

Automated Cause Analysis of Latency Outliers Using System-Level Dependency Graphs

Sneh Patel
sp1800@brocku.ca

Brock University
St. Catharines, ON Canada
L2S 3A1

Brendan Park
bp18ul@brocku.ca

Brock University
St. Catharines, ON Canada
L2S 3A1

Naser Ezzati-Jivan
nezzati@brocku.ca

Brock University
St. Catharines, ON Canada
L2S 3A1

Quentin Fournier
quentin.fournier@polymtl.ca

Polytechnique Montréal
Montreal, QC Canada
H3T 1J4

Abstract—Detecting performance issues and identifying their root causes in the runtime is a challenging task. Typically, developers use methods such as logging and tracing to identify bottlenecks. These solutions are, however, not ideal as they are time-consuming and require manual effort. In this paper, we propose a method to automate the task of detecting latency outliers using system-level traces and then comparing them to identify the root cause(s). Our method makes use of dependency graphs to show internal interactions between threads and system resources. With these graphs, one can pinpoint where performance issues occur. However, a single trace can be composed of a large number of requests, each generating one graph. To automate the task of identifying outliers within the dataset, we use machine learning density-based models and statistical calculations such as Z -score. Our evaluation shows an accuracy greater than 97% on outlier detection, making them appropriate for in-production servers and industry-level use cases.

Index Terms—Performance evaluation; latency outliers; outlier detection; root cause analysis; execution tracing; dependency graphs.

I. INTRODUCTION

Performance issues such as unexpected latency can negatively affect a program. During the development phase, debuggers and profilers help detect and resolve potential performance issues originating from bugs or poor design choices. However, after the application has been released, new variables that the developer may not have expected, such as different workloads or a new piece of hardware, may cause performance issues. Investigating latency issues of a released application and identifying their potential root causes is known to be a difficult task. One method often used for performance issues is execution trace-based runtime analysis.

Analyses are conducted either on-CPU or off-CPU depending on the type of issue [1, 2]. When the analysis is conducted on-CPU, any threads running on the CPU can be analyzed. As a consequence, on-CPU analysis is only capable of detecting issues that occur during the execution time on the CPU. Such analysis is limited as it cannot detect performance issues related to thread scheduling and thread interactions. These performance anomalies can be detected with an off-CPU analysis. Notably, off-CPU analysis can detect issues related to hardware, I/O, network timers, and thread-waiting.

For instance, let us consider a Java program that performs numerical computations with integers read from a local text file. In this situation, the on-CPU analysis can detect performance issues related to the calculation, while the off-CPU analysis can detect performance issues associated with reading the file. The proposed method in this paper is about analyzing latency outliers using off-CPU analysis. We collect runtime execution data by tracing the operating system kernel. We then convert the collected trace data into dependency graphs to model the internal structure of an execution. We base our latency and root cause analyses on these graphs.

A trace represents all the functions and operations accessed at the layer traced, either the kernel or the user space, during a program's execution. Each time that a tracepoint¹ is encountered during runtime, an event is generated and added to the trace. An event is composed of a name, a timestamp, and possibly many other arguments. For example, the following event (`thread1`, `syscall_open`, `file1`, `cpu0`, `t1`) corresponds to a file open system call, running on core 0, called on thread 1 at time `t1`. The Linux kernel already contains read-to-use tracepoints, making kernel-level tracing straightforward. Furthermore, kernel trace data includes all the interactions of a program with other threads/processes running simultaneously on the system and the system resources accessed during the execution of a program. This makes kernel tracing an excellent data source to investigate execution issues such as performance and latency issues. However, given the size of the collected trace data, they can be hard to analyze and may introduce a high overhead to the system. Therefore, there is a real need to introduce an efficient and automatic analysis tool for massive traces such as those collected at the kernel level.

In this paper, we propose a method to meet that demand. Our method automates the task of finding latency outliers in system-level traces. Additionally, the method is able to identify the root cause of latency issues by comparing individual outliers with clusters of normal executions to find shared and distinct behaviours between them. Automating the task of identifying outliers reduces the manual effort required to analyze a trace dataset. To automate the task, we first

¹A tracepoint is a tracing macro added to the source code or binary code

filter out the parts of the trace that do not correspond to the program. Using the filtered trace data, we then construct Waiting-Dependency Graphs (DepGraphs) [1], displaying the internal interactions between threads and system resources. Next, we use graph embedding techniques to convert the DepGraphs into fixed-size vectors compatible with most off-the-shelf machine learning algorithms. We use the vectors in density-based clustering models to detect outliers. Finally, we use our proposed merging and comparison algorithms to find the areas contributing to the latency issues. Our work in this paper expands upon what was presented by Ezzati-Jivan et al. [1]. The authors discuss converting system-level traces into DepGraphs to reduce the effort required to analyze raw trace data. In this paper, our main contribution is the automation of analyzing DepGraphs for the outlier detection as well as clustering them to use in discovering the root cause of the outliers.

The remainder of this paper is organized as follows. Section II reviews the related work. Section III explains our proposed method for outlier detection and root cause analysis. Section IV provides an evaluation regarding the automated analysis along with a use case. Section V summarizes the results and concludes this paper.

II. RELATED WORK

Performance is an essential consideration for many developers when designing their programs. For this reason, developers often spend a significant fraction of their time investigating latency issues and optimizing their code. This section looks at some of the existing methods designed to detect and analyze performance issues.

The first work is GAPP [3] which uses weighted criticality measures to identify different serialization bottlenecks in parallel Linux applications caused by an imbalance or shared resource contention. GAPP can reveal a wide range of bottleneck-related performance issues, and the authors show how the extended Berkeley Packet Filter (eBPF) tracer can make a lightweight-based profiler. When used in combination with the bottleneck detector, it will pinpoint the lines of code responsible for the bottleneck. However, profilers are known to be ineffective because they operate by averaging the metrics, which may hide outliers [4].

Another method for bottleneck detection is using a causal profiling technique. These techniques discover bottlenecks and display the effects on the program were it to be optimized. COZ [5] is a causal profiler that uses a virtual speed up to perform causal profiling without actually optimizing the program's code. However, limitations to causal profiling do exist. For instance, causal profiling only applies to user space as it affects only the thread level. To address this issue, Ahn et al. [6] introduced SCOZ, a system-wide causal profiler. Nevertheless, SCOZ is also a profiler and thus ineffective when it comes to outliers.

