

Scalable Architecture for Anomaly Detection and Visualization in Power Generating Assets

Paras Jain^{*}, Chirag Tailor^{*}, Sam Ford^{*}, Liexiao (Richard) Ding[‡], Michael Phillips[‡],
Fang (Cherry) Liu[§], Nagi Gebraeel[‡], Duen Horng (Polo) Chau^{*}

^{*}College of Computing

[‡]H. Milton Stewart School of Industrial & Systems Engineering

[§]Partnership for an Advanced Computing Environment

Georgia Institute of Technology

Atlanta, Georgia, USA.

Email: {paras, chirag.tailor, sford100, richard.ding, mphilips68, polo}@gatech.edu,
fang.liu@oit.gatech.edu, nagi.gebraeel@isye.gatech.edu

Abstract—Power-generating assets (e.g., jet engines, gas turbines) are often instrumented with tens to hundreds of sensors for monitoring physical and performance degradation. Anomaly detection algorithms highlight deviations from predetermined benchmarks with the goal of detecting incipient faults.

We are developing an integrated system to address three key challenges within analyzing sensor data from power-generating assets: (1) difficulty in ingesting and analyzing data from large numbers of machines; (2) prevalence of false alarms generated by anomaly detection algorithms resulting in unnecessary downtime and maintenance; and (3) lack of an integrated visualization that helps users understand and explore the flagged anomalies and relevant sensor context in the energy domain.

We present preliminary results and our key findings in addressing these challenges. Our system’s scalable event ingestion framework, based on OpenTSDB, ingests nearly 400,000 sensor data samples per seconds using a 30 machine cluster. To reduce false alarm rates, we leverage the False Discovery Rate (FDR) algorithm which significantly reduces the number of false alarms. Our visualization tool presents the anomalies and associated content flagged by the FDR algorithm to inform users and practitioners in their decision making process.

We believe our integrated platform will help reduce maintenance costs significantly while increasing asset lifespan. We are working to extend our system on multiple fronts, such as scaling to more data and more compute nodes (70 in total).

Keywords—False discovery rate; visualization; OpenTSDB; power asset; energy sensor

I. INTRODUCTION

To improve public safety, modern power generating assets (e.g., jet engines, gas turbines) are instrumented with hundreds of sensors to monitor physical performance degradation. Such sensors, such as temperature or pressure, is installed with the goal of measuring potential signals of asset failure. Given the scale of data produced, it is impossible for humans to directly monitor every signal. Instead, this monitoring process can be automated by studying possible deviations from pre-specified benchmarks, with the goal of detecting incipient faults.

Anomaly detection involves defining a pattern in observations that represent normal behavior and declaring observations that do not belong to that region as anomalies. A multitude of detection algorithms and techniques have been developed and

commercialized over the years, many of which have been applied in the manufacturing domain [1], [2] for what has become known as Statistical Process Control (SPC).

In large scale settings involving thousands of power generating assets where each asset is monitored by a large number of sensors, the problem of false alarms becomes a significant challenge. False alarms can be very costly — for example, 50% of replaced parts in aircraft are classified as “no fault found” [3]. In energy and aerospace domains, reducing rates of false alarms can measurably reduce lifetime maintenance costs for power generating assets.

Our goal is to develop an integrated system that reduces false alarms in multi-stream condition monitoring of power generating assets using a scalable analytics architecture that ingests, stores and analyzes large amounts of sensor data, and interactively visualize the computation results to enhance user understanding and advance decision making capabilities.

Our ongoing work’s contributions are:

- We present a scalable event ingestion and storage architecture that can handle 399,000 sensor samples per second with a 30 node storage cluster.
- We adapt the FDR algorithm to the energy domain in order to reduce the rate of falsely identified anomalies.
- We demonstrate a visualization tool that enables interactive exploration of power-generating asset sensor data with associated anomalies.

II. SYSTEM OVERVIEW

Our proposed architecture consists of three key components: the anomaly detection algorithm, streaming sensor data ingestion, and interactive visualization. Figure 1 summarizes our envisioned system architecture.

