

# Moving Metric Detection and Alerting System at eBay

**Zezhong Zhang, Keyu Nie and Ted Tao Yuan**

eBay Inc., San Jose CA 95125, USA  
zezzhang, knie, teyuan@ebay.com

## Abstract

At eBay, there are thousands of product health metrics for different domain teams to monitor. We built a two-phase alerting system to notify users with actionable alerts based on anomaly detection and alert retrieval. In the first phase, we developed an efficient anomaly detection algorithm, called Moving Metric Detector (MMD), to identify potential alerts among metrics with distribution agnostic criteria. In the second alert retrieval phase, we built additional logic with feedbacks to select valid actionable alerts with point-wise ranking model and business rules. Compared with other trend and seasonality decomposition methods, our decomposer is faster and better to detect anomalies in unsupervised cases. Our two-phase approach dramatically improves alert precision and avoids alert spamming in eBay production.

## Introduction

Over the era of big data, large companies have the desire of building automated alerting system to continuously monitor the health of their sites and applications. Currently, many solutions are proposed to serve their needs, including Alibaba (Xu et al. 2018), Amazon (Guha et al. 2016), Anodot (Toledano et al. 2018), Baidu (Sun et al. 2018), AT&T (Yan et al. 2012), Facebook (Taylor and Letham 2018), Google (Shipmon et al. 2017), LinkedIn, Microsoft (Ren et al. 2019), Twitter (Hochenbaum, Vallis, and Kejariwal 2017), and Yahoo (Laptev, Amizadeh, and Flint 2015).

At eBay, product owners monitor various metrics every day. Each metric potentially contains thousands of time series with different dimension aspects. To effectively monitor these millions of time series and alert automatically, we developed a two-phase approach with anomaly detection and alert retrieval to build eBay automated alerting system.

In our practice, there are three challenges for building the automated alerting system:

- Lack of true alert labels: it is demanding for domain analysts to label past anomalies in each time series for machine to learn,
- Scalability and efficiency: we need to detect millions of time series every day and alert in time,

- Avoid alert spamming: flooded alerts do not help users since not all finding outliers are valid alerts to all users.

To tackle the first two challenges, a fast and robust time series anomaly detection algorithm (MMD) is developed. It firstly decomposes a metric time series to "normal" patterns (e.g. trend, seasonality) and a noise/residual part. To gauge the noise level as well as filter out potential anomalies without assuming particular distribution of the noise part, we use Chebyshev's Inequality to determine the anomaly criteria that controls the overall false positive detection rate (type I error outside "normal" range). At the end of the anomaly detection phase, we obtained a list of time series with anomalies that are potentially used to alert users.

Secondly, to increase precision of the potential alerts and reduce alert spamming, we designed an alert retrieval framework that takes input from the anomaly detection phase, combined with additional rules and policies, to generate the final ranked or prioritized alert list to notify users.

In recent years, many anomaly detectors are proposed to tackle above issues, such as Isolation Forest (Liu, Ting, and Zhou 2008), Argus (Yan et al. 2012), OC-SVM (Amer, Goldstein, and Abdennadher 2013), EGAD (Laptev, Amizadeh, and Flint 2015), Opprentice (Liu et al. 2015), VAE (An and Cho 2015), RRFCF (Guha et al. 2016), Donut (Xu et al. 2018), SR-CNN (Ren et al. 2019) and ATAD (Zhang et al. 2019). Our time-series anomaly detection focuses more on statistical model as Makridakis, Spiliotis, and Assimakopoulos (2018) pointed out that the performance of ML models is lower to that of statistical methods in time-series area. Besides, many existing anomaly detection methods may cause alert spam since they only focus on anomalies without considering the second phase to retrieve valid alerts.

In this paper, we first explain our two-phase approach, especially the efficient and robust moving metric detector; then we discuss the evaluation of alert result and production performance; and conclude with some discussion.

## Approach

Let's assume a business metric is measured daily and represented as a time series. Anomaly detection is to determine if the last observation of the metric is significantly different from what it is supposed to be given the past measurements

of the time series. In Figure 1, we illustrate the overall alert detection and alert notification pipeline. The anomaly detection phase comprises of time series signal decomposition and statistical test for anomaly. The alert retrieval phase is to filter out valid alert from recently detected anomalies found in the previous phase.



Figure 1: The Two-Phase Alerting System.

