

# Effectiveness of Tree-based Ensembles for Anomaly Discovery: Insights, Batch and Streaming Active Learning

Shubhomoy Das

SHUBHOMOY.DAS@WSU.EDU

Md Rakibul Islam

MDRAKIBUL.ISLAM@WSU.EDU

Nitthilan Kannappan Jayakodi

N.KANNAPPANJAYAKODI@WSU.EDU

Janardhan Rao Doppa

JANA.DOPPA@WSU.EDU

*School of EECS, Washington State University, Pullman, WA 99164, USA*

## Abstract

Anomaly detection (AD) task corresponds to identifying the true anomalies among a given set of data instances. AD algorithms score the data instances and produce a ranked list of candidate anomalies. The ranked list of anomalies is then analyzed by a human to discover the true anomalies. Ensemble of tree-based anomaly detectors trained in an unsupervised manner and scoring based on uniform weights for ensembles are shown to work well in practice. However, the manual process of analysis can be laborious for the human analyst when the number of false-positives is very high. Therefore, in many real-world AD applications including computer security and fraud prevention, the anomaly detector must be configurable by the human analyst to minimize the effort on false positives. One important way to configure the detector is by providing true labels (nominal or anomaly) for a few instances. Recent work on active anomaly discovery has shown that greedily querying the top-scoring instance and tuning the weights of ensemble detectors based on label feedback allows us to quickly discover true anomalies.

This paper makes four main contributions to improve the state-of-the-art in anomaly discovery using tree-based ensembles. First, we provide an important insight that explains the practical successes of unsupervised tree-based ensembles and active learning based on greedy query selection strategy. We also present empirical results on real-world data to support our insights and theoretical analysis to support active learning. Second, we develop a novel batch active learning algorithm to improve the diversity of discovered anomalies based on a formalism called compact description to describe the discovered anomalies. Third, we develop a novel active learning algorithm to handle streaming data setting. We present a data drift detection algorithm that not only detects the drift robustly, but also allows us to take corrective actions to adapt the anomaly detector in a principled manner. Fourth, we present extensive experiments to evaluate our insights and our tree-based active anomaly discovery algorithms in both batch and streaming data settings. Our results show that active learning allows us to discover significantly more anomalies than state-of-the-art unsupervised baselines, our batch active learning algorithm discovers diverse anomalies, and our algorithms under the streaming-data setup are competitive with the batch setup.

## 1. Introduction

We consider the problem of *anomaly detection* (AD), where the goal is to detect unusual but interesting data (referred to as *anomalies*) among the regular data (referred to as *nominals*). This problem has many real-world applications including credit card transactions, medical diagnostics, computer security, etc., where anomalies point to the presence of phenomena such as fraud, disease, malpractices, or novelty which could have a large impact on the

domain, making their timely detection crucial. Anomaly detection poses severe challenges that are not seen in traditional learning problems. First, anomalies are significantly fewer in number than nominal instances. Second, unlike classification problems, there may not be a hard decision boundary to separate anomalies and nominals.

Typically, anomaly detection algorithms train models to compute scores for all instances, and report instances which receive the highest scores as anomalies. The candidate anomalies in the form of top-ranked data instances are analyzed by a human analyst to identify the true anomalies. Prior work has shown that a class of ensemble anomaly detectors trained in an unsupervised manner and scoring based on uniform weights work very well (Das, Wong, Dietterich, Fern, & Emmott, 2016). Algorithms based on tree-based ensembles such as isolation forest (IFOR) consistently perform well on real-world data (Emmott, Das, Dietterich, Fern, & Wong, 2015). However, purely unsupervised algorithms are usually ineffective in the following scenarios: a) The underlying model assumptions are not suited to the real-world task (Chandola, Banerjee, & Kumar, 2009); b) The data volume is so high that even a detector with low false-positive rate reports a very large absolute number of potential anomalies; c) The nature/type of anomalies is not known beforehand (unknown unknowns) and the user wants to explore the data by directing the detector with feedback in the spirit of collaborative problem solving; and d) the same algorithm is able to identify multiple types of anomalies but only a subset of them are relevant for a particular domain. Consider the following simple example. Let us assume that there are three features (say A, B, and C) and we use a generic unsupervised detector which reports instances which are outside 2-standard deviation along each axis as anomalies. However, if this rule is relevant only for features A and B for a particular domain, then the purely unsupervised algorithm would be less useful for this domain as it would report too many false positives. As a concrete real-world scenario, assume that in a computer network security application, the features A, B, and C correspond to *total bytes transferred over the network*, *no. of unique IP addresses connected to*, and *no. of logins in off-business hours* respectively for a server. With work from home mandates, the last feature might become irrelevant. Unless the detector is adjusted to account for this, the number of false positives could be high.

If the candidate set of anomalies contains many false-positives, it will significantly increase the effort of the human analyst in discovering true anomalies. To overcome the above drawbacks of unsupervised AD algorithms, recent work has explored active learning framework to configure the anomaly detector by the human analyst to minimize the overall effort on false-positives (Das et al., 2016; Siddiqui, Fern, Dietterich, Wright, Theriault, & Archer, 2018). In this framework, a human analyst provides label feedback (nominal or anomaly) for one or more data instances during each round of interaction. This feedback is used by the anomaly detector to change the scoring mechanism of data instances towards the goal of increasing the true anomalies appearing at the top of the ranked list. Specifically, prior work employs a greedy query selection strategy by selecting the top-scoring unlabeled instance for feedback. Similar to unsupervised anomaly detection ensembles, prior work showed that active learning using tree-based ensembles performs better (Das et al., 2016; Siddiqui et al., 2018).

The effectiveness of tree-based ensembles for anomaly discovery in both unsupervised and active-learning settings raise two fundamental questions: **(Q1)** Why does the average score across ensemble members perform best in most cases (Chiang & Yeh, 2015) instead

of other score combination strategies (e.g., min, max, median etc.); and **(Q2)** Why does the greedy query selection strategy for active learning almost always perform best? In this paper, we investigate these two fundamental questions related to tree-based ensembles and provide an important insight that explains their effectiveness for anomaly discovery. We also provide empirical evidence on real-world data to support our intuitive understanding and theoretical analysis to support active anomaly discovery. Prior work on active learning for anomaly discovery has some important shortcomings. First, algorithmic work on enhancing the diversity of discovered anomalies is lacking (Görnitz, Kloft, Rieck, & Brefeld, 2013). Second, most algorithms are designed to handle batch data well, but there are few principled algorithms to handle streaming data setting that arises in many real-world applications. To fill this important gap in active anomaly discovery, we exploit the inherent strengths of tree-based ensembles and propose novel batch and streaming active learning algorithms.

**Contributions.** The main contribution of this paper is to study and evaluate anomaly discovery algorithms based on tree-based ensembles in both unsupervised and active learning settings. Specific contributions include:

- We present an important insight into how tree-based anomaly detector ensembles are naturally suited for active learning, and why the greedy querying strategy of seeking labels for instances with the highest anomaly scores is efficient. We also provide theoretical analysis to support active anomaly discovery using tree-based ensembles.
- A novel formalism called compact description (CD) is developed to describe the discovered anomalies using tree-based ensembles. We develop batch active learning algorithms based on CD to improve the diversity of discovered anomalies.
- To handle streaming data setting, we develop a novel algorithm to robustly detect drift in data streams and design associated algorithms to adapt the anomaly detector on-the-fly in a principled manner.
- We present extensive empirical evidence in support of our insights and algorithms on several benchmark datasets. The results show the efficacy of the proposed active learning algorithms in both batch and streaming data settings. Our results show that in addition to discovering significantly more anomalies than state-of-the-art unsupervised baselines, our learning algorithm under the streaming-data setup is competitive with the batch setup. Our code and data are publicly available<sup>1</sup>.

**Outline of the Paper.** The remainder of the paper is organized as follows. In Section 2, we discuss the prior work related to this paper. We introduce our problem setup and give a high-level overview of the generic human-in-the-loop learning framework for active anomaly discovery in Section 3. In Section 4, we describe the necessary background on tree-based anomaly detector ensembles and provide intuition along with empirical evidence and theoretical analysis to explain their suitability for active anomaly discovery. In Section 5 and 6, we present a series of algorithms to support the framework for batch data and streaming data settings respectively. Section 7 presents our experimental results and finally Section 8 provides summary and directions for future work.