Data storage is non-trivial with large scale deployments of assets generating huge volumes of data. Our system will need to eventually scale to process hundreds of thousands of sensor samples per second. After evaluating candidate storage solutions, we chose to build our platform over the Open Time Series Database (OpenTSDB) [4] to store both streaming sensor data as well as flagged anomalies (Section III). OpenTSDB

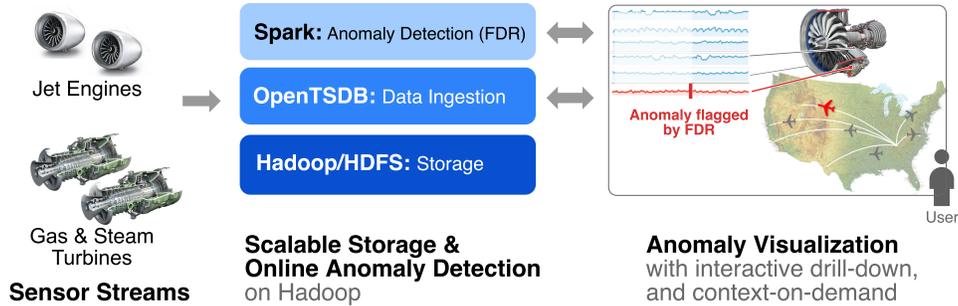


Fig. 1. Overview of our system architecture, with a scalable Hadoop-based storage and anomaly detection backend, and an interactive visual frontend to aid user understanding.

is supported by HBase [5] as the underlying distributed file storage system.

The choice of anomaly detection algorithm is central to the project. A survey [6] of techniques for anomaly detection divides traditional algorithms into several categories: rule-based systems, statistical techniques, spectral techniques and neural networks. As the primary criteria for anomaly detection is to control the rate of false positives, we specifically chose the False Discovery Rate algorithm [7], [8] (described in section IV). Offline evaluation of the anomaly detection algorithm currently executes in the Spark framework [9] in batch mode. Given its rich distributed matrix computation libraries, Spark is a natural choice for evaluating the FDR algorithm.

Both the online anomaly detection and streaming sensor data ingestion components run on the Data Science Platform (DSP) [10] built at Georgia Tech College of Computing. The online anomaly detection component runs on a 44 nodes HDFS/Hadoop/Spark cluster, while the streaming sensor data ingestion runs on a 32 nodes HDFS/HBase/OpenTSDB cluster.

Our visualization tool allows for operators to interactively explore large volumes of time-series data (section V). By enriching sensor data with potential anomalies and integrated real-time analytics, the platform can serve as a powerful control center to monitor faults across large sensor networks.

A. Evaluation Dataset

Real datasets from industry partners contain actual sensor readings from power generating assets like gas turbine and jet engines. However, these datasets often include sensitive information, and currently are not available for off-site evaluation. As our very first goals of developing the system are to investigate the algorithmic scalability, visualization capabilities, and hardware platform requirements, access to a real dataset is not critical to our current stage of investigation. Therefore, we generated a dataset for training and evaluation of the algorithm. This allows measuring the exact degree to which FDR reduces false alarm rates while allowing us to verify the algorithms ability to detect various classes of injected faults.

The training dataset contains 100 simulated units, each with 1000 sensors (on order with the 3000 sensors in the Siemens

SGT5-8000H gas turbine [11]). We modeled three primary categories of faults:

- Pure random noise for comparison
- Pure random noise plus gradual degradation signal
- Pure random noise plus sharp shift

Injected faults are correlated across sensors which allows measuring the algorithm’s response to deviations across multiple signals.

III. SCALABLE DATA INGESTION & STORAGE

Processing and storing sensor samples is non-trivial at the scale of a production deployment. We expect our system will need to ingest and analyze at least 100,000 sensor samples per second (based on estimation in subsection II-A, assuming each sensor will generate data at 1Hz).

This project utilizes OpenTSDB [4], an open-source, scalable, time series database which leverages HBase [5], an Apache top level project inspired by Google’s BigTable [12], to manage data in a distributed manner and provide horizontal scalability. We chose OpenTSDB because it allows us to easily horizontally scale out our system to more storage nodes, while maintaining a stable, linear scaleup in streaming ingestion (as we shall describe in Section III-B).

A. Streaming Event Ingestion and Storage using OpenTSDB

For the purposes of this preliminary work, a distributed system of 32 nodes was deployed running HDFS and HBase. HDFS was set up with one NameNode (co-running Hbase master), one Secondary NameNode, one Hbase backup master and 29 DataNodes. HBase is configured with one HMaster, one BackupHMaster, and 29 RegionServers that communicate through the built-in Apache Zookeeper [13] coordination service. Each node is also running an instance of a TSD Daemon for time series data writing and querying.

