

Detecting Anomalous Packets in Network Transfers: Investigations using PCA, Autoencoder and Isolation Forest in TCP Packets

Mariam Kiran*, Cong Wang[†], George Papadimitriou[‡], Anirban Mandal[†], Ewa Deelman[‡]

*Lawrence Berkeley National Laboratory, CA, USA

[†]RENCI, University of North Carolina at Chapel Hill, NC, USA

[‡]University of Southern California, Information Sciences Institute, Marina del Rey, CA, USA

Abstract—Large-scale scientific workflows rely heavily on high-performance file transfers. These transfers require strict quality parameters such as guaranteed bandwidth, no packet loss or data duplication. To have successful file transfers, methods such as predetermined thresholds and statistical analysis need to be done to determine abnormal patterns. Network administrators routinely monitor and analyze network data for diagnosing and alleviating these, making decisions based on their experience. However, as networks grow and become complex, monitoring large data files and quickly processing them, makes it improbable to identify errors and rectify these.

Abnormal file transfers have been classified by simply setting alert thresholds, via tools such as PerfSonar and TCP statistics (Tstat). This paper investigates the feasibility of unsupervised feature extraction methods for identifying network anomaly patterns with three unsupervised classification methods - Principal Component Analysis (PCA), Autoencoder and Isolation Forest. We collect file transfer statistics from two experiment sets - synthetic iPerf generated traffic and 1000 Genome workflow runs, with synthetically introduced anomalies. Our results show that while PCA and a simple Autoencoder finds it difficult to detect clusters, the tree-variant isolation forest is able to identify anomalous packets by breaking down TCP traces into tree classes early.

Index Terms—TCP anomalies, file transfer, unsupervised feature extraction, PCA, autoencoders, isolation forest

I. INTRODUCTION

Today’s high-performance networked systems are highly complex, consisting of many heterogeneous hardware devices and software working together. These systems are essential to people’s daily lives, from mobile connectivity to inter-cloud communications, from checking emails to running complex scientific workflows on distributed cloud infrastructures. Reliable and high throughput connectivity, security and 99.9% availability are just some of the guarantees that are warranted from them. These networking infrastructures support diverse user requirements and are constantly under pressure to provide end-to-end connectivity that is reliable, with high performance, minimum packet loss, and sometimes providing dedicated links for certain data-intensive applications.

Performance and reliability of high-speed networks are affected by many factors, such as congestion, packet loss, end-host I/O performance, end-host network tuning, and more. The complexity of data movement also brings increased potential

for failures and performance problems that need to be detected and mitigated.

Majority of data transfers rely on the TCP protocol to guarantee packet delivery with minimum loss. Network researchers have studied TCP performance by performing anomaly detection and exploring reliable packet delivery using monitoring tools such as PerfSonar [1], sFlow [2], TCP statistics [3] and more [4]. In addition to employing simple statistical analysis, multivariate machine learning techniques have been used to find outliers in traffic statistics, hardware/software failures in security software, and improving utilization [5].

Usually, identification of abnormal transfers is done through classification or clustering methods. Examples include using Naive Bayes Theorem, Support Vector Machines, Random Forests or rule-based approaches to detect traffic anomalies [6], [7]. But using these methods requires large amounts of data about packet numbers, flow direction, IP addresses, QoS values, latency or file traces to characterize normal traffic behavior. These methods also require prior and domain knowledge to help classify what is “normal”. However, networks are extremely dynamic and what is normal for one, may not be true for others. Distributed, heterogeneous and complex nature of end-to-end networked systems makes it difficult to collect labeled data about anomalies and failures. Additionally, dynamic characteristics of the platform-induced anomalies often manifest themselves in unknown manners. Hence, lack of labeled data sets makes it impossible to appropriately train and validate any machine learning models for detecting transfer anomalies. Some researchers have used formula-based methods to understand packet behaviors or history-based methods [7]. However, again these methods would work only under certain conditions, and fail when network systems change or data set becomes too big to process expeditiously.

In this work, we analyze anomalous network transfers by utilizing data collected using Tstat [3], which is a tool to collect TCP traces for transfers. Figure 1 shows how TCP works. The server sends a packet to a client with a packet load (in bytes), when the client receives this, it sends an “acknowledgment” (ACK) back. The total time for receiving the ACK is recorded as round trip time (RTT). The TCP protocol also uses time windows to allow servers to wait for




Figure 1: Round Trip Time (RTT) measured in the three way communication in TCP.

the ACK, before deciding to resend the packet again. The longer the wait, the window size grows over time, allowing the protocol to adjust its waiting time. This not only affects the total RTT recorded but also builds up retransmissions and overflowing flows in both TCP directions.

The TCP protocol is designed to ensure all packets are delivered reliably. It does this by tracking various packet information, for example window timeouts, packet numbers, RTT and more, to calculate if loss has occurred. If yes, it triggers the host to resend the packets again. Collectively, Tstat traces contain about 150 variables per packet on both server and client sides.

Anomalous network transfers have been classified into three major groups - packet loss, packet duplication and retransmissions [8], [9]. The retransmissions could be triggered by lack of client acknowledgments or some other error in the link. The TCP window size measures the waiting time before TCP sends the packet again. For instance, this may cause more retransmissions and eventually congestion on the link, causing packet loss eventually. Current approaches classify anomalous behaviors using thresholds and prior known formulas, but due to TCP complexity, it is difficult to understand root causes and relevant features/variables.