Similar to our proposed method, Inagaki et al. [7] uses thread graphs to detect layered bottlenecks. They presented a method to build a model that detects layered bottlenecks

by profiling threads and their dependencies in the target system. This method is also capable of analyzing third-party components. They detect layered bottlenecks by hierarchically traversing the path with the most significant threads in their thread dependency graph. One limitation that remains in this work is the manual effort required from the user, which our method seeks to reduce.

Noble and Cook [8] introduces two new techniques for graph-based anomaly detection. They present an algorithm called Subdue that can find repetitive patterns (substructures) within graphs. For their first approach, the authors use Subdue to find patterns within the dataset, and then any substructure that is opposite of the patterns is declared to be an anomaly. For their second approach, the authors separate the graph into distinct structures and compare them to each other to find anomalous subgraphs. The limitation of these methods is that Subdue can be very biased towards smaller graphs.

Lane and Brodley [9] propose machine learning algorithms to detect anomalous behaviour within their program. Their system learns a user profile and subsequently applies it to detect anomalous behaviours. The program compares the current behaviours with the user profiles using a similarity measure. To determine whether the current behaviour is abnormal, input tokens are divided into fixed-size segments and compared using a similarity measure. Using this technique, the authors aim to detect intrusion attacks done by an outsider automatically.

Critical paths [10] can also be utilized to resolve performance problems. A critical path can show the state of threads' executions at different points of time to the user. A drawback of critical paths is that they only show a limited view of the execution rather than the complete process. For example, the critical path of a thread that is waiting for another thread to finish would not display the details of that other thread. Dependency graphs, proposed by Ezzati-Jivan et al. [1] and adapted in our work, can show all thread interactions and help identify bottlenecks.

A common factor with existing tools for performance anomaly detection is that they require some degree of manual labour, making the analysis time-consuming and tedious. Our proposed method uses Waiting-Dependency Graphs (DepGraphs) to illustrate the interactions between different threads and system resources (i.e., disk, CPU, network, etc.). Unlike Ezzati-Jivan et al. [1], however, we aim to avoid manual analysis of these DepGraphs. Instead, we want to automate the process of detecting outliers within a dataset of DepGraphs. To do so, we utilize machine learning algorithms over the DepGraph data to identify outliers. After detecting outliers, we can compare them with a group of normal DepGraphs to identify the root causes of the anomalies.

III. METHODOLOGY

This section provides an in-depth explanation of the proposed anomaly detection and root cause identification method. Using a kernel-level trace dataset, we extracted multiple requests to convert them into DepGraphs. Each DepGraph shows an overall view of the request's internal behaviours, which can

then pinpoint latency issues. However, the analysis becomes time-consuming as the number of DepGraphs increases. Our proposed approach converts the DepGraphs into graph embeddings that are later used in machine learning clustering and outlier detection algorithms to detect and investigate outliers automatically. After identifying the outliers, we compare them with clusters of normal executions to pinpoint the exact reasons behind the issues. In the following paragraphs, we explain each step in more detail.

A. System-Level Data Collection

We use a low-overhead open-source Linux tracing tool called the Linux Tracing Toolkit: next generation (LTTng [11]) to collect kernel-level trace data from a Linux operating system. The collected trace includes the full execution details of all active threads of the system. A thread may handle several tasks in its entire lifetime. We are, however, interested only in some aspects of a group of active threads. We first extract the requests (e.g., web requests) handled by a web server thread (or a group of similar threads, say Apache threads) from its kernel execution trace data, and then we construct a DepGraph for each request. This paper relies on the algorithm presented by Ezzati-Jivan et al. [1] to extract the requests and construct the DepGraphs from trace data. Although this paper focuses on web requests, our method is generic enough to apply to any execution (e.g., a button click, a compiler event, a graphical-view rendering, or a network service).

B. Waiting-Dependency Graphs

A Waiting-Dependency Graph (DepGraph) [1] is a directed graph whose nodes are labelled with the name of the thread, system call, or other execution states and with their runtime duration in milliseconds. In this graph, edges indicate the percentage of time the source node spends waiting for the destination node. DepGraphs show a visual representation of a request's internal execution breakdown. Including the aggregated duration of a thread running on the CPU and the aggregated duration of a thread waiting for a system resource (e.g., disk, network, timer, CPU).

In previous works, developers would use critical paths to observe how long each thread's state takes [4]. A critical path is a tool to display the execution states of a given thread at any particular time. However, a critical path cannot generally show why a thread is in a waiting state or how long a thread has been waiting for a resource. One would need to do a manual investigation to find why a thread is in a waiting state. If only a few threads exist investigating them is an easy task to do. However, in a real-world scenario, there are usually a large number of threads, making them tedious to investigate. In contrast, DepGraphs display the overall breakdown of waiting dependencies that cause latency issues. In essence, DepGraphs are the aggregation of all interacting threads' critical paths that contribute to handling a request. Consider, for example, a web server thread that is communicating with a database thread to handle a web request. In this case, the DepGraph of the web request is constructed by merging the critical paths of the two

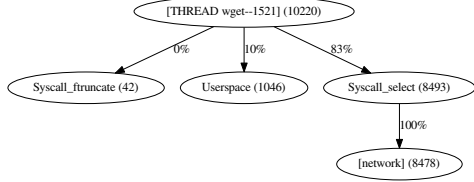
threads. For more information regarding the construction of DepGraphs we refer the reader to the original paper by Ezzati-Jivan et al. [1].

DepGraphs allow finding latency issues and their root causes by comparing the graphs of the different normal and abnormal requests executions. To illustrate how DepGraphs can be used to find anomalies' root cause, let us consider the example of a web server that runs a PHP script to print all the table records from a MySQL database. The script repeats this database access three times. The command-line web client `wget` was used to send two consecutive requests to the server. In this example, one would expect the second call to be faster due to client and server caching mechanisms. However, we observed the opposite to be true. For this scenario, we have made four DepGraphs (Figure 1): one for each client and server execution. Note that all edges with less than 3% have been removed for readability since the edges with low values provide no significant latency to the execution. In Figure 1a, the DepGraph of the client execution of the first request shows that 83% of the `wget` thread's execution was spent waiting for the network. In Figure 1b, the DepGraph of the first request's server execution, we observe what occurred on the server during that time. Note that although `syscall_ftruncate` takes less than 3% of the time, we preserved it for comparison with the second request. Figure 1b shows that the majority of the server time was spent waiting for the MySQL database to handle the queries. For the second request, from the server's DepGraph shown in Figure 1d, we can see that the database requests only take 700ms (compared to the 6298ms of the first request). Such a decrease in response time indicates that database files/queries are (probably) cached in memory from the first call, making the second database call much more efficient. Therefore, the web/database server is not the bottleneck. The DepGraph of the second request's client execution (Figure 1c) explains why the second request is slower than the first request. The latency is mainly due to a new state in the execution of the `wget` client called "Syscall_ftruncate" where the `wget` waits for the disk to write something while already busy serving tasks from two other threads (`[kworker/7:1H--307]` and `[ltnng-consumerd--17665]`). This explains why the total time of the second request has increased by almost three times compared to the first request.

