

Griffon: Reasoning about Job Anomalies with Unlabeled Data in Cloud-based Platforms

Liqun Shao
Microsoft

Yiwen Zhu
Microsoft

Siqi Liu*
University of Pittsburgh, PA, USA

Abhiram Eswaran
Microsoft

Kristin Lieber
Microsoft

Janhavi Mahajan
Microsoft

Minsoo Thigpen
Microsoft

Sudhir Darbha
Microsoft

Subru Krishnan
Microsoft

Soundar Srinivasan
Microsoft

Carlo Curino
Microsoft

Konstantinos Karanasos
Microsoft

ABSTRACT

Microsoft’s internal big data analytics platform is comprised of hundreds of thousands of machines, serving over half a million jobs daily, from thousands of users. The majority of these jobs are recurring and are crucial for the company’s operation. Although administrators spend significant effort tuning system performance, some jobs inevitably experience slowdowns, i.e., their execution time degrades over previous runs. Currently, the investigation of such slowdowns is a labor-intensive and error-prone process, which costs Microsoft significant human and machine resources, and which negatively impacts several lines of business.

In this work, we present Griffon, a system we built and have deployed in our production analytics clusters since last year to automatically discover the root cause of job slowdowns. Most existing solutions rely on labeled data (i.e., resolved incidents with labeled reasons for job slowdowns), which is in most practical scenarios non-existent or non-trivial to acquire. Others rely on time-series analysis of individual metrics that do not target specific jobs holistically. In contrast, in Griffon we cast the problem to a corresponding regression one that predicts the runtime of a job, and we show how the relative contributions of the features used to train our interpretable model can be exploited to rank the potential causes of job slowdowns. Evaluated over historical incidents, we show that Griffon discovers slowdown causes that are consistent with the ones validated by domain-expert engineers in a fraction of the time required by them.

*Work done while the author was at Microsoft.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SoCC '19, November 20–23, 2019, Santa Cruz, CA, USA

© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6973-2/19/11...\$15.00
<https://doi.org/10.1145/3357223.3362716>

CCS CONCEPTS

• **Computing methodologies** → **Causal reasoning and diagnostics**; Anomaly detection; Classification and regression trees; • **Mathematics of computing** → *Regression analysis*; • **Computer systems organization** → *Cloud computing*.

KEYWORDS

Job anomaly, job slowdown, anomaly detection, root cause analysis, unlabeled data, analytics job, big data analytics cluster, reasoning and diagnostics

ACM Reference Format:

Liqun Shao, Yiwen Zhu, Siqi Liu, Abhiram Eswaran, Kristin Lieber, Janhavi Mahajan, Minsoo Thigpen, Sudhir Darbha, Subru Krishnan, Soundar Srinivasan, Carlo Curino, and Konstantinos Karanasos. 2019. Griffon: Reasoning about Job Anomalies with Unlabeled Data in Cloud-based Platforms. In *ACM Symposium on Cloud Computing (SoCC '19), November*

20–23, 2019, Santa Cruz, CA, USA. ACM, New York, NY, USA, 12 pages.
<https://doi.org/10.1145/3357223.3362716>

1 INTRODUCTION

Microsoft operates one of the biggest data lakes worldwide for its big data analytics needs [9]. It is comprised of several clusters for a total of over 250k machines and receives approximately half a million jobs daily that process exabytes of data on behalf of thousands of users across the organization. The majority of these jobs are recurring and several of them are critical services for the company. Hence, administrators and users put significant effort in tuning the system and the jobs to optimize their performance. Nevertheless, some jobs inevitably experience slowdowns in their execution time (i.e., they take longer to complete than their previous occurrences) due to either system-induced (e.g., upgrades in the execution environment, network issues, hotspots in the cluster) or user-induced reasons (e.g., changes in job scripts, increase in data consumed).

Such *job slowdowns* can have a catastrophic impact to the company. In fact, runtime predictability is often considered more important than pure job performance in recurring production jobs [17]. First, several jobs are interdependent, that is, the output of a job

might be consumed by multiple other jobs [7]. Thus, the slowdown of the first job can have a cascading effect on all other dependent jobs, impacting vital services across the company. Second, some business-critical jobs are associated with deadlines in the form of service-level objectives (SLOs). Missing those SLOs can result in substantial financial penalties on the order of millions of dollars.

Despite the importance of promptly resolving such incidents, the current approach remains largely manual. Job slowdowns are signaled either through tickets raised by customers or by missed deadlines (for jobs with SLOs). In either case, a slow, labor-intensive process of error triaging and root-cause analysis must be initiated as the causes are usually far from the effects. In particular, on-call engineers manually investigate causes of job slowdowns by analyzing hundreds of logs and system traces through a complex monitoring dashboard. Despite the existence of detailed metrics, it can sometimes take several hours to resolve an incident. This bottleneck costs millions of dollars in engineering time wasted on investigation and in job SLO violations, and results in degraded user experience.

In this work, we present Griffon, the system we built and have deployed in our production big data analytics infrastructure to automatically discover the main factors causing a job’s runtime deviation through the use of machine learning. Griffon greatly improves the situation described above. First, it helps users automatically find user-induced causes of their job slowdowns and prevents them from raising tickets that are “false alarms” to system administrators. Second, in case of actual infrastructure issues, it directs administrators towards the most probable causes for a job slowdown and allows early elimination of factors unrelated to the slowdown. Third, by observing slowdowns in jobs submitted for testing purposes, administrators can resolve system issues before they affect user jobs.

Existing related works have used detection methods such as classification and clustering to perform analysis of anomalies in cloud computing [1, 25]. However, to analyze anomalies, these methods rely on labeled data, e.g., data from existing incidents that associate jobs with their slowdown causes. Such labeled training data in production cloud systems are extremely hard to obtain and can also be erroneous. A few approaches do consider unlabeled data, but rely either on time-series analysis or restrict their focus to machine or VM behavior [5, 8, 10, 15, 30, 31, 34]. In contrast, we focus on job instances that span several hundreds of machines but only during the lifetime of the job. As a result, existing techniques are frequently not applicable to identify root causes of job slowdowns.

Unlike existing efforts, Griffon employs an interpretable regression model to predict job runtime, which it leverages to suggest reasons for runtime deviations. In particular, Griffon relies on: (1) telemetry data that we collect at various levels of abstraction in our clusters (i.e., at the job, machine, and cluster level); and (2) the fact that the majority of our analytics jobs are recurring, i.e., similar jobs are executed at regular intervals (e.g., every hour, day, or week) [9]. Combining these two, Griffon uses recurring jobs as historical data to build models for predicting the runtime of future occurrences of these jobs, using the collected metrics. Then, based on the relative contribution of each metric/feature to the runtime of a job that experienced a slowdown, we emit a list of possible causes for the slowdown, ranked by their importance.




Figure 1: Running tasks in one of Microsoft’s production analytics clusters, comprised of tens of thousands of machines.

