Clustered Traces, Metrics, and Component-Associated Attribute Values. Filed: 12 December 2020. Application No.: US 17/119,462. Patent No.: US 11,416,364 B2. Granted: 16 August 2022.

Author Contributions: Conceptualization, A.P., A.H. and N.B.; Data curation, E.D. and K.P.; Formal analysis, A.P. and A.H.; Investigation, A.P. and A.H.; Methodology, A.P. and A.H.; Project administration, N.B.; Resources, N.B.; Software, E.D. and K.P.; Supervision, N.B.; Validation, A.P. and A.H.; Visualization, E.D. and K.P.; Writing—original draft, A.P. and A.H.; Writing—review and editing, A.P., A.H., E.D., K.P. and N.B. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by ADVANCE Research Grants from the Foundation for Armenian Science and Technology.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are available on request from the corresponding author. The data are not publicly available due to privacy.

Acknowledgments: We sincerely appreciate the reviewers' commitment, exceptional dedication, and invaluable contributions. Their efforts and insightful feedback have greatly enriched the quality and depth of this paper.

Conflicts of Interest: Author Edgar Davtyan was employed by the company Picsart. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest. The funders had no role in the design of the study, in the collection, analyses, or interpretation of data, in the writing of the manuscript, or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
AIOps	AI for IT Operations
API	Application Programming Interface
AWS	Amazon Web Services
DST	Dempster–Shafer Theory
IT	Information Technologies
MAD	Median Absolute Deviation
ML	Machine Learning
MTTD	Mean Time to Detect
MTTR	Mean Time to Repair
SRE	Site Reliability Engineer
DT	Distributed Tracing
XAI	Explainable Artificial Intelligence

References

1. Parker, A.; Spoonhower, D.; Mace, J.; Sigelman, B.; Isaacs, R. *Distributed Tracing in Practice: Instrumenting, Analyzing, and Debugging Microservices*; O'Reilly Media, Incorporated: Sebastopol, CA, USA, 2020.
2. Shkuro, Y. *Mastering Distributed Tracing: Analyzing Performance in Microservices and Complex Systems*; Packt Publishing: Birmingham, UK, 2019.
3. Opentracing. What Is Distributed Tracing? 2019. Available online: <https://opentracing.io/docs/overview/what-is-tracing/> (accessed on 26 January 2021).
4. Cai, Z.; Li, W.; Zhu, W.; Liu, L.; Yang, B. A real-time trace-level root-cause diagnosis system in Alibaba datacenters. *IEEE Access* **2019**, *7*, 142692–142702. [CrossRef]
5. Liu, D.; He, C.; Peng, X.; Lin, F.; Zhang, C.; Gong, S.; Li, Z.; Ou, J.; Wu, Z. MicroHECL: High-Efficient Root Cause Localization in Large-Scale Microservice Systems. In Proceedings of the 2021 IEEE/ACM 43rd International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP), Madrid, Spain, 25–28 May 2021; pp. 338–347.
6. Poghosyan, A.V.; Harutyunyan, A.N.; Grigoryan, N.M.; Pang, C. Root Cause Analysis of Application Performance Degradations via Distributed Tracing. In Proceedings of the Third CODASSCA Workshop, Yerevan, Armenia: Collaborative Technologies and Data Science in Artificial Intelligence Applications, Yerevan, Armenia, 23–26 August 2022; pp. 27–31.
7. Poghosyan, A.; Harutyunyan, A.; Grigoryan, N.; Pang, C. Distributed Tracing for Troubleshooting of Native Cloud Applications via Rule-Induction Systems. *JUCS J. Univers. Comput. Sci.* **2023**, *29*, 1274–1297. [CrossRef]
8. Distributed Tracing—Past, Present and Future. 2023. Available online: <https://www.zerok.ai/post/distributed-tracing-past-present-future> (accessed on 25 June 2024).
9. Young, T.; Parker, A. *Learning OpenTelemetry*; O'Reilly Media: Sebastopol, CA, USA, 2024.
10. Cotroneo, D.; De Simone, L.; Liguori, P.; Natella, R. Run-time failure detection via non-intrusive event analysis in a large-scale cloud computing platform. *J. Syst. Softw.* **2023**, *198*, 111611. [CrossRef]
11. Zhang, X.; Lin, Q.; Xu, Y.; Qin, S.; Zhang, H.; Qiao, B.; Dang, Y.; Yang, X.; Cheng, Q.; Chintalapati, M.; et al. Cross-dataset Time Series Anomaly Detection for Cloud Systems. In Proceedings of the 2019 USENIX Annual Technical Conference (USENIX ATC 19), Renton, WA, USA, 10–12 July 2019; pp. 1063–1076.

