
Improve black-box sequential anomaly detector relevancy with limited user feedback

Luyang Kong¹ Lifan Chen¹ Ming Chen¹ Parminder Bhatia¹ Laurent Callot¹

Abstract

Anomaly detectors are often designed to catch statistical anomalies. End-users typically do not have interest in all of the detected outliers, but only those relevant to their application. Given an existing black-box sequential anomaly detector, this paper proposes a method to improve its user relevancy using a small number of human feedback. As our first contribution, the method is agnostic to the detector: it only assumes access to its anomaly scores, without requirement on any additional information inside it. Inspired by a fact that anomalies are of different types, our approach identifies these types and utilizes user feedback to assign relevancy to types. This relevancy score, as our second contribution, is used to adjust the subsequent anomaly selection process. Empirical results on synthetic and real-world datasets show that our approach yields significant improvements on precision and recall over a range of anomaly detectors.

1. INTRODUCTION

Sequential anomaly detection is widely used in various fields, including cyber-security (Teoh et al., 2018), scientific discovery (Sabeti & Høst-Madsen, 2019), operational metrics monitoring (Shi et al., 2018), etc. The goal is to detect unlikely and rare events that are both statistically abnormal and are relevant to user’s specific application.

Most anomaly detectors provides a ranked list of statistical anomalies. End-users are typically not interested in all statistical anomalies but only in those that are relevant to their specific application (Das et al., 2016). As a result, extra manual efforts are needed to locate the subset of interest. Sometimes users need to expand the subset by finding

more points with similar characteristics, which can be tedious. One simple example is monitoring the time series of revenue: although sharp increase anomalies and decrease anomalies are all statistical anomalies, a risk management professional may only care about unexpected decreases in revenue. This phenomenon, which is referred to as the “relevancy challenge” in this paper, lies in the center of a successful anomaly detection system for real world problems.

We tackle the relevancy challenge by collecting a small number of feedback from end users. Users can choose to provide feedback on any previously reported anomaly. In our setting, feedback could either be “positive”, if an anomaly is relevant to their application, or “negative” otherwise. The feedback thus collected is used to adjust future anomaly selection process in order to improve relevancy.

One important assumption in this work is that the *anomaly detector is a black-box*. Only anomaly scores are observed. This is motivated by a fact that, in many scenarios, improving the statistical anomaly detector is orthogonal to improving the relevancy of reported anomalies, since the user’s relevance function is domain specific while the definition of generic statistical anomaly is not. Decoupling these two development processes is critical for generalization and customization of an end-to-end system. To the best of our knowledge, this paper is the first to assume a given black-box statistical anomaly detector, and tries to improve its relevancy using limited user feedback.

Motivated by real world applications, we assume user feedback are generated in *sequential batches*. Within each batch, anomalies are detected for the users to provide feedback. Note that such *sequential batch* setting is different from *pure online* anomaly detection setting where detection happens in real-time for each data point. In comparison, the *sequential batch* setting generates detection for one batch at a time. We argue that sequential batch setting is a more general setting for feedback based anomaly detection system, especially for users who monitor a system and provide feedback on a certain frequency e.g. daily/hourly basis.

Our method is inspired by a simple observation: statistical anomalies are of different types (Vercruyssen et al., 2018), only a subset is relevant to user’s application. It starts with

¹Amazon Research. Correspondence to: Laurent Callot <lcallot@amazon.com>.

first identifying the types using unsupervised clustering. **Second**, each type (or cluster) is assigned a relevancy score, according to the user feedback previously reported within that cluster. **third**, we prioritize statistical anomalies from high relevancy cluster and de-prioritize the low relevancy ones during future anomaly selection.

The proposed framework is evaluated on a synthetic data set as well as a real world anomaly detection challenge. We document that our approach consistently improves accuracy of four different base detectors.

The remainder of this article is organized as follows. Section 2 reviews related work. Our main approach is discussed in Section 3. Section 4 describes the data used in our work, followed by the experiment results and case reviews in Section 5. Section 6 concludes.

2. RELATED WORK

While the paper belongs to the broad fields of active learning or semi-supervised learning, we narrow our review of related works to the sub-field of active anomaly detection (AAD).

A common setup in AAD is to tightly couple a specific anomaly detector method with an online optimization algorithm. Siddiqui et al. (2018) is tied with isolation forest Liu et al. (2008), whose anomaly score is a weighted leaf node value. It uses online optimization to update the weight to reflect the feedback. (Das et al., 2017) is specifically tied with tree-based methods. Das et al. (2016) is tied with the Loda algorithm (Pevný, 2016). (Siddiqui et al., 2019) is also tied with isolation forest.