---

1. [https://github.com/shubhomoydas/ad\\_examples](https://github.com/shubhomoydas/ad_examples)

## 2. Related Work

**Unsupervised anomaly detection algorithms** are trained without labeled data and have assumptions baked into the model about what defines an anomaly or a nominal (Breunig, Kriegel, Ng, & Sander, 2000; Liu, Ting, & Zhou, 2008; Pevný, 2016; Emmott et al., 2015). They typically cannot change their behavior to correct for the false positives after they have been deployed. Ensembles of unsupervised anomaly detectors (Aggarwal & Sathe, 2017) try to guard against the bias induced by a single detector by incorporating decisions from multiple detectors. The potential advantage of ensembles is that when the data is seen from multiple views by more than one detector, the set of anomalies reported by their joint consensus is more reliable and has fewer false positives. Different methods of creating ensembles include the collection of heterogeneous detectors (Ted, Goldberg, Memory, Young, Rees, Pierce, Huang, Reardon, Bader, Chow, et al., 2013), feature bagging (Lazarevic & Kumar, 2005), varying the parameters of existing detectors such as the number of clusters in a Gaussian Mixture Model (Emmott et al., 2015), sub-sampling, bootstrap aggregation etc. Some unsupervised detectors such as LODA (Pevný, 2016) are not specifically designed as ensembles but may be treated as such because of their internal structure. Isolation forest (IFOR) (Liu et al., 2008) is a state-of-the-art tree-based ensemble anomaly detector. Recently, some prior work looked at the problem of ensemble selection by transfer learning (Campos, Zimek, & Meira, 2018), ensemble score aggregation methods (Zhao & Hryniewicki, 2019) by looking at the locality of a data point, approximated neighborhood creation techniques and combine their scores (Kirner, Schubert, & Zimek, 2017) as outlier scores, feature selection process in outlier detection (Cheng, Wang, Liu, & Li, 2020) (Mícenková, McWilliams, & Assent, 2014). Zero++ (Pang, Ting, Albrecht, & Jin, 2016) is an unsupervised approach that works only for categorical data. It is similar to IFOR, but there is no statistically significant difference in detection accuracy between Zero++ and IFOR for numerical data that is transformed to categorical data.

**Active learning** corresponds to the setup where the learning algorithm can selectively query a human analyst for labels of input instances to improve its prediction accuracy. The overall goal is to minimize the number of queries to reach the target performance. There is a significant body of work on both theoretical analysis (Freund, Seung, Shamir, & Tishby, 1997; Balcan, Broder, & Zhang, 2007; Balcan & Feldman, 2015; Monteleoni, 2006; Dasgupta, Kalai, & Monteleoni, 2009; Yan & Zhang, 2017; Cao & Tsang, 2022; Cao, Liu, & Tsang, 2023) and applications (Settles, 2012) of active learning.

**Active learning for anomaly detection** has been explored by researchers in several earlier works (Almgren & Jonsson, 2004; Abe, Zadrozny, & Langford, 2006; Stokes, Platt, Kravis, & Shilman, 2008; He & Carbonell, 2007; Pichara & Soto, 2011; Görnitz et al., 2013). In this setting, the human analyst provides feedback to the algorithm on true labels (nominal or anomaly). If the algorithm makes wrong predictions, it updates its model parameters to be consistent with the analyst’s feedback. Interest in this area has gained prominence in recent years due to significant rise in the volume of data in real-world settings, which made reduction in false-positives much more critical. Most of the popular unsupervised anomaly detection algorithms have been adapted to the active learning setting. Some of these methods are based on ensembles and support streaming data (Veeramachaneni, Arnaldo, Korrapati, Bassias, & Li, 2016; Stokes et al., 2008) but maintain separate models

for anomalies and nominals internally, and do not exploit the inherent strength of the ensemble in a systematic manner. (Nissim, Cohen, Moskovitch, Shabtai, Edry, Bar-Ad, & Elovici, 2014) incorporates user feedback into SVM, a classification algorithm. (Das et al., 2016) and (Das, Wong, Fern, Dietterich, & Siddiqui, 2017), based on LODA and IFOR respectively, re-weight the components of the ensemble with every user feedback in a manner similar to our work, however, they do not support streaming data and discovery of diverse anomalies. (Siddiqui et al., 2018) updates the same model as (Das et al., 2017) in an online incremental manner when incorporating user feedback.

(Ikeda, Ishibashi, Nakano, Watanabe, Tajiri, & Kawahara, 2018) employs an autoencoder as the base anomaly detector and adds a penalty term for the false positives (i.e., normal data that are incorrectly reported as anomalies). The user examines the anomalies reported by the algorithm and labels the false-positives. The model is then retrained with the new labeled data in an online manner. (Pimentel, Monteiro, Veloso, & Ziviani, 2020) adds a classifier layer on top of any unsupervised deep network based anomaly detector and updates the parameters of the deep network based on human feedback. (Trittenbach & Böhm, 2019) extends SVDD, a one-class classifier model for outlier detection, to multiple sub-spaces. The multiple sub-spaces are intended to help with interpretability of the outliers and support better query strategies for active learning. (Kim, Kim, Yu, & Choi, 2023) learns an adaptive boundary using a deep SVDD model by searching for locations having a similar number of normal and abnormal samples. This boundary is iteratively adjusted by getting labeling feedback from users with an uncertainty based query strategy. Deep semi-supervised anomaly detection techniques have also been used to incorporate feedback from users (Ruff, Vandermeulen, Göernitz, Deecke, Siddiqui, Binder, Müller, & Kloft, 2018; Ruff, Vandermeulen, Görnitz, Binder, Müller, Müller, & Kloft, 2019; Li, Qiu, Kloft, Smyth, Mandt, & Rudolph, 2023). The performance gain for such techniques on classical tabular datasets can be quite low in comparison to performance on complex raw data such as images. SOEL (Li et al., 2023) employs a loss function that has opposing effects on normal and anomalous data points. Using this loss, SOEL infers latent labels  $\{anomaly, normal\}$  for each training example using a semi-supervised EM-style algorithm. In this work, we present a comparison of our algorithm with SOEL in Section 7.3.1.

Our proposed algorithmic instantiation of HiLAD for batch data setting (HiLAD-Batch) and recent work on feedback-guided AD via online optimization (**Feedback-guided Online**) (Siddiqui et al., 2018) both build upon the same tree-based model proposed by (Das et al., 2017). Therefore, there is no fundamental difference in their performance (7.3). The uniform initialization of weights employed by both HiLAD-Batch and **Feedback-guided Online** plays a critical role in their overall effectiveness. **Feedback-guided Online** also adopts the greedy query selection strategy mainly because it works well in practice. In this work, we present the fundamental reason why greedy query selection strategy is label-efficient for human-in-the-loop learning with tree-based anomaly detection ensembles.

Anomaly detection for streaming data setting has many real-life applications. Under this setup, data is sent continuously as a stream to the anomaly detector. (Gupta, Gao, Aggarwal, & Han, 2014) present a broad overview of this setup and categorize the outlier detection techniques into three sub-categories: (a) evolving prediction models, (b) distance based outlier detection for sliding windows, (c) methods for high-dimensional data streams. Our proposed approach falls under the class of distance based outlier detection in a sliding

Table 1: Notational convention

Symbol	Description
$\mathbf{D}$	Dataset
$x_i$	A data instance
$y_i$	A label (hidden)
$\mathcal{E}$	Ensemble
$z_m$	Anomaly score from $m^{\text{th}}$ ensemble member
$\mathbf{H}$	Score matrix of all unlabeled data instances
$\mathbf{H}_+$	Score matrix of all positively labeled data instances
$\mathbf{H}_-$	Score matrix of all negatively labeled data instances
$\mathbf{w}$	Weight vector
$B$	Query budget
$\mathcal{L}$	Labeled dataset
$\tau$	Initial guess on the fraction of anomalous data instances
$q$	Score for $\tau^{\text{th}}$ ranked anomaly

window. Existing approaches can detect either local anomalies (Breunig et al., 2000; Salehi, Leckie, Bezdek, Vaithianathan, & Zhang, 2016; Na, Kim, & Yu, 2018) or global anomalies (Angiulli & Fassetti, 2007; Yang, Rundensteiner, & Ward, 2009; Kontaki, Gounaris, Papadopoulos, Tsihclas, & Manolopoulos, 2011). (Guha, Mishra, Roy, & Schrijvers, 2016) proposed a tree-based streaming anomaly detector. The key idea is to create a sketch (or summary) of the data using a random cut forest and to report a data point as anomaly when the displacement score is high. However, unlike our work, none of these prior streaming anomaly detection methods incorporate human feedback.

### 3. Problem Setup and Human-in-the-Loop Learning Framework

We are given a dataset  $\mathbf{D} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ , where  $\mathbf{x}_i \in \mathbb{R}^d$  is a data instance that is associated with a hidden label  $y_i \in \{-1, +1\}$ . Instances labeled  $+1$  represent the *anomaly* class and are at most a small fraction  $\tau$  of all instances. The label  $-1$  represents the *nominal* class. We also assume the availability of an ensemble  $\mathcal{E}$  of  $m$  anomaly detectors which assigns scores  $\mathbf{z} = \{z^1, \dots, z^m\}$  to each instance  $\mathbf{x}$  such that instances labeled  $+1$  tend to have scores higher than instances labeled  $-1$ . Our framework is applicable for any ensemble of detectors, but we specifically study the instantiation for tree-based ensembles because of their beneficial properties. We denote the ensemble score matrix for all unlabeled instances by  $\mathbf{H}$ . The score matrix for the set of instances labeled  $+1$  is denoted by  $\mathbf{H}_+$ , and the matrix for those labeled  $-1$  is denoted by  $\mathbf{H}_-$ .