12. Abad, C.; Taylor, J.; Sengul, C.; Yurcik, W.; Zhou, Y.; Rowe, K. Log correlation for intrusion detection: A proof of concept. In Proceedings of the 19th Annual Computer Security Applications Conference, Las Vegas, NV, USA, 8–12 December 2003; pp. 255–264.
13. Suriadi, S.; Ouyang, C.; van der Aalst, W.; ter Hofstede, A. Root cause analysis with enriched process logs. In Proceedings of the Business Process Management Workshops, International Workshop on Business Process Intelligence (BPI 2012), Tallin, Estonia, 3–6 September 2012; pp. 174–186.
14. BigPanda. Incident Management. 2020. Available online: <https://docs.bigpanda.io/docs/incident-management> (accessed on 26 January 2021).
15. Josefsson, T. Root-Cause Analysis through Machine Learning in the Cloud. Master's Thesis, Uppsala Universitet, Uppsala, Sweden, 2017.
16. Tak, B.; Tao, S.; Yang, L.; Zhu, C.; Ruan, Y. LOGAN: Problem diagnosis in the cloud using log-based reference models. In Proceedings of the 2016 IEEE International Conference on Cloud Engineering (IC2E), Berlin, Germany, 4–8 April 2016; pp. 62–67.
17. Mi, H.; Wang, H.; Zhou, Y.; Lyu, M.R.; Cai, H. Localizing root causes of performance anomalies in cloud computing systems by analyzing request trace logs. *Sci. China Inf. Sci.* **2012**, *55*, 2757–2773. [CrossRef]
18. Poghosyan, A.; Harutyunyan, A.; Grigoryan, N.; Pang, C.; Oganessian, G.; Ghazaryan, S.; Hovhannisyan, N. An Enterprise Time Series Forecasting System for Cloud Applications Using Transfer Learning. *Sensors* **2021**, *21*, 1590. [CrossRef] [PubMed]
19. Barredo Arrieta, A.; Díaz-Rodríguez, N.; Del Ser, J.; Benetot, A.; Tabik, S.; Barbado, A.; Garcia, S.; Gil-Lopez, S.; Molina, D.; Benjamins, R.; et al. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf. Fusion* **2020**, *58*, 82–115. [CrossRef]
20. Cohen, W.W. Fast Effective Rule Induction. In Proceedings of the Twelfth International Conference on Machine Learning, Tahoe City, CA, USA, 9–12 July 1995; Morgan Kaufmann: San Francisco, CA, USA, 1995; pp. 115–123.
21. Quinlan, J.R. *C4.5: Programs for Machine Learning*; Elsevier: Amsterdam, The Netherlands, 2014.
22. Jaeger: Sampling. 2024. Available online: <https://www.jaegertracing.io/docs/1.55/sampling/> (accessed on 25 June 2024).
23. Anomaly Detection in Zipkin Trace Data. 2024. Available online: <https://engineering.salesforce.com/anomaly-detection-in-zipkin-trace-data-87c8a2ded8a1/> (accessed on 25 June 2024).
24. LightStep: Sampling, Verbosity, and the Case for (Much) Broader Applications of Distributed Tracing. 2024. Available online: <https://medium.com/lightstephq/sampling-verbosity-and-the-case-for-much-broader-applications-of-distributed-tracing-f3500a174c17> (accessed on 25 June 2024).
25. Datadog: Trace Sampling Use Cases. 2024. Available online: https://docs.datadoghq.com/tracing/guide/ingestion_sampling_use_cases/ (accessed on 25 June 2024).
26. Partial Trace Sampling: A New Approach to Distributed Trace Sampling. 2024. Available online: <https://engineering.dynatrace.com/blog/partial-trace-sampling-a-new-approach-to-distributed-trace-sampling/> (accessed on 25 June 2024).
27. New Relic: Technical Distributed Tracing Details. 2024. Available online: <https://docs.newrelic.com/docs/distributed-tracing/concepts/how-new-relic-distributed-tracing-works/#sampling> (accessed on 25 June 2024).