Our contributions in this paper are the following:

- (1) We present an end-to-end ranking system to identify the root causes of job slowdowns without human-labeled data.
- (2) We show how an interpretable regression model can be used to reason about job slowdowns.
- (3) We experimentally compare various models in terms of accuracy, scalability in the model size and number of jobs, and generalizability to jobs not seen before by the system.
- (4) Griffon is deployed in our clusters and is used by our engineers. Early indications show that slowdown causes generated by Griffon are closely correlated to causes validated by domain experts. At the same time, Griffon drops the time of investigation by orders of magnitude compared to the existing manual process.

The rest of the paper is organized as follows. Section 2 provides details on our production environment and the relevance of the problem we focus on. Section 3 gives an overview of Griffon. Section 4 describes our anomaly reasoning algorithm, while Section 5 discusses feature engineering and data collection. Section 6 provides details on Griffon’s deployment, and Section 7 presents the results of our experimental evaluation. Section 8 discusses related work, and Section 9 provides our concluding remarks.

2 BACKGROUND ON OUR ENVIRONMENT

In this section, we provide background on the characteristics of our analytics clusters to give a sense of the scale of the problem we target (Section 2.1). Then, we describe the current state of affairs in finding the reasons for a job’s slowdown (Section 2.2).

2.1 Cluster Characteristics

At Microsoft we operate a massive data infrastructure, powering our internal analytics processing [9]. This infrastructure consists of several clusters, each comprised of tens of thousands of machines—Table 1 highlights some details of our clusters’ scale. To make the situation even more challenging, our cluster environments are also heterogeneous, including several generations of machines.

Tens of thousands of users submit hundreds of thousands of jobs to these clusters daily. Each job is a directed-acyclic graph (DAG) of operators (which we term *stages*), and each stage consists of several



Figure 2: The critical path of execution for two occurrences of the same job. The top one completed in 57 mins, while the bottom one took 88 mins to complete. Each line of each occurrence corresponds to a job stage’s start/finish time. Just by looking at this visualization tool, one might think that the top job is the errant job (due to some long stages), while in fact the bottom one is the one diverging from the normal behavior.

Dimension	Description	Size
Daily Data I/O	Total bytes processed daily	>1EB
Fleet Size	Number of servers in the fleet	>250k
Cluster Size	Number of servers per cluster	>50k

Table 1: Microsoft cluster environments.

tasks [35]. Each task gets executed in a cluster’s machine (and each machine runs several tasks in parallel). Figure 1 depicts the number of running tasks in one of our clusters over the course of a week. At each moment in time there are between 200k–300k tasks running.

Given this extreme scale and complexity, job slowdowns are quite common. Manually investigating such slowdowns, as we explain in the following section is a painful and time-consuming effort.

2.2 Manual Job Slow-down Investigation

We now describe how Microsoft engineers used to approach job slowdowns before Griffon got deployed in our clusters. Figure 2 shows two occurrences of the same job. The top one corresponds to its regular execution, taking 57 mins to complete. The bottom one experiences a slow down with a completion time of 88 mins. The figure visualizes how long the various stages of the job take to execute (although several stages might run in parallel, this tool shows the ones in the critical path, as those determine the job’s execution time).

To investigate this slowdown, an engineer will typically start by looking at the visualization tool of Figure 2, trying to detect the stages that seem abnormal. Note, however, that in this example, the “regular” top occurrence is the one that seems to have longer stages. Therefore, this tool is of limited use. Next, the engineer will have to manually combine several other tools and system files to get

more information about the job and the system during the time this job was executed. Given the scale of the system and the amount of metrics collected, this process can take a considerable amount of time to complete. Multiply this by the number of slowdowns and one can easily see the significant opportunity in saving engineering time and improving user experience by speeding up this process. Note also that only a few engineers have the knowledge to perform this manual analysis.

3 SYSTEM OVERVIEW

Griffon’s goal is to find the causes for job runtime degradations in our big data analytics clusters. A central requirement we set when we started designing Griffon was to not rely on labeled data, i.e., there should be no need for existing slowdown instances associated with their causes. Acquiring labeled data in a complex infrastructure like ours that has been operating for years is very hard. Labeling data requires both infrastructure support and engineering training. Most importantly, it would require significant engineering time to label years-worth of historical data to get to a meaningful training set.

Each job in our clusters is associated with a set of telemetry data that we already collect for monitoring and debugging purposes (e.g., number of tasks, size of input data, load of machines the job was executed on—see Section 5 for details), some of which contributed to the job’s slowdown. Instead of finding a subset of the slowdown causes (as a system that relies on labeled data would do), Griffon *ranks* the causes (i.e., the features) in the order they affected the deviation of the job’s runtime from its expected runtime, and then suggests the top causes to the users. A formal description of the problem and our approach for solving it is presented in Section 4.

Griffon’s architecture is depicted in Figure 3. It consists of two pipelines: the offline training and the online prediction, which we detail below.

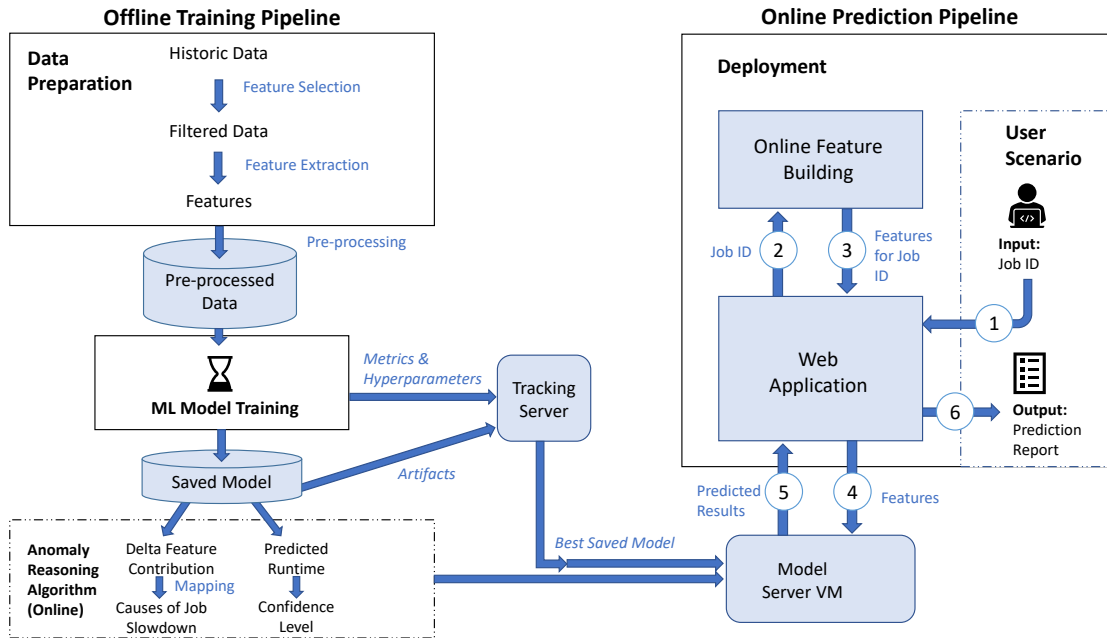


Figure 3: System architecture.




Figure 4: Griffon's output for the slowdown of the job depicted in Figure 2.

