COMPUTER SCIENCE
& ENGINEERING
TEXAS A&M UNIVERSITY

Data Analytics at Texas A&M (DATA Lab)

Meta-AAD: Active Anomaly Detection with Deep Reinforcement Learning

Daochen Zha, Kwei-Herng Lai, Mingyang Wan and Xia Hu

Department of Computer Science and Engineering, Texas A&M University

Emails: {daochen.zha, khlai037, w1996, xiahu}@tamu.edu

Code: https://github.com/daochenzha/Meta-AAD

What is anomaly detection?

- *Goal:* ***Identify the data objects or behaviors that significantly deviate from the majority.***
- *Applications:* ***Fraud detection, cybersecurity attack detection, medical diagnosis, etc.***
- *Challenges:* ***High false-positive rate. Lots of false alarms.***
- *Why high false-positive rate?* ***Most algorithms are unsupervised with assumptions on the anomaly patterns. There is usually discrepancy between the assumptions and real world applications.***

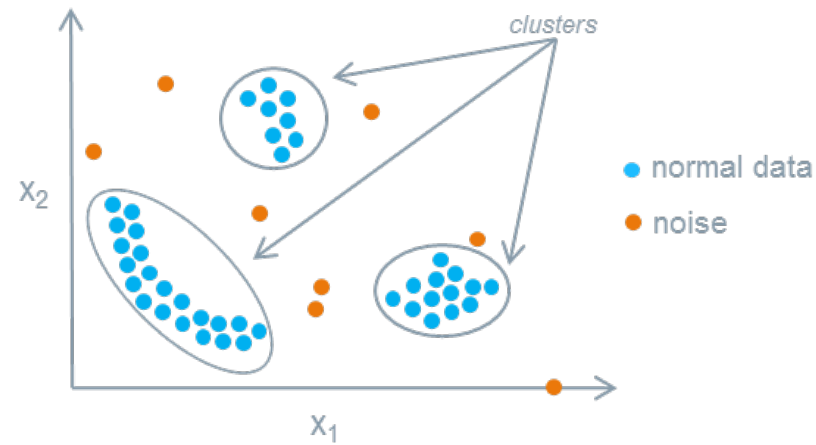

Source: https://developer.mindsphere.io/apis/analytics-anomalydetection/api-anomalydetection-overview.html

Active Anomaly Detection (AAD)

- ***Main Idea:*** ***Correcting the assumptions with feedback from the experts.***
- ***Human-in-the-Loop:*** ***(1) Select a query; (2) label the query; (3) adjust scores; (4) go to (1)***