We setup a scoring function  $\text{Score}(\mathbf{x})$  to score data instances, which is parameterized by  $\mathbf{w}$  and ensemble of detectors  $\mathcal{E}$ . For example, with tree-based ensembles, we consider a linear model with weights  $\mathbf{w} \in \mathbb{R}^m$  that will be used to combine the scores of  $m$  anomaly detectors as follows:  $\text{Score}(\mathbf{x}) = \mathbf{w} \cdot \mathbf{z}$ , where  $\mathbf{z} \in \mathbb{R}^m$  corresponds to the scores from anomaly detectors for instance  $\mathbf{x}$ . This linear hyperplane separates anomalies from nominals. We will denote the optimal weight vector by  $\mathbf{w}^*$ .

The generic human-in-the-loop learning framework HiLAD assumes the availability of an analyst who can provide the true label for any instance in an interactive loop as shown in Figure 1. In each iteration of the learning loop, we perform the following steps: 1) Select one or more unlabeled instances from the input dataset  $\mathbf{D}$  according to a query selection strategy; 2) Query the human analyst for labels of selected instances; and 3) Update the weights of the scoring function based on the aggregate set of labeled and unlabeled instances. The overall goal of the framework is to learn optimal weights for maximizing the number of true anomalies shown to the human analyst.

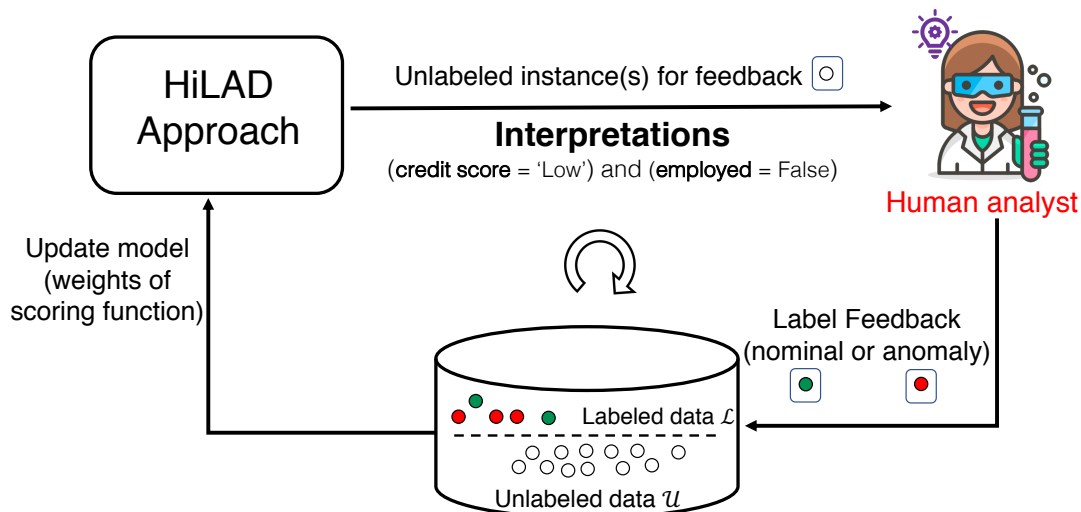


Figure 1: High-level overview of the human-in-the-loop learning framework for anomaly detection. Our goal is to maximize the number of true anomalies presented to the analyst.

---

**Algorithm 1** HiLAD Framework
 

---

**Input:** Unlabeled dataset  $\mathcal{U}$ , Feedback budget  $B$ , and Anomaly detection ensembles  $\mathcal{E}$  of size  $m$

Initialize weights  $\mathbf{w} \in \mathbb{R}^m$  of the linear scoring function  $\text{Score}(\mathbf{x})$

Initialize  $\mathcal{L} \leftarrow \emptyset$  // Labeled dataset

**while**  $|\mathcal{L}| < B$  **do**

$\mathbf{X}_{\text{query}} \leftarrow$  one or more instances selected from  $\mathcal{U}$  using query selection strategy

    Get label feedback for instances in  $\mathbf{X}_{\text{query}}$  from the human analyst by showing interpretations of anomalous instances

    Add newly labeled examples to  $\mathcal{L}$

    Update weights  $\mathbf{w}$  to be consistent with aggregate labeled data  $\mathcal{L}$

    Update unlabeled set:  $\mathcal{U} \leftarrow \mathcal{U} \setminus \mathbf{X}_{\text{query}}$

**end while**

**return** discovered anomalous instances from  $\mathcal{L}$

---

We provide algorithmic solutions for key elements of this human-in-the-loop framework:

- Initializing the parameters  $\mathbf{w}$  of the scoring function  $\text{Score}(\mathbf{x})$  based on a key insight for tree-based anomaly detection ensembles.
- Query selection strategies to improve the label-efficiency of learning.
- Updating the weights of scoring function based on label feedback.
- Updating ensemble members as needed to support the streaming data setting.

## 4. Suitability of Tree-based Ensembles for Human-in-the-Loop Learning

In this section, we first provide the background on tree-based anomaly detector ensembles and describe *Isolation Forest (IFOR)* in more depth as we employ it for our experimental evaluation. Subsequently, we provide intuition on how tree-based ensembles are naturally suited for human-in-the-loop learning and show empirical evidence using isolation forest.

### 4.1 Tree-based Anomaly Detection Ensembles

Ensemble of tree-based anomaly detectors have several attractive properties that make them an ideal candidate for human-in-the-loop learning:

- They can be employed to construct large ensembles inexpensively.
- Treating the nodes of tree as ensemble members allows us to both focus our feedback on fine-grained subspaces<sup>2</sup> as well as increase the *capacity* of the model in terms of separating anomalies (positives) from nominals (negatives) using complex class boundaries.
- Since some of the tree-based models such as *Isolation Forest (IFOR)* (Liu et al., 2008), *HS Trees (HST)* (Tan, Ting, & Liu, 2011), and *RS Forest (RSF)* (Wu, Zhang, Fan, Edwards, & Yu, 2014) are state-of-the-art unsupervised detectors (Emmott et al., 2015; Domingues, Filippone, Michiardi, & Zouaoui, 2018), it is a significant gain if their performance can be improved with minimal label feedback.

In this work, we will focus mainly on IFOR because it performed best across all datasets (Emmott et al., 2015). However, we also present results on HST and RSF wherever applicable.

**Isolation Forest (IFOR)** comprises of an ensemble of *isolation* trees. Each tree partitions the original feature space at random by recursively splitting an unlabeled dataset. At every tree-node, first a feature is selected at random, and then a split point for that feature is sampled uniformly at random (Figure 2). This partitioning operation is carried out until every instance reaches its own leaf node. The key idea is that anomalous instances, which are generally isolated in the feature space, reach the leaf nodes faster by this partitioning strategy than nominal instances which belong to denser regions (Figure 3). Hence, the path

---

2. In this article, subspace stands for location encompassed by a leaf node of the tree, i.e., it is a rectangle for 2D space. All the leaf nodes for a tree cover the complete data space. This not to be confused with subspaces from Linear Algebra.



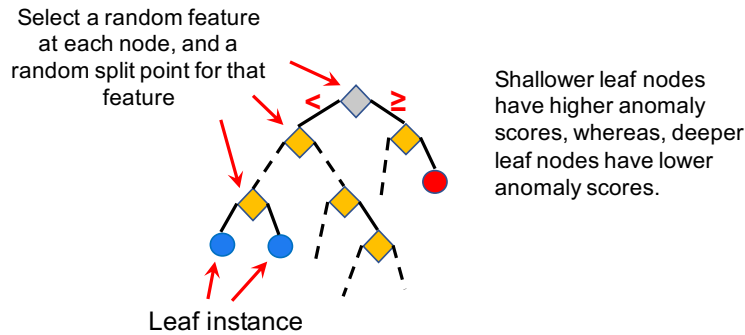


Figure 2: Illustration of Isolation tree from (Das et al., 2017).

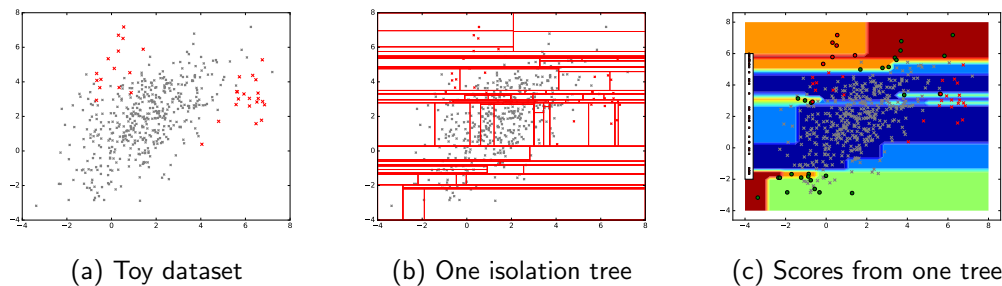


Figure 3: Illustration of Isolation Tree on simple data. **(a)** Toy dataset (Das et al., 2017). **(b)** A single isolation tree for the Toy dataset. **(c)** Regions having deeper red belong to leaf nodes which have shorter path lengths from the root and correspondingly, higher anomaly scores. Regions having deeper blue correspond to longer path lengths and lower anomaly scores.

from the root node is shorter to the leaves of anomalous instances when compared to the leaves of nominal instances. This path length is assigned as the unnormalized score for an instance by an isolation tree. After training an IFOR with  $T$  trees, we extract the leaf nodes as the members of the ensemble. Such members could number in the thousands (typically 4000 – 7000 when  $T = 100$ ). Assume that a leaf is at depth  $l$  from the root. If an instance belongs to the partition defined by the leaf, it gets assigned a score  $-l$  by the leaf, else 0. As a result, anomalous instances receive higher scores on average than nominal instances. Since every instance belongs to only a few leaf nodes (equal to  $T$ ), the score vectors are *sparse* resulting in low memory and computational costs.