Training The offline training process uses various metrics that we collect at the job, machine, and cluster level to generate a model that will be able to predict the runtime of a job given these features. Our anomaly reasoning algorithm will use this model to rank the features that contributed to the job's slowdown. The training process involves the following steps: (1) data preparation, i.e., collect and clean the data from different sources in the cluster at regular intervals; (2) feature engineering, i.e., extract raw features from the collected data, create new ones, and choose the ones that we will be using to train the model; and (3) model generation, i.e., the creation of the model, including hyper-parameter tuning, training, and model evaluation. The data preparation is detailed in Section 5, whereas the model details are in Section 4.

The generated model is stored in the Tracking Server, which tracks model runs and stores performance metrics and hyper-parameter values (see Section 6).

Prediction The online prediction pipeline provides an API that takes as input the ID of a job that was executed and experienced a slowdown. This API is exposed to the users through a web application. Then the Online Feature Building gathers the metrics associated with that job and provides the data to the Model Server where the prediction model and the anomaly reasoning algorithm are deployed. The output is a report with ranked reasons for job slowdown. In addition, the system provides a confidence level in the results (see Section 4.4).

High confidence means that users can rely on the output of the system. On the other hand, low confidence means that the metrics used are not sufficient to explain the job slowdown. The output of the system is useful even in the latter case, because users can rule out these metrics and focus their investigation in other areas.

Figure 4 shows Griffon's output for the job slowdown of Figure 2. In this case, Griffon suggests with high confidence that the increase in data written by the job is the main reason for its slowdown. As this is a user-induced reason and not a problem with the system, the corresponding ticket can be closed without further investigation.

4 ANOMALY REASONING ALGORITHM

In this section, we describe our algorithm for reasoning about anomalies. First, we formally define the problem (Section 4.1) and discuss model interpretability (Section 4.2). Then we describe the interpretable tree-based model that Griffon uses for determining the reasons for a job's slowdown (Section 4.3) and its associated confidence level (Section 4.4).

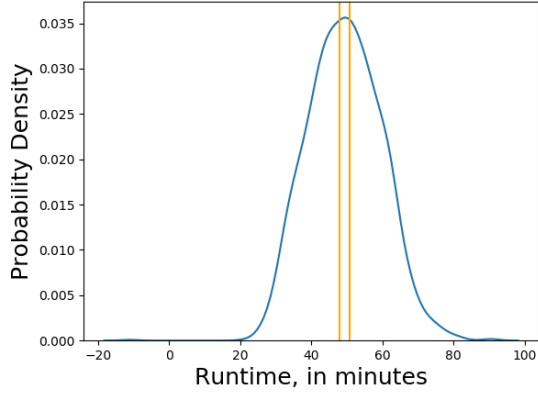


Figure 5: Example of baseline selection for a job template.

4.1 Problem Statement

We consider a set of jobs that have already been executed. Hence, we know the runtime of each job. Through our collected metrics, we also know the values of the features that we are interested in (see Section 5.2 for feature selection).

The majority of the jobs submitted in our clusters are analytics jobs¹ that are recurring, i.e., they are submitted at regular intervals (typically hourly, daily, or weekly) [9]. We use the notion of **job template** to refer to each of these recurring jobs. Jobs belonging to the same template have very similar scripts with minor differences, e.g., to access the latest data.

We also define the **baseline** of a job to be its “expected” runtime, given the runtime of the other jobs that belong to the same template. In practice, we use a set of jobs whose runtime falls between the 45th and 55th percentile for that template. A benefit of using a percentile measure is that we avoid outliers. Therefore “slow jobs” in our training data will not affect the baseline set. Figure 5 shows the runtime distribution for one job template. The data we use for baseline selection falls between the two orange lines.

For jobs that belong to templates with no previous occurrences, we use the baseline of jobs with similar characteristics (in data size and performed operations). Similarly, we define the baseline of various features of a job to be their expected value, given the jobs of the corresponding template.

Let $\mathbf{x} \in \mathbb{R}^p$ be the p -dimensional features of a job, $y \in \mathbb{R}$ be the job’s runtime, and $\hat{y} \in \mathbb{R}$ its predicted runtime. Let y^β and \mathbf{x}^β be the baseline of the runtime and the features, respectively. We define the problem as follows: for each job, lacking human labeled reasons, with features \mathbf{x} and runtime y , predict the rank of different features based on their influence on the deviation of \hat{y} from \hat{y}^β .

Note that we could use the actual runtimes for both the job and its baseline, instead of their predicted values. However, as we show below, using the predicted runtimes allows us to decompose a job’s runtime to the contributions of its features. This is key to identify reasons for a job’s slowdown.

¹Analytics jobs are executed using Scope, an internal SQL-like distributed query engine that enables processing of petabytes of data per job [35].

4.2 Interpretable Model

Consider a machine learning model that is trained to predict the runtime y of a job using a set of features \mathbf{x} . That alone would be a standard regression problem. However, in our setting we want to use such a runtime prediction to find the features that contribute the most to a job’s slowdown, that is, to the runtime’s deviation from the job’s baseline y^β . To this end, we need an *interpretable* regression model for the job’s runtime.

We define a regression model to be interpretable, if the output of the model can be expressed as the sum of contributions of each of the model’s features:

$$\hat{y} = c + \sum_{k=1}^K f c_k \quad (1)$$

where c is a constant and $f c_k$ is the contribution of feature x_k to the prediction.

Similarly, for a baseline job, let \hat{y}^β be the predicted runtime based on the same model using the baseline features \mathbf{x}_k^β . We can decompose the model prediction as:

$$\hat{y}^\beta = c + \sum_{k=1}^K \bar{f} c_k^\beta \quad (2)$$

where $f c_k^\beta$ is the contribution of feature x_k^β to the prediction. If we have multiple baseline jobs, the mean of \hat{y}^β and $f c_k^\beta$ across all the baseline jobs will be used:

$$\bar{\hat{y}}^\beta = c + \sum_{k=1}^K \bar{f} c_k^\beta \quad (3)$$

where $\bar{\hat{y}}^\beta$ is the average predicted runtime based on the same model and $\bar{f} c_k^\beta$ the average contribution of feature x_k for all the baseline jobs.

In our setting, we can quantify the **delta feature contribution** $\Delta f c_k$ of each feature x_k to the deviation of \hat{y} from $\bar{\hat{y}}^\beta$ as:

$$\hat{y} - \bar{\hat{y}}^\beta = \sum_{k=1}^K (f c_k - \bar{f} c_k^\beta) = \sum_{k=1}^K \Delta f c_k \quad (4)$$

Being able to quantify the contributions of each feature to a job’s slowdown allows us to rank the features in order of importance, which is the goal of Griffon.

Note that proving that our method *always* finds the right reasons for a job’s slowdown would require us to be able to determine the exact contribution of each feature to the job’s runtime. This is notoriously hard in complex distributed systems, comprising tens of thousands of machines, such as our big data analytics clusters that Griffon targets. Even if we consider the same exact job, its behavior can vary greatly each time it gets executed, based on various factors: the specific machines (often belonging to different hardware generations) that are chosen by the cluster’s resource manager to execute the job’s (possibly thousands of) tasks [9], the machine load and possible hardware problems during execution (e.g., network congestion, failed disks), tasks of other jobs sharing the same machines (the “noisy neighbor” problem [14]), and so on. Therefore, we cannot *guarantee* that Griffon will always find

the right reasons for a slowdown. Nevertheless, our experimental results (see Section 7) show that in most practical scenarios, Griffon does find the same reasons that were identified by expert engineers. In parallel, we quantify the level of trust that a user should have on a prediction by defining the confidence level for a prediction (see Section 4.4).