OpenTSDB organizes time series data into metrics and allows for the assignment of multiple tags per metric. The tags provide unique identifiers for querying data and allow the data to be compartmentalized into sub series. The simulated data generated for this project is stored into a metric called “energy” with tags for “unit ID” and “sensor ID”. Associated metadata for a particular unit and sensor is stored though

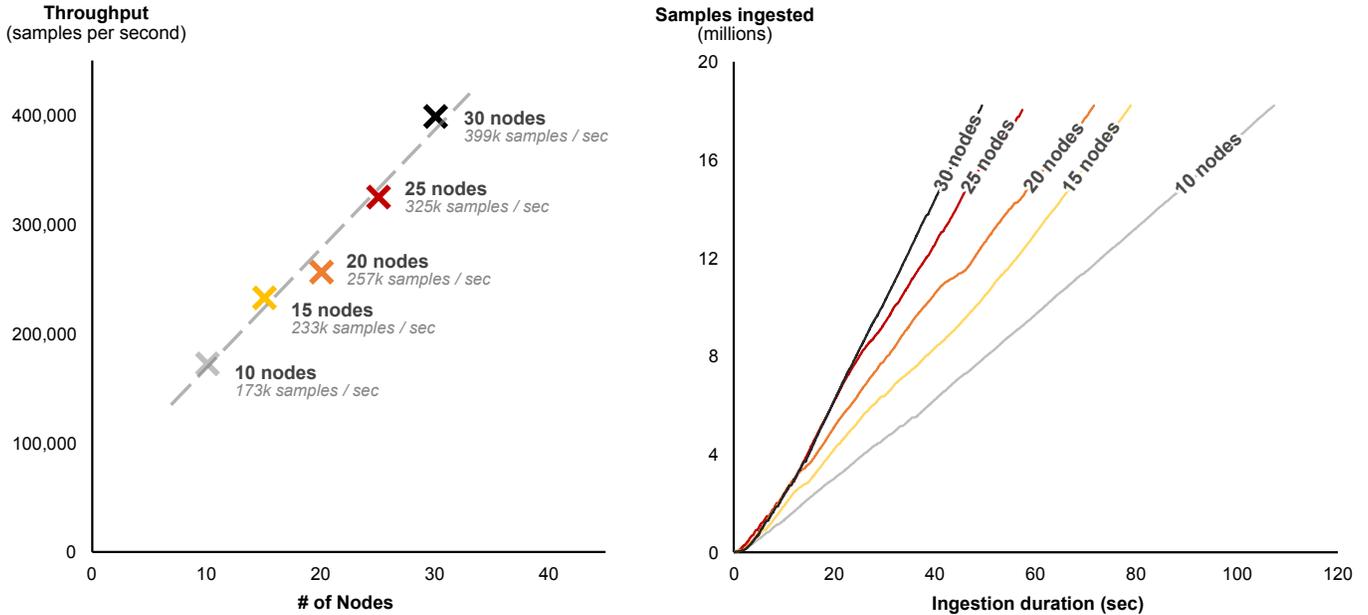


Fig. 2. **Left:** Ingestion rates achieved as the ingestion framework was ramped up from 10 machines to 30 machines. The system scales linearly, with each added machine increasing throughput by 11K samples per second on average. **Right:** The line graph of sensor samples ingested versus the ingestion duration shows a constant and stable ingestion rate for each configuration of the framework.

configuration in JSON outside the time series database, which allows configuration updates to be managed in version control.

For storing data, the TSD Daemon takes a metric, timestamp, data value, and tag identifiers as input and produces an entry to be written to an HBase table. First, the key is generated from a binary encoding of the metric, timestamp, and tag values. Then, the TSD Daemon submits an RPC call to HBase which distributes the write based upon the key value. The writes for similar keys are grouped onto the same Regionserver by HBase.

B. Preliminary Results & Key Findings

Linear Scale-up. We tested our event ingestion framework with the evaluation dataset (described in Section II-A, consisting of 100 assets with 1000 sensors each). Figure 2a shows the ingestion rate scales up linearly with the number of machines. Each machine runs one HBase RegionServer instance and one OpenTSDB daemon. The ingestion rate reaches 300,000 samples per second with 30 machines, while maintaining a stable ingestion speed, shown in Figure 2b.