This paper departs from using prior knowledge, and uses unsupervised feature extraction to learn normal and abnormal features. Based on three techniques - Principle Component Analysis (PCA), Autoencoders and Isolation Forest - we build and train classifiers to help identify anomalous transfers. The goal of this study is to understand what can be used to classify TCP abnormal behavior and the feasibility of unsupervised classification techniques for doing this.

Approach and Contributions: We approach the problem of anomaly detection by performing unsupervised feature extraction on TCP traces collected from simple transfers between two nodes (iPerf [10] transfer) and from transfers in a real scientific workflow (1000 Genome workflow [11]), to find unique characteristics from normal transfers and from transfers with synthetically introduced anomalies (e.g. packet loss, packet duplication and synthetic reordering).

To inform the machine learning models with some labeled data sets, the experiments are set up to generate data in a real network environment. File transfers using the iPerf tool and a workflow manager are set up on the ExoGENI net-

worked cloud testbed [12], where the traces in these controlled conditions, are used to inform common features in a reliable transfer. The purpose of doing two kinds of transfers is (1) to compare the results of workflow transfer versus iPerf transfer experiments, and (2) to learn common features of normal transfers given whatever the infrastructure setup is.

By leveraging innovations in machine learning and existing research on TCP protocols, we make the following contributions in this paper:

- We develop and test multiple unsupervised feature extraction methods, including Principle Component Analysis, Isolation Forest and Autoencoders, to study TCP traces and deduce anomalous features.
- We perform experiments on a real testbed (ExoGENI) with iPerf transfers and real workflow network transfers. Each experiment was run for one hour with synthetic anomalies like packet loss, duplication, and reordering. We extract key features from TCP statistics to understand packet behavior to improve transfer performance.
- We test anomaly detection methods using a set of synthetically generated network traffic using the iPerf tool and a real scientific workflow called the 1000 Genome workflow, which produces a reference for human variation, having reconstructed the genomes of 2,504 individuals across 26 different populations [11].
- We perform in-depth analysis of issues and data dependencies of the algorithms, laying the foundation for further research on feature extraction and TCP performance.

Motivation for this work: This work is unique to TCP anomaly detection and relevant for understanding transfer performance. With increasing network complexity and transfer sizes, it is imperative to find classification techniques that can detect bad transfers as soon as they happen. Exploring multiple techniques, we can determine what issues are faced while building these classifiers and whether such tools can be built to work in tandem with monitoring tools.

The rest of this paper is organized as follows: Section 2 discusses the problem formulation, as well as TCP, data preparation and machine learning techniques. Section 3 provides evaluations in two scenarios - iPerf transfer experiments and 1000 Genome workflow experiments. This section also discusses how the training data set is collected and features extracted. Section 4 provides a discussion of the feature extraction methods. Section 5 presents related work, and finally, Section 6 concludes the paper.

II. PROBLEM FORMULATION

Networks have a significant impact on scientific workflow performance. Gaikwad et al. [13] show how science workflows can suffer on multiple levels of hardware infrastructure, software, middleware, application and workflows reacting to anomalous events during execution. To monitor these behaviors, researchers deploy measurement tools (e.g. perfSONAR, Tstat, netflow, sflow, to name a few), monitoring various properties like usage traffic, router health, link performance, and more. Each of these tools focuses on different aspects:

No.	Experiment type (runs for 1 hour)	Tstat Sample
1	Normal Traffic	152
2	No flow	120
3	1% Loss	120
4	5% Loss	120
5	1% Packet Duplication	60
6	5% Packet Duplication	60
7	25% - 50% Packet reordering	50
8	50% - 50% Packet reordering	60

Table I: Scenario 1: iPerf Transfer Experiments

No.	Experiment type	Tstat Sample
1	Normal workflow runs	1475
2	1% Loss	1563
3	2% Loss	1636
4	3% Loss	1948
5	1% Packet Duplication	1574
6	5% Packet Duplication	1527
7	25% - 50% Packet reordering	1491
8	50% - 50% Packet reordering	1592

Table II: Scenario 2: 1000 Genome Workflow Experiments

PerfSONAR actively probes packet loss, delay, jitter and utilization, whereas netflow records packet data with IP addresses, bytes sent and the hops taken. In this paper, we primarily focus on TCP statistics to determine features for file transfers.

The Tstat tool records TCP statistics such as client/server IP address, packets sent/received, end-to-end RTT, packets observed in a TCP window, payload and more [3]. Mellia et al. [8] discuss TCP packet anomalies collected at the source. Other works such as [14], [15] capture web behavior by mining Tstat to understand passive and active network attacks. However, they use statistical thresholds and some prior knowledge to help determine if an anomaly is present.

A. What makes a good data set for feature extraction?

Machine learning techniques need correct and large data sets for training purposes. A better training set can greatly improve the accuracy of the machine learning model. In this paper, we use two sets of network experiments: (1) controlled experiment with two nodes generating synthetic iPerf traffic, and (2) real scientific workflow experiment using the 1000 Genome workflow software [11]. Traces are recorded for successful transfers and anomalous transfers with packet loss, duplication and reordering packets. These labeled traces are then fed to unsupervised feature extraction methods to extract features.

Table I and II document the traces collected per experiment run. Both experiments, have different network topology setups and different transfers, as seen in the initial throughput recorded (Figure 2).