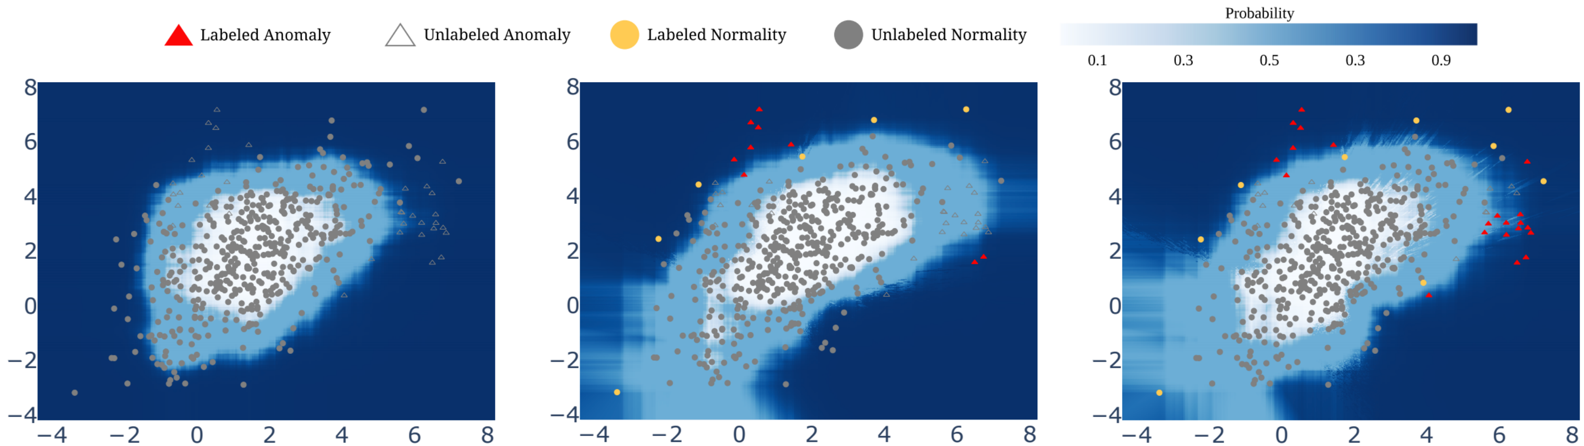

(a) Initial state (b) 15 queries (c) 30 queries

- ***Observation:*** ***The decision boundary evolves with more and more queries.***

Problem Statement & Related Work

- *Given the dataset* $\mathbf{X}$ and a budget T, in each iteration, we aim to select one instance $\mathbf{X_i}$ from $\mathbf{X}$ for query. An analyst will give a label y_i to indicate whether it is anomalous or not.
- **Objective:** Maximize the number of discovered anomalies when the budget T is used up. i.e., within T queries.

Existing solutions:

- *Active Anomaly Detection (AAD) [1]: state-of-the-art method based on node re-weighting*
- *Feedback-Guided Isolation Forest [2]: active anomaly detector via online optimization*
- *OJRANK [3]: re-rank the instances and select top-1 as feedback*

Observation: ***They focus on making top-1 instance anomalous, but not long-term performance.***

[1] Das, Shubhomoy, et al. Incorporating feedback into tree-based anomaly detection. arXiv:1708.09441 (2017).
[2] Siddiqui, Md Amran, et al. Feedback-guided anomaly discovery via online optimization. KDD, 2018.
[3] Lamba, Hemank, and Leman Akoglu. Learning on-the-job to re-rank anomalies from top-1 feedback. SDM, 2019.

Meta-AAD: Optimizing the Performance with RL

- **Motivation:** RL can inherently balance long-term and short-term rewards.
- **Challenges:** (1) Huge decision space; (2) RL is not sample efficient.

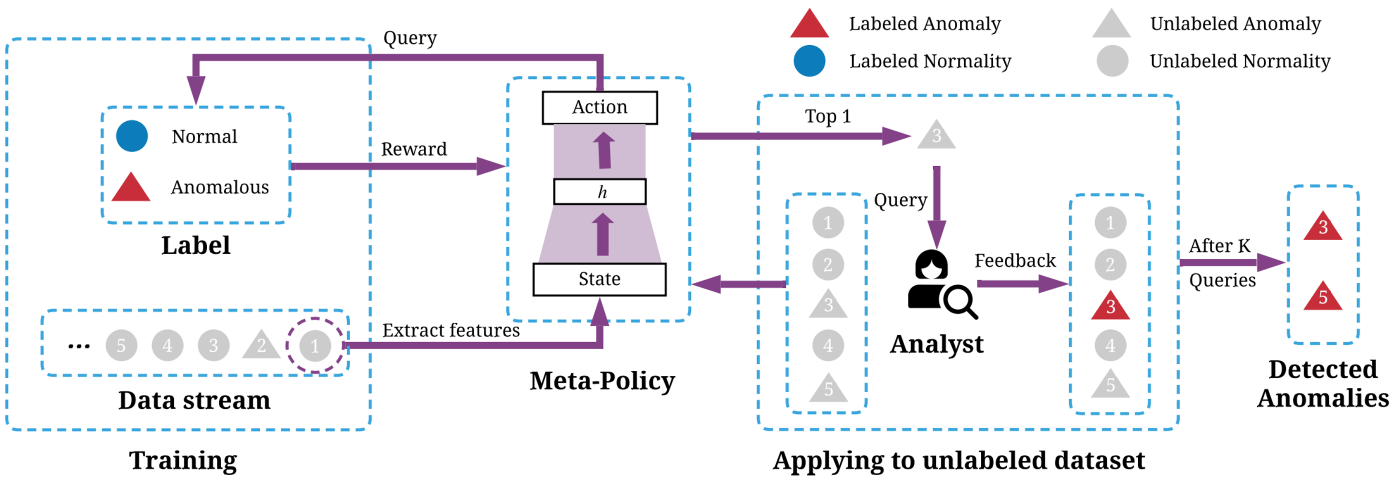

- **How it works?** (1) Train in a streaming manner on labeled data. (2) Transfer to unlabeled data.