OpenTSDB Key Design. One obstacle encountered early in the ingestion process was an issue with these writes not being distributed across all the HBase Regionserver efficiently. Since sequential data values share the same metric and similar timestamp values, their binary encoded keys are also similar resulting in the RPC calls being sent to the same HBase Regionserver. To combat this, the binary key encodings were salted with an additional uniformly randomly generated byte at the beginning to create unique keys for chronological data values. Additionally, HBase regions were manually split to ensure each region handled an equal proportion of the writes. Salting the keys allowed for the full utilization of all

the deployed HBase Regionserver and provided a dramatic increase to the ingestion rate.

Buffering Requests for Backpressure and Scalability. Another obstacle encountered while working with OpenTSDB and HBase was frequent crashes of Regionserver due to overloaded RPC Queues. Initially, the cause of these crashes was attributed to HBase having no means of providing back pressure to RPC calls made by OpenTSDB. To remedy this, we built a reverse proxy to buffer requests to OpenTSDB in order to limit the number of concurrent requests. Compaction was also disabled on OpenTSDB to reduce RPC calls to HBase. This proxy also serves to increase ingestion throughput by load-balancing traffic to multiple ingestion processes. Ingestion throughput scales horizontally by distributing the requests to the OpenTSDB nodes via a round-robin fashion.

IV. FLAGGING ANOMALIES WITH LOW FALSE ALARM RATES

A key component in our system is the algorithm to flag potential anomalies. Given the large expense caused by erroneously flagged faults, we aim to use an anomaly detection algorithm that balances identifying the majority of true faults while also controlling the rate of false alarms.

From a statistical standpoint, anomaly detection amounts to performing a hypothesis test on sample observations to detect possible shifts in the mean of the sampling distribution. Rejection of the null hypothesis implies that there is significant evidence to conclude that the distribution has indeed changed. A common mistake committed in hypothesis testing is to reject the null hypothesis when it is actually true (type I error).

In our setting, a type I error amounts to a false alarm, i.e., an equipment is classified as faulty when in reality there was

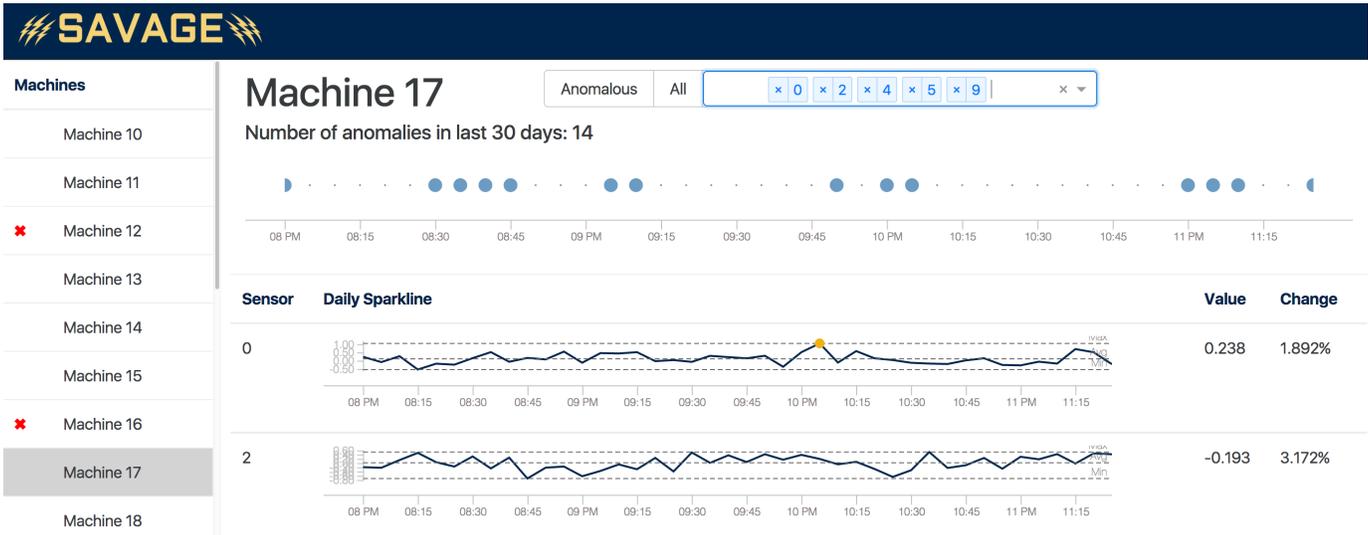


Fig. 3. A web application shows sensor readings across all machine and sensors. For machine 17, the timeline at the top visualizes correlated anomalous events across all sensors. A detailed view of each sensor’s time series with associated anomalies is displayed at the bottom.