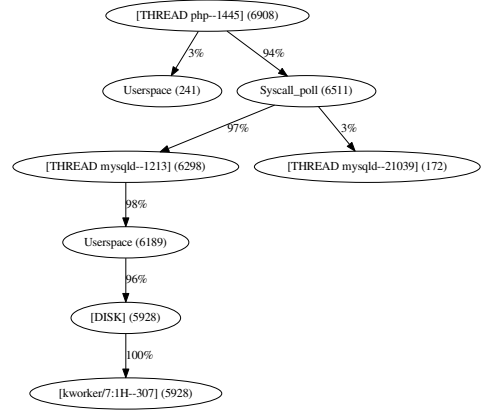
As illustrated by the example, the analysis of DepGraphs, although very useful, is manual and requires significant human labour to identify issues. In order to reduce the manual labour behind analyzing these DepGraphs, we convert them into fixed-size vectors and analyze them with machine learning algorithms.

C. Graph2Vec

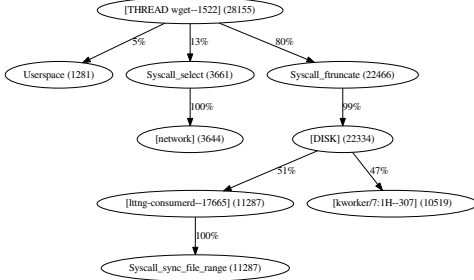
The proposed method is to analyze DepGraphs with machine learning algorithms automatically. Most off-the-shelf machine learning algorithms require real-value vectors as input; thus, one must embed the graphs to use them as input. Graph embedding is the approach of transforming nodes, edges, and other graph features into a vector space.



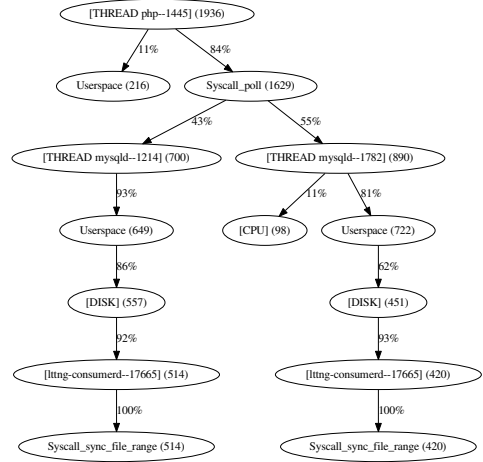
(a) Client DepGraph of the first request.



(b) Server DepGraph of the first request.



(c) Client DepGraph of the second request.



(d) Server DepGraph of the second request.

Fig. 1: The Waiting-Dependency Graphs of two `wget` requests handled by PHP web server. (Top) The first request behave as expected. (Bottom) The second request is slower than the first one, although MySQL cached the query.

For our proposed method, we used a technique introduced by Narayanan et al. [12] called Graph2Vec. Given a set of unidirectional graphs $\mathbf{G} = \{G_1 \dots G_n\}$, Graph2Vec is able to learn graph embeddings of any given size that contains information pertaining to the graph’s topology. Given a dataset of graphs, Graph2Vec attempts to explore all the nodes of their sub-graph up to a certain degree. The technique follows Doc2Vec [13] Skip-gram training process. It first maps random numerical values to the vector spaces, and then using the Weisfeiler-Lehman (WL) relabeling strategy, they are trained and refined over several epochs.

To show an example of converting a DepGraph into a graph embedding using Graph2Vec, we converted the DepGraph of the previous example, shown in Figure 1d. Figure 2 displays the input edges and nodes and the corresponding embedding output. It is worth noting that similar graphs have similar vector representations, and identical graphs have similar but not identical vector representations since embeddings are randomly initialized. This shortcoming is not a concern for

our approach since the strong similarity is enough. Additionally, Graph2Vec requires graphs with undirected edges, while DepGraphs have directed edges. Therefore, to use Graph2Vec with our dataset, we need to convert all the directed edges into undirected edges, losing some information in the process. Nonetheless, Graph2Vec is suitable for our method since it preserves most of the relevant information.

D. Automatic Analysis

As explained in the previous section III-C, a graph embedding is learned for each DepGraph in the kernel trace dataset. Using outlier detection algorithms, we can automate the process of finding outliers by clustering and analyzing similar DepGraphs. After identifying the outliers, we use a comparison algorithm to compare each outlier to a cluster of normal executions to find the differences, hence the cause of the given outlier.

1) *Outlier Detection*: The definition of an outlier as provided by Hawkins [14] is: “an observation which deviates so

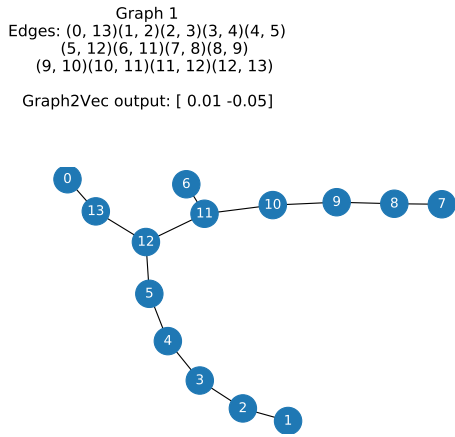


Fig. 2: Converted the DepGraph shown in Figure 1d into an embedding using Graph2Vec.

much from the other observations as to arouse suspicions that it was generated by a different mechanism.”