28. OpenTelemetry Trace Sampling. 2024. Available online: <https://docs.appdynamics.com/observability/cisco-cloud-observability/en/application-performance-monitoring/opentelemetry-trace-sampling> (accessed on 25 June 2024).
29. When to Sample. 2024. Available online: <https://docs.honeycomb.io/manage-data-volume/sample/guidelines/> (accessed on 25 June 2024).
30. An Introduction to Trace Sampling with Grafana Tempo and Grafana Agent. 2024. Available online: <https://grafana.com/blog/2022/05/11/an-introduction-to-trace-sampling-with-grafana-tempo-and-grafana-agent/> (accessed on 25 June 2024).
31. Application Performance Monitoring: Transaction Sampling. 2024. Available online: <https://www.elastic.co/guide/en/observability/current/apm-sampling.html> (accessed on 25 June 2024).
32. Las-Casas, P.; Papakerashvili, G.; Anand, V.; Mace, J. Sifter: Scalable Sampling for Distributed Traces, without Feature Engineering. In Proceedings of the ACM Symposium on Cloud Computing, New York, NY, USA, 20–23 November 2019; pp. 312–324.
33. Thereska, E.; Salmon, B.; Strunk, J.; Wachs, M.; Abd-El-Malek, M.; Lopez, J.; Ganger, G.R. Stardust: Tracking activity in a distributed storage system. In Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS'06), Saint-Malo, France, 26–30 June 2006.
34. Sambasivan, R.R.; Zheng, A.X.; Rosa, M.D.; Krevat, E.; Whitman, S.; Stroucken, M.; Wang, W.; Xu, L.; Ganger, G.R. Diagnosing performance changes by comparing request flows. In Proceedings of the 8th USENIX Symposium on Networked Systems Design and Implementation (NSDI), Boston, MA, USA, 1 April–30 March 2011.
35. Fonseca, R.; Porter, G.; Katz, R.H.; Shenker, S.; Stoica, I. X-trace: A pervasive network tracing framework. In Proceedings of the 4th USENIX Conference on Networked Systems Design & Implementation, Cambridge, MA, USA, 11–13 April 2007; p. 20.
36. Sigelman, B.H.; Barroso, L.A.; Burrows, M.; Stephenson, P.; Plakal, M.; Beaver, D.; Jaspan, S.; Shanbhag, C. *Dapper, a Large-Scale Distributed Systems Tracing Infrastructure*; Technical Report; Google, Inc.: Menlo Park, CA, USA, 2010.
37. Kaldor, J.; Mace, J.; Bejda, M.; Gao, E.; Kuropatwa, W.; O'Neill, J.; Ong, K.W.; Schaller, B.; Shan, P.; Viscomi, B.; et al. Canopy: An End-to-End Performance Tracing And Analysis System. In Proceedings of the 26th Symposium on Operating Systems Principles, Shanghai, China, 28–31 October 2017; pp. 34–50.
38. OpenTelemetry. 2024. Available online: <https://opentelemetry.io/> (accessed on 25 June 2024).

39. Las-Casas, P.; Mace, J.; Guedes, D.; Fonseca, R. Weighted Sampling of Execution Traces: Capturing More Needles and Less Hay. In Proceedings of the ACM Symposium on Cloud Computing, Carlsbad, CA, USA, 11–13 October 2018; pp. 326–332.
40. Google Cloud Observability: Trace Sampling. 2024. Available online: <https://cloud.google.com/trace/docs/trace-sampling> (accessed on 25 June 2024).
41. OpenCensus: Sampling. 2024. Available online: <https://opencensus.io/tracing/sampling/> (accessed on 25 June 2024).
42. Azure Monitor: Sampling in Application Insights. 2024. Available online: <https://learn.microsoft.com/en-us/azure/azure-monitor/app/sampling-classic-api> (accessed on 25 June 2024).
43. He, S.; Feng, B.; Li, L.; Zhang, X.; Kang, Y.; Lin, Q.; Rajmohan, S.; Zhang, D. STEAM: Observability-Preserving Trace Sampling. In Proceedings of the FSE’23 Industry, San Francisco, CA, USA, 3–9 December 2023.
44. AWS: Advanced Sampling Using ADOT. 2024. Available online: <https://aws-otel.github.io/docs/getting-started/advanced-sampling#best-practices-for-advanced-sampling> (accessed on 25 June 2024).