**HST** and **RSF** apply different node splitting criteria than IFOR, and compute the anomaly scores on the basis of the sample counts and densities at the nodes. We apply log-transform to the leaf-level scores so that their unsupervised performance remains similar to the original model and yet improves with feedback. The trees in HST and RSF have a fixed depth which needs to be larger in order to improve the accuracy. In contrast, trees in IFOR have adaptive depth and most anomalous subspaces are shallow. Larger depths are associated with smaller subspaces, which are shared by fewer instances. As a result, feedback on any individual instance gets passed on to very few instances in HST and RSF,




Figure 4: Illustration of differences among different tree-based ensembles. The red rectangles show the union of the 5 most anomalous subspaces across each of the 15 most anomalous instances (blue). These subspaces have the highest influence in propagating feedback across instances through gradient-based learning under our model. HST has fixed depth which needs to be high for accuracy (recommended 15 (Tan et al., 2011)). IFOR has adaptive height and most anomalous subspaces are shallow. Higher depths are associated with smaller subspaces which are shared by fewer instances. As a result, feedback on any individual instance gets passed on to many other instances in IFOR, but to fewer instances in HST. RSF has similar behavior as HST. We set the depth for HST (and RSF (Wu et al., 2014)) to 8 (Figure 4c) in our experiments in order to balance accuracy and feedback efficiency.

but to much more number of instances in IFOR. Therefore, it is more efficient to incorporate feedback in IFOR than it is in HST or RSF (see Figure 4).

### 4.2 Beneficial Properties of Tree-based Anomaly Detector Ensembles

In this section, we provide intuition for the effectiveness of unsupervised tree-based anomaly detector ensembles and list their beneficial properties for human-in-the-loop learning framework. Our arguments are motivated by the active learning theory for standard classification.




Figure 5: Illustration of candidate score distributions from an ensemble in 2D. The two axes represent two different ensemble members. (a) **C1** represents the common case where both anomaly detectors want to score anomalous data points higher, (b) **C2** illustrates how active learning helps the model to learn the slight angle deviation,  $\theta$ , (c) **C3** is specifically for IFOR case where the anomalous data points has smaller path length. So, anomalous data points will be located at two extremes where the path lengths are smallest and can be separated by the non-homogeneous hyperplane.

Without loss of generality, we assume that all scores from the members of tree-based ensemble of anomaly detectors are normalized (i.e., they lie in  $[-1, 1]$  or  $[0, 1]$ ), with higher scores implying more anomalous. For the following discussion,  $\mathbf{w}_{unif} \in \mathbb{R}^m$  represents a vector of equal values, and  $\|\mathbf{w}_{unif}\| = 1$ . Figure 5 illustrates a few possible distributions of normalized scores from the ensemble members in 2D. In each figure, there are two ensemble members and each member tends to assign higher scores to anomalies than to nominals such that, on *average*, anomalies tend to have higher scores than nominals. Let  $\mathbf{z}$  represent the vector of anomaly scores for an instance and let  $m$  be the number of ensemble members. Then, the average score for the same instance can be computed as  $\text{Score} = \frac{1}{m} \sum_{i=1}^m z_i = \frac{1}{\sqrt{m}} (\mathbf{w}_{unif} \cdot \mathbf{z}) \propto \mathbf{w}_{unif} \cdot \mathbf{z}$ . We observe that under this setup, points in the score space closest to the top right region of the score space have the highest average scores since they are most aligned with the direction of  $\mathbf{w}_{unif}$ . When ensemble members are “good”, they assign higher scores to anomalies and push them to this region of the score space. The most general case is illustrated in case **C1** in Figure 5a. This behavior makes it easier to separate anomalies from nominals by a hyperplane. Moreover, it also sets a prior on the position and alignment of the separating hyperplane (as shown by the red solid line). Having priors on the hyperplane helps determine the region of label uncertainty (anomaly vs. nominal) even without any supervision, and this, in theory, strongly motivates active learning using the human analyst in the loop.

Most theoretical research on active learning for *classification* (Kearns, 1998; Balcan et al., 2007; Kalai, Klivans, Mansour, & Servedio, 2008; Dasgupta et al., 2009; Balcan & Feldman, 2015) makes simplifying assumptions such as uniform data distribution over a unit sphere and with homogeneous (i.e., passing through the origin) hyperplanes. However, for *anomaly detection*, arguably, the idealized setup is closer to case **C2** (Figure 5b), where *non-homogeneous* decision boundaries are more important. More importantly, in order for the theory to be relevant for anomaly detection, the assumptions behind it should be realizable with multi-dimensional real-world data and a competitive anomaly detector. We present empirical evidence (Section 7) which shows that scores from the state-of-the-art *Isolation Forest* (IFOR) detector are distributed in a similar manner as case **C3** (Figure 5c). **C3** and **C2** are similar in theory (for active learning) because both involve searching for the optimum non-homogeneous decision boundary. In all cases, the common theme is that when the ensemble members are ideal, then *the scores of true anomalies tend to lie in the farthest possible location in the positive direction of the uniform weight vector  $\mathbf{w}_{unif}$  by design*. Consequently, the average score for an instance across all ensemble members works well for anomaly detection. However, not all ensemble members are ideal in practice, and the true weight vector ( $\mathbf{w}^*$ ) is displaced by an angle  $\theta$  from  $\mathbf{w}_{unif}$ . Figure 6c shows an illustration of this scenario on a *Toy* dataset. In large datasets, even a small misalignment between  $\mathbf{w}_{unif}$  and  $\mathbf{w}^*$  results in many false positives. While the performance of ensemble on the basis of the AUC metric may be high, the detector could still be impractical for use by analysts.

The property, that the misalignment is usually small, can be leveraged by active learning to learn the optimal weights efficiently using a small amount of label feedback. To understand this, observe that the top-ranked instances are close to the decision boundary and are therefore, in the uncertainty region. The key idea is to design a hyperplane that passes through the uncertainty region which then allows us to select query instances by uncertainty sampling. Selecting instances on which the model is uncertain for labeling is

efficient for active learning (Cohn, Atlas, & Ladner, 1994; Balcan et al., 2007; Cao et al., 2023). Specifically, greedily selecting instances with the highest scores is first of all more likely to reveal anomalies (i.e., true positives), and even if the selected instance is nominal (i.e., false positive), it still helps in learning the decision boundary efficiently. This is an important insight and has significant practical implications. We present the label complexity of this greedy (but efficient) active learning algorithm for anomaly detection below.

**Label complexity of active learning with fixed weights:** Our trained ensemble of two anomaly detectors assign scores uniformly distributed over a unit sphere (Fig 5). Since we have assumed that  $\tau$  fraction of the data is anomalous, marginal density of anomalies is then  $\frac{\tau}{2\pi}$ . We also have a pool of unlabeled data. For simplicity, we assume that anomalies have a Gaussian distribution as a function of the angle  $\omega \in [-\pi, \pi]$  from  $\mathbf{u}$  on the unit sphere:  $f(y = +1|\omega) \sim \mathcal{N}(0, \sigma^2)$ . We query an analyst for the labels of only the top  $\beta$  fraction of instances sorted on scores, i.e.,  $\{\mathbf{z} : \arccos(\mathbf{w} \cdot \mathbf{z}) \leq \beta\pi\}$ . Let the angle between  $\mathbf{u}$  and  $\mathbf{w}_{unif}$  be  $\theta$ . Let  $\beta_l$  and  $\beta_r$  be the minimum and maximum of the values  $\{\theta - \beta\pi, \theta + \beta\pi\}$  respectively. As per our query selection strategy, the probability of labeling a true anomaly is then  $p_\theta = \int_{\beta_l}^{\beta_r} f(y = +1|\omega)d\omega = (\Phi(\frac{\beta_r}{\sigma}) - \Phi(\frac{\beta_l}{\sigma}))$ , where  $\Phi(\cdot)$  is the c.d.f of the Gaussian distribution. If  $|\theta|$  is large, i.e.,  $\mathbf{w}_{unif}$  and  $\mathbf{u}$  are not “close”, then  $p_\theta$  will be small.