There are also works on active learning for *online* anomaly detection. Varadarajan et al. (2017) provides a framework for anomaly detection in video analytics field. Raginsky et al. (2009) and Horn & Willett (2011) incorporates expert feedback for dynamic threshold control. They employ a two step approach: a filtering step, which assigns log-likelihoods to each observation, followed by a hedging step, which sets a dynamic threshold of the anomaly. While dynamic threshold control is another important topic for anomaly detection, it is relatively orthogonal to learning user preference over different anomaly types, which is the main focus of this paper.

Vercruyssen et al. (2018) formulates this problem with a semi-supervised clustering approach. It uses COP K-means with cannot-link constraint to incorporate user feedback. The final anomaly score is a function of clustering results including cluster size and each point’s distance with labeled anomalies. Thus, it is also tied with its specific anomaly detection algorithm, lacking the capacity of handling generic black-box detector.

3. METHODOLOGY

Assume we have a time-series $V^t = (v_1, v_2, \dots, v_t)$, a black-box base anomaly detector generates scores $P^t = (p_1, p_2, \dots, p_t)$ where $p_i \in [0, 1)$, with its detection threshold $\tau_a \in (0, 1)$ such that observations with $p_i > \tau_a$ are classified as statistical anomalies.

Detected anomalies are shown to the end-user who then provides positive or negative feedback on a subset of them. We denote V_+^t to the set of positively reviewed points, and V_-^t the set of negatively reviewed ones.

3.1. CLUSTERING

The first step of our framework is clustering time points with high p_i .

Candidate Selection: we define the clustering candidate population by giving certain threshold $\tau^c \in [0, 1]$:

$$S^c := \{p_i | \forall i \in [1, 2, \dots, t] : p_i > \tau^c\}. \quad (1)$$

Recall p_i denotes the anomaly scores. If $\tau^c = 0$, all the points are included in the clustering process.

Context Collection. Consider the original time point v_t at time t , we define the context vector with length m at time t as

$$h_t := (v_{t-m}, \dots, v_t). \quad (2)$$

We can collect the context set S^c of anomaly candidate set S^c :

$$S^c := \{h_t | \forall z_t \in S^c\}. \quad (3)$$

Sub-sequence Embedding. We create a low-dimensional embedding of sub-sequence contexts in S^c , such that we could obtain the contextual feature map of each time point:

$$S^E := \{\vec{e}_t = f_E(h_t) | h_t \in S^c\}. \quad (4)$$

Here f_E is a sequence embedding function. In this work, we use a sequence auto-encoder with bi-directional LSTM as encoder and decoder, illustrated in Fig 1. We argue that auto-encoder is very suitable for this task because it can extract vector representation of a sequence efficiently in a totally unsupervised fashion.

Clustering. With the feature map S^E generated above, we can apply clustering algorithms on top of it:

$$S^C := \{c_i = g(\vec{e}_i), c_i \in [1, \dots, k] | \vec{e}_i \in S^E\} \quad (5)$$

g is our clustering function. k is number of clusters. Specifically, we choose K-means in our framework, thus k is a hyper-parameter. We provide rationale on this choice in latter chapters.

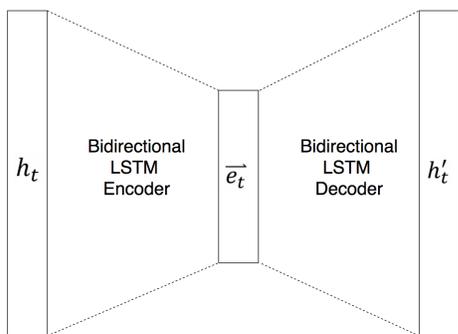


Figure 1. Our sequence embedding module: auto-encoder with Bi-directional LSTM as encoder and decoder

3.2. CLUSTER RELEVANCY

The second main step of our approach is to assign each cluster a relevancy score based on the concentration of feedback and representing the degree to which each cluster is relevant to the users’ preferences.

Cluster distribution. Consider the distribution \vec{d}_c over clusters on anomaly candidate set S^c , the distribution \vec{d}^+ over clusters on intersection of S^c and positive labels and the distribution \vec{d}^- on intersection of S^c and negative labels.