Outliers may correspond to noise in the dataset, unusual patterns or behaviours of interest, or anomalies such as latency, intrusions, and bugs. In our paper, all the outliers detected are latency-related issues. We consider a latency outlier to be an execution/request with a high execution/response time. Most of the requests in our dataset have a runtime lower than 200ms (97.8% of the dataset). Any request that has a runtime above 200ms is considered an outlier. Furthermore, any request with unusual internal nodes is considered an outlier, even if the execution time is lower than 200ms. After all, they represent a system resource, system call, or request that has an unusual behaviour compared to the other requests in the dataset. Using the outliers we are able to observe the potential reason for the performance issue, since it displays the internal threads/processes with high runtime or it might show unnecessary threads/process that the request is using, making it slow.

From our dataset, we show the DepGraph in Figure 3a as an outlier compared to the DepGraph in Figure 3b. Figure 3a has a runtime of 710ms and also contains elements that are unusual compared to other DepGraphs in the dataset, such as `thread Xorg`. Meanwhile, Figure 3b is a normal request since it only has a runtime of 78ms, and all of the elements are similar to other DepGraphs in the dataset. Later on, in Section III-D2, we display the use of a comparison algorithm to compare the outlier shown in Figure 3a with the other normal DepGraphs to find the differences between the two and find the root cause of the outlier.

We use three different models to detect outliers: Z -score, distance-based model, and density-based model.

The Z -score approach indicates how many standard deviations away a point lies from the sample’s mean value. We decided to use Z -score because it is an effective method to detect outliers on a Gaussian distribution [15] and because it is easy to implement. The limitation of the Z -score is that it

may provide inaccurate results with smaller datasets. However, the sample size would not usually be an issue for our use case since, for accurate Z -score calculations, the sample size should be greater than 30. Z -score uses the mean and standard deviation of the dataset to calculate its value. Outliers can highly influence these values. Therefore, datasets with too many outliers may also yield inaccurate results.

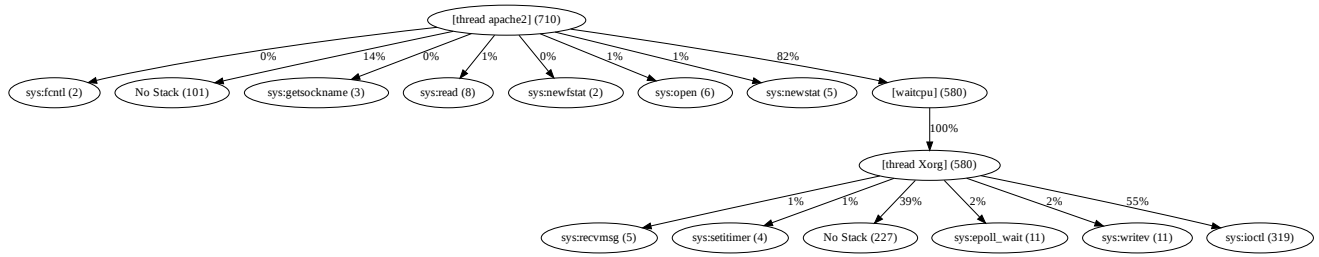
For the distance-based model, we used the k -nearest neighbour (k -NN) algorithm. A data point is defined as an outlier if the distance between it and its k -nearest neighbour is greater than a predefined threshold. While the distance-based models are easy to understand and implement, they are sensitive to the distance metric.

Density-based models assume that outliers are points that lie in a sparsely populated region of space. We consider the Density-Based Spatial Clustering of Applications with Noise (DBSCAN [16]) and Ordering Points To Identify the Clustering Structure (OPTICS [17]) methods. Although these methods were designed for clustering, they are also able to identify outliers as a byproduct. DBSCAN and OPTICS work similarly to each other. Both algorithms require two parameters: ϵ and m . ϵ is the maximum distance between two data points for them to be neighbours, and m is the minimum number of points required to create a cluster. The algorithms start by selecting a random point x , then count the points in its ϵ -neighbourhood. If the number of points is greater than m , DBSCAN creates a new cluster with x and its neighbours, whereas OPTICS creates an ordered list based on the reach-ability distance by keeping a priority queue. If the above condition is not satisfied, DBSCAN and OPTICS both consider x to be an outlier. However, the neighbours of x are considered and added to the cluster in DBSCAN, and OPTICS will add it to its ordered list. After this process, OPTICS and DBSCAN choose a new data point. When all data points have been considered, DBSCAN and OPTICS terminate. It is important to note that outliers declared by DBSCAN and OPTICS are only points that lie in a low-density region.

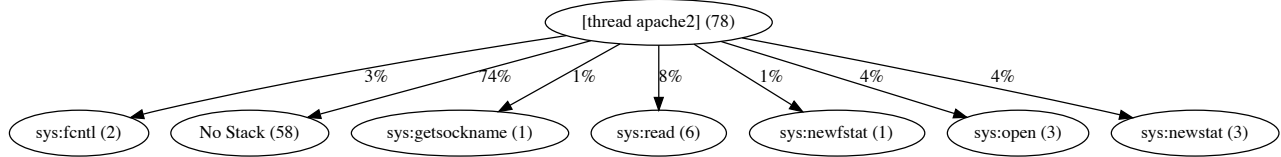
Figure 4a and Figure 4b show a visualization of the DepGraphs clustering with OPTICS and DBSCAN. From these figures, we observe that both methods cluster points in a similar manner, as they were able to get the same clusters and outliers.

Our reasoning behind using DBSCAN and OPTICS is that we can detect outliers with great accuracy and cluster the rest of the dataset into likewise groups. The clusters created by these algorithms are later used in Section III-D2. In Section III-D2 we show our comparison method to compare the outliers with the cluster of normal executions to find the difference between the two and possibly identify the root cause of an outlier.

Evaluating the outlier detection is challenging since it is an unsupervised problem. In the past, synthetically generated datasets or a few rare aspects of a real dataset were often used as proxies to evaluate unsupervised algorithms. Since external criteria such as known labels have restricted requirements, one might look for internal criteria for outlier validation. However,



(a) This DepGraph is considered to be an outlier because it contains novel elements compared to other DepGraphs in our dataset. Also it is one of the only DepGraph with that big of a request time.



(b) This DepGraph is considered to be a normal request, since it has a runtime under 200ms and it contains similar elements as the other clustered DepGraphs.

Fig. 3: Displays an example of a normal request and an outlier to differentiate between the two.