**Proposition 4.1.** *Let  $\delta \in [0, 1]$ . For the 2D case, the number of labels needed to learn the decision boundary with probability  $(1 - \delta)$  with pool-based active learning is  $T = O(\log(\frac{1}{\sigma})\frac{1}{p_\theta} \log(\frac{1}{\delta}))$ .*

*Proof:* Essentially, we need to locate the mean of the Gaussian distribution. This can be done through binary search because the density is highest at the mean and decreases continuously as we move away from it. For the search, we estimate the density of anomalies at  $O(\log(\frac{1}{\sigma}))$  locations (for accuracy to one s.d.). For any estimate, at least one anomaly must be sampled. When the angle is  $\theta$  (i.e.,  $\mathbf{w} = \mathbf{w}_{unif}$ ) at the initial location, then the number of samples required is, with probability  $(1 - \delta)$ :  $(1 - p_\theta)^l \leq \delta \Rightarrow l \geq \frac{1}{p_\theta} \log(\frac{1}{\delta})$ . The number of samples at the initial location sets the baseline estimate; therefore, the total number of samples is  $T = O(\log(\frac{1}{\sigma})\frac{1}{p_\theta} \log(\frac{1}{\delta}))$ .  $\square$

**Remark.** If we ignore the initial location of search, the number of labels required would be  $O(\log(\frac{1}{\sigma})\frac{2\pi}{\tau} \log(\frac{1}{\delta}))$ . Thus, initializing at  $\mathbf{w}_{unif}$  helps when  $p_\theta$  is larger than  $\frac{\tau}{2\pi}$ .

**Label complexity of active learning with varying weights:** In this case, we select one data instance from the top ranked  $\beta$  fraction of instances at each time step for querying, and then adjust the weights  $\mathbf{w}$  in response to the label received from the analyst. We denote the weight vector and its angular displacement from  $\mathbf{u}$  at each time step  $t$  by  $\mathbf{w}_t$  and  $\theta_t$  respectively, starting with  $\mathbf{w}_1 = \mathbf{w}_{unif}$  and  $\theta_1 = \theta$ .

**Proposition 4.2.** *Let the algorithm  $\mathcal{A}$  update the weights with each new label at time  $t$  while learning the decision boundary such that  $\theta_{t+1} \leq \theta_t$ . If the label complexity of algorithm  $\mathcal{A}$  is  $T'$ , then  $T' \leq T$  (where  $T$  is defined in Proposition 4.1).*

*Proof:* As in the previous case, we sample at  $O(\log(\frac{1}{\sigma}))$  locations for the binary search. Let  $l'$  be the number of samples required at the initial location of the binary search for density estimation, and set  $\theta_1 = \theta$ . Now,  $l'$  sets the baseline for the number of samples in each round of the search. Since, by assumption,  $\theta_{t+1} \leq \theta_t$  at each label iteration  $t$ ,  $p_{\theta_{t+1}} \geq p_{\theta_t}$ .




Figure 6: Isolation Tree. (a) Toy dataset (Das et al., 2017) which will be used as the running example through out the text to illustrate the ideas. Red points are anomalies and black points are nominals. (b) Anomaly scores assigned by IFOR to the Toy dataset. (c) Histogram distribution of the angles between score vectors from IFOR and  $\mathbf{w}_{unif}$ . The red and green histograms show the angle distributions for anomalies and nominals respectively. Since the red histograms are closer to the left, anomalies are aligned closer to  $\mathbf{w}_{unif}$ .

For any estimate, at least one anomaly needs to be selected with probability  $(1 - \delta)$ ; i.e.,  $(1 - p_{\theta_1}) \dots (1 - p_{\theta_{l'}}) \leq \delta$ . Since,  $(1 - p_{\theta})^{l'} \geq (1 - p_{\theta_1}) \dots (1 - p_{\theta_{l'}})$ ,  $\exists l \geq l'$  s.t.  $(1 - p_{\theta})^l \leq \delta$ . Therefore,  $T' = O(\log(\frac{1}{\delta})l') \leq O(\log(\frac{1}{\delta})l) = T$ .  $\square$

When the number of members in the ensemble is  $m$ , then  $p_{\theta}$  will be a function of  $m$  and the number of locations for binary search will be  $O(m \log(\frac{1}{\sigma}))$ .

**Summary.** Tree-based anomaly detector ensembles have several beneficial properties for human-in-the-loop learning for anomaly discovery:

- Uniform prior over weights helps in improving the label-efficiency of discovering true anomalies. Indeed, in our experiments, we demonstrate the effectiveness of uniform weights over random weights for initialization in discovering anomalies. We also show the histogram distribution of the angles between score vectors from IFOR and uniform weights  $\mathbf{w}_{unif}$  to demonstrate that anomalies are aligned closer to  $\mathbf{w}_{unif}$ .
- With uniform prior over weights, the greedy strategy of querying the labels for top ranked instances is efficient, and is therefore a good yardstick for evaluating the performance of other querying strategies as well. This point will be particularly significant when we evaluate a different querying strategy to enhance the diversity of discovered anomalies as part of this work.
- Learning the weights of anomaly detector ensembles with human-in-the-loop learning that generalizes to unseen data helps in limited-memory or streaming data settings. Indeed, our experimental evaluation corroborates this hypothesis.

## 5. Algorithmic Instantiation of HiLAD for Batch Data Setting

In this section, we describe a series of algorithms to instantiate the generic human-in-the-loop learning framework HiLAD for batch data setting (HiLAD-Batch). First, we present a novel formalism called *compact description* that describes groups of instances compactly

using a tree-based model, and apply it to improve the diversity of instances selected for labeling (Section 5.1). Second, we describe an algorithm to update the weights of the scoring function based on label feedback in the *batch* setting, where the entire data is available at the outset (Section 5.2).

### 5.1 Compact Description for Diversified Querying and Interpretability

In this section, we first describe the compact description formalism to describe a group of instances. Subsequently, we propose algorithms for selecting diverse instances for querying and to generate succinct interpretable rules using compact description.

**Compact Description (CD).** The tree-based model assigns a weight and an anomaly score to each leaf (i.e., subspace). We denote the vector of leaf-level anomaly scores by  $\mathbf{d}$ , and the overall anomaly scores of the subspaces (corresponding to the leaf-nodes) by  $\mathbf{a} = [a_1, \dots, a_m] = \mathbf{w} \circ \mathbf{d}$ , where  $\circ$  denotes element-wise product operation. The score  $a_i$  provides a good measure of the *relevance* of the  $i$ -th subspace. This relevance for each subspace is determined automatically through the label feedback. Our goal is to select a small subset of the most relevant and “compact” (by volume) subspaces which together contain all the instances in a group that we want to describe. We treat this problem as a specific instance of the *set covering* problem. We illustrate this idea on a synthetic dataset in Figure 7. This approach can be potentially interpreted as a form of non-parametric clustering.




Figure 7: Top 30 subspaces ranked by  $\mathbf{w} \circ \mathbf{d}$  (relevance). Red points are anomalies. (a) shows the top 30 most relevant subspaces (w.r.t their *anomalousness*) without any feedback. We can see that initially these subspaces simply correspond to the exterior regions of the dataset. Our human-in-the-loop anomaly detection approach **learns the true relevance of subspaces** via label feedback. (b) shows that after incorporating the labels of 35 instances, the subspaces around the labeled anomalies have emerged as the most relevant. (c) shows the set of **important** subspaces which compactly cover all labeled anomalies. These were computed by solving Equation 1. Note that the compact subspaces only cover anomalies that were discovered in the 35 feedback iterations. Anomalies which were not detected are likely to fall outside these compact subspaces.

Let  $\mathcal{Z}$  be the set of instances that we want to describe, where  $|\mathcal{Z}| = p$ . For example,  $\mathcal{Z}$  could correspond to the set of anomalous instances discovered by our human-in-the-loop learning approach. Let  $s_i$  be the  $\delta$  most relevant subspaces (i.e., leaf nodes) which contain  $\mathbf{z}_i \in \mathcal{Z}, i = 1, \dots, p$ . Let  $\mathcal{S} = s_1 \cup \dots \cup s_p$  and  $|\mathcal{S}| = k$ . Denote the *volumes* of the subspaces in  $\mathcal{S}$  by the vector  $\mathbf{v} \in \mathbb{R}^k$ . Suppose  $\mathbf{x} \in \{0, 1\}^k$  is a binary vector which contains 1 in

locations corresponding to the subspaces in  $\mathcal{S}$  which are included in the covering set, and 0 otherwise. Let  $\mathbf{u}_i \in \{0, 1\}^k$  denote a vector for each instance  $\mathbf{z}_i \in \mathcal{Z}$  which contains 1 in all locations corresponding to subspaces in  $s_i$ . Let  $\mathbf{U} = [\mathbf{u}_1^T, \dots, \mathbf{u}_n^T]^T$ . A compact set of subspaces  $\mathbf{S}^*$  which contains (i.e., describes) all the candidate instances can be computed using the optimization formulation in Equation 1. We employ an off-the-shelf ILP solver (CVX-OPT) to solve this problem.

$$\begin{aligned} \mathbf{S}^* &= \arg \min_{\mathbf{x} \in \{0, 1\}^k} \mathbf{x} \cdot \mathbf{v} \\ \text{s.t. } &\mathbf{U} \cdot \mathbf{x} \geq \mathbf{1} \quad (\text{where } \mathbf{1} \text{ is a column vector of } p \text{ 1's}) \end{aligned} \tag{1}$$

*Applications of Compact Description.* Compact descriptions have multiple uses including:

- Discovery of diverse classes of anomalies very quickly by querying instances from different subspaces of the description.
- Improved interpretability of anomalous instances. We assume that in a practical setting, the analyst(s) will be presented with instances along with their corresponding description(s). Additional information can be derived from the description and shown to the analyst (e.g., number of instances in each compact subspace), which can help prioritize the analysis.