Cluster relevancy. Cluster relevancy vector is a function of the three distribution vectors above. It is formulated as:

$$\vec{r} = \frac{\exp((\vec{d}^+ - \vec{d}^-)/\vec{d}^c)}{\exp((\vec{d}^- - \vec{d}^+/\vec{d}^c)} = \exp\left(\frac{\vec{d}^+ - \vec{d}^-}{\vec{d}^c}\right)$$

Intuitively, the numerator represents positive feedback’s *distribution deviation* from the overall population, while the denominator is that of negative feedback. The exponential transformation is used for smoothing purposes, since \vec{r} will be used for scaling in the following steps. In practice, we usually set an upper and lower bound (U, L) on \vec{r} for regularization.

3.3. ANOMALY SELECTION ADJUSTMENT

In the last step, the relevancy score \vec{r} is used to adjust the anomaly selection process in next incoming batch.

Suppose next batch’s data comes in as

$$V^{t:T} = (v_t, v_{t+1}, \dots, v_T),$$

with base detector scores

$$P^{t:T} = (p_t, p_{t+1}, \dots, p_T).$$

Following a similar process as section 3.1, we conduct inference step on new data batch:

- Candidate selection $S_{t:T}^c$
- Context collection $S_{t:T}^e$.
- Apply learned embedding function f_E^l to get embedding $S_{t:T}^E$.
- Apply learned clustering function g to get cluster assignment $S_{t:T}^C$.

The base anomaly detector will classify all the observations $S_{t:T}^b$ whose scores are larger than arbitrary threshold τ^a as anomalies:

$$S_{t:T}^b := \{v_i | \forall p_i > \tau^a, i = t, t+1, \dots, T\}.$$

We count their occurrence across clusters $S_{t:T}^C$:

$$\vec{N}^b = [n_1, n_2, \dots, n_k]^T.$$

The *adjusted cluster-wise anomaly size* is then formulated as:

$$\vec{N}^{adj} = \vec{N}^b \odot \vec{r}, \quad (6)$$

where \odot is point-wise multiplication.

With each element in \vec{N}^{adj} rounded down to integer, it represents our *estimated number of relevant anomalies* in each cluster.

We then select the corresponding number of anomalies per cluster according to \vec{N}^{adj} . Inside each individual cluster, ranking is still produced by the base anomaly detector scores.

One key advantage of this adjustment approach is: \vec{r} only impacts the number of anomalies per cluster, while the original anomaly ranking inside each cluster is kept. This property, to its maximum extent, preserves the knowledge from the original base detector.

The procedure described is repeated for every new coming batch, as is illustrated in the pseudo-code of Human-In-The-Loop(HITL) Algorithm 1. We show in the following sections that, even with a limited feedback budget, our procedure can help improve black-box anomaly detectors’ performance consistently.

4. DATASETS

We test our framework on two datasets: a simulated dataset, and a real world dataset.

KPI dataset is a labeled anomaly detection dataset released by the AIOPS competition (kpi, 2017). It contains a set of KPI metrics (as time-series) with human annotated anomaly

Algorithm 1 HITL for black-box anomaly detector

Input: Sequential data batch P^i , base anomaly detector
for $i = 1, 2, 3, \dots$ **do**
 Candidate selection S_i^c
 Context collection S_i^e
 if $i \geq 2$ **then**
 Apply clustering and sub-sequence embedding
 Anomaly adjustment based on \vec{r}_{i-1}
 end if
 Collect feedback from end-users
 Train sub-sequence embedding function f_E^i
 Train clustering function g^i
 Calculate cluster relevancy \vec{r}_i
end for

labels from companies including Tencent, eBay, Baidu, Alibaba. Example KPI metrics include the order placement rate on certain website or the CPU usage for hosts, etc. There are 29 time-series in the dataset and it contains around 3M data points with 80k labeled true anomalies. We split all the time-series into two parts, the first half is used as training (when applicable) and the second half is used as testing.

Synthetic dataset contains 100 time series, each generated by 200,000 iid gaussian random variables simulating seconds level time series. Around 0.3% time points are injected anomalies. Half of the anomalies is random sharp increase added to original value while the other half is sharp decrease. However, only the decrease half is labeled as anomaly which simulates user preference. Ideally, a successful anomaly detection system should be able to identify this preference with a small number of feedback.

5. EXPERIMENTS

5.1. BASE ANOMALY DETECTOR

IID model is the simplest model which assumes the data point is generated from an *i.i.d.* Gaussian distribution, parameterized by the sample mean and sample variance of the data. The two-sided tail probability of observations on this Gaussian distribution are used to calculate the anomaly score.

Holt-winters (Chatfield, 1978) is an exponential smoothing-based forecasting models where the prediction distribution is used to calculate the tail probability of observations which constitute the anomaly scores.