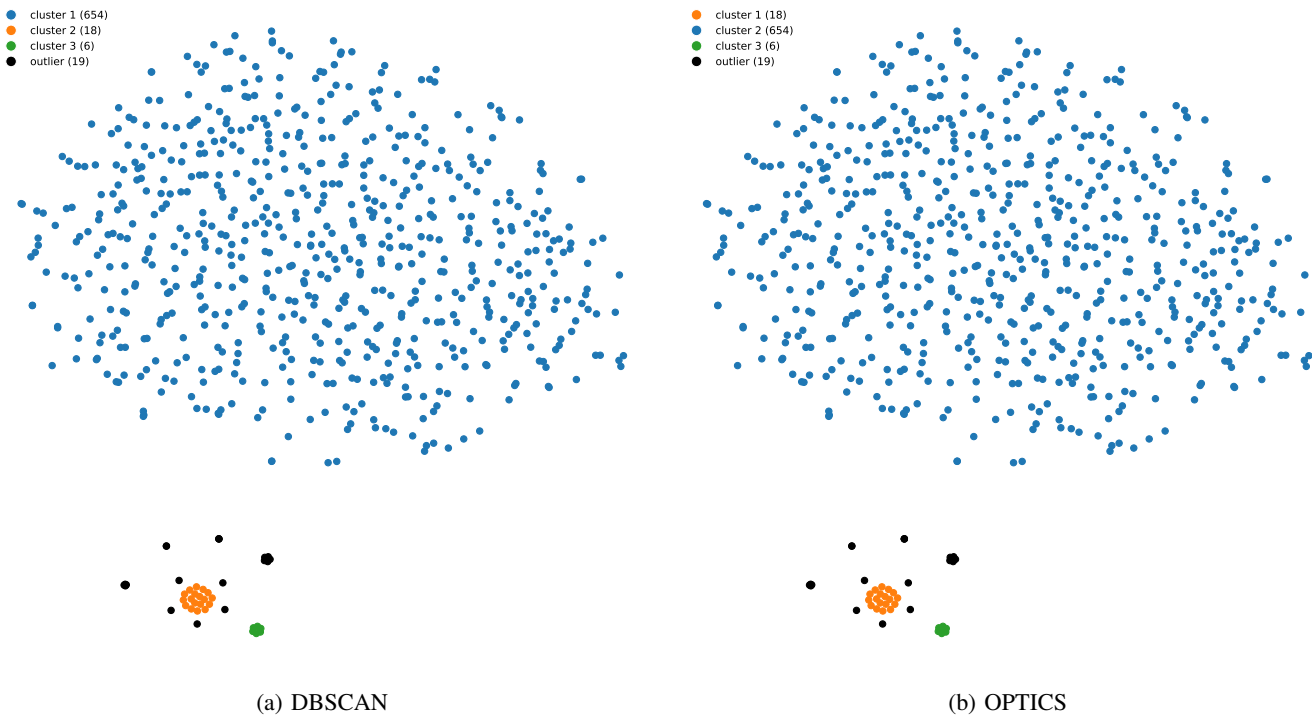


Fig. 4: Results of two different clustering methods. The visualization is obtained by projecting the output of Graph2Vec in 2D using t-SNE [18]. (a) DBSCAN is a clustering method that can produce outliers as a byproduct of clustering. It is not built for outlier detection, however, it is often used for that. (b) OPTICS produces similar results to DBSCAN's since they are both density-based methods.

internal criteria are hardly ever used in outlier analysis because they are known to be flawed [19]. Internal criteria's flaw

is that since outliers are rare by nature, a certain validity measure might favour one outlier detection algorithm over others. Therefore most of the validity measures used in outlier analysis tend to be external measures.

Typically, unsupervised algorithms are evaluated with internal measures such as the value of a loss function on the training set. Outlier detection algorithms may also be evaluated with ground truths retrieved from synthetic or real datasets. In that case, one may see the outlier detection task as a supervised problem. Typically, soft methods return a probability of belonging to the positive class, and a simple decision threshold is used to make the classification (outlier or normal). Then, standard metrics such as accuracy, precision, and recall may be computed.

An outlier detection algorithm’s primary goal is to declare all positive outliers and ensure as few false positives and false negatives as possible. In order to achieve this in a real-world scenario, one may do a grid search on the decision threshold value² or a random search³.

The dataset used for this paper did not contain the ground truths. As a solution, we considered slow requests as the only outliers, although some fast requests could also contain unusual behaviours or patterns. The binary labels were generated using a response time threshold set by an expert and used to evaluate the different outlier detection methods. The threshold was 200ms since the majority of the dataset was under 200ms. Based on this threshold value 2.5% of the dataset was declared as an outlier.

TABLE I: Accuracy, precision, recall of each outlier detection method used.

	DBSCAN	OPTICS	<i>k</i> -NN	Z-SCORE
# of outliers	19	19	17	20
Accuracy (%)	97.7	97.7	97.1	98.1
Precision (%)	47.4	47.7	35.3	55.0
Recall (%)	60	60	40.0	73.3
F1 (%)	52.9	52.9	37.5	62.9

Table I displays the accuracy, precision, recall, and f1 score of each outlier detection method considered (DBSCAN, OPTICS, Z-Score, and *k*-means). Each of these metrics requires two parameters ground truths and predictions. While the ground truth remains the same for each method, the predictions are different depending on the outlier detection method.

The results obtained from this evaluation methodology should be considered with caution since detected outliers may correspond to anomalies unrelated to latency, in which case we acknowledge the prediction to be erroneous. The accuracy measures the ratio of correctly predicted samples, both normal and outliers, while the precision and recall only consider the outliers. The f1 score is a harmonic mean of precision and

²A grid search exhaustively considers all parameter combinations for the set of values specified.

³A random search samples a given number of parameter combinations from the distributions specified.

recall. The table allows us to determine which outlier detection methods were successful in detecting slow requests.

2) *Comparing DepGraphs*: So far, using the proposed outlier detection methods, we have been able to find anomalies and group similar DepGraphs. Now using these groups and outliers, we would like to find the root cause of these anomalies. This is achieved by first merging all the DepGraphs in each cluster. By merging all the DepGraphs in a cluster, we create a new DepGraph representing the entire cluster. We can then use the comparison algorithm to discover the differences between the outliers and the merged graphs.