In this work, we present empirical results on improving query diversity and also compare with another state-of-the-art algorithm that extracts interpretations (Wang, Rudin, Doshi-Velez, Liu, Klampfl, & MacNeille, 2016).

---

**Algorithm 2** Select-Diverse ( $\mathbf{X}, b, n$ )

---

**Input:** Unlabeled dataset  $\mathbf{X}$ , # instances to select  $b$ , # candidate instances  $n$  ( $n \geq b$ )  
 Let  $\mathcal{Z} = n$  top-ranked instances as candidates  $\subseteq \mathbf{X}$  (blue points in Figure 8a)  
 Let  $\mathbf{S}^*$  = subspaces with Equation 1 that contain  $\mathcal{Z}$  (rectangles in Figures 8b and 8c)  
 Set  $\mathbf{Q} = \emptyset$   
**while**  $|\mathbf{Q}| < b$  **do**  
   Let  $\mathbf{x}$  = instance with highest anomaly score  $\in \mathcal{Z}$  s.t.  $\mathbf{x}$  has minimal  
     overlapping regions in  $\mathbf{S}^*$  with instances in  $\mathbf{Q}$   
   Set  $\mathbf{Q} = \mathbf{Q} \cup \{\mathbf{x}\}$  (green circles in Figure 8c)  
   Set  $\mathcal{Z} = \mathcal{Z} \setminus \{\mathbf{x}\}$   
**end while**  
**return**  $\mathbf{Q}$

---

**Diversity-based Query Selection Strategy.** In Section 4, we reasoned that the greedy strategy of selecting top-scored instances (referred as **Select-Top**) for labeling is efficient. However, this strategy might lack diversity in the types of instances presented to the human analyst. It is likely that different types of instances belong to different subspaces in the original feature space. Our proposed strategy (**Select-Diverse**), which is described next, is intended to increase the diversity by employing tree-based ensembles and compact description to select groups of instances from subspaces that have minimum overlap.

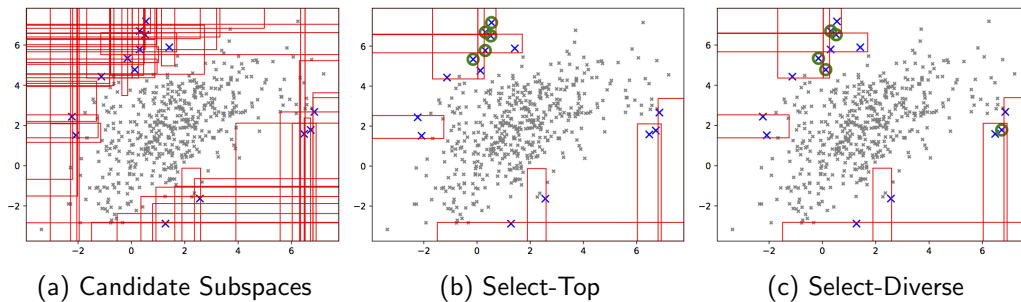


Figure 8: Illustration of compact description and diversity of selected instances for labeling using IFOR. Most anomalous 15 instances (blue checks) are selected as the query candidates. The red rectangles in (a) form the union of the  $\delta$  ( $= 5$  works well in practice) most *relevant* subspaces across each of the query candidates. (b) and (c) show the most “compact” set of subspaces which together cover all the query candidates. (b) shows the most anomalous 5 instances (green circles) selected by the greedy **Select-Top** strategy. (c) shows the 5 “diverse” instances (green circles) selected by **Select-Diverse**.

Assume that the human analyst can label a batch of  $b$  instances, where  $b > 1$ , in each feedback iteration. Algorithm 2 employs the compact description to achieve this diversity. The algorithm first selects  $n$  ( $> b$ ) top-ranked anomalous instances  $\mathcal{Z}$  and computes the corresponding compact description (small set of subspaces  $\mathbf{S}^*$ ). Subsequently, performs an iterative selection of  $b$  instances from  $\mathcal{Z}$  by minimizing the overlap in the corresponding subspaces from  $\mathbf{S}^*$ . Figure 8 provides an illustration comparing **Select-Top** and **Select-Diverse** query selection strategies.

## 5.2 Algorithmic Approach to Update Weights of Scoring Function

In this section, we provide an algorithm to update weights of the scoring function for HiLAD-Batch instantiation for batch setting: the entire data  $\mathbf{D}$  is available at the outset.

Recall that our scoring function is of the following form:  $\text{Score}(\mathbf{x}) = \mathbf{w} \cdot \mathbf{z}$ , where  $\mathbf{z} \in \mathbb{R}^m$  corresponds to the scores from anomaly detectors for instance  $\mathbf{x}$ . We extend the AAD approach (based on LODA projections) (Das et al., 2016) to update the weights for tree-based models. AAD makes the following assumptions: (1)  $\tau$  fraction of instances (i.e.,  $n\tau$ ) are anomalous, and (2) Anomalies should lie above the optimal hyperplane while nominals should lie below. AAD tries to satisfy these assumptions by enforcing constraints on the labeled examples while learning the weights of the hyperplane. If the anomalies are rare and we set  $\tau$  to a small value, then the two assumptions make it more likely that the hyperplane will pass through the region of uncertainty. Our previous discussion then suggests that the optimal hyperplane can now be learned efficiently by greedily asking the analyst to label the most anomalous instance in each feedback iteration. We simplify the AAD formulation with a more scalable unconstrained optimization objective. Crucially, the ensemble weights are updated with an intent to maintain the hyperplane in the region of uncertainty through the entire budget  $B$ . The HiLAD-Batch learning approach is presented in Algorithm 3 and depends on only one hyper-parameter  $\tau$ .



**Algorithm 3** HiLAD-Batch ( $B, \mathbf{w}^{(0)}, \mathbf{H}, \mathbf{H}_+, \mathbf{H}_-$ )

---

**Input:** Query budget  $B$ , initial weights  $\mathbf{w}^{(0)}$ , unlabeled instances  $\mathbf{H}$ , labeled instances  $\mathbf{H}_+$  and  $\mathbf{H}_-$

Set  $t = 0$

**while**  $t \leq B$  **do**

  Set  $t = t + 1$

  Set  $\mathbf{a} = \mathbf{H} \cdot \mathbf{w}$  (i.e.,  $\mathbf{a}$  is the vector of anomaly scores)

  Let  $\mathbf{q} = \mathbf{z}_i$ , where  $i = \arg \max_i (a_i)$

  Get  $y_i \in \{-1, +1\}$  for  $\mathbf{q}$  from analyst

**if**  $y_i = +1$  **then**

    Set  $\mathbf{H}_+ = \{\mathbf{z}_i\} \cup \mathbf{H}_+$

**else**

    Set  $\mathbf{H}_- = \{\mathbf{z}_i\} \cup \mathbf{H}_-$

**end if**

14: Set  $\mathbf{H} = \mathbf{H} \setminus \mathbf{z}_i$

$\mathbf{w}^{(t)} = \text{learn new weights; normalize } \|\mathbf{w}^{(t)}\| = 1$

**end while**

**return**  $\mathbf{w}^{(t)}, \mathbf{H}, \mathbf{H}_+, \mathbf{H}_-$

---

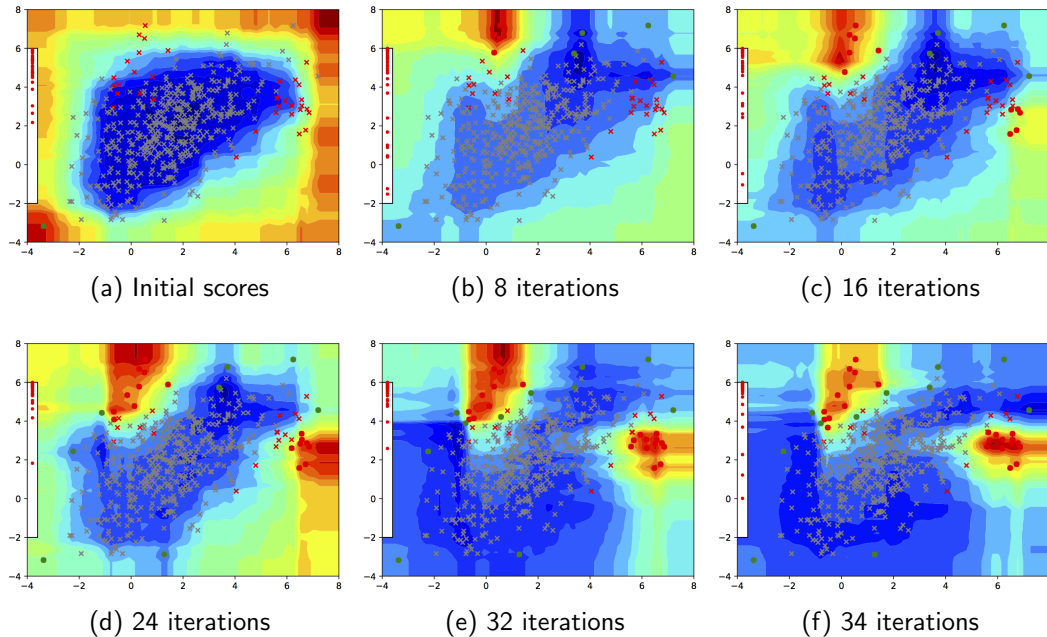


Figure 9: Score contours across 34 HiLAD-Batch feedback iterations on the Toy dataset (Figure 6a).