Model choice We considered various model categories to predict the runtime of a job, namely Linear Regression (LR), Random Forest (RF), Gradient Boosted Tree (GBT), and Deep Neural Network (DNN). Our main requirements were that the model be interpretable and that it offers good accuracy.

A linear model can be expressed as $y = \alpha + \sum_{k=1}^K \beta_k x_k$, where $\alpha, \beta_k \in \mathbb{R}$. It is trivial to show that it satisfies the interpretability criterion of Eq. 1. However, as we show in Section 7.2, the accuracy is worse than that of the other models. The DNN has acceptable accuracy, but its interpretability is hard to establish. The GBT is less accurate than RF. The RF model exhibited the best accuracy in our experiments and therefore, is our model of choice in Griffon. In the next section, we describe an appropriate tree interpreter that reformulates a tree-based model to a linear form, so that we can use it to rank feature contributions to job slowdowns.

Note that when training our models, we considered both a *global* and *per-template* models. In the former case, we train a single unified model to predict runtime using jobs of all templates together in the training set. In the latter, we train one model per job template utilizing training data drawn exclusively from jobs that belong to the particular template. In Section 7, we compare the two approaches in terms of accuracy, scalability, and generalizability.

4.3 Interpretable Random Forest

In a Random Forest (RF) model, for each tree, in order to make a prediction, we traverse a path from the root of the tree to a leaf. This path consists of a series of decisions based on the model's features. Assuming there are M nodes on the path, each node separates the feature space into two, given a feature x_k and a threshold t_k : the one child node corresponds to $x_k \leq t_k$, the other to $x_k > t_k$. In other words, from the root node where all the samples reside, a partition based on feature x_k and threshold t_k thus separates the data samples to the two children that correspond to smaller feature spaces.

Consider a tree j of the model and a node $m \in j$ that is partitioned from its sibling based on feature x_k . Let $\bar{y}_{m,j}$ be the mean target value for all samples that reside on node m . Then the contribution of feature x_k to the final prediction due to this partitioning is calculated as:

$$\Delta_m \text{contrib}_j(x, k) = (\bar{y}_{m,j} - \bar{y}_{m-1,j}) I_j(m, k) \quad (5)$$

for $2 < m \leq M$, where node $m-1$ is m 's parent. $I_j(m, k)$ equals to 1 if the partitioning at node $m-1$ involves feature x_k for tree j or 0 otherwise. The number of samples that reside on each node becomes smaller and smaller by traversing the path, as the feature space gets smaller. The contribution of x_k to the final prediction can be calculated as the sum of all $\Delta_m \text{contrib}_j(x, k)$:

$$\text{contrib}_j(x, k) = \sum_{m=2}^M \Delta_m \text{contrib}_j(x, k) \quad (6)$$

$$= \sum_{m=2}^M (\bar{y}_{m,j} - \bar{y}_{m-1,j}) I_j(m, k) \quad (7)$$

The prediction of the target value from this tree is $\bar{y}_{M,j}$ and can be expressed using the sum of all features' contributions along the path:

$$y_j = \bar{y}_{M,j} = \bar{y}_{1,j} + \sum_{m=2}^M \bar{y}_{m,j} - \bar{y}_{m-1,j} \quad (8)$$

$$= c_j + \sum_{m=2}^M \sum_{k=1}^K \Delta_m \text{contrib}_j(x, k) \quad (9)$$

where c_j is the full sample mean. *TreeInterpreter* [28] combines the results of all trees in our Random Forest by taking the sum of the contribution from each tree. Thus, each prediction is decomposed into a sum of contributions from the features, as follows:

$$y = \frac{1}{J} \sum_{j=1}^J c_j + \sum_{k=1}^K \left(\frac{1}{J} \sum_{j=1}^J \text{contrib}_j(x, k) \right) \quad (10)$$

where J is the number of trees, c_j is the full sample mean for each j^{th} tree, and K is the number of features involved.

Refers to Eq. 10, using $c = \frac{1}{J} \sum_{j=1}^J c_j$ for the average runtime across the whole training set and $f c_k = \frac{1}{J} \sum_{j=1}^J \text{contrib}_j(x, k)$ for the contribution of feature x_k to the predicted runtime, we get to Eq. 1, which shows that our Random Forest model meets the interpretability criterion. Therefore, it can be used to detect reasons for job slowdowns in Griffon.

4.4 Confidence Level

The confidence level shows how reliable is the prediction made by our model for the contribution of each feature to a job's slowdown. We consider two factors that affect our model confidence: (1) the relative error in predicting the runtime of the job (by comparing the model prediction with the actual runtime of the job);² (2) the confidence intervals estimated by the random forest [20].

The relative error is defined as following:

$$\text{error_rate} = \frac{|\text{predicted_runtime} - \text{actual_runtime}|}{\text{actual_runtime}} \quad (11)$$

We use two thresholds, t_1 and t_2 , for the relative error, as explained below.

The confidence interval of the random forest method is estimated based on the prediction of each decision tree, $y_j, \forall j \in \{1, 2, 3, \dots, J\}$. We take the p^{th} and $(100 - p)^{\text{th}}$ percentile of the distribution of y_j . If the final prediction y is within this range, we

²We also considered taking into account the prediction errors for the baseline jobs in the definition of confidence. However, we decided to not include them, as we observed that the prediction for baseline jobs is very accurate (with a Mean Absolute Ratio Error of 2.2%). This is due to the fact that a large part of the baseline jobs are included in the training set of the prediction models. Therefore, it is expected that predictions especially for these jobs to be very accurate.

consider the prediction to have low variance, since the predictions from all trees are consistent.

We define three confidence levels as follows:

- High** The prediction y is within the range of p^{th} and $(100 - p)^{\text{th}}$ percentile of y_j , and the relative error is lower than threshold t_1 ;
- Medium** The relative error is between t_1 and t_2 ;
- Low** Other scenarios.

Parameters p , t_1 and t_2 are tuned as hyper parameters using validation data. High confidence means our model can reliably predict the slowdown reasons. Low confidence means the reasons are likely to fall outside the metrics we used. In the API, the user will be presented the level of confidence. A low confidence indicates more investigation will be needed. However, even in this case, as the model has examined many metrics, the DRI can focus their investigation in other areas.

5 DATA PREPARATION AND FEATURE ENGINEERING

5.1 Data Preparation

In our clusters, we keep hundreds of metrics for monitoring, reporting, and troubleshooting purposes, which result in petabytes of logs and metrics per day. Moreover, the features we are interested in are scattered both *physically* (in different files across our hundreds of thousands of machines) and *logically* (we need to process and combine different files to generate features). To perform the required data preparation and extract features out of the data, we use Scope [35], which provides a SQL-like language and can support our scale.