Algorithm 1 Method used to merge the cluster

```

1: procedure MERGE(file)
2:   Data ← formatData(path, function(file)) ▷
   Collect all the executions and make sure, there are no
   loose nodes
3:   max = dictionary()
4:   min = dictionary()
5:   count = dictionary()
6:   size = dictionary()
7:   for each execution in data do
8:     for each node in the given execution do ▷ Add
   the node if and only if its distinct else add the value
9:       if nodePath is not in counts then
10:        count[nodePath] = 1
11:      else
12:        count[nodePath] = count[nodePath] + 1
13:      end if
14:      if nodePath is not in min then
15:        min[nodePath] = nodePath[min]
16:      else if nodePath[min] < min[nodePath] then
17:        min[nodePath] = nodePath[min]
18:      end if
19:      if nodePath is not in max then
20:        max[nodePath] = nodePath[max]
21:      else if nodePath[max] > max[nodePath] then
22:        max[nodePath] = nodePath[max]
23:      end if
24:      if nodePath is not in sizes then
25:        sizes[nodePath] = nodePath[size]
26:      else
27:        sizes[nodePath] = sizes[nodePath] +
   nodePath[size]
28:      end if
29:    end for
30:  end for
31:  CREATEGRAPH(size, count, min, max) ▷ Pass on
   the new distinct nodes to make a new DepGraph
32: end procedure

```

Algorithm 1 displays the method used to merge the clusters. As input, it takes a file comprising of the DepGraphs to merge. After retrieving the file, the algorithm formats the data by extracting all the information of each execution and ensuring no nodes without a parent exist. Then, the algorithm analyzes

each node from each execution (line 7). If the node is unseen, the dictionary `counts` would add it; otherwise, it would increase the count of that node, indicating that the node is repeated (lines 9-13). The algorithm also records the minimum and maximum values of all similar nodes. Therefore if a node’s size is less than the current size, it will be updated as the new minimum. Similarly, if a node’s size is greater than the current size, it will be updated as the new maximum. If it is a distinct node, its value becomes the maximum and the minimum (lines 14-23). Furthermore, the algorithm stores the sizes of each distinct node. If the node already exists, it adds its value with the new node’s value (lines 24-28). Using the `size`, `count`, `min`, and `max` (line 31) dictionaries, the algorithm is able to make a `DepGraph` that has the `id` and `name` of all distinct nodes, the cumulative size of all repeated nodes, the number of repeated nodes (`count`), the `max` and `min` values from all the common nodes, and the `path`.

We then use the comparison algorithm shown in Algorithm 2 to compare the merged graph (i.e., the representative graph of the normal cluster) and the `DepGraph` of the outlier request. As input, the algorithm takes two files, each containing a `DepGraph` to compare. In our case, one `DepGraph` represented the clusters, and the other represented the outlier. After retrieving both files, the algorithm extracts all the nodes from both executions and merged their data (`counts`, `max`, `min`, and `size`) using a similar method shown in Algorithm 1. Next, the algorithm gets the total number of repetitions for each node from both executions (lines 12-22). Using the `totCount`, the algorithm calculates the mean count of each node in the first execution and calculates the standard deviation of the first `DepGraph` (`leftSds`). Later it repeats the process and calculates the mean count of each node in the second execution (lines 23-30). With the two means, the algorithm can find the differences between the means, which is needed to find the similarities between the two `DepGraphs`. The algorithm iterates over each difference in `diffMeans` and calculates the difference between the standard deviation of the means. If the two nodes are present in both the `DepGraphs`, the algorithm uses `getBoldness` method to compare each node’s count. Depending on the standard deviation, it can determine the level of boldness. The higher the standard deviation, the higher the difference between the counts of the two nodes. We use boldness to show the differences in the value of the corresponding nodes of the given graphs. The `getBoldness` method uses five different boldness levels. The larger the value difference, the higher the boldness level of the edges. The algorithm can also show distinct nodes that may only appear in one of the two executions. If the node is only present in the first `DepGraph`, it will have a dashed edge, and if the node is only present in the second `DepGraph`, it will have a dotted edge. The resulting graph will depict the exact differences (i.e., value differences between corresponding nodes and existence of a node in only one tree) between the two given `DepGraphs`, useful in identifying the root cause(s) of the outlier request.

To display an example of using the merge and comparison algorithms to find root causes of anomalies, we will compare

the outlier shown in Figure 3a and the normal `DepGraph` in Figure 3b. Using OPTICS, we can cluster the normal `DepGraph` in cluster 2 and identify the outlier. Before comparing the cluster with the outlier, we first need to merge the entire cluster into one `DepGraph` since the comparison algorithm can only take two `DepGraphs` as input. Using the Algorithm 1, we merged cluster 2 into a single representative `DepGraph` with all distinct nodes. Figure 5 displays the final `DepGraph` after merging all 654 (normal) `DepGraphs` in the cluster. In this figure, each node has a name, the total cumulative size of all the same nodes, the maximum and minimum of all the same nodes (i.e., the execution/run time of that node), and the number of repeated nodes. For example, in the `DepGraph`, the node `thread apache2` shows a count of 654, which means that 654 `DepGraphs` had the same node in their `DepGraphs`. Out of the 654 nodes, the highest value was the runtime of 306ms, and the lowest value was the runtime of 60ms. The node also shows a cumulative size of 69,791, which allows computing the percentage of time spent by the thread in the node. With the single merged `DepGraph` that represents cluster 2, we can now compare it with the outlier in Figure 3a using the comparison Algorithm 2. Using the Algorithm 2, we get the comparison graph shown in Figure 6.

The comparison graph shown in Figure 6 depicts all distinct and common nodes from both of the `DepGraphs`. Nodes with dotted edges are from the first `DepGraph` (representing cluster 2), and nodes with dashed edges are from the second `DepGraph` (outlier). If an edge is a simple straight line (possibly with boldness level), it represents that the node is present in both the `DepGraphs`. In Figure 6, we can observe that most nodes are present in both `DepGraphs`. However, `waitcpu` is only present in the second `DepGraph` (outlier). By investigating the `waitcpu`, we can identify nodes (e.g., other active thread(s), which is the thread `xorg` here) that take up the majority of the CPU time while the main thread is waiting to get the CPU. In other words, this dash-lined node from the difference graph shows that the CPU was utilized by `xorg` thread, and consequently, it is not available for the `apache2` thread to handle the web request causing it to get blocked and wait longer than usual to get the CPU. We can conclude that this is one of the reasons why this request is an outlier. There might be, however, some other reasons that can be seen from the difference graph.

IV. EVALUATION

A. Computational Cost

1) *Setup*: The trace used for our example above was collected using LTTng 2.11 on a workstation equipped with 32 GB of RAM and a quad-core Intel® Core™ i7-6700K CPU @ 4.00 GHz. The operating system was Linux Ubuntu 16.04.6 LTS with the 64-bit kernel 4.15.0-62. The local hard disk reference 7200 RPM WDC WD1003FZEX-0 stored the trace data and trace analysis results.