B. Data preparation

1) *Processing skewed iPerf transfers*: iPerf sometimes randomly sends very large packets in one flow. This creates various throughput discrepancies and affects the overall sample reliability. To prevent this, large transfers were removed using k-means clustering. These clusters help group the very large

transfers as shown in (Figure 3), ensuring the generation of correct data samples.

2) *Tstat logs*: Tstat collects logs where each record corresponds to a single TCP flow. It collects data in both directions: (1) Client-to-Server (C2S) and (2) Server-to-Client (S2C), while each Tstat record contains approximately 150 variables. Among these variables there are details on the source and destination IP addresses, ports, number of bytes transferred, completion times, congestion window sizes and more. Each TCP connection is established by the first SYN segment and ends either when a FIN segment is observed or the default timeout of 10s is reached. Tstat discards all connections for which the three-way handshake is not properly seen in the `log_tcp_nocomplete` log file, while it logs the rest of them in the `log_tcp_complete` file. A comprehensive list of Tstat variables can be found in [16].

C. Unsupervised feature extraction algorithms

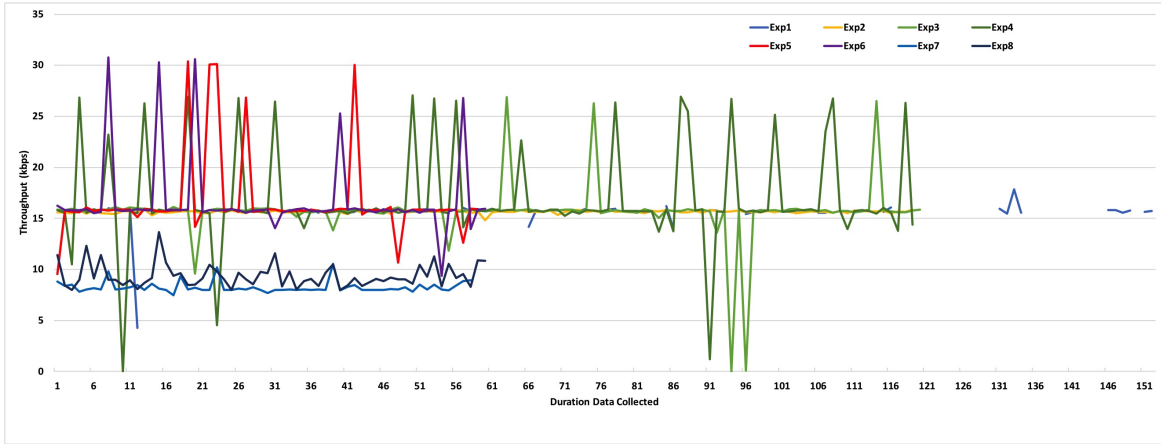
A feature is a property of a data sample. Average, mean, median and standard deviation can also be features. In TCP transfer, each recorded detail can be a feature which can be recognized for successful transfers, or “normal”, or packet loss, as “bad” features. Unsupervised feature extraction can be used to identify these features from a trace log data.

1) *Principal Component Analysis*: Principal Component Analysis (PCA) [17] is an unsupervised feature extraction algorithm that can be used to reduce dimensions and extract vectors summarizing data properties and correlation among variables. PCA depends on linear combinations and constructs new features (e.g. eigenvectors) that summarize variation among data variables. These eigenvectors can help find the most influential variables (out of 150) and also reduce the data dimension for faster processing. In this paper, we use PCA to extract eigenvectors as features of the packet transfer. The k^{th} principal component of a data vector $x(i)$ can be given a score $tk(i) = x(i) \otimes w(k)$ in a transformed coordinates space. The corresponding vector in this space is $x(i) \otimes w(k)$, where $w(k)$ is the k^{th} eigenvector of $X^T X$.


The transfers are labeled per scenario and fed into the PCA algorithm. To visualize these eigenvectors, we can plot these to identify if two clear clusters are formed.

2) *Autoencoders*: Autoencoders [18] belong to the class of deep learning algorithms and are used for data or image compression and decompression. Autoencoders take training data input, compress it to reduce the dimensions into a compressed representation. These are then decoded from the compressed version to find the original data set, by calculating a loss between the result and the original input. Therefore, the algorithm removes noise (or randomness) in the input data and learns key features only. Autoencoders are often used in image recognition and text analysis, where they can successfully compress features into learned sets.

Autoencoders take a high-dimensional space (e.g. 150x100 (150 variables and 100 transfers)) and reduce it to vectors of size 15000. We train the autoencoder to minimize the loss rate between the original input and output data. The reduced



(a) iPerf Transfer Experiments.



(b) 1000 Genome Workflow.

Figure 2: Throughput (in one direction) across iPerf Transfer and 1000 Genome Workflow Experiments.




Figure 3: K-means to identify bad iPerf transfers.

dimensions can then be visualized on a 2D space to visualize clusters formed. Autoencoders consist of two parts: (1) the encoder ϕ and (2) the decoder ψ of the input data set X , represented by:

$$\phi : X \rightarrow F \quad (1)$$

$$\psi : F \rightarrow X \quad (2)$$

$$\phi, \psi = \operatorname{argmin} \|X - (\psi \circ \phi)X\|^2 \quad (3)$$

In our case, the transfers are labeled based on their experiment number and fed to Autoencoder to reduce the dimensions and extract features.