Feature extraction occurs both during the offline training phase and the online prediction (see Figure 3). Given that data freshness is not an issue for training (we do not need data of the current day), we use data that becomes available daily in our clusters and includes years-worth of historical data. For prediction, we need to collect features only for a single job (or a group of jobs), but latest data is required, as a user might want to debug their job that just finished. Thus, we use different data sources that allow us to access data within minutes from when they are produced (but that do not allow access to historical data, so cannot be used for training).

5.2 Feature Engineering and Selection

In collaboration with domain experts, we selected a subset of the features we collect to train our models in Griffon, based on what could potentially impact the runtime of a job. As already discussed, a job can experience a slowdown compared to its previous occurrences, due to either user-induced or system-induced reasons. User-induced reasons can be captured by metrics collected at the job-level, whereas metrics related to system-induced reasons can be split to either machine-level or cluster-level, as detailed below.

- Job-level** These are metrics collected for each job. Griffon currently uses approximately 15 such features, including data read within and across racks, data written, data skewness metrics, job priority, execution DAG features (e.g., number of stages and tasks), and user information.

- Machine-level** These are metrics collected at the machines used during the execution of a job, such as CPU load, allocation delays, I/O reads/writes. Due to the challenges of collecting such features and correlating them with each job, in production Griffon currently uses only a few of them, but we are working on adding more.

- Cluster-level** These relate to the cluster environment when a job was executed. Examples of the ones we use are job queuing times, number of failed and revoked vertices, and execution environment version.

Challenges One of the biggest challenges with feature engineering is the correlation between features, which often results in high variance in the model prediction. While certain highly correlated features (> 0.95) were removed, feature data was preserved to the greatest extent, because correlated features may indicate different problems with a job. For instance, input size and input size per task have a correlation of 0.9. However, the former might indicate that the slowness reason was more data, whereas the latter might indicate data skew. Fortunately, random forest models, as the one Griffon relies on, are by design robust to correlated features thanks to the way they are constructed. First, each tree of the forest is developed based on a small subset of the overall features. This automatically reduces the likelihood of having correlated features within the same tree. Second, once a set of features is picked for a tree, this tree is constructed using a decision tree learner algorithm. This algorithm typically works top-down, by choosing a feature at each step that “best” splits the dataset. Given that correlated features would split the dataset in a similar way, once a feature is picked, its correlated features are less likely to be picked.

6 TRACKING AND DEPLOYMENT

Model Tracking Machine learning is an iterative process, and reproducibility and versioning are crucial to productionalize machine learning models. To track model history in Griffon, we use MLflow [24] in our Tracking Server (see Figure 3), which is deployed as an Azure Linux VM. For each model, we track logging parameters, code versions, metrics, and model artifacts.

Model Serving In order to make available to end users a model that we trained and stored in the Tracking Server, we use one of MLflow’s “flavors”³ to build an Azure Machine Learning image out of that model. Then we deploy this image as an Azure Container Instance, using the Azure ML Service’s [23] in the Model Server (see Figure 3).

The models that are deployed in the Model Server are made available to end users through a web application, which exposes a scoring API. The web application runs on a flask web server [13].

Model Monitoring It is critical to monitor the model performance and retrain models if they go stale. To ensure this, we have provisioned our pipeline to allow single click retraining. Decoupling data sources from the training pipeline helps to easily refresh our data, retrain, and deploy the model with minimal impact to end users.

³An MLflow flavor is a convention that deployment tools use to understand the model.




Figure 6: Runtime distribution for the jobs of different templates.

7 EVALUATION

We now present our experimental evaluation for Griffon. In Section 7.1 we discuss Griffon’s effectiveness in finding the actual causes of job slowdowns. In Section 7.2 we compare different machine learning models for training, whereas Section 7.3 studies the scalability of different models as the number of job templates increases. In Section 7.4 we compare the model performance when training our model with an increasing number of jobs per template. In Section 7.5 we show early evidence that Griffon can be applicable in different domains.

We carried out our experiments on Windows using Python 3.7. We used a machine with eight 2.90 GHz processors and 64 GB RAM for the experiments in Sections 7.1 and 7.2, and a high-memory Virtual Machine (VM) for the scalability experiments of Section 7.3.

The job and feature data is obtained from the Microsoft production clusters, as described in Section 5.1. We use historical job data with different job templates, as described in Section 5.2, over a period of three months.

7.1 Validation Results

Working with domain experts at Microsoft, we picked a set of job templates that are considered important for our production clusters (SLO critical), and trained Griffon based on those. Note that the runtime distribution of the jobs of different templates varies significantly, which poses extra challenges for the runtime prediction based on machine learning models. Figure 6 shows the runtime distributions for five of the templates that we used.

From these templates, we then randomly picked seven jobs that experienced slowdowns (five from these templates and two from different templates), and compared the causes for slowdowns that were identified by the experts with those suggested by Griffon. For these jobs, Table 2 shows the reasons identified by the experts and Griffon (with their ranking), Griffon’s confidence level, and if the job belonged to one of the templates used for training the model (in-t). We use R_x to denote the reason that Griffon predicted for a job’s slowdown with rank x . For readability, we show only the

reasons that were common between Griffon and experts and use R_x variables for the rest.

The results in Table 2 show that the reasons generated by Griffon are highly correlated with the reasons manually validated by our domain expert engineers. For job 1, the top predicted reason is the same as the manually validated reason with high confidence. For jobs 2 and 4, our system predicted the validated reason in the top 5 slowdown reasons, which is consistent with the confidence level medium. For job 3, our model does not identify the same reasons as the experts, also consistent with low confidence. Adding more features as planned (e.g., additional machine-level features) will allow us to improve the model’s prediction capability and minimize such low-confidence cases. Jobs 6, 7 and 8 show the robustness of our model: although Griffon was not trained using these job templates, it can still find the correct reasons with high confidence by using knowledge gathered by other similar job templates. Importantly, we observed no misleading predictions, i.e., there were no cases where Griffon predicted wrong slowdown reasons with high confidence. This means that even predictions with low confidence, which are usually due to large deviations between the predicted and the actual runtimes of a job, can be useful in ruling out the currently used features from the investigation.

7.2 Picking the Right Model

As discussed in Section 4, we experimented with various categories of models for the job runtime prediction, including Linear Regression (LR), Random Forest (RF), Gradient Boosted Trees (GBT) and Deep Neural Networks (DNN) with two hidden layers without hyper parameter tuning. For each of these categories, we consider both a global and per-template models.