Recurrent neural network, or RNN represents the class of models which employs deep learning to generate forecasts. It uses a similar scheme as classical models to calculate the anomaly score based on the tail probability, yet the

forecasts are calculated by neural networks. Specifically, in our experiments we choose the DeepAR (Salinas et al., 2019) architecture.

Random Cut Forest (Guha et al., 2016), or RCF represents the class of ensemble methods (similar to isolation forest), which estimates the density of data points directly by a forest of random cut trees. Each tree randomly cuts high dimensional data points into sub-spaces. The number of cuts is required to isolate a point is proportional to its estimated density. The depth of a data point in a tree is a measure of the rarity of this point, which yields an anomaly score.

The above four anomaly detectors are typically used in the field of statistical anomaly detection, yet each of them uses a different methodology. We expect they are representative and are able to serve the purpose of ‘black box’ detector to some extent.

5.2. HYPER-PARAMETERS

Silhouette methods (Rousseeuw, 1987) were used to determine the number of clusters, with the upper limit of five to avoid overfitting in sparse feedback setting. We choose the auto-encoder’s LSTM hidden size to be 20. Experiments show that the method is not very sensitive to this parameter.

5.3. PERFORMANCE

Our human-in-the-loop(HITL) framework is evaluated with four base anomaly detectors mentioned above. We compare the “base detector only” version, versus “HITL augmented” version. For HITL, only 20 feedback(10 positive, 10 negative) are randomly sampled from the ground truth labels per batch. In KPI dataset, each batch contains 1500(mins) of data which roughly equals to one day’s length. After a total of five batches, we stop collecting more feedback and calculate precision, recall, F1 score for both datasets. Results are shown in Table 1 and Table 2. Consistent improvements are observed across all four methods on both datasets.

5.4. CASE REVIEWS

We conducted a qualitative analysis, focused on the learned cluster relevancy \vec{r} . Two major typologies emerge.

In the **First Type**, clear disparity exists between \vec{d}_c and \vec{d}^+ , \vec{d}^- . This indicates strong user preference on some clusters over the others, which makes anomaly selection adjustment necessary. In one example with RCF as base detector:

$$\vec{d}^c = \begin{bmatrix} 0.1 \\ 0.4 \\ 0.1 \\ 0.2 \\ 0.2 \end{bmatrix} \quad \vec{d}^+ = \begin{bmatrix} 0 \\ 1.0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad \vec{d}^- = \begin{bmatrix} 0.1 \\ 0.2 \\ 0 \\ 0.3 \\ 0.4 \end{bmatrix} \Rightarrow \vec{r} = \begin{bmatrix} 0.4 \\ 2.0 \\ 0.2 \\ 0.2 \\ 0.1 \end{bmatrix}$$

Table 1. KPI dataset results comparison between base anomaly detector only versus its combination with human-in-the-loop(HITL) framework

Base Detector	Base Only			Base w/ HITL		
	F1	Precision	Recall	F1	Precision	Recall
IID	0.26	0.39	0.39	0.31	0.48	0.39
HW	0.26	0.32	0.35	0.32	0.40	0.38
RCF	0.39	0.36	0.60	0.43	0.41	0.62
RNN	0.27	0.38	0.37	0.34	0.45	0.39

Table 2. Simulated dataset results comparison between base anomaly detector only versus its combination with human-in-the-loop(HITL) framework

Base Detector	Base Only			Base w/ HITL		
	F1	Precision	Recall	F1	Precision	Recall
IID	0.38	0.39	0.63	0.55	0.44	0.74
HW	0.52	0.42	0.70	0.59	0.47	0.79
RCF	0.11	0.09	0.14	0.16	0.13	0.21
RNN	0.58	0.47	0.76	0.65	0.56	0.78

This means the model manages to learn from the feedback that the second cluster, where all positive feedback is concentrated on, is probably much more important to the user than the others. For this example, F1 improves from 27%(base only) to 43%(base w/ HITL) after only 20 feedback. We take a look at cluster 1, and observe that it mainly consists of sharp decreases in low-value region.

Second Type is where \vec{d}^+ , \vec{d}^- , \vec{d}_c are similar in value. In this scenario, our model makes relatively small impact on the existing selected anomalies.

6. Conclusion

In this paper, we propose a general method to improve the relevancy of the anomalies generated by any black-box statistical anomaly detector based on sparse user feedback. The method employs the clustering of anomalies in an embedding space and computes a relevancy score based on the concentrations of negative and positive feedback within each cluster. With limited number of feedback required, empirical results show the method improves the precision and recall over a range of different anomaly detectors significantly on both synthetic and real-world datasets.