2) *Data*: The trace data was collected using the LTTng tracer and comprised 697 requests. The trace contains the execution of an Apache web server handling a PHP web

Algorithm 2 Method to compare the merged cluster DepGraph and the outlier.

```
1: procedure MERGE(file1, file2)
2:   data1  $\leftarrow$  formatData(file1)
3:   data2  $\leftarrow$  formatData(file2)
4:   totCount = dictionary()
5:   means1 = dictionary()
6:   means2 = dictionary()
7:   leftCounts = []
8:   meanDiffSd = dictionary()
9:   diffMeans = dictionary()
10:  Merge the counts, max, min and size for the data1 and data2            $\triangleright$  For loop on line 7 in Merge algorithm 1
11:  for each node in execution1 do
12:    if node[path] in totCount then
13:      totCount[node[path]] = totCount[node[path]] + node[count]
14:    else totCount[node[path]] = node[count]
15:    end if
16:  end for
17:  for each node in execution2 do
18:    if node[path] in totCount then
19:      totCount[node[path]] = totCount[node[path]] + node[count]
20:    else totCount[node[path]] = node[count]
21:    end if
22:  end for
23:                                      $\triangleright$  mean of each node count is count/totalcount for that node
24:  for node in execution1 do
25:    means1[node[path]] = node[count]/totCount[node[path]]
26:    leftCounts.append(node[count])
27:  end for
28:  leftSds = standardDeviation(leftCounts)
29:  for node in execution2 do
30:    means2[node[path]] = node[count]/totCount[node[path]]
31:  end for
32:  for mean in means1 do
33:    if mean is in means2 then
34:      diffMeans[means] = means[mean] - means[mean]
35:    end if
36:  end for
37:  for diff in diffMeans do
38:    if leftSds != None then
39:      meanDiffSd[diff] = diffMeans[diff]/leftSds
40:    else
41:      meanDiffSd[diff] = infinity
42:    end if
43:  end for
44:  for x in meanDiffsd do
45:    boldness[x] = getBold(meanDiffSd[x])
46:  end for
47:  CREATEGRAPH(Mergedmin, Mergedmax, Mergedcount, Mergedsize, boldness)
48:                                      $\triangleright$  since the comparison graph has all nodes from both depgraphs it needs the merged data we got at line 10
49: end procedure
```

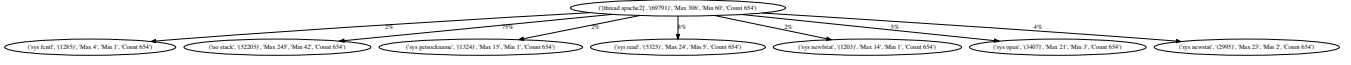


Fig. 5: This DepGraph is a result of merging all the DepGraphs in cluster 2 from the OPTICS clustering method.

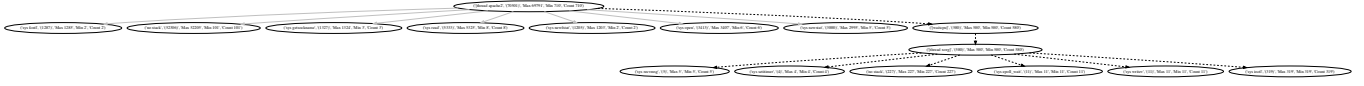


Fig. 6: This is a comparison graph that compares the cluster 2 from OPTICS and the outlier shown in Figure 3a

application. The internal execution of each request is reflected in a distinct DepGraph. Therefore, there are 697 DepGraphs in the dataset. Our observation shows that 97.8% of all requests have a runtime between 0 and 200ms (considered as normal executions), as shown in Figure 7. Therefore, for this dataset, we consider any request and corresponding DepGraph that has a runtime below 200ms to be a normal request and the rest to be slow requests. Normal requests can still contain anomalies that are unrelated to latency issues. More details of the dataset and the constructed DepGraphs can be found in Ezzati-Jivan et al. [1].

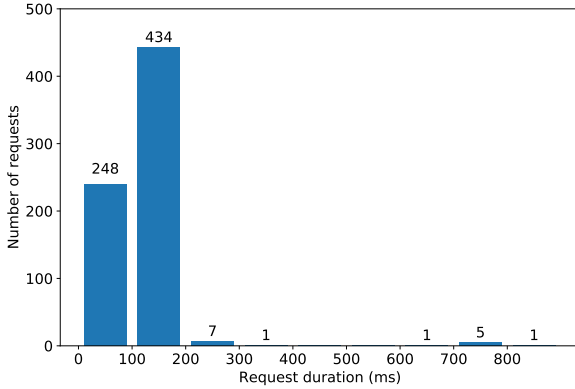


Fig. 7: Histogram of the request duration in ms.

3) *Tracing and DepGraph Construction Cost*: In the worst-case scenario of tracing the Linux kernel where all kernel tracepoints are enabled, which is not the case of our method, the overhead imposed by tracing is 42%. In our method, when collecting the trace to construct DepGraphs, only system calls and events needed to extract the DepGraphs are enabled. Our analyses show that the overhead imposed by the tracing never exceeds 10.1%. After obtaining the trace, the DepGraphs were created. When creating the DepGraphs, one needs to analyze the trace, which is done offline after collecting it. Therefore it does not have any overhead on the application and is not considered in our time analysis.

4) *Automated-Analysis*: Table II displays the time taken for various tasks in Section III-D. The table shows the time it took to load the dataset, make the graph embeddings, and

run each outlier detection and clustering method. From that, we observed that DBSCAN took the least amount of time to cluster and detect outliers. The table also shows how long it took to merge a cluster into one DepGraph and compare it with an outlier. For this table, we measured the time taken to merge cluster 2 and the time taken to compare it with the outlier shown in Figure 3a. Here we notice that the majority of the time was spent constructing the DepGraphs and embedding, whereas everything else took less than a second to do.

The overhead imposed by our method is modest compared to the cost of tracing and certainly negligible compared to the time required by a human operator to investigate each request manually.