3) *Isolation Forest*: Isolation Forest [19] is a relatively new, unsupervised classification algorithm that is commonly used for outlier detection in high-dimensional data sets. This algorithm is based on the fact that anomalies are data points that are usually fewer in number and have different distribution patterns from the normal data. The algorithm constructs a separation by first creating isolation trees, or random decision trees, and then calculates the score as the path lengths to isolate observations. It isolates these observations by randomly selecting a feature and then randomly selecting a split value between the maximum and minimum values of the selected feature. Usually, random partitioning produces noticeably shorter paths for anomalies. Hence, when a forest of random trees collectively produces shorter path lengths for particular samples, they are highly likely to be anomalies.

With X as input data, L as a number of trees and N as training samples, the Isolation Forest algorithm is shown in

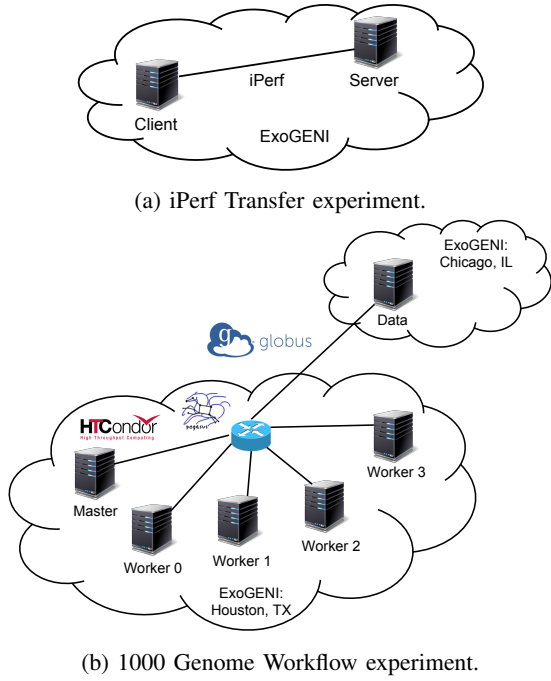


Figure 4: Experiment set up for iPerf transfers and 1000 Genome workflow.

Algorithm 1 Random forest

```

Initialize Forest = ( )
Set  $i$  tree height  $h = \text{ceiling}(\log_2 N)$ 
for  $i=1$  to  $L$  do
   $X = \text{sample}(X, N)$ 
  Forest =  $iTree(X, 0, h)$ 
end for
return Forest

```

Algorithm 1. In our scenario, we train the Isolation Forest algorithm using the data collected from normal workflow transfers. The other experiments are then fed into the algorithm as test data, to check if the Isolation Forest can predict whether the corresponding network flow is normal or anomalous.


III. EXPERIMENTAL EVALUATION

In this section, we discuss the evaluation of our anomaly detection classifiers.

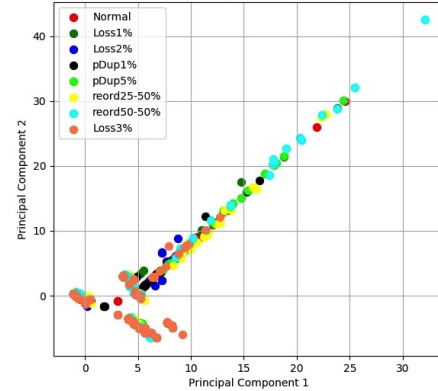
A. Experiment setup

In the iPerf transfer experiment (Figure 4a), we set up two nodes on ExoGENI testbed, each node has 4 cores and 12 GB RAM. The two nodes are connected via a dedicated 500 Mbps network link. We use iPerf to generate TCP traffic and use Linux TC [20] to introduce synthetic network anomalies, like packet loss, duplication and reordering.

For the second experiment, we set up an HTCondor cluster with 4 worker nodes on ExoGENI, (shown in Figure 4b). The bandwidth is set to 500 Mbps for all links and we use the 1000 Genome workflow composed of five different tasks:



(a) iPerf transfer experiments.



(b) 1000 Genome Workflow.

Figure 5: Feature extraction using PCA.

(1) *individuals* - task fetches and parses the data from the 1000 genomes project per chromosome; (2) *populations* - The task fetches and parses five super populations (African, Mixed American, East Asian, European, and South Asian), and a set of all individuals; (3) *sifting* - task computes the SIFT scores of all of the SNPs (single nucleotide polymorphisms) variants, as computed by the Variant Effect Predictor; (4) *pair overlap mutations* - task measures the overlap in mutations (SNPs) among pairs of individuals; and (5) *frequency overlap mutations* - task calculates the frequency of overlapping mutations across subsamples of certain individuals.

B. Extracting TCP Features

Here, we present the results of testing the TCP traces with the different techniques.

1) *PCA - Dimension Reduction for Identifying Variance Features*: PCA is a commonly used dimension reduction method allowing large sets of variables to be reduced to smaller sets containing most information. The new dimensions are sectioned into principal components, where the first component accounts for the major variability in the data, with remaining components containing the rest of the variability. It also allows to select a subset of variables, which have high correlation in the data set.

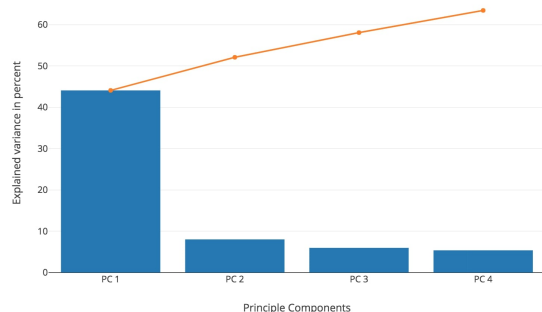


Figure 6: Optimal PCA clusters in 1000 Genome workflow based on variance.