nothing wrong. One of the key aspects of type I errors is that they tend to increase as the number of hypothesis tests increases. In our context, this translates to higher false alarm rates as the number of sensors increases. For example, for a single sensor with an allowable $\alpha = 0.05$, the probability of making at least one false alarm is 5%. However, if we increase the number of sensors to 10 sensors each with $\alpha = 0.05$, that probability jumps to 40%, i.e., $1 - (1 - \alpha)^{10} = 0.4$.

Traditionally, false alarms in multi-inference studies were controlled using family-wise error rate (FWER). FWER focuses on controlling the probability of committing any type I error when performing multiple hypothesis tests by applying a correction to a family of inferences. A popular example is the Bonferroni correction [14] where for m hypotheses each having a probability α of committing type I error, then the corrected probability for the family would be α/m . In other words, reject all hypotheses with p -values $\leq \alpha/m$. One drawback of this approach was that it provided much less detection power and was overly conservative.

FDR was first introduced by Benjamini and Hochberg in 1995 [7], [8]. The goal was to reduce false alarms in multiple inferences (hypothesis tests) in clinical trials. Compared to FWER, FDR was designed to control the expected proportion of type I errors.

The underlying premise of FDR is that when multiple tested hypotheses are rejected, it is more important to control the proportion of errors (wrong rejections) than it is a single erroneous rejection. As we wish to balance the rate of type I and type II errors, FDR is a promising option for a choice of algorithm for preliminary testing.

Production latency and throughput requirements necessitates that the chosen anomaly detection algorithm be scalable. The FDR algorithm is easy to scale using distributing matrix computation provided by Spark’s MLlib. This property makes FDR more attractive than more complex algorithms.

A. Preliminary Results

Our implementation of the FDR algorithm is composed of two parts — an offline training component and an online evaluation component. Offline training occurs in Spark, running in batch mode. Spark’s MLlib [15] provides an implementation of distributed matrix factorization, which allows our offline training system to scale to large numbers of sensors. Once training is complete, we can evaluate for anomalies at a rate of 939,000 sensor samples per second on average.

In offline training, model estimation of each sensor on each unit begins by calculating the covariance matrix of each data set. Singular Value Decomposition is then performed on each covariance matrix to obtain the mean and variance. Results from the decomposition are cached to HDFS. Evaluation is thereby relatively fast requiring a single matrix multiplication per iteration. Results from online evaluation are reported back to OpenTSDB for use by the integrated visualization tool. The current system can deal with one machine at a time and we plan to utilize concurrency of Spark to scale up workload.

V. ANOMALY VISUALIZATION

Advanced anomaly detection without commensurate visualization presents limited value for operators. Our platform includes an interactive visualization tool that equips users to quickly respond to flagged anomalies. Our tool integrates A) live sensor data, B) highlighted anomalies, and C) real-time system analytics into a single control center for users to monitor and react to events in a network of power-generating assets.

A. Preliminary Results

Our tool provides an overview of the overall health of the network of power-generating assets. It helps users explore and understand the context surrounding the flagged anomalies. By using the FDR anomaly detection algorithm, we avoid unnecessarily notifying users of false alarms.

Analytics summarize global system status across a large deployment of power-generating assets. By selectively surfacing the most concerning anomalies, we allow users to focus only on what is important. Anomalous machines are clearly marked in a single list on the left of Figure 3, allowing operators to quickly check system health.

Power-generating asset faults often result in correlated anomalies across multiple sensors (e.g., pressure and temperature). Our tool displays a summary of detected correlated anomalies in a single timeline at the top of Figure 3.

Drill-down capabilities enable users to quickly examine details about a fault with necessary context. Clicking on a particular event in the machine-wide status timeline at the top of the application displays detailed time series for relevant sensors at the bottom of the page. This allows operators to quickly view local context surrounding an anomaly with visualizations for all relevant sensors. Operators can filter the list of displayed sensor time-series to only anomalous sensors, all sensors or some custom subset of sensors.