45. Solé, M.; Muntés-Mulero, V.; Rana, A.I.; Estrada, G. Survey on models and techniques for root-cause analysis. *arXiv* **2017**, arXiv:1701.08546.
46. Harutyunyan, A.N.; Poghosyan, A.V.; Grigoryan, N.M.; Hovhannisyan, N.A.; Kushmerick, N. On machine learning approaches for automated log management. *JUCS J. Univers. Comput. Sci.* **2019**, *25*, 925–945.
47. Poghosyan, A.; Harutyunyan, A.; Grigoryan, N.; Kushmerick, N. Incident Management for Explainable and Automated Root Cause Analysis in Cloud Data Centers. *JUCS J. Univers. Comput. Sci.* **2021**, *27*, 1152–1173. [[CrossRef](#)]
48. Poghosyan, A.V.; Harutyunyan, A.N.; Grigoryan, N.M. Managing cloud infrastructures by a multi-layer data analytics. In Proceedings of the 2016 IEEE International Conference on Autonomic Computing, ICAC 2016, Wuerzburg, Germany, 17–22 July 2016; Kounev, S., Giese, H., Liu, J., Eds.; IEEE Computer Society: Piscataway, NJ, USA, 2016; pp. 351–356.
49. Marvasti, M.A.; Poghosyan, A.V.; Harutyunyan, A.N.; Grigoryan, N.M. Pattern detection in unstructured data: An experience for a virtualized IT infrastructure. In Proceedings of the 2013 IFIP/IEEE International Symposium on Integrated Network Management, IM 2013, Ghent, Belgium, 27–31 May 2013; Turck, F.D., Diao, Y., Hong, C.S., Medhi, D., Sadre, R., Eds.; IEEE: Piscataway, NJ, USA, 2013; pp. 1048–1053.
50. Reynolds, P.; Killian, C.E.; Wiener, J.L.; Mogul, J.C.; Shah, M.A.; Vahdat, A. Pip: Detecting the Unexpected in Distributed Systems. In Proceedings of the Symposium on Networked Systems Design and Implementation, San Jose, CA, USA, 8–10 May 2007.
51. Harutyunyan, A.; Poghosyan, A.; Harutyunyan, L.; Aghajanyan, N.; Bunariyan, T.; Vinck, A.H. Challenges and Experiences in Designing Interpretable KPI-diagnostics for Cloud Applications. *JUCS J. Univers. Comput. Sci.* **2023**, *29*, 1298–1318. [[CrossRef](#)]
52. Fürnkranz, J.; Gamberger, D.; Lavrač, N. *Foundations of Rule Learning*; Cognitive Technologies; Springer: Berlin/Heidelberg, Germany, 2012; pp. xviii+334.
53. Fürnkranz, J.; Kliegr, T. A brief overview of rule learning. In Proceedings of the Rule Technologies: Foundations, Tools, and Applications, Berlin, Germany, 2–5 August 2015; Bassiliades, N., Gottlob, G., Sadri, F., Paschke, A., Roman, D., Eds.; Springer: Cham, Switzerland, 2015; pp. 54–69.
54. Fürnkranz, J. Pruning Algorithms for Rule Learning. *Mach. Learn.* **1997**, *27*, 139–172. [[CrossRef](#)]
55. Fürnkranz, J.; Widmer, G. Incremental reduced error pruning. In Proceedings of the 11th International Conference on Machine Learning, New Brunswick, NJ, USA, 10–13 July 1994; Morgan Kaufmann: San Francisco, CA, USA, 1994; pp. 70–77.
56. Hühn, J.; Hüllermeier, E. FURIA: An algorithm for unordered fuzzy rule induction. *Data Min. Knowl. Discov.* **2009**, *19*, 293–319. [[CrossRef](#)]
57. Lin, F.; Muzumdar, K.; Laptev, N.P.; Curelea, M.V.; Lee, S.; Sankar, S. Fast Dimensional Analysis for Root Cause Investigation in a Large-Scale Service Environment. *Proc. ACM Meas. Anal. Comput. Syst.* **2020**, *4*, 31. [[CrossRef](#)]
58. Lee, W.; Stolfo, S.J. Data Mining Approaches for Intrusion Detection. In Proceedings of the 7th Conference on USENIX Security Symposium, San Antonio, TX, USA, 26–29 January 1998; Volume 7, p. 6.