We first define the hinge loss  $\ell(q, \mathbf{w}; (\mathbf{z}_i, y_i))$  in Equation 2 that penalizes the model when anomalies are assigned scores lower than  $q$  and nominals higher. Equation 3 then formulates the optimization problem for learning the optimal weights in Line 14 of the batch algorithm

(Algorithm 3). Figure 9 illustrates how HiLAD-Batch changes the anomaly score contours across feedback iterations on the synthetic dataset using an isolation forest with 100 trees.

$$\ell(q, \mathbf{w}; (\mathbf{z}_i, y_i)) = \begin{cases} 0 & \mathbf{w} \cdot \mathbf{z}_i \geq q \text{ and } y_i = +1 \\ 0 & \mathbf{w} \cdot \mathbf{z}_i < q \text{ and } y_i = -1 \\ (q - \mathbf{w} \cdot \mathbf{z}_i) & \mathbf{w} \cdot \mathbf{z}_i < q \text{ and } y_i = +1 \\ (\mathbf{w} \cdot \mathbf{z}_i - q) & \mathbf{w} \cdot \mathbf{z}_i \geq q \text{ and } y_i = -1 \end{cases} \quad (2)$$

$\lambda^{(t)}$  determines the influence of the prior. For the batch setup, we set  $\lambda^{(t)} = \frac{0.5}{|\mathbf{H}_+| + |\mathbf{H}_-|}$  such that the prior becomes less important as more instances are labeled. When there are no labeled instances,  $\lambda^{(t)}$  is set to  $\frac{1}{2}$ . The third and fourth terms of Equation 3 encourage the scores of anomalies in  $\mathbf{H}_+$  to be higher than that of  $\mathbf{z}_\tau^{(t-1)}$  (the  $n\tau$ -th ranked instance from the previous iteration), and the scores of nominals in  $\mathbf{H}_-$  to be lower than that of  $\mathbf{z}_\tau^{(t-1)}$ . We employ gradient descent to learn the optimal weights  $\mathbf{w}$  in Equation 3. Our prior knowledge that  $\mathbf{w}_{unif}$  is a good prior provides a good initialization for gradient descent. We later show empirically that  $\mathbf{w}_{unif}$  is a better starting point than random weights.

$$\begin{aligned} \mathbf{w}^{(t)} = \arg \min_{\mathbf{w}} \sum_{s \in \{-, +\}} & \left( \frac{1}{|\mathbf{H}_s|} \sum_{\mathbf{z}_i \in \mathbf{H}_s} \ell(\hat{q}_\tau(\mathbf{w}^{(t-1)}), \mathbf{w}; (\mathbf{z}_i, y_i)) \right. \\ & \left. + \frac{1}{|\mathbf{H}_s|} \sum_{\mathbf{z}_i \in \mathbf{H}_s} \ell(\mathbf{z}_\tau^{(t-1)} \cdot \mathbf{w}, \mathbf{w}; (\mathbf{z}_i, y_i)) \right) \\ & + \lambda^{(t)} \|\mathbf{w} - \mathbf{w}_{unif}\|^2 \end{aligned} \quad (3)$$

where,  $\mathbf{w}_{unif} = [\frac{1}{\sqrt{m}}, \dots, \frac{1}{\sqrt{m}}]^T$ , and,

$\mathbf{z}_\tau^{(t-1)}$  and  $\hat{q}_\tau(\mathbf{w}^{(t-1)})$  are computed by ranking anomaly scores with  $\mathbf{w} = \mathbf{w}^{(t-1)}$

## 6. Algorithmic Instantiation of HiLAD for Streaming Data Setting

In this section, we describe algorithms to support human-in-the-loop anomaly detection using tree-based ensembles in the *streaming* data setting (HiLAD-Stream), where the data comes as a continuous stream.

In the streaming setting, we assume that the data is input to the algorithm continuously in *windows* of size  $K$  and is potentially unlimited. The HiLAD-Stream instantiation is shown in Algorithm 4. Initially, we train all the members of the ensemble with the first window of data. When a new window of data arrives, the underlying tree model is updated as follows: in case the model is an HST or RSF, only the node counts are updated while keeping the tree structures and weights unchanged; whereas, if the model is an IFOR, a subset of the current set of trees is replaced as shown in `Update-Model` (Algorithm 4). The updated

---

**Algorithm 4** HiLAD-Stream ( $K, B, Q, \mathcal{E}^{(0)}, \mathbf{X}_0, \mathbf{w}^{(0)}, \alpha_{KL}$ )

---

**Input:** Stream window size  $K$ , total query budget  $B$ ,  
queries per window  $Q$ , anomaly detector ensemble  $\mathcal{E}^{(0)}$ ,  
initial instances  $\mathbf{X}_0$  (used to create  $\mathcal{E}^{(0)}$ ), initial weights  $\mathbf{w}^{(0)}$ ,  
significance level  $\alpha_{KL}$

Set  $\mathbf{H} = \mathbf{H}_+ = \mathbf{H}_- = \emptyset$   
// initialize KL-divergence baselines  
Set  $q_{KL}^{(0)} = \text{Get-KL-Threshold}(\mathbf{X}_0, \mathcal{E}^{(0)}, \alpha_{KL}, 10)$   
Set  $\mathcal{P}^{(0)} = \text{Get-Ensemble-Distribution}(\mathbf{X}_0, \mathcal{E}^{(0)})$

Set  $t = 0$   
**while** <stream is not empty> **do**  
  Set  $t = t + 1$   
  Set  $\mathbf{X}_t = K$  data instances from stream  
  Set  $\mathbf{H}_t = \text{transform } \mathbf{X}_t \text{ to ensemble features}$

  // **Update-Model** either updates node counts (e.g., for HST and RSF),  
  // or replaces a fraction of the oldest trees in  $\mathcal{E}$  with new  
  // ones constructed using  $\mathbf{X}_t$  (e.g., for IFOR)  
  Set  $\mathcal{E}^{(t)}, q_{KL}^{(t)}, \mathcal{P}^{(t)} = \text{Update-Model}(\mathbf{X}_t, \mathcal{E}^{(t-1)}, q_{KL}^{(t-1)}, \mathcal{P}^{(t-1)}, \alpha_{KL})$

  // **Merge-and-Retain**( $\mathbf{w}, \mathbf{H}, K$ ) retains  $K$  most anomalous instances in  $\mathbf{H}$   
  Set  $\mathbf{H} = \text{Merge-and-Retain}(\mathbf{w}^{(t-1)}, \{\mathbf{H} \cup \mathbf{H}_t\}, K)$   
  Set  $\mathbf{w}^{(t)}, \mathbf{H}, \mathbf{H}_+, \mathbf{H}_- = \text{HiLAD-Batch}(Q, \mathbf{w}^{(t-1)}, \mathbf{H}, \mathbf{H}_+, \mathbf{H}_-)$

**end while**

---

model is then employed to determine which unlabeled instances to retain in memory, and which to “forget”. This step, referred to as **Merge-and-Retain**, applies the simple strategy of retaining only the most anomalous instances among those in the memory and in the current window, and discarding the rest. Next, the weights are fine-tuned with analyst feedback through human-in-the-loop learning loop similar to the batch setting with a small budget  $Q$ . Finally, the next window of data is read, and the process is repeated until the stream is empty or the total budget  $B$  is exhausted. *In the rest of this section, we will assume that the underlying tree model is IFOR.* When we replace a tree in **Update-Model**, its leaf nodes and corresponding weights get discarded. On the other hand, adding a new tree implies adding all its leaf nodes with weights initialized to a default value  $v$ . We first set  $v = \frac{1}{\sqrt{m'}}$  where  $m'$  is the total number of leaves in the new model, and then re-normalize the updated  $\mathbf{w}$  to unit length. The HiLAD instantiation for streaming data is presented in Algorithm 4. In all the HiLAD-Stream experiments, we set the number of queries per window  $Q = 20$ , and  $\lambda^{(t)} = \frac{1}{2}$ .