Mainly the two principal components can be used to create a linear combination of variables based on the maximum variance among them. The second linear combination explains the maximum proportion of the remaining variance. PCA produces eigenvectors which plot the variance and weight correlations explaining the weighted component directions. By applying PCA on our two datasets of approx. 150 variables, we are able to extract useful information about them.

Figure 5 shows the 2D PCA plots of both iPerf transfers and 1000 Genome workflow. Reducing the feature dimensions, Figure 5a shows that all normal transfers lie along a diagonal of the two components, while all transfers with anomalies have a higher variance in the second principal component. However, this behavior is not observed in the 1000 Genome workflow. Here all good and bad transfers, all lie along the diagonal of the two principal components (Figure 5b). Therefore the variance in the data is less distinct in the 1000 Genome workflow transfers.

Further, PCA clusters can be analyzed to study how many clusters contain the most variance. Figure 6 shows that in the 1000 Genome workflow, about 60% of the variance is captured in the first two components, justifying the 2D plots. However, unique variance features are still not detected.

PCA is also able to extract the most relevant variables out of the 150 variables we collected to describe the transfer. These influential variables are reduced into 26 variables that include Round Trip time (RTT), Return time Observed (RTO), congestion window details (including size, scaling and segment sizes), which are known as the most important variables in other TCP research [8].

2) *Autoencoders - Compressing and Decompressing to Find Features:* Autoencoders use artificial neural networks to learn unique data features by first compressing the data into lower dimensions and then reconstructing the compressed data into the original data that are as close to the original set as possible. Using backpropagation, most autoencoder architectures extract features of the input distribution as a *good representation* of the data and remove any noise by computing the loss between the reconstructed output and original input. Trained on linear combinations, autoencoders behave very similar to PCA, where the weights of the hidden layers represent the




Figure 7: Autoencoder architecture constructed to extract TCP features.

variance among the principal components [18]. However, these weights are not the same as principal components. In our experiments, we encoded the TCP traces into a three-layer fully connected dense neural network, using one layer to encode the traces, one hidden layer and one output layer to decode them. Figure 7 shows the layer architecture of the autoencoder used. This autoencoder had only three layers—layer 1 for encoder using RELU activation to process the input data, two neuron layer for learning representations and, finally third layer for the decoder that uses the SIGMOID activation function.


During the training phase, the autoencoders learn a 2-dimension feature vector, extracted as learned features. These are presented in Figure 8.

Comparing to PCA results, the autoencoders also reduce the TCP dimension. however, no clear clusters are formed for good and bad file transfers. While autoencoder have given successful results in images, for text data they are more complicated to work with and require additional layers and training.


3) *Isolation Forest - Detecting anomalies through tree classification and scoring:* Isolation forest works by building a tree and classifying data into known nodes on a normal data set. Test data are then parsed using the same data assigning scores of how close the data point lies on the tree. This score can be used to identify data points which lie far from the tree classes (or anomalies) or not. Figures (9, 10, 11) and (12, 13, 14) show the results of isolation forest on the two experiments. In Figure 9a, the no-flow experiment shows that there is some TCP noise in the system, which are random monitoring tool transfers recorded in the link¹. As packet loss is increased, most transfers are not detected as bad, which is also the case with packet duplication. In packet reordering, the anomaly score is quite high and catches most of the transfers as bad, giving an error rate of 0.05% (Figure 15a).

In the 1000 Genome workflow (Figure 12, 13, 14), the error rate is slightly better, as it is able to predict more bad transfers. However, it is still unable to extract key features between the experiments with loss and duplication.

¹PerfSonar tool generates TCP transfers for its monitoring purposes.




(a) iPerf Transfer Experiment.




(b) 1000 Genome workflow.


Figure 8: Feature extraction using Autoencoder.



(a) No Flow.




(b) Packet loss 1%.

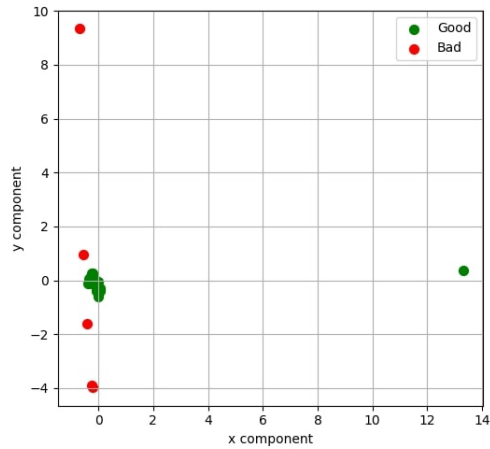


(c) Packet Loss 5%.

Figure 9: Isolation Forest identifying good and bad transfers in iPerf transfer experiments with packet loss.




(a) Packet Duplication 1%.




(b) Packet Duplication 5%.

Figure 10: Isolation Forest identifying good and bad transfers in iPerf transfer experiments with packet duplication.