## Anomaly Detection Phase

Our basic assumption is that metric time series data should contain mostly "normal" observations than "abnormal" observations, and we do not require the abnormal data being labeled. In order to identify anomalies, the general scheme is to learn the "normal" patterns in the time series and use it as prediction to test if the last observation is within the "normal" metric value range.

**Time Series Decomposer** As key part of anomaly detection phase, we decompose each time series to trend, seasonality and residual parts. There are existing time series decomposition implementations, however, they are either sensitive to outliers because using moving average, e.g. classical decomposition (Kendall and Stuart 1983), or with high computation cost of iterations, e.g. STL (Cleveland et al. 1990). To extract "normal" patterns from time series containing sparse anomalies in the past but without labels, we developed our moving metric decomposer using median to extract robust trend and seasonality in the time series.

For a given time series  $X$ , we first do a rough estimate of trend ( $L$ ) using symmetric moving average with a sliding time window ( $w$ ), which is the number of observations per cycle in the given time-series  $X$ . To estimate the  $w$  for each metric, we applied a one-time signal frequency estimation using ESPRIT method (Roy and Kailath 1989) since the frequency of each metric rarely changes for almost all cases in eBay. Using the trend  $L$ , we can discover seasonality and reduce part of the variance in  $X$ .

Let's define trend-removed time series as  $L' = X - L$ . To extract a robust seasonality ( $S$ ), we use the median of historical periodic values of trend-removed time series  $L'$ , instead of simple averaging which is sensitive to outliers.

$$S_t = \text{median}(L'_{t \pm iw} | i \leq \tau, i \in \mathbb{N})$$

where  $w$  is the number of observations per cycle,  $i$  is the number of cycles away from observation at  $t$  and  $\tau$  is the number of cycles near the observation at  $t$  that is used to estimate seasonality  $S_t$ . In our case, we set  $\tau$  as length of the input time series  $X$  that contains all the observations.

With extracted seasonality  $S$ , we then estimate trend once again from the seasonality-removed series as  $S' = X - S$ . We ignore the first estimate of trend  $L$  and recompute

trend from  $S'$  as  $T$  in the following. Using right alignment to calculate rolling median in a sliding time window ( $w$ ),

$$T_t^f = \text{median}(S'_{t-j} | j \leq w, j \in \mathbb{N})$$

where  $j$  is the number of observations previous to observation at  $t$ . The  $f$  indicates it is right alignment.

The right alignment rolling method naturally lags behind current observations, so we estimated a potential bias term using median in the following formula to model the lagging effect. Hence, the extracted trend contains two terms,

$$T = T^f + \text{median}(S' - T^f).$$

Compared with STL (Cleveland et al. 1990), which iteratively estimates trend and seasonality with inner loop and outer loop to moderate outlier impact on trend and seasonality, our method first computes an approximated trend  $L$  and extract seasonality  $S$ , and then extract the final trend  $T$  with median as a robust replacement of average. Our method requires less computation with better and robust performance.

**Determine "Normal" Range of a Metric** With time series decomposer, given time series  $X$  is decomposed into trend ( $T$ ), seasonality ( $S$ ) and residual ( $R$ ), which is defined as  $R = X - T - S$ . We detect anomaly as the most recent observation in a time series whose value differs significantly from the "normal" prediction. To detect anomaly, we have to calculate the "normal" range from the input time series  $X$ .

In many cases, the noise residual  $R$  is assumed to be normally distributed. However, in our approach, we do not assume "normal" residuals distribution, i.e., whether it is Gaussian or not. By using Chebyshev's Inequality, for any real number  $k > 0$ ,

$$P(|R - \hat{\mu}| \geq k\hat{\sigma}) \leq \frac{1}{k^2}$$

where the  $\hat{\mu}$  is the expected value and  $\hat{\sigma}$  is the standard deviation of the residual  $R$ . Using the above inequality formula, we can define false positive criteria to differentiate anomaly values outside of "normal" region.

The distribution-free "normal" range of  $X$  can be calculated with a given expected probability  $p$  of seeing an anomaly (Amidan, Ferryman, and Cooley 2005),

$$T + S + \hat{\mu} \pm k\hat{\sigma}$$

where  $k = 1/\sqrt{p}$ .

To estimate the  $\hat{\mu}$  and  $\hat{\sigma}$  for "normal" range, Hampel (1974) pointed out that mean and standard deviation are sensitive to outliers and Leys et al. (2013) suggested using median and median absolute deviation (MAD) as robust replacements, as follow:

$$\begin{aligned} \hat{\mu} &= \text{median}(R) \\ \hat{\sigma} &= b \times \text{median}(|R_t - \text{median}(R)|) \end{aligned}$$

where the robust standard deviation  $\hat{\sigma}$  is estimated by scaling MAD with a constant  $b$  (commonly set as 1.4826).