B. Limitation

The proposed method described in this paper is heavily reliant on Graph2Vec used to get our fixed-size representation of a trace. However, it has a few limitations. When training the Skip-gram model to produce the embeddings, it is primarily reliant on the embedding size and the hyper-parameters `wl_iterations` and `epochs`. Therefore when we increase the epochs and iterations, the time it takes to train the model also increases. If a significant portion of the graphs used to train the model is large, that can also increase the time it takes to train the model. Also, to train the model to embed, a minimum number of graphs are needed. If there is less than the minimum threshold, we run into a small sample problem, resulting in overfitting.

Another limitation our proposed method has is that it can only detect off-CPU issues and cannot detect any issues that happen at the on-CPU level. Since our proposed method focuses on finding latency issues within dependency graphs, this is not an issue. However, one cannot use our method to detect issues within their code (on-CPU analysis). To do so, we would need to get data from on-CPU profilers. Since there are no interactions between kernel space and user space, one cannot get information about the user space functions from the kernel level.

V. CONCLUSION AND FUTURE WORK

Our proposed method is able to automatically detect outliers and their root causes from within a system-level trace data. Using our method, we remove the need to manually examine all the trace data and the abstract data (such the dependency

TABLE II: Time taken by each task.

		Time (s)
Data Processing	Load DepGraphs	0.076
	DepGraph Embeddings w/ Graph2Vec	11.608
Cluster	DBSCAN	0.020
	OPTICS	0.325
Outlier Detection	DBSCAN	0.017
	OPTICS	0.275
	k -NN	0.075
	Z-Score	0.025
Merging clusters	Load the cluster	0.027
	Merge DepGraphs in cluster	0.033
	Construct DepGraphs	32.391
Compare outlier with cluster	Load the two DepGraphs	0.001
	Merge both execution for details	0.001
	Compare the two DepGraphs	0.002
	Construct comparison graph	0.672

graphs) extracted from them. Furthermore, with our approach to compare the declared outliers with the normal executions, we can identify unusual activities contributing to the latency. This method makes the process of analyzing the software execution trace more efficient for developers, as automating the process reduces the time taken for debugging and gets the task done to find potential bottlenecks in the program. Additionally, having an automated process removes the potential for human error in the analysis. For future work, we would like to extend our approach by reducing the graph embedding size while still retaining the level of accuracy (or potentially increasing it). By reducing the dimension of the vectors, we can train the models faster while increasing the number of dependency graphs.

REFERENCES

- [1] N. Ezzati-Jivan, Q. Fournier, M. R. Dagenais, and A. Hamou-Lhadj, “Depgraph: Localizing performance bottlenecks in multi-core applications using waiting dependency graphs and software tracing,” in *2020 IEEE 20th International Working Conference on Source Code Analysis and Manipulation (SCAM)*. IEEE, 2020, pp. 149–159.
- [2] F. Zhou, Y. Gan, S. Ma, and Y. Wang, “wperf: Generic off-cpu analysis to identify bottleneck waiting events,” in *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*. Carlsbad, CA: USENIX Association, Oct. 2018, pp. 527–543. [Online]. Available: <https://www.usenix.org/conference/osdi18/presentation/zhou>
- [3] R. Nair and T. Field, “Gapp: A fast profiler for detecting serialization bottlenecks in parallel linux applications,” in *Proceedings of the ACM/SPEC International Conference on Performance Engineering*, 2020, pp. 257–264.
- [4] Q. Fournier, N. Ezzati-jivan, D. Aloise, and M. R. Dagenais, “Automatic cause detection of performance problems in web applications,” in *2019 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW)*. IEEE, 2019, pp. 398–405.
- [5] C. Curtsinger and E. D. Berger, “Coz: Finding code that counts with causal profiling,” in *Proceedings of the 25th Symposium on Operating Systems Principles*, 2015, pp. 184–197.
- [6] M. Ahn, D. Kim, T. Nam, and J. Jeong, “Scoz: A system-wide causal profiler for multicore systems,” *Software: Practice and Experience*, 2020.
- [7] T. Inagaki, Y. Ueda, T. Nakaike, and M. Ohara, “Profile-based detection of layered bottlenecks,” in *Proceedings of the 2019 ACM/SPEC International Conference on Performance Engineering*, 2019, pp. 197–208.
- [8] C. C. Noble and D. J. Cook, “Graph-based anomaly detection,” in *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD ’03. New York, NY, USA: Association for Computing Machinery, 2003, p. 631–636. [Online]. Available: <https://doi.org/10.1145/956750.956831>
- [9] T. Lane and C. E. Brodley, “An application of machine learning to anomaly detection,” in *Proceedings of the 20th National Information Systems Security Conference*, vol. 377. Baltimore, USA, 1997, pp. 366–380.
- [10] F. Doray and M. Dagenais, “Diagnosing performance variations by comparing multi-level execution traces,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 28, no. 2, pp. 462–474, 2016.
- [11] M. Desnoyers and M. R. Dagenais, “The lttng tracer: A low impact performance and behavior monitor for gnu/linux,” in *OLS (Ottawa Linux Symposium) 2006*, 2006, pp. 209–224.
- [12] A. Narayanan, M. Chandramohan, R. Venkatesan, L. Chen, Y. Liu, and S. Jaiswal, “graph2vec: Learning distributed representations of graphs,” *CoRR*, vol. abs/1707.05005, 2017. [Online]. Available: <http://arxiv.org/abs/1707.05005>
- [13] Q. V. Le and T. Mikolov, “Distributed representations of sentences and documents,” *CoRR*, vol. abs/1405.4053,

2014. [Online]. Available: <http://arxiv.org/abs/1405.4053>
- [14] D. M. Hawkins, *Identification of outliers*. Springer, 1980, vol. 11.
- [15] S. Santoyo, “A brief overview of outlier detection techniques,” Nov. 2017. [Online]. Available: <https://towardsdatascience.com/a-brief-overview-of-outlier-detection-techniques-1e0b2c19e561>
- [16] M. Ester, H.-P. Kriegel, J. Sander, X. Xu *et al.*, “A density-based algorithm for discovering clusters in large spatial databases with noise.” in *Kdd*, vol. 96, no. 34, 1996, pp. 226–231.
- [17] M. Ester, “Density-based clustering.” *Data Clustering: Algorithms and Applications*, vol. 31, p. 111, 2013.
- [18] L. van der Maaten and G. Hinton, “Visualizing data using t-sne,” *Journal of Machine Learning Research*, vol. 9, no. 86, pp. 2579–2605, 2008. [Online]. Available: <http://jmlr.org/papers/v9/vandermaaten08a.html>
- [19] C. C. Aggarwal, *Data mining: the textbook*. Springer, 2015.