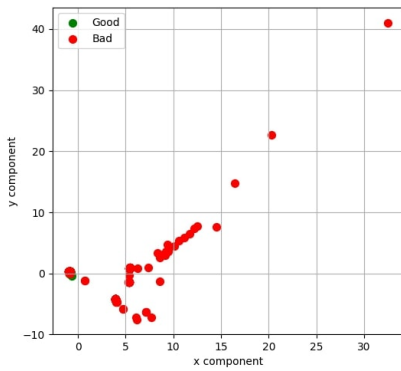


(a) Packet reordering 25%-50%.




(b) Packet reordering 50%-50%.


Figure 11: Isolation Forest identifying good and bad transfers in iPerf transfer experiments with packet reordering.



(a) Packet Loss 1%.




(b) Packet loss 2%.




(c) Packet Loss 3%.

Figure 12: Isolation Forest identifying good and bad transfers in 1000 Genome Workflow in packet loss.




(a) Packet Duplication 1%.




(b) Packet Duplication 5%.

Figure 13: Isolation Forest identifying good and bad transfers in 1000 Genome Workflow in packet duplication.

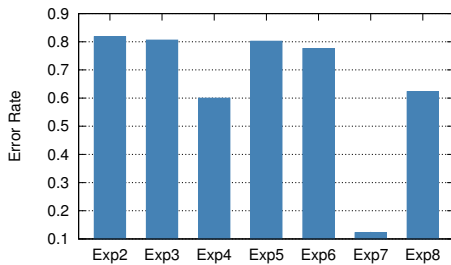


(a) Packet reordering 25%-50%. (b) Packet reordering 50%-50%.

Figure 14: Isolation Forest identifying good and bad transfers in 1000 Genome Workflow in packet reordering.



(a) iPerf transfer experiment.



(b) 1000 Genome Workflow.

Figure 15: Isolation Forest Prediction Error Rate.

IV. DISCUSSIONS

In this section, we discuss the lessons learned from building unsupervised feature extraction classifiers.

A. Transfer learning does not work in networks

Transfer learning can be used to generate training data from smaller sets of experiments to use machine learning models to learn from the representative sets (or simulation-based scenarios) about the real-world problems, where data is often difficult to collect. For example, using simulators to learn real-world driving conditions help improve the models about certain road conditions, for which data does not exist at present [18].

Most machine learning approaches assume that the training and the test data belong to the same distribution. However, in situations where labeled data is generated via transfer learning, it is often that machine learning models trained on controlled experiments cannot recognize reasonable conclusions and features in real-world examples. This implies that test data is a different data distribution than the training data. In traditional transfer learning, the labeled and unlabeled data, both come from the same class to aid with classification [21]. However, in situations where this is not possible, [22] talks about using a learner which can calculate the difference between training and test data, to improve the overall predictions.

This behavior was also witnessed in both network transfer experiments. Any features learned in the iPerf transfers were not detected in the 1000 Genome workflow experiments, which means transfer learning cannot be used here to learn unique features with different network topologies.

B. PCA, as a feature extraction, is not enough

Dimension reduction is a useful technique to extract higher-level features of a data set. PCA is a useful technique which caters to this by using variance matrix, covariance matrix, eigenvector and eigenvalues. However, for our application, we see that PCA is highly dependent on the numerical data. Thus, these cannot be general vector spaces, that can be applied in all situations.

The feature extraction method allows us to reduce the dimension space of 150 TCP variables to 26 key variables, including Round trip time (RTT), Return time observed (RTO), segment size and TCP window details. Additionally, both scenarios have different network topologies. The iPerf transfers had two nodes versus the 1000 genome workflow experiment used ExoGENI sites in Jacksonville, FL and Chicago, IL. This makes the RTT in the workflow to be much higher than the iPerf experiment. Because PCA is dependent on variables, this completely threw off the results and was not able to classify good or bad transfers.

C. Autoencoders cannot capture behavior data through compressed features

Similar to PCA, autoencoders capture the learned weights of the hidden layers to determine the best reconstruction

ratio from input to output. Results show that although these do not capture variance, the learned representations are not sufficient to identify unique clusters of good and bad transfer characteristics.

D. Isolation features show that high-level relationships exist

Our experiments are able to capture packet reordering cases which shows that decision trees are able to build some high-level relationships among the data set. The anomaly score to reordered packets is very high and extracting this feature can help researchers to look for particular relationships. Further work in this area is needed to extract these features separately.

V. RELATED WORK

Computing workflows leverage a collection of many elements including compute, storage and networks. In a recent paper, [23] presented a framework that used machine learning to predict workflow performance and forecast workflow behavior. Presenting results on independent workflows like on the XSEDE SDSC Comet cluster, the paper highlights the dependency of multiple components when trying to predict overall performance. Although most workflow performance improvement techniques have looked at distributed computing to improve performance of computing jobs, networks are becoming overloaded and are prone to breakdowns affecting overall performance. Additionally, [13] presented an anomaly detection method based on autoregression and time-series prediction to understand how workflow anomalies can be detected on networked cloud infrastructures. This highlights the different experimental setups and the anomalies that can affect a healthy workflow based on network topologies.

Jiang et al. [24] used network topology and destination data for detecting anomalous behaviors. This method was able to identify ill behaving devices and users. Similarly, [25] proposed using entropy methods to extract multiple traffic features about flows. While entropy was successful in detecting low-volume anomalies, this method was unable to detect anomalies in high-volume data. In these approaches, the feature extraction method was dependent on the current traffic's nature, which can vary dynamically in a large-scale distributed network system. Therefore, these features were not sufficient to identify common anomalies in abnormal file transfers.

