# *Möbius: Online Anomaly Detection and Diagnosis*

Jian Zhang, Richard Wydrowski, Zeqiang Wang, Sai Sankalp Arrabolu, Keiji Kanazawa, Lech Gudalewicz, Hong Gao, Roman Batoukov, Souren Aghajanyan
Microsoft Universal Store Data and Analytics Team
proj_mobius@microsoft.com

Kenneth Tran
Microsoft Deep Learning Research Group
ktran@microsoft.com

## ABSTRACT

Project Möbius targets to develop online anomaly detection solution and diagnosis tools for time series data at scale. It is collaboration between Data and Analytics (DnA) team of Microsoft Universal Store and Deep Learning Research Group. Solution developed in this project has been integrated into Xpert system intelligence platform that monitors millions of cores in services and billions of Windows devices, to help to shorten Time-To-Detect (TTD) and Time-To-Mitigate (TTM) system outage. This paper describes design and algorithms used in the Möbius solution.

## KEYWORDS

anomaly detection, machine learning, time series, Streaming Least Squares, Holt-Winters, Mahalanobis distance

## 1 INTRODUCTION

[1] For large scale online store such as Microsoft Store, every single minute of system downtime especially those impact purchase flow could potentially lead to revenue loss. Historically it was hard for engineers to define alert threshold accurately because of diversity of workload and volume of metrics to monitor. Inaccurate alerts cause delay in Time-To-Detect (TTD) and Time-To-Mitigate (TTM) system outage and cost more engineering resources to investigate.

Project Möbius targets to provide platform support for online anomaly detection and diagnosis for time series data at scale. Our contributions are: First, A collection of algorithms, each of which targeting different types of workload. Second, algorithm recommendation and assisted parameter tuning to make alert setup and calibration simple and easy. In addition, diagnosis scripts are provided to facilitate post-detection investigation. Alert annotation enables engineers to provide feedback to alerts received and can be used as labeled data to re-train the algorithms. The solution has been integrated into Xpert system intelligence platform, which monitors millions of cores in services and billions of Windows devices. Please reference Figure 1 for its system design.

## 2 ALGORITHMS

The solution leverages a collection of detection algorithms such as proprietary adaptive anomaly detection based on Holt-Winter [2] and Streaming Least Square (SLS) [3] as shown in Figure 2 and also introduced as below.

[1]Project Möbius

*KDD'2018, August, 2018, El London, UK*
2018. ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

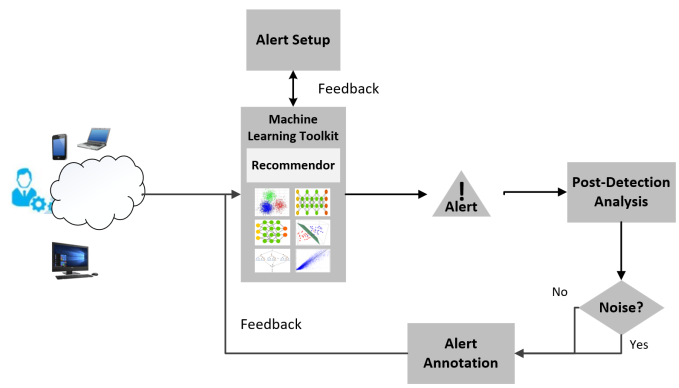


**Figure 1: Möbius Anomaly Detection System Design**

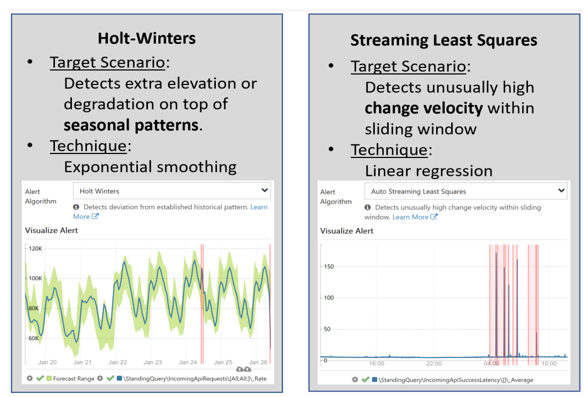


**Figure 2: Möbius Anomaly Detection Algorithms**

### 2.1 Streaming Least Squares

Streaming Least Squares (SLS) is an open source algorithm developed by MSR (Microsoft Research). It transforms from Ordinary Least Squares (OLS) with improved efficiency. For each sliding window, it runs OLS and computes its regression residual. The regression residual is reported as SLS score. The results are sorted by SLS Score and classified into levels, which are formed on the basis that anomalies in the same level have similar scores, in similar fashion of K-Means in one dimension. The classification also provides the automatic SLS threshold used to determine the level of each anomaly. Please reference [3] for detail explanation.