Especially with our decomposition method, we have the  $\hat{\mu} = \text{median}(R) = 0$  due to the extra added estimated bias by  $\text{median}(S' - T^f)$ . In our model, the  $k$  of the "normal"

range can be interpreted as the number of standard deviations away from the prediction ( $T + S$ ). With Chebyshev’s Inequality, the “normal” range with  $p = 0.01$  would be interpreted as at most 1% probability of exceed the range with  $k = 10$ . In this case, we identify anomaly if the value is beyond 10 standard deviation away from prediction.

### Alert Retrieval Phase

Not all anomalies are valid actionable alerts for all users. In eBay practice, we can only send 10 alerts for each user every day, otherwise it will cause alert spamming. In the alert retrieval phase, we developed a ranking model and filtering rules and polices to retrieve valid alerts.

**Ranking of Anomalies** To retrieve valid alerts and increase precision, we created a point-wise anomaly ranking algorithm with generalized linear ranking model to sort valid alerts from identified anomalies. The ranking score has several components, i.e.,

- Anomaly deviation severity feature computed using outputs from the previous phase as  $f_d = |R - \hat{\mu}|/\hat{\sigma}$ .
- User defined 4-level priority (e.g. P1, P2, P3, P4) for each metric as  $f_p$  (one-hot encoding).
- Granularity of time series dimension aspect values as  $f_g$ , which is the number of dimensions that not rolled up in metric dimension hierarchy. For example, “Country: US” has granularity  $f_g = 1$  and “Country: US, Device: PC” has granularity  $f_g = 2$ .

With above features, our ranking model is

$$g(p) = w_d f_d + w_p f_p + w_g f_g$$

where  $w_p$ ,  $w_g$  and  $w_d$  are the weight vector that can be learned from users feedbacks. The  $g(p)$  is the ranking score function considering both anomaly severity and importance.

**Retrieval Logic and Rules** Some anomalies that we detected are true and valid alerts to users, however, they are related and pointing to the same issue. From user perspective, it is one kind of alert spam if flooding very similar alerts to user. As a business requirement to consider the alerts diversity, our retrieval mechanism combined the diversity of the alerts and ranking score with following two steps:

- For each metric, detected anomalies are ranked by the ranking score  $g(p)$ . We take the top anomaly from the list and compare the next highest ranked anomaly with the top one by computing the absolute Pearson correlation coefficient, if the coefficient is less than 0.9, which implies different time series pattern, we include it in the final alerts, if not we compare the next one on the list and so on. Overall, we select two alerts for each metric.
- Retrieve at most 10 alerts across metrics and rank them with score  $g(p)$  from selected anomalies in each metric.

To reduce spam, many business rules are proposed from our user, such as stop alerting duplicated/similar anomalies in recent  $k$  days and alerting after  $k$  days continuous exceed interval, which archived good result in eBay.

## Validation and Production Performance

To evaluate our approach, we crowdsourced alert labels from product owners and domain analysts in eBay for a selected set of metrics to collect alert feedbacks (valid alert or not). Also, on our production environment, every alert that we sent out, are triaged and labeled as a valid alert if it is a truly useful and actionable issue with further investigation by the domain analysts.

For the crowdsourced labels, we showed the time series to some colleagues that labeled them independently. Here are the details of our crowdsourced labeling setup:

- Each person had the same amount (15) of time series to look at out of a pool of 164 different time-series.
- Each person was told to label all points that shall trigger alert for each 7-month (from Jan to July) daily time series.
- While random, the allocation of time series was designed in such a way that each metric was seen by at least two people.

At the end of above alert crowdsourcing collection, we received total 321 alert feedbacks from 38 different product owners and domain analysts in eBay.

### Crowdsourcing Labels Analysis

It turned out that people, when looking at data independently have very differing views about what constituted a valid alert, despite the fact that they were all given the exact same, clear definition of an alert. In Figure 2, in approximately over half of the crowdsourced alert labels there was no common agreement (agreement ratio  $\leq 0.50$ ) as to whether a given data point was an alert or not. Here, the agreement ratio of one alert is defined as the percentage of people who labeled the observation as an alert over the people who saw the time-series.




Figure 2: Histogram of Agreement Ratio.