Further, [5] also performed a classification of anomalies based on application identification. Using feature selection techniques for finding optimal flow attributes, the algorithm was able to predict accurate anomalous behavior up to 86%.

Traffic anomalies can also be a result of malicious user or faulty device behavior and can lead to network disruption. Often they can show a degradation of networks, as investigated by [26] detecting slow file transfers as symptoms of failures. The authors detect slow transfers, based on a performance model and on the observed distribution of file transfer rates. The authors find that both methods found different results in different workflow transfers, based on the nature of the transfer.

Just analyzing file transfer, researchers [27] have been investigating anomalies in GridFTP transfers, by creating controlled experiments and injecting faults such as disk injection, latency perturbations and loss, to measure effects on performance. Although this study is focused on GridFTP and statistical methods, it discusses the nature of how anomalies can affect healthy transfers. But et al. [28] showed that just monitoring RTT is not enough for studying network issues. Applying the Jacobson's algorithm, the authors showed the effect of minimizing packet transfer times on TCP stream jitters, which showed additional factors affecting network behaviors. Luca et al. [29] proposed a heuristic classification approach to identify possible anomalies by quantifying the sensitivity of certain parameters on flow measurements. Although based on established literature data, this technique showed that RTT had a high dependence on current traffic load that affected prediction accuracy.

Feature extraction can also be used to investigate user behavior with traffic models [30]. Zhang et al. [31] described anomaly detection as a significant step for network intrusion detection systems (NIDS) to work, but highlighted that these systems are based on known or supervised data sets. The authors recognize that a more robust unsupervised feature extraction method is needed, with features learned from real network data sets, to make the NIDS systems more reliable.

In this paper, we focus on using simple unsupervised feature extraction, based on domain knowledge, to extract certain characteristics from known transfer datasets. Using feature extraction and dimension reduction we build relationships among sensitive parameters such as congestion and availability with a transfer file type. These feature filters can be leveraged for any future machine learning methods. With the goal for understanding packet loss, congestion and impact on end-to-end performance. Our results show positive results with isolation forest in identifying packet reordering features, where PCA and autoencoders have failed. Further investigation of these experiments and future work are discussed in the next section.

VI. CONCLUSION

Feature extraction is a powerful technique to identify behavior characteristics, which can be used to identify anomalous transfers as they happen. In this paper, we used a variety of unsupervised feature extraction, where algorithms were fed only training data sets, to learn and extract what they thought was a feature to predict an anomaly. Using PCA we were able to find clear clusters of normal and abnormal transfer, but this was not seen for the 1000 genome workflow experiment. Alternatively, autoencoders completely failed in extracting clusters on normal or abnormal transfers in both experiments. Isolation Forest, being trained on the normal transfers, was able to give better results, but sometimes did not work (Figure 15).

The results have given a number of unique conclusions. Firstly, the nature of the two scenarios was completely different. The iPerf transfer experiment was run on two VMs that

existed on the same rack, whereas the 1000 Genome workflow involved geographically distributed nodes, Jacksonville and Chicago respectively. The added RTT in the second scenario affected the performance of the PCA. Second, isolation forest is worth exploring further with more extensive experiments across multiple workflows. Third, we were able to identify out of 150 variable only 26 are the most influencing TCP statistics that could give good results of their behaviors - belonging to RTT, RTO, segment size and window details.

In the future, we plan to extend these experiments to build a real-time classifier that can predict anomalous transfers as they happen. We will be extending the isolation forest result with more loss, duplication and reordering experiments, to find key relationships to reduce error rates as experiments have more anomalies in them.

Acknowledgments. This work was funded by DOE contract number #DE-SC0012636M, "Panorama 360: Performance Data Capture and Analysis for End-to-end Scientific Workflows". Any opinions, findings, or conclusions in this article are those of the authors and do not necessarily reflect the views of the sponsor.