The visualization tool is a web application that is available on both desktop and mobile devices. Mobile access allows technicians to explore pertinent sensor data while performing maintenance on a particular machine in the field.

VI. CONCLUSION & ONGOING WORK

We present our preliminary work and key findings in developing an integrated system to address three key challenges within analyzing sensor data from power-generating assets: (1) difficulty in ingesting and analyzing data from large numbers of machines; (2) prevalence of false alarms generated by anomaly detection algorithms resulting in unnecessary downtime and maintenance; and (3) lack of an integrated visualization that helps users understand and explore the flagged anomalies and relevant sensor context in the energy domain. Our system can currently ingest and analyze 399,000 sensor samples per second while running on a 30 node cluster. The system visualizes sensor data in an interactive web application which presents potential anomalies with the associated context surrounding the event.

Ongoing work for the project includes: experimenting with increasing storage nodes to further scale up throughput, migrating our anomaly detection implementation to Spark Streaming [16] for online training, and evaluating our system with domain users through our collaboration with industry partners like General Electric (GE) to test the system on their datasets.

ACKNOWLEDGMENT

This work is supported in part by the Strategic Energy Institute (SEI) at Georgia Tech, and NSF grants IIS-1563816, TWC-1526254, IIS-1217559. We also thank Yahoo! for their generous 200-machine donation. We thank Will Powell on hardware support for our system.

REFERENCES

- [1] R. Isermann, *Fault-diagnosis applications : model-based condition monitoring: actuators, drives, machinery, plants, sensors, and fault-tolerant systems*. Berlin Heidelberg New York: Springer, 2011.
- [2] L. Chiang, *Fault Detection and Diagnosis in Industrial Systems*. London: Springer London, 2001.
- [3] P. Sderholm, "A system view of the no fault found (nff) phenomenon," *Reliability Engineering & System Safety*, vol. 92, no. 1, pp. 1 – 14, 2007. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0951832005002139>
- [4] "Opentsdb : The scalable time series database, <http://opentsdb.net/>," 2016.
- [5] "Apache hbase : A distributed, scalable, big data store, <https://hbase.apache.org/>," 2016.
- [6] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM computing surveys (CSUR)*, vol. 41, no. 3, p. 15, 2009.
- [7] Y. Benjamini and Y. Hochberg, "Controlling the false discovery rate: a practical and powerful approach to multiple testing," *Journal of the royal statistical society. Series B (Methodological)*, pp. 289–300, 1995.
- [8] Y. Benjamini and D. Yekutieli, "The control of the false discovery rate in multiple testing under dependency," *Annals of statistics*, pp. 1165–1188, 2001.
- [9] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica, "Spark: cluster computing with working sets."
- [10] F. C. Liu, F. Shen, D. H. Chau, N. Bright, and M. Belgin, "Building a research data science platform from industrial machines," *3rd Workshop on Advances in Software and Hardware for Big Data to Knowledge Discovery (ASH) co-located with IEEE Big Data Conference*, 2016.
- [11] P. Ratliff, P. Garbett, and W. Fischer, "The new siemens gas turbine sgt5-8000h for more customer benefit," *VGB powertech*, vol. 87, no. 9, pp. 128–132, 2007.
- [12] F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber, "Bigtable: A distributed storage system for structured data," *ACM Transactions on Computer Systems (TOCS)*, vol. 26, no. 2, p. 4, 2008.
- [13] P. Hunt, M. Konar, F. P. Junqueira, and B. Reed, "Zookeeper: Wait-free coordination for internet-scale systems." in *USENIX Annual Technical Conference*, vol. 8, 2010, p. 9.
- [14] O. J. Dunn, "Multiple comparisons among means," *Journal of the American Statistical Association*, vol. 56, no. 293, pp. 52–64, 1961.
- [15] X. Meng, J. Bradley, E. Sparks, S. Venkataraman, D. Liu, J. Freeman, D. Tsai, M. Amde, S. Owen *et al.*, "Mllib: Machine learning in apache spark," 2016.
- [16] M. Zaharia, T. Das, H. Li, T. Hunter, S. Shenker, and I. Stoica, "Discretized streams: Fault-tolerant streaming computation at scale," in *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*. ACM, 2013, pp. 423–438.