We use Mean Absolute Ratio Error (MARE) as a metric to evaluate each model’s accuracy. As the runtime distribution varies significantly across different job templates, we normalize the estimation error by the baseline runtime of each job template (see Section 4.1), calculating the average runtime per job template in the training data:

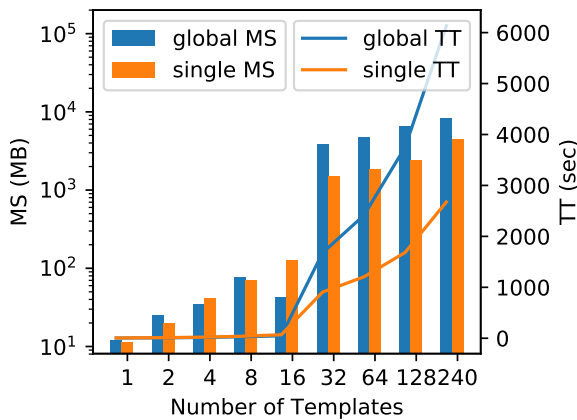
$$\text{MARE} = \frac{1}{n} \sum_{i=1}^n \left| \frac{y_{\text{test}_i} - \hat{y}_{\text{test}_i}}{y_{\text{test}_i}^{\beta}} \right| \quad (12)$$

where n is the number of jobs in the testing data, y_{test_i} , \hat{y}_{test_i} , and $y_{\text{test}_i}^{\beta}$ are the predicted, actual, and baseline runtime from testing data, respectively.

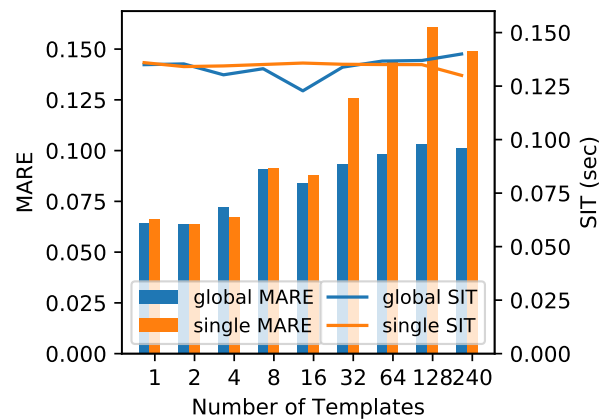
Table 3 shows the results of MARE scores for the four model categories. Random forest performs best in terms of accuracy, both for the global and per-template model. Given its high accuracy and interpretability (as discussed in Section 4.2), RF is the approach we use in our production Griffon deployments. Moreover, we observe that the global model tracks closely the performance of the per-template model, while allowing to reason about jobs that we have not sufficiently encountered previously. In the next section, we demonstrate that the global model scales much better than the per-template models with an increasing number of templates. Hence, we use the global model in production.

Table 2: Result subset validated by engineers.

Job	Griffon's Predicted List of Ranked Reasons	Engineer Validated Reason	Confidence Level	in-t
1	[Input Size, R_2 , R_3]	Input size	High	Yes
2	[R_1 , R_2 , R_3 , Revocation, R_5]	Revocation	Medium	Yes
3	[R_1 , R_2 , R_3 , R_4 , R_5 , R_6]	Framework issue	Low	Yes
4	[R_1 , R_2 , R_3 , R_4 , High compute hours]	High compute hours	Medium	Yes
5	[Time skew, R_2 , R_3 , R_4]	Time skew	High	Yes
6	[High compute hours, R_2 , R_3]	High compute hours	High	No
7	[R_1 , Usable machine count, R_3 , R_4]	Usable machine count	High	No
8	[High compute hours, R_2]	High compute hours	High	No



(a) MS and TT



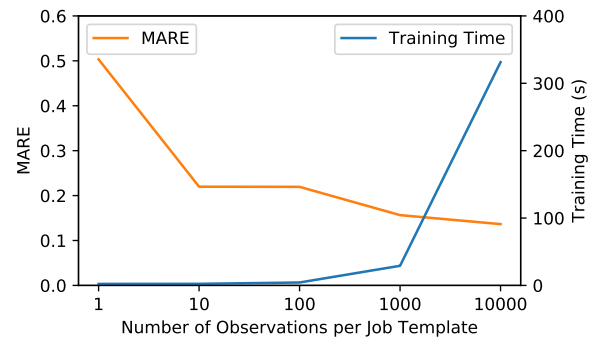
(b) MARE and SIT

Figure 7: Scalability of global vs. per-template models**Table 3: MARE scores for runtime predictions by LR, RF, GBT, and DNN with global and per-template models (lower is better).**

	LR	RF	GBT	DNN
Per-Template Model	0.186	0.116	0.124	0.146
Global Model	0.235	0.121	0.277	0.353

7.3 Scalability of Global vs. Per-template Models

We assess the scalability of the global and per-template models by training them with an increasing number of job templates, as shown in Figure 7. We added job templates incrementally to the training set and repeated each experiment 10 times. We use 5-fold cross validation for hyper parameter tuning with random grid search. With more job templates, the training time (TT) and the model size (MS) for both global and per-template models increased. We observe that the MARE for global model is better than per-template

**Figure 8: Model performance (in terms of MARE and training time) with varying training sample sizes.**

models when training with a larger number of job templates. This can be attributed for the most part to the larger sample size for the global model, as a result of the unification of the training set. In contrast, when training the per-template models with 240 job

Table 4: Contribution of each feature to the high gas mileage of American-made Ford Granada compared to other American cars.

Feature	Ford Granada	Mean	Baseline	Delta FC (in mpg)
Year	81	76	75	2.07
Horsepower	88	104	106	1.07
Weight	3060	2978	3224	0.7
Displacement	200	194	221	0.04
Acceleration	17.1	15.5	16.1	-0.19
Cylinders	6	5.5	5.8	-1.4

templates, many templates only have a small number of samples and the MARE is high. We also report the prediction time on a single job, i.e. the single inference time (SIT). The SIT didn't increase with the model size, which is important to deliver a real-time experience to Griffon's users.

Overall, the above experiments demonstrate the scalability of the global model for cloud-scale training. At the other extreme, a template-specific model suffers from lack of training data and the ability to generalize to new (unseen) templates. As part of our future work, we plan to cluster job templates using unsupervised machine learning methods and train "semi-global" models that take into account multiple job templates that share similar characteristics to strike a balance between the two approaches.

7.4 Varying Size of Training Data

In order to determine the impact of the number of jobs per template on our model performance, we retrain our global model assuming that we only have a limited number of observations for each job template. In particular, we train our model with n observations per job template, where $n \in \{1, 10, 100, 1000, 10000\}$. The MARE and the training time are reported in Figure 8. We used 22 job templates and, for each run of the experiment, a random sample of n observations were selected as the training data. When $n \geq 10$, the MARE dropped below 20%, and the prediction accuracy continues to improve for larger sample sizes, although less significantly. Note that the training time increased exponentially to the sample size.

7.5 Model Generalization

The baseline approach for interpretation described in this paper allows job runtime prediction results to be interpreted and compared to a set of similar jobs. Data in the real world frequently looks similar to our dataset: Gaussian mixture distributions of a target variable are commonly encountered.

This section presents an example of employing Griffon on another dataset in another field. A classic dataset from statistics is the "Auto" dataset of gas mileage, which is represented well by a linear superposition of three Gaussians by region of origin: American, European and Japanese. Manufacturers might be interested in understanding what factors drive higher gas mileage in American cars relative to other American cars. Here gas mileage is the equivalent to job runtime.

Table 4 summarizes the delta contribution of each feature (FC) to the high gas mileage of American-made Ford Granada compared to other American cars based on our anomaly reasoning algorithm, described in Section 4. We observe that "Year" and "Horsepower" contribute the most to high gas mileage, while "Weight" and "Displacement" make marginal contributions. "Acceleration" and "Cylinders" contribute to low gas mileage.