To understand the lack of agreement situation, we talk with domain analysts to build business rules to reduce spamming. For example, some domain’s analysts do not expect alerts on just the first day exceed interval since there service exist auto-recovery mechanism, however, they are interested in the 3 days continuous exceed interval, so we employed the “alerting after  $k$  days continuous exceed interval” rule.

## Validation of Anomaly Detection

To keep only good quality data, we introduced majority vote rule (agreement ratio  $> 0.5$ ) over the crowdsourced labels, therefore excluding a fair share of data points that were ties. Then, we compared our time series decomposition method with classical method and STL on anomaly detection task using our distribution-free “normal” range.

In order to weigh in recall more than precision during anomaly detection phase, we choose F2-Score as measurement since it focuses on recalling hundreds of potential alert candidates out of millions. We tuned model parameters with grid search that maximize the F2-Score for each anomaly detection method using different time series decomposer with data during Jan-May 2018, and ran the tuned model over an unseen testing set (Jun-Jul 2018), which we had held aside.

Method	Recall	Precision	F2-Score	Processing Time
Classical	64.86%	28.92%	0.519	252.1ms
STL	59.46%	37.93%	0.534	156.8ms
Ours	72.97%	38.03%	0.616	90.3ms

Table 1: Performance comparison on crowdsourcing dataset.

Besides, we compared speed by using average processing time for every 100 time-series (no parallel computing enabled). The implementation of classical method and STL are the *decompose* and *stl* functions in R Package *stats*. In Table 1, our decomposer shows better F2-Score performance over eBay Metrics with much faster processing speed.

## Production Performance

Currently, our alert system monitors several million time-series and process hundreds of gigabytes of data every day. we implemented the two-phase solution on our private cloud as several micro-services, such as the data-driven anomaly detector, alerting rules tagger and ranking scorer. In figure 3, Function as a Service (FaaS) framework is adopted to scale up these micro-services as  $\lambda$  functions in the anomaly detection phase. Then, the alert retrieval phase leveraged the tags and scores generated in the previous phase to filter and rank anomalies according to domain analysts’ settings.




Figure 3: Architecture with Function as a Service Cluster.

With the 15 servers (each has 8 VCPUs and 16GB RAM) FaaS cluster, Our anomaly detector can handle at least 6.4 thousand time-series per minutes, and we can easily scale up the system by adding additional hardware if needed.

To evaluate our alert result, we built an evaluation system to collect alert feedbacks that we sent out. Here, we com-

Approach	Weeks	#Alert	#Valid Alert	Precision
First Phase Only	10	164	67	40.85%
Two Phases	9	118	106	89.83%

Table 2: Production performance on valid alert precision.

pared the anomaly detection phase only with the two-phase approach. As in Table 2, the extra alert retrieval phase dramatically improved the precision of valid alert.

## Conclusion

To automate product health moving metric alert process at eBay, we developed a two-phase approach to identify anomalies and retrieve valid alerts for different domain users. At anomaly detection phase, we developed the Moving Metric Detector (MMD), which contains a fast and robust time-series decomposition algorithm with better F2-Score performance compared with classical method and STL on anomaly detection task. We leverage Chebyshev’s Inequality to determine a distribution-free “normal” range. To avoid alert spamming and improve alert diversity, we designed the alert retrieval phase with a point-wise ranking model and business rules. Our alerting system is adopted by users across eBay and proved to be an effective way to early alarm issues for business needs.

## Acknowledgments

We thank Rong Song, Christine Wu, Nathan Ni, Zhixuan Jia for their insights and expertise that greatly assisted the research, and continued support by eBay infrastructure organization. We would also like to show our gratitude to Giorgio Ballardin, Woody Zhou, Liren Sun, Jianpeng Xu and Jiahui Ruan for their stimulating discussions.

## References

- Amer, M.; Goldstein, M.; and Abdennadher, S. 2013. Enhancing one-class support vector machines for unsupervised anomaly detection. In *Proceedings of the ACM SIGKDD Workshop on Outlier Detection and Description*, 8–15. ACM.
- Amidan, B. G.; Ferryman, T. A.; and Cooley, S. K. 2005. Data outlier detection using the chebyshev theorem. In *2005 IEEE Aerospace Conference*, 3814–3819. IEEE.
- An, J., and Cho, S. 2015. Variational autoencoder based anomaly detection using reconstruction probability. *Special Lecture on IE 2(1)*.
- Cleveland, R. B.; Cleveland, W. S.; McRae, J. E.; and Terpenning, I. 1990. Stl: a seasonal-trend decomposition procedure based on loess. *Journal of official statistics* 6(1):3–73.
- Guha, S.; Mishra, N.; Roy, G.; and Schrijvers, O. 2016. Robust random cut forest based anomaly detection on streams. In *International conference on machine learning*, 2712–2721.
- Hampel, F. R. 1974. The influence curve and its role in robust estimation. *Journal of the american statistical association* 69(346):383–393.