Training of Meta-Policy

- **State:** meta-features of an instance, including the unsupervised anomaly scores, distance to the labeled anomalies, and the distance to the labeled normalities.
- **Action:** 1 for query; 0 for not query.
- **Reward:** a positive reward of 1 if the queried instance is indeed anomaly; a negative reward 0f -0.1 if it is not; a reward of 0 if not queried.

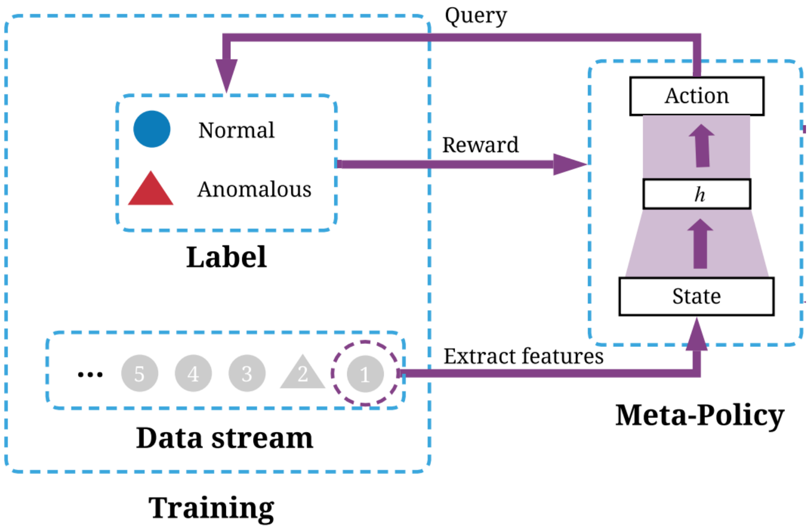

Applying Meta-Policy to Unlabeled data

- **Step 1:** Select the instance with the highest probability of taking action 1 for query.
- **Step 2:** The analyst gives the label.
- **Step 3:** The meta-features are adjusted according to the feedback.
- **Step 4:** Repeat step 1 to 3 until budget is used up.

Nice Properties:

(1) RL models long-term performance.

(2) The policy can be directly transferred.

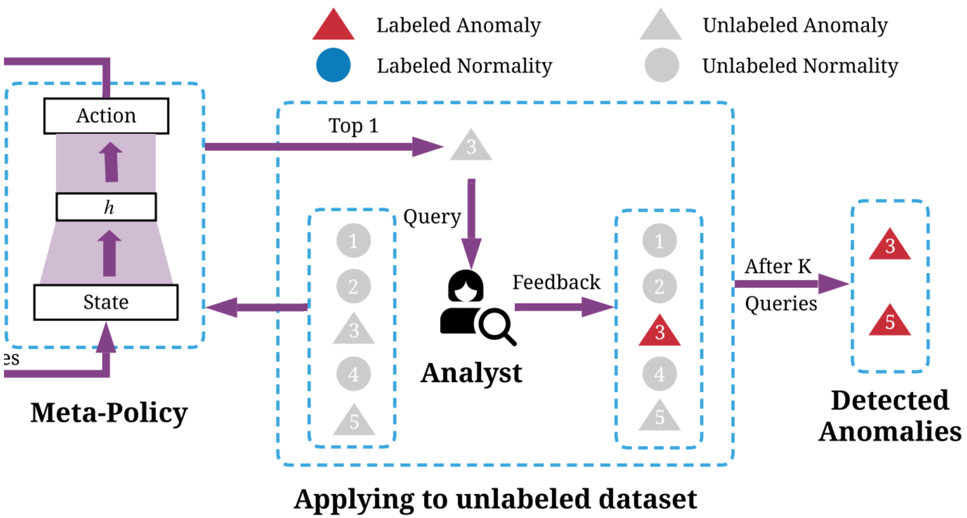

Basslines and Datasets

- **AAD:** Active Anomaly Detection [1] is a state-of-the-art method based on node re-weighting.
- **FIF:** Feedback-Guided Isolation Forest [2] is a recently proposed active anomaly detector via online optimization.
- **SSDO:** Semi-Supervised Detection of Outliers [3] also use label information.
- **Unsupervised:** We use Isolation Forest (IF) [4] as an unsupervised baseline.