8 RELATED WORK

Anomaly detection [4] refers to the problem of finding patterns in data that do not conform to expected behavior. In contrast, *anomaly reasoning*, which is the purpose of this work, encompasses recognizing, interpreting, and reacting to unfamiliar objects or familiar objects appearing in unexpected contexts.

Anomaly reasoning is particularly important for large-scale systems, as it is not possible to manually track all machines and applications at scale. Below we discuss the main categories of works that are related to Griffon.

Interpretable models Building an effective anomaly system requires both interpretable models and reasoning algorithms. Several efforts have focused on interpreting the results of machine learning models. Their goal is to provide proper explanation about how or why the algorithm produces a specific prediction and to identify interactions between features and estimation results [11, 16, 29]. In Griffon, we use an interpretable random forest model to rank a job's slowdown reasons given the contribution of various features. Similar methods to interpret a tree model can be generalized to boosting algorithms [33].

Anomaly reasoning with labeled data Existing related works have used detection methods such as classification and clustering to perform analysis of anomalies in cloud computing. A detailed survey of those work can be seen in [1, 25]. For example, a fault detection and isolation system based on k-nearest neighbor has been proposed to rank machines in order of their anomalous behavior [3]. Other works have used a hybrid of SVM classification and k-medoids clustering to detect intrusions of the network [6]. An anomaly-based clustering method has also been suggested to detect failures in general production systems [12].

The downside of these approaches is that they require labeled data. Such data is hard to acquire in many settings, including ours. In the context of an infrastructure that has been operating for many years, labeling data requires infrastructure support and, most importantly, training a large number of engineers to add labels when resolving anomalies. Moreover, it is almost impossible to perform labeling for the years-worth of historical data.

Anomaly reasoning with unlabeled data A few approaches have considered unlabeled data, but focus either on time-series analysis or are restricted to machine or VM metrics [5, 8, 10, 15, 30, 31, 34]. For instance, an anomaly detection and reasoning system has been proposed to detect security problems during VM Live Migration [34]. This work is based on time series data related to resource utilization statistics, e.g., file read/write, system call, CPU usage. PREDICTIVE PERFORMANCE ANOMALY PREVENTION (PRE-PARE) [31] predicts performance anomalies using a 2-dependent Markov model and a classifier based on system-level metrics, such

as CPU, memory, network traffic. Another related work developed an Unsupervised Behavior Learning (UBL) system to capture the anomalies and infer their causes [10]. To circumvent the need for labeling data, Self Organizing Map (SOM) has been suggested to model the system behavior, and deviations are used for the anomaly detection [18].

Similar to Griffon, those methods estimate the contribution of each attribute to the anomaly and provides information about which system-level metrics to look into. However, they require time-dependent series of data to capture the anomalies. In contrast, we focus on job instances that span several hundreds of machines but only during the lifetime of the job. Thus, our features are neither time series nor machine-centric (although we do employ some system-level data to examine the system’s impact on a particular job’s execution).

Other works in anomaly reasoning aim to pinpoint the faulty components of a system by tracing the system’s activities [2, 21, 22, 26]. The methods rely more on the estimation of the time series’ change point and the propagation pattern or the execution graph. However, those methods require significant domain knowledge and are hard to generalize.

To the best of our knowledge, Griffon is the first anomaly reasoning system to be deployed at this scale in production to identify the causes of job slowdowns in analytics clusters. Unlike existing approaches, it follows a job-centric approach and does not rely neither on labeled data nor on time series analysis.

Weak supervision Acknowledging the difficulty of hand-labeling data but also the importance of having labeled data, weak supervision attempts to strike a balance between supervised and unsupervised learning. Recent systems, such as Snorkel [27], bypass the problem of hand-labeling data by instead providing labeling functions, which label data programmatically. It would be interesting to see how a system like Griffon can use such techniques to encode knowledge of domain expert engineers as functions to further improve its anomaly detection capabilities.

9 CONCLUSION & FUTURE WORK

We presented Griffon, a system that we built and have deployed in production to detect the causes of job slowdowns in Microsoft’s big data analytics clusters, consisting of hundreds of thousands of machines. Griffon does not require labeled data to perform anomaly reasoning. Instead, it uses an interpretable machine learning model to predict the runtime of a job that has experienced a slowdown. Using this model, we can determine the contribution of each feature in the deviation of the job’s runtime compared to previous normal executions of the job (or of jobs with similar characteristics).

Our evaluation results using historical incidents showed that Griffon discovers the same slowdown reasons that were detected by domain expert engineers. We also compared various categories of models and showed that a global (i.e., trained over all jobs) random forest model strikes a good balance between accuracy, training time, model size, and generalization capabilities.

Towards data-driven decisions Griffon is part of our bigger vision towards employing data-driven decisions to optimize various

aspects of our systems. Taking Griffon’s capabilities a step further, knowing the job slowdown reasons allows us to automatically tune the system to avoid such slowdowns in the future. This may include both system parameters, such as dynamically setting the number of running tasks per machine, and application parameters, such as the degree of parallelism for each stage of a job. Moreover, such parameter autotuning does not have to be constrained to job slowdowns—we can use it to automatically and dynamically set various parameters in our systems to improve their performance.

Furthermore, although Griffon currently targets our internal analytics clusters, the above techniques can be applied to other environments, such as various public Azure services, including the Azure SQL and HDInsight offerings. Similar data-driven decisions are increasingly applied in various companies [19, 32].

ACKNOWLEDGMENTS

We would like to thank Microsoft’s Big Data analytics team, and in particular Sarvesh Sakalanaga, Vamsi Kuppa, Sankar Nemani, David Olix, and Anusha Chavva, for their invaluable help in bringing Griffon to production. We also thank Anupam Upadhyay for many stimulating discussions that helped shape this effort, as well as Panagiotis Garefalakis for his initial work on visualizing cluster metrics. Finally, we thank Mathias Lécuyer, our shepherd, and the anonymous reviewers for their insightful feedback and suggestions.