Acknowledgements

We thank all the team members in AWS AI Lab for meaningful discussions. We thank Tim Januschowski, Bing Xi-ang and Taha Kass-Hout for organizing the collaboration. Moreover, we thank Kristjan Arumae for providing valuable feedbacks.

References

- AI-Ops KPI Challenge Dataset. http://iops.ai/competition_detail/?competition_id=5&flag=1, 2017. Accessed: 2020-05-27.
- Chatfield, C. The holt-winters forecasting procedure. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 27(3):264–279, 1978.
- Das, S., Wong, W., Dietterich, T., Fern, A., and Emmott, A. Incorporating expert feedback into active anomaly discovery. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*, pp. 853–858, Dec 2016. doi: 10.1109/ICDM.2016.0102.
- Das, S., Wong, W., Fern, A., Dietterich, T. G., and Siddiqui, M. A. Incorporating feedback into tree-based anomaly detection. *CoRR*, abs/1708.09441, 2017. URL <http://arxiv.org/abs/1708.09441>.
- Guha, S., Mishra, N., Roy, G., and Schrijvers, O. Robust random cut forest based anomaly detection on streams. In *International conference on machine learning*, pp. 2712–2721, 2016.
- Horn, C. and Willett, R. Online anomaly detection with expert system feedback in social networks. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1936–1939, 2011.
- Liu, F. T., Ting, K. M., and Zhou, Z. Isolation forest. In *2008 Eighth IEEE International Conference on Data Mining*, pp. 413–422, 2008.

- Pevný, T. Loda: Lightweight on-line detector of anomalies. *Mach. Learn.*, 102(2):275–304, February 2016. ISSN 0885-6125. doi: 10.1007/s10994-015-5521-0. URL <https://doi.org/10.1007/s10994-015-5521-0>.
- Raginsky, M., Marcia, R. F., Silva, J. G., and Willett, R. Sequential anomaly detection in the presence of noise and limited feedback. *CoRR*, abs/0911.2904, 2009. URL <http://arxiv.org/abs/0911.2904>.
- Rousseeuw, P. J. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53 – 65, 1987. ISSN 0377-0427. doi: [https://doi.org/10.1016/0377-0427\(87\)90125-7](https://doi.org/10.1016/0377-0427(87)90125-7). URL <http://www.sciencedirect.com/science/article/pii/0377042787901257>.
- Sabeti, E. and Høst-Madsen, A. Data discovery and anomaly detection using atypicality for real-valued data. *Entropy*, 21(3), 2019. ISSN 1099-4300. doi: 10.3390/e21030219. URL <https://www.mdpi.com/1099-4300/21/3/219>.
- Salinas, D., Flunkert, V., Gasthaus, J., and Januschowski, T. Deepar: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 2019.
- Shi, J., He, G., and Liu, X. Anomaly detection for key performance indicators through machine learning. In *2018 International Conference on Network Infrastructure and Digital Content (IC-NIDC)*, pp. 1–5, 2018.
- Siddiqui, M. A., Fern, A., Dietterich, T. G., Wright, R., Theriault, A., and Archer, D. W. Feedback-guided anomaly discovery via online optimization. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery Data Mining, KDD '18*, pp. 2200–2209, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450355520. doi: 10.1145/3219819.3220083. URL <https://doi.org/10.1145/3219819.3220083>.
- Siddiqui, M. A., Stokes, J. W., Seifert, C., Argyle, E., McCann, R., Neil, J., and Carroll, J. Detecting cyber attacks using anomaly detection with explanations and expert feedback. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2872–2876, 2019.
- Teoh, T. T., Chiew, G., Franco, E. J., Ng, P. C., Benjamin, M. P., and Goh, Y. J. Anomaly detection in cyber security attacks on networks using mlp deep learning. In *2018 International Conference on Smart Computing and Electronic Enterprise (ICSCEE)*, pp. 1–5, 2018.
- Varadarajan, J., Subramanian, R., Ahuja, N., Moulin, P., and Odobez, J. Active online anomaly detection using dirichlet process mixture model and gaussian process classification. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 615–623, 2017.
- Vercruyssen, V., Meert, W., Verbruggen, G., Maes, K., Bäumer, R., and Davis, J. Semi-supervised anomaly detection with an application to water analytics. In *2018 IEEE International Conference on Data Mining (ICDM)*, pp. 527–536, 2018.