- **Datasets:** We 24 real-world datasets from ODDS [5].

[1] Das, Shubhomoy, et al. Incorporating feedback into tree-based anomaly detection. arXiv:1708.09441 (2017).
[2] Siddiqui, Md Amran, et al. Feedback-guided anomaly discovery via online optimization. KDD, 2018.
[3] Vercruyssen, Vincent, et al. Semi-supervised anomaly detection with an application to water analytics. ICDM, 2018.
[4] Liu, Fei Tony, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. ICDM, 2008
[5] http://odds.cs.stonybrook.edu/

How does Meta-AAD perform on Benchmarks?

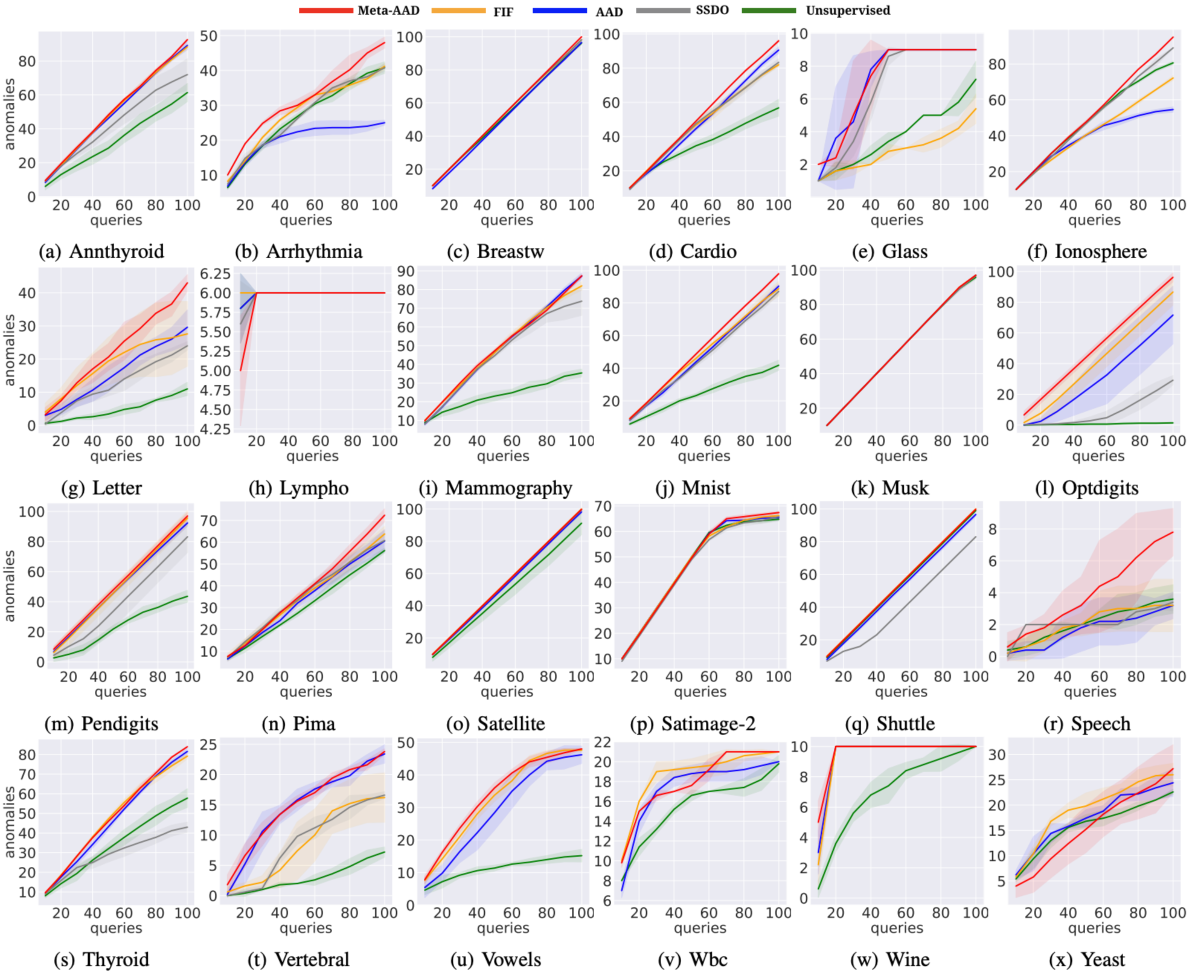

Method	**20**	**40**	**60**	**80**	**100**
unsupervised [2]	4.188▲	4.146▲	4.167▲	4.333▲	4.375▲
SSDO [18]	3.312▲	3.396▲	3.500▲	3.625▲	3.438▲
AAD [5]	3.229▲	3.208▲	3.271▲	3.167▲	3.104▲
FIF [6]	2.208	2.333	2.312	2.396▲	2.708▲
Meta-AAD	**2.062**	**1.917**	**1.750**	**1.479**	**1.375**
Improvement	0.146	0.416	0.562	0.917	1.333

▲ Meta-AAD is significantly better than the baseline w.r.t. the Wilcoxon signed rank test ($p < 0.01$).

Observations:

(1) Meta-AAD outperforms baselines.

(2) Meta-AAD has stronger performance in the long-term

Ablation Study

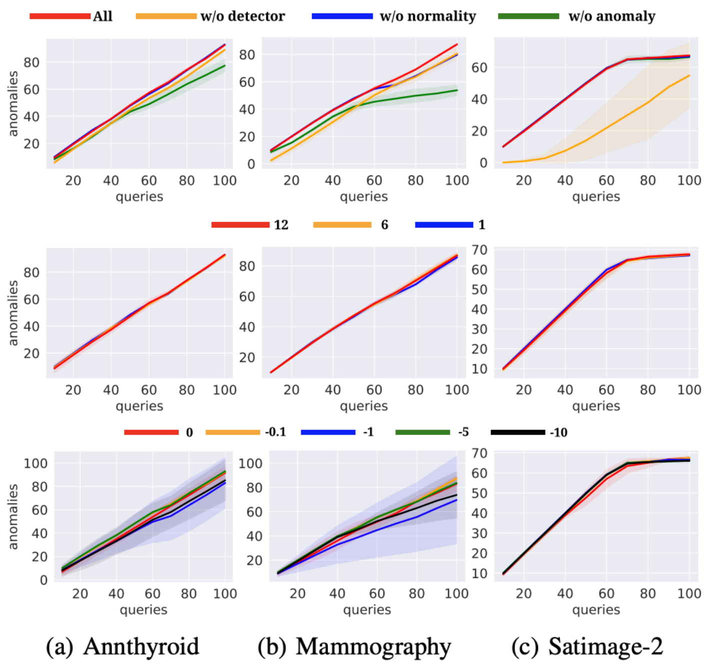

Fig. 4: Ablation study of Meta-AAD. We show the learning curves on Annthyroid, Mammography, Satimage-2 by dropping different features (top row), using different number of training datasets (mid row), and using different negative rewards for a missed query.

Observations:

(1) All the proposed meta-features are helpful

(2) One dataset is enough for training, which suggests the meta-policy is indeed transferable.

(3) The negative reward can not be too large nor too small.

Efficiency & Sensitivity Analysis

Fig. 5: The average discovered anomalies across all the datasets given 100 queries with respect to the number of training steps (left) and different γ values (right).

Observations:

(1) It usually takes less than 2 minutes for training with one core on a PC.

(2) The hyperparameter γ can balance long-term and short-term performance.

Takeaways

Some insights:

(1) With very few labels, active learning can effectively correct the anomaly detector and boost the performance.

(2) The active learning strategy is transferable and can easily deployed.

Our contributions:

(1) We propose a practical framework, called Meta-AAD, which optimizes the performance of active anomaly detection with deep reinforcement learning.

(2) Extensive experiments are presented to validate our framework.

(3) We open-source the code and all the datasets to facilitate future research: *https://github.com/daochenzha/Meta-AAD*

Acknowledgement

- *DATA Lab and collaborators*

COMPUTER SCIENCE & ENGINEERING
TEXAS A&M UNIVERSITY

DATA LAB
TAMU

Data Analytics at Texas A&M (DATA Lab)

- *National Science Foundation (NSF)*
- *Everyone watching this video!*