59. Helmer, G.; Wong, J.; Honavar, V.; Miller, L. Intelligent agents for intrusion detection. In Proceedings of the IEEE Information Technology Conference, Syracuse, NY, USA, 3 June 1998; Springer: Berlin/Heidelberg, Germany, 1998; pp. 121–124.
60. Helmer, G.; Wong, J.S.; Honavar, V.; Miller, L. Automated discovery of concise predictive rules for intrusion detection. *J. Syst. Softw.* **2002**, *60*, 165–175. [[CrossRef](#)]
61. Mannila, H.; Toivonen, H.; Verkamo, A.I. Discovering Frequent Episodes in Sequences Extended Abstract. In Proceedings of the First International Conference on Knowledge Discovery and Data Mining, Montreal, QC, Canada, 20–21 August 1995; AAAI Press: Washington, DC, USA, 1995; pp. 210–215.
62. Liu, H.; Motoda, H. Perspectives of Feature Selection. In *Feature Selection for Knowledge Discovery and Data Mining*; Springer: Boston, MA, USA, 1998; pp. 17–41.
63. John, G.H.; Kohavi, R.; Pfleger, K. Irrelevant features and the subset selection problem. In Proceedings of the Machine Learning: Proceedings of the 11th International Conference, New Brunswick, NJ, USA, 10–13 July 1994; Morgan Kaufmann: Burlington, MA, USA, 1994; pp. 121–129.
64. Agrawal, R.; Imieliński, T.; Swami, A. Mining association rules between sets of items in large databases. In Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data, Washington, DC, USA, 26–28 May 1993; pp. 207–216.
65. Shafer, G. *A Mathematical Theory of Evidence*; Princeton University Press: Princeton, NJ, USA, 1976.
66. Peñafiel, S.; Baloian, N.; Sanson, H.; Pino, J.A. Applying Dempster–Shafer theory for developing a flexible, accurate and interpretable classifier. *Expert Syst. Appl.* **2020**, *148*, 113262. [[CrossRef](#)]

67. Poghosyan, A.; Harutyunyan, A.; Davtyan, E.; Petrosyan, K.; Baloian, N. A Study on Automated Problem Troubleshooting in Cloud Environments with Rule Induction and Verification. *Appl. Sci.* **2024**, *14*, 1047. [[CrossRef](#)]
68. Chen, Z.; Jiang, Z.; Su, Y.; Lyu, M.R.; Zheng, Z. TraceMesh: Scalable and Streaming Sampling for Distributed Traces. *arXiv* **2024**, arXiv:2406.06975.
69. Gias, A.U.; Gao, Y.; Sheldon, M.; Perusquía, J.A.; O'Brien, O.; Casale, G. SampleHST: Efficient On-the-Fly Selection of Distributed Traces. *arXiv* **2022**, arXiv:2210.04595.
70. Huang, Z.; Chen, P.; Yu, G.; Chen, H.; Zheng, Z. Sieve: Attention-based Sampling of End-to-End Trace Data in Distributed Microservice Systems. In Proceedings of the 2021 IEEE International Conference on Web Services (ICWS), Virtual, 5–11 September 2021; pp. 436–446.
71. Zhou, T.; Zhang, C.; Peng, X.; Yan, Z.; Li, P.; Liang, J.; Zheng, H.; Zheng, W.; Deng, Y. TraceStream: Anomalous Service Localization based on Trace Stream Clustering with Online Feedback. In Proceedings of the 2023 IEEE 34th International Symposium on Software Reliability Engineering (ISSRE), Florence, Italy, 9–12 October 2023; pp. 601–611.
72. Lundberg, S.; Lee, S.I. A Unified Approach to Interpreting Model Predictions. *arXiv* **2017**, arXiv:1705.07874.
73. Ribeiro, M.T.; Singh, S.; Guestrin, C. “Why Should I Trust You?": Explaining the Predictions of Any Classifier. *arXiv* **2016**, arXiv:1602.04938.
74. Leys, C.; Ley, C.; Klein, O.; Bernard, P.; Licata, L. Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *J. Exp. Soc. Psychol.* **2013**, *49*, 764–766. [[CrossRef](#)]
75. Dunning, T. The t-digest: Efficient estimates of distributions. *Softw. Impacts* **2021**, *7*, 100049. [[CrossRef](#)]
76. Witten, I.H.; Frank, E.; Hall, M.A.; Pal, C.J. *Practical Machine Learning Tools and Techniques*; Morgan Kaufmann: Burlington, MA, USA, 2005.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.