REFERENCES

- [1] Shikha Agrawal and Jitendra Agrawal. 2015. Survey on anomaly detection using data mining techniques. *Procedia Computer Science* 60 (2015), 708–713.
- [2] Marcos K Aguilera, Jeffrey C Mogul, Janet L Wiener, Patrick Reynolds, and Athicha Muthitacharoen. 2003. Performance debugging for distributed systems of black boxes. In *ACM SIGOPS Operating Systems Review*, Vol. 37. ACM, 74–89.
- [3] Kanishka Bhaduri, Kamalika Das, and Bryan L Matthews. 2011. Detecting abnormal machine characteristics in cloud infrastructures. In *2011 IEEE 11th International Conference on Data Mining Workshops*. IEEE, 137–144.
- [4] Varun Chandola, Arindam Banerjee, and Vipin Kumar. 2009. Anomaly detection: A survey. *ACM Comput. Surv.* 41 (2009), 15:1–15:58.
- [5] Ludmila Cherkasova, Kivanc Ozonat, Ningfang Mi, Julie Symons, and Evgenia Smirni. 2009. Automated anomaly detection and performance modeling of enterprise applications. *ACM Transactions on Computer Systems (TOCS)* 27, 3 (2009), 6.
- [6] Roshan Chitrakar and Chuanhe Huang. 2012. Anomaly based intrusion detection using hybrid learning approach of combining k-medoids clustering and naive bayes classification. In *2012 8th International Conference on Wireless Communications, Networking and Mobile Computing*. IEEE, 1–5.
- [7] Andrew Chung, Carlo Curino, Subru Krishnan, Konstantinos Karanasos, Greg Ganger, and Panagiotis Garefalakis. 2019. Peering through the dark: an Owl’s view of inter-job dependencies and jobs’ impact in shared clusters. In *SIGMOD*.
- [8] Ira Cohen, Jeffrey S Chase, Moises Goldszmidt, Terence Kelly, and Julie Symons. 2004. Correlating Instrumentation Data to System States: A Building Block for Automated Diagnosis and Control. In *OSDI*, Vol. 4. 16–16.
- [9] Carlo Curino, Subru Krishnan, Konstantinos Karanasos, Sriram Rao, Giovanni M. Fumarola, Botong Huang, Kishore Chaliparambil, Arun Suresh, Young Chen, Solom Heddaya, Roni Burd, Sarvesh Sakalanaga, Chris Douglas, Bill Ramsey, and Raghu Ramakrishnan. 2019. Hydra: a federated resource manager for data-center scale analytics. In *NSDI*.
- [10] Daniel Joseph Dean, Hiep Nguyen, and Xiaohui Gu. 2012. Ubl: Unsupervised behavior learning for predicting performance anomalies in virtualized cloud systems. In *Proceedings of the 9th international conference on Autonomic computing*. ACM, 191–200.
- [11] Filip Karlo Došilović, Mario Brčić, and Nikica Hlupić. 2018. Explainable artificial intelligence: A survey. In *2018 41st International convention on information and communication technology, electronics and microelectronics (MIPRO)*. IEEE, 0210–0215.
- [12] Songyun Duan, Shivnath Babu, and Kamesh Munagala. 2009. Fa: A system for automating failure diagnosis. In *2009 IEEE 25th International Conference on Data Engineering*. IEEE, 1012–1023.
- [13] Flask. 2019. *Flask - A Python Microframework*.

- [14] Panagiotis Garefalakis, Konstantinos Karanasos, Peter R Pietzuch, Arun Suresh, and Sriram Rao. 2018. Medea: scheduling of long running applications in shared production clusters. *EuroSys* (2018).
- [15] Xiaohui Gu and Haixun Wang. 2009. Online anomaly prediction for robust cluster systems. In *2009 IEEE 25th International Conference on Data Engineering*. IEEE, 1000–1011.
- [16] David Gunning. 2017. Explainable artificial intelligence (xai). *Defense Advanced Research Projects Agency (DARPA), nd Web* (2017).
- [17] Sangeetha Abdu Jyothi, Carlo Curino, Ishai Menache, Shravan Matthur Narayana-murthy, Alexey Tumanov, Jonathan Yaniv, Ruslan Mavlyutov, Iñigo Goiri, Subru Krishnan, Janardhan Kulkarni, and Sriram Rao. 2016. Morpheus: Towards Automated SLOs for Enterprise Clusters. In *OSDI*.
- [18] Teuvo Kohonen. 2012. *Self-organizing maps*. Vol. 30. Springer Science & Business Media.
- [19] linkedinds 2019. An Introduction to AI at LinkedIn. <https://engineering.linkedin.com/blog/2018/10/an-introduction-to-ai-at-linkedin>.
- [20] Nicolai Meinshausen. 2006. Quantile Regression Forests. *Journal of Machine Learning Research* 7 (2006), 983–999.
- [21] Haibo Mi, Huaimin Wang, Gang Yin, Hua Cai, Qi Zhou, and Tingtao Sun. 2012. Performance problems diagnosis in cloud computing systems by mining request trace logs. In *2012 IEEE Network Operations and Management Symposium*. IEEE, 893–899.
- [22] Haibo Mi, Huaimin Wang, Gang Yin, Hua Cai, Qi Zhou, Tingtao Sun, and Yangfan Zhou. 2011. Magnifier: Online detection of performance problems in large-scale cloud computing systems. In *2011 IEEE International Conference on Services Computing*. IEEE, 418–425.
- [23] Microsoft. 2019. *Azure Machine Learning Service - Build, train, and deploy models from the cloud to the edge*.
- [24] MLflow. 2019. *MLflow - A platform for machine learning lifecycle*.
- [25] Chirag Modi, Dhiren Patel, Bhavesh Borisaniya, Hiren Patel, Avi Patel, and Muttukrishnan Rajarajan. 2013. A survey of intrusion detection techniques in cloud. *Journal of network and computer applications* 36, 1 (2013), 42–57.
- [26] Hiep Nguyen, Yongmin Tan, and Xiaohui Gu. 2011. Pal: Propagation-aware anomaly localization for cloud hosted distributed applications. In *Managing Large-scale Systems via the Analysis of System Logs and the Application of Machine Learning Techniques*. ACM, 1.
- [27] Alexander Ratner, Stephen H. Bach, Henry R. Ehrenberg, Jason Alan Fries, Sen Wu, and Christopher Ré. 2017. Snorkel: Rapid Training Data Creation with Weak Supervision. *PVLDB* (2017).
- [28] Ando Saabas. 2018. *TreeInterpreter*.
- [29] Wojciech Samek, Thomas Wiegand, and Klaus-Robert Müller. 2017. Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models. *arXiv preprint arXiv:1708.08296* (2017).
- [30] Yongmin Tan, Xiaohui Gu, and Haixun Wang. 2010. Adaptive system anomaly prediction for large-scale hosting infrastructures. In *Proceedings of the 29th ACM SIGACT-SIGOPS symposium on Principles of distributed computing*. ACM, 173–182.
- [31] Yongmin Tan, Hiep Nguyen, Zhiming Shen, Xiaohui Gu, Chitra Venkatramani, and Deepak Rajan. 2012. Prepare: Predictive performance anomaly prevention for virtualized cloud systems. In *2012 IEEE 32nd International Conference on Distributed Computing Systems*. IEEE, 285–294.
- [32] uberds 2019. How Uber Organizes Around Machine Learning. <https://urlzs.com/J4Rk9>.
- [33] Soeren H Welling, Hanne HF Refsgaard, Per B Brockhoff, and Line H Clemmensen. 2016. Forest floor visualizations of random forests. *arXiv preprint arXiv:1605.09196* (2016).
- [34] Qiannan Zhang, Yafei Wu, Tian Huang, and Yongxin Zhu. 2013. An intelligent anomaly detection and reasoning scheme for VM live migration via cloud data mining. In *2013 IEEE 25th International Conference on Tools with Artificial Intelligence*. IEEE, 412–419.
- [35] Jingren Zhou, Nicolas Bruno, Ming-Chuan Wu, Per-Åke Larson, Ronnie Chaiken, and Darren Shakib. 2012. SCOPE: parallel databases meet MapReduce. *VLDB J.* 21, 5 (2012), 611–636.