REFERENCES

- [1] A. Hanemann, J. W. Boote, E. L. Boyd, J. Durand, L. Kudarimoti, R. Łapacz, D. M. Swany, S. Trocha, and J. Zurawski, "PerfSONAR: A service oriented architecture for multi-domain network monitoring," in *Service-Oriented Computing - ICSOC 2005*, B. Benatallah, F. Casati, and P. Traverso, Eds., 2005.
- [2] E. Jasinska, "sflow, i can feel your traffic," in *Amsterdam Internet Exchange (AMS-IX)*, 2006.
- [3] M. Mellia, "Tcp statistic and analysis tool," *IEEE Network*, vol. 16, 2002.
- [4] R. Hofstede, P. Celeda, B. Trammell, I. Drago, R. Sadre, A. Sperotto, and A. Pras, "Flow monitoring explained: From packet capture to data analysis with netflow and ipfix," in *IEEE Communications Surveys Tutorials*. IEEE Communications Society, 2014, pp. 2037–2064.
- [5] S. Zander, T. Nguyen, and G. Armitage, "Automated traffic classification and application identification using machine learning," in *Proceedings of the The IEEE Conference on Local Computer Networks 30th Anniversary*, ser. LCN, 2005, pp. 250–257.
- [6] N. Bansal and R. Kaushal, "Unusual internet traffic detection at network edge," *International Conference on Computing and Network Communications (CoCoNet)*, 2015.
- [7] M. Mirza, J. Sommers, P. Barford, and X. Zhu, "A machine learning approach to tcp throughput prediction," *IEEE/ACM Transactions on Networking*, vol. 18, no. 4, pp. 1026–1039, Aug 2010.
- [8] M. Mellia, M. Meo, L. Muscariello, and D. Rossi, "Passive analysis of tcp anomalies," *Computer Networks*, vol. 52, pp. 663–2676, 2008.
- [9] P. Casas, P. Fiandino, S. Wassermann, S. Traverso, A. D'Alconzo, E. Tego, F. Matera, and M. Mellia, "Unveiling network and service performance degradation in the wild with mplane," *IEEE Communications Magazine*, vol. 54, pp. 71–79, 2016.
- [10] Iperf, <https://iperf.fr/>.
- [11] 1000 Genomes Project Consortium, "A global reference for human genetic variation," *Nature*, vol. 526, no. 7571, pp. 68–74, 2012.
- [12] I. Baldine, Y. Xin, A. Mandal, P. Ruth, C. Heerman, and J. Chase, "Ex-ogeni: A multi-domain infrastructure-as-a-service testbed," in *Testbeds and Research Infrastructure. Development of Networks and Communities*, Berlin, Heidelberg, 2012.
- [13] P. Gaikwad, A. Mandal, P. Ruth, G. Juve, D. Krol, and E. Deelman, "Anomaly detection for scientific workflow applications on networked clouds," *International Conference on High Performance Computing Simulation*, 2016.
- [14] L. Vassio, F. Figuereido, A. Paula, C. da Silva, M. Mellia, and J. Almeida, "Mining and modeling web trajectories from passive traces," *IEEE Bigtable*, December 2017.
- [15] M. Trevisan, I. Drago, and M. Mellia, "Measuring web speed from passive traces," *ACM, IRTF ISOC Applied Networking Research Workshop 2018 (ANRW 18)*, July 2018.
- [16] A. Finamore, M. Mellia, M. Meo, M. M. Munafo, P. D. Torino, and D. Rossi, "Experiences of internet traffic monitoring with tstat," *IEEE Network*, vol. 25, no. 3, pp. 8–14, May 2011.
- [17] I. Jolliffe, "Principal component analysis," *International Encyclopedia of Statistical Science*, pp. 1094–1096, 2011.
- [18] Y. Bengio, "Learning deep architectures for AI," *Foundations and Trends in Machine Learning*, 2009.
- [19] F. Liu, K. T. Ming, and Z.-H. Zhou, "Isolation forest," in *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining*, ser. ICDM, 2008, pp. 413–422.
- [20] B. Hubert, T. Graf, G. Maxwell, R. van Mook, M. van Oosterhout, P. Schroeder, J. Spaans, and P. Larroy, "Linux advanced routing & traffic control," in *Ottawa Linux Symposium*, vol. 213, 2002.
- [21] R. Raina, A. Battle, H. Lee, B. Packer, and A. Y. Ng, "Self-taught learning: Transfer learning from unlabeled data," in *Proceedings of the 24th International Conference on Machine Learning*, ser. ICML, 2007, pp. 759–766.
- [22] W. Dai, Q. Yang, G.-R. Xue, and Y. Yu, "Boosting for transfer learning," in *Proceedings of the 24th International Conference on Machine Learning*, ser. ICML, 2007, pp. 193–200.
- [23] A. Singh, A. Rao, S. Purawat, and I. Altintas, "A machine learning approach for modular workflow performance prediction," in *Proceedings of the 12th Workshop on Workflows in Support of Large-Scale Science*, ser. WORKS, 2017, pp. 7:1–7:11.
- [24] D. Jiang, Z. Xu, P. Zhang, and T. Zhu, "A transform domain-based anomaly detection approach to network-wide traffic," *J. Netw. Comput. Appl.*, vol. 40, no. C, pp. 292–306, Apr. 2014.
- [25] A. Lakhina, M. Crovella, and C. Diot, "Mining anomalies using traffic feature distributions," in *Proceedings of the 2005 conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM)*, pp. 217–228, 2005.
- [26] M. Yang, X. Liu, W. Kroeger, A. Sim, and K. Wu, "Identifying anomalous file transfer events in lcls workflow," in *Proceedings of the 1st International Workshop on Autonomous Infrastructure for Science*, ser. AI-Science, 2018, pp. 7:1–7:4.
- [27] D. Gunter, B. L. Tierney, A. Brown, M. Swany, J. Bresnahan, and J. M. Schopf, "Log summarization and anomaly detection for troubleshooting distributed systems," in *Proceedings of the 8th IEEE/ACM International Conference on Grid Computing*, ser. GRID, 2007, pp. 226–234.
- [28] J. But, U. Keller, D. Kennedy, and G. Armitage, "Passive tcp stream estimation of rtt and jitter parameters," in *The IEEE Conference on Local Computer Networks (LCN)*, Nov 2005.
- [29] L. Muscariello, M. Mellia, and M. Meo, "Passive identification and analysis of tcp anomalies," *Distributed Cooperative Laboratories: Networking, Instrumentation, and Measurements*, 2006.
- [30] D. Rossi, M. Mellia, and C. Casetti, "User patience and the web: a hands-on investigation," *Global Telecommunications Conference*, 2003.
- [31] J. Zhang and M. Zulkernine, "Anomaly based network intrusion detection with unsupervised outlier detection," *IEEE Communications*, 2006.