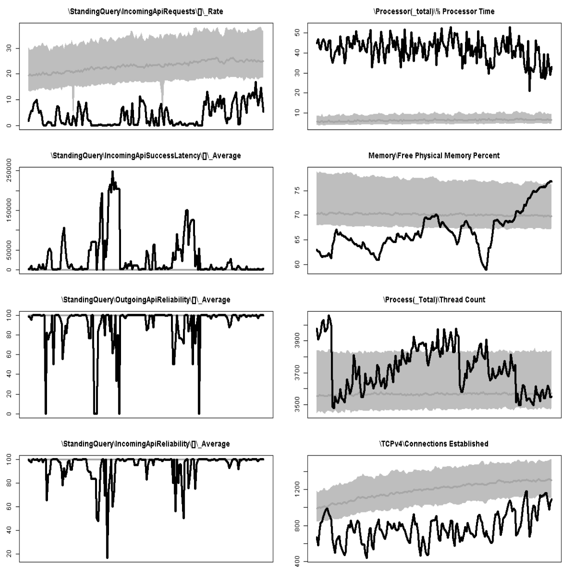


**Figure 3: Möbius Outlier Detection**

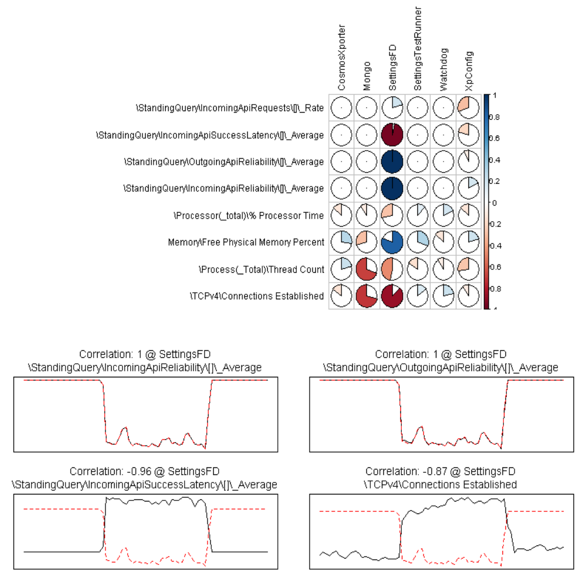


**Figure 4: Möbius Correlation Scan**

### 2.2 Adaptive Anomaly Detection Scheme Based on Holt-Winters

Holt-Winters algorithm is based on exponential smoothing and has been well established for decades as a good way to predict values in a time series with seasonal components. The Holt-Winter based adaptive anomaly detection scheme used in the solution supports adaptivity by weighing recent history more heavily than the distant past in both the prediction and estimation of error. The results for the algorithm is visualized as the time series pictured in the context of the tolerance bands. If the series crosses the bands, an anomaly is declared.

### 2.3 Algorithm Recommendation

It is hard to have an algorithm that works best for all types of workload. The algorithm recommender is boot strapped with recommendation based on knowledge of human experts as shown in Figure 1 and crowd sourced to users, who have options to experiment with different algorithms and choose algorithm that works best for the workload. Overtime, signals of algorithms selection and matching workload characteristics are collected and analyzed to optimize algorithm and algorithm recommendation.

## 3 DIAGNOSIS TOOLS

A collection of diagnosis tools are provided to support post-detection analysis such as machine outlier detection and correlation scan.

Machine outlier detection is used to identify servers, which behave significantly differently from the rest of its peers in the same machine role and same data center, using Mahalanobis distance [1] as shown in Figure 3. This helps to narrow down scope of investigation.

Correlation scan is used to identify QoS (Quality of Service) counters of upstream and downstream services, which have high correlation to the service alerts under investigation as shown in an example in Figure 4. This provides additional information to on call engineers to facilitate investigation.

## 4 DEMO PLAN

The following demo will be provided in the poster session:

- Anomaly alert setup and calibration
- Diagnosis tools

Please reference Github repository for SLS in [3] for sample demo.

## 5 SUMMARY

The anomaly detection and diagnosis solutions developed in project Möbius have been integrated into Xpert system intelligence platform, which monitors millions of cores in services and billions of Windows devices. Feedback from engineers of Universal Store states that Möbius enables them to set up alert for counters which they have hard time to define threshold for before and to catch early signals of system outage.

## REFERENCES

[1] Mahalanobis Prasanta Chandra. 1936. On the generalized distance in statistics. In *Proceedings of the National Institute of Sciences of India*, Vol. 2. 49–55.
[2] Chris Chatfield. 1978. *The holt-winters forecasting procedure.* 264–279.
[3] Kenneth Tran. [n. d.]. Streaming Least Square Explained. ([n. d.]). https://github.com/MSRDL/Deep4Cast/tree/master/AnomalyDetection The source code, demo, and introduction of Streaming Least Squares algorithm.