- Hochenbaum, J.; Vallis, O. S.; and Kejariwal, A. 2017. Automatic anomaly detection in the cloud via statistical learning. *arXiv preprint arXiv:1704.07706*.
- Kendall, M. G., and Stuart, A. 1983. The advanced theory of statistics volume iii. *Charles Griffin and Company Limited, London and High Wycombe* 242.
- Laptev, N.; Amizadeh, S.; and Flint, I. 2015. Generic and scalable framework for automated time-series anomaly detection. In *Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.
- Leys, C.; Ley, C.; Klein, O.; Bernard, P.; and Licata, L. 2013. Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology* 49(4):764–766.
- Liu, D.; Zhao, Y.; Xu, H.; Sun, Y.; Pei, D.; Luo, J.; Jing, X.; and Feng, M. 2015. Opprentice: Towards practical and automatic anomaly detection through machine learning. In *Proceedings of the 2015 Internet Measurement Conference*, 211–224. ACM.
- Liu, F. T.; Ting, K. M.; and Zhou, Z.-H. 2008. Isolation forest. In *2008 Eighth IEEE International Conference on Data Mining*, 413–422. IEEE.
- Makridakis, S.; Spiliotis, E.; and Assimakopoulos, V. 2018. Statistical and machine learning forecasting methods: Concerns and ways forward. *PloS one* 13(3):e0194889.
- Ren, H.; Xu, B.; Wang, Y.; Yi, C.; Huang, C.; Kou, X.; Xing, T.; Yang, M.; Tong, J.; and Zhang, Q. 2019. Time-series anomaly detection service at microsoft. *arXiv preprint arXiv:1906.03821*.
- Roy, R., and Kailath, T. 1989. Esprit-estimation of signal parameters via rotational invariance techniques. *IEEE Transactions on acoustics, speech, and signal processing* 37(7):984–995.
- Shipmon, D.; Gurevitch, J.; Piselli, P. M.; and Edwards, S. 2017. Time series anomaly detection: Detection of anomalous drops with limited features and sparse examples in noisy periodic data. Technical report, Google Inc.
- Sun, Y.; Zhao, Y.; Su, Y.; Liu, D.; Nie, X.; Meng, Y.; Cheng, S.; Pei, D.; Zhang, S.; Qu, X.; et al. 2018. Hotspot: Anomaly localization for additive kpis with multi-dimensional attributes. *IEEE Access* 6:10909–10923.
- Taylor, S. J., and Letham, B. 2018. Forecasting at scale. *The American Statistician* 72(1):37–45.
- Toledano, M.; Cohen, I.; Ben-Simhon, Y.; and Tadeski, I. 2018. Real-time anomaly detection system for time series at scale. In *KDD 2017 Workshop on Anomaly Detection in Finance*, 56–65.
- Xu, H.; Chen, W.; Zhao, N.; Li, Z.; Bu, J.; Li, Z.; Liu, Y.; Zhao, Y.; Pei, D.; Feng, Y.; Chen, J.; Wang, Z.; and Qiao, H. 2018. Unsupervised anomaly detection via variational auto-encoder for seasonal kpis in web applications. In *Proceedings of the 2018 World Wide Web Conference, WWW '18*, 187–196. Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee.
- Yan, H.; Flavel, A.; Ge, Z.; Gerber, A.; Massey, D.; Papadopoulos, C.; Shah, H.; and Yates, J. 2012. Argus: End-to-end service anomaly detection and localization from an isp's point of view. In *2012 Proceedings IEEE INFOCOM*, 2756–2760. IEEE.
- Zhang, X.; Lin, Q.; Xu, Y.; Qin, S.; Zhang, H.; Qiao, B.; Dang, Y.; Yang, X.; Cheng, Q.; Chintalapati, M.; Wu, Y.; Hsieh, K.; Sui, K.; Meng, X.; Xu, Y.; Zhang, W.; Shen, F.; and Zhang, D. 2019. Cross-dataset time series anomaly detection for cloud systems. In *2019 USENIX Annual Technical Conference (USENIX ATC 19)*, 1063–1076. Renton, WA: USENIX Association.