HiLAD-Stream approach can be employed in two different situations: (1) **limited memory with no concept drift**, and (2) **streaming data with concept drift**. The type of situation determines how fast the model needs to be updated in **Update-Model**. If there

---

**Algorithm 5** Update-Model ( $\mathbf{X}, \mathcal{E}, q_{KL}, \mathcal{P}, \alpha_{KL}$ )

---

**Input:** Instances  $\mathbf{X}$ , anomaly detector ensemble  $\mathcal{E}$ ,  
current KL threshold  $q_{KL}$ , baseline distributions  $\mathcal{P}$ ,  
significance level  $\alpha_{KL}$   
Set  $T$  = number of trees in  $\mathcal{E}$   
Set  $\mathcal{Q} = \text{Get-Ensemble-Distribution}(\mathbf{X}, \mathcal{E})$   
Initialize  $\mathbf{KL} = \mathbf{0} \in \mathbb{R}^T$   
**for**  $t \in 1..T$  **do**  
    Set  $KL_t = D_{KL}(\mathcal{P}_t || \mathcal{Q}_t)$   
**end for**  
Set  $S = \{t : KL_t > q_{KL}\}$   
**if**  $|S| < 2\alpha_{KL}T$  **then**  
    // the number of trees with divergence is not significant  
    **return**  $\mathcal{E}, q_{KL}, \mathcal{P}$   
**end if**  
Set  $\mathcal{E}' =$  replace trees in  $\mathcal{E}$  whose indexes are in  $S$ , with new trees trained using  $\mathbf{X}$   
// Recompute threshold and baseline distributions  
Set  $q'_{KL} = \text{Get-KL-Threshold}(\mathbf{X}, \mathcal{E}', \alpha_{KL}, 10)$   
Set  $\mathcal{P}' = \text{Get-Ensemble-Distribution}(\mathbf{X}, \mathcal{E}')$   
**return**  $\mathcal{E}', q'_{KL}, \mathcal{P}'$

---

is no concept drift, we need not update the model at all. If there is a large change in the distribution of data from one window to the next, then a large fraction of members need to be replaced. When we replace a member tree in our tree-based model, all its corresponding nodes along with their learned weights have to be discarded. Thus, some of the “knowledge” is lost with the model update. In general, it is hard to determine the true rate of drift in the data. One approach is to replace, in the **Update-Model** step, a reasonable number (e.g. 20%) of older ensemble members with new members trained on new data. Although this ad hoc approach often works well in practice, a more principled approach is preferable.

**Drift Detection Algorithm.** Algorithm 5 presents a principled methodology that employs KL-divergence (denoted by  $D_{KL}$ ) to determine which trees should be replaced. The set of all leaf nodes in a tree are treated as a set of histogram bins which are then used to estimate the data distribution. We denote the total number of trees in the model by  $T$ , and the  $t$ -th tree by  $\mathcal{T}_t$ . When  $\mathcal{T}_t$  is initially created with the first window of data, the data from the same window is also used to initialize the *baseline* distribution for  $\mathcal{T}_t$ , denoted by  $\mathbf{p}_t$  (**Get-Ensemble-Distribution** in Algorithm 7). After computing the baseline distributions for each tree, we estimate the  $D_{KL}$  threshold  $q_{KL}$  at the  $\alpha_{KL}$  (typically 0.05) significance level by sub-sampling (**Get-KL-Threshold** in Algorithm 6). When a new window is read, we first use it to compute the new distribution  $\mathbf{q}_t$  (for  $\mathcal{T}_t$ ). Next, if  $\mathbf{q}_t$  differs *significantly* from  $\mathbf{p}_t$  (i.e.,  $D_{KL}(\mathbf{p}_t || \mathbf{q}_t) > q_{KL}$ ) for at least  $2\alpha_{KL}T$  trees, then we replace all such trees with new ones created using the data from the new window. Finally, if any tree in the forest is replaced, then the baseline densities for all trees are recomputed with the data in the new window.

---

**Algorithm 6** Get-KL-Threshold ( $\mathbf{X}, \mathcal{E}, \alpha_{KL}, n$ )

---

**Input:** Instances  $\mathbf{X}$ , anomaly detector ensemble  $\mathcal{E}$ ,  
significance level  $\alpha_{KL}$ ,  
repetitions of KL-divergence computations  $n$   
Set  $T =$  number of trees in  $\mathcal{E}$   
Initialize  $\mathbf{KL} = \mathbf{0} \in \mathbb{R}^T$  // mean KL-divergence for each tree  
**for**  $i$  in  $1 \cdots n$  **do**  
  Partition  $\mathbf{X} = \{\mathbf{A}, \mathbf{B}\}$  at random s.t.  $\mathbf{X} = \mathbf{A} \cup \mathbf{B}$  and  $|\mathbf{A}| \approx |\mathbf{B}|$   
  **for**  $t \in 1 \cdots T$  **do**  
    Let  $\mathcal{T}_t = t$ -th tree in  $\mathcal{E}$   
    Set  $\mathbf{p}_A = \text{Get-Tree-Distribution}(\mathbf{A}, \mathcal{T}_t)$   
    Set  $\mathbf{p}_B = \text{Get-Tree-Distribution}(\mathbf{B}, \mathcal{T}_t)$   
    Set  $KL_t = KL_t + D_{KL}(\mathbf{p}_A || \mathbf{p}_B)$   
  **end for**  
**end for**  
Set  $\mathbf{KL} = \frac{\mathbf{KL}}{n}$  // average the values  
Set  $q_{KL} = (1 - \alpha_{KL}) \times 100$  quantile value in  $\mathbf{KL}$   
**return**  $q_{KL}$

---



---

**Algorithm 7** Get-Ensemble-Distribution ( $\mathbf{X}, \mathcal{E}$ )

---

**Input:** Instances  $\mathbf{X}$ , ensemble  $\mathcal{E}$   
Set  $T =$  number of trees in  $\mathcal{E}$   
Set  $\mathcal{P} = \emptyset$   
**for**  $t \in 1 \cdots T$  **do**  
  Let  $\mathcal{T}_t = t$ -th tree in  $\mathcal{E}$   
  Set  $\mathbf{p}_t = \text{Get-Tree-Distribution}(\mathbf{X}, \mathcal{T}_t)$   
  Set  $\mathcal{P} = \mathcal{P} \cup \mathbf{p}_t$   
**end for**  
**return**  $\mathcal{P}$

---



---

**Algorithm 8** Get-Tree-Distribution ( $\mathbf{X}, \mathcal{T}$ )

---

**Input:** Instances  $\mathbf{X}$ , tree  $\mathcal{T}$   
Set  $\mathbf{p} =$  distribution of instances in  $\mathbf{X}$  at the leaves of  $\mathcal{T}$   
**return**  $\mathbf{p}$

---

## 7. Experiments and Results

**Datasets.** We evaluate our human-in-the-loop learning framework on ten publicly available benchmark datasets ((Woods, Doss, Bowyer, Solka, Priebe, & Kegelmeyer, 1993),(Ditzler & Polikar, 2013),(Harries & of New South Wales., 1999), UCI(Dheeru & Karra Taniskidou, 2017)) listed in Table 2. The anomaly classes in *Electricity* and *Weather* were down-sampled to be 5% of the total.

Table 2: Description of benchmark datasets used in our experiments.

Dataset	Nominal Class	Anomaly Class	Total	Dims	Anomalies(%)
Abalone	8, 9, 10	3, 21	1920	9	29 (1.5%)
ANN-Thyroid-1v3	3	1	3251	21	73 (2.25%)
Cardiotocography	1 (Normal)	3 (Pathological)	1700	22	45 (2.65%)
KDD-Cup-99	<i>'normal'</i>	<i>'u2r', 'probe'</i>	63009	91	2416 (3.83%)
Mammography	-1	+1	11183	6	260 (2.32%)
Shuttle	1	2, 3, 5, 6, 7	12345	9	867 (7.02%)
Yeast	<i>CYT, NUC, MIT</i>	<i>ERL, POX, VAC</i>	1191	8	55 (4.6%)
Covtype	2	4	286048	54	2747 (0.9%)
Electricity	<i>DOWN</i>	<i>UP</i>	27447	13	1372 (5%)
Weather	<i>No Rain</i>	<i>Rain</i>	13117	8	656 (5%)

**Evaluation Methodology.** For each variant of the human-in-the-loop framework HiLAD, we plot the percentage of the total number of anomalies shown to the analyst versus the number of instances queried; this is the most relevant metric for an analyst in any real-world application. A higher plot means the algorithmic instantiation of the framework is better in terms of discovering anomalies. All results presented are averaged over 10 different runs and the error-bars represent 95% confidence intervals.

## 7.1 Human-in-the-Loop Learning for Batch Data Setting

**Experimental Setup.** All versions of HiLAD-Batch employ IFOR with the number of trees  $T = 100$  and subsample size 256. The initial starting weights are denoted by  $\mathbf{w}^{(0)}$ . We normalize the score vector for each instance to unit length such that the score vectors lie on a unit sphere. This normalization helps adhere to the discussion in Section 4, but is otherwise unnecessary. Figure 10 shows that  $\mathbf{w}_{unif}$  tends to have a smaller angular separation from the normalized IFOR score vectors of anomalies than from those of nominals. This holds true for most of our datasets (Table 2). *Weather* is a hard dataset for all anomaly detectors (Wu et al., 2014), as reflected in its angular distribution in Figure 10i. In all our experiments, **Unsupervised Baseline** shows the number of anomalies detected without any feedback, i.e., using the uniform ensemble weights  $\mathbf{w}_{unif}$ ; **HiLAD-Batch (No Prior - Unif)** and **HiLAD-Batch (No Prior - Rand)** impose no priors on the model, and start human-in-the-loop learning with  $\mathbf{w}^{(0)}$  set to  $\mathbf{w}_{unif}$  and a random vector respectively; **HiLAD-Batch** sets  $\mathbf{w}_{unif}$  as prior, and starts with  $\mathbf{w}^{(0)} = \mathbf{w}_{unif}$ . For HST, we present two sets of results with batch input only: **HST-Batch** with original settings ( $T = 25$ , depth=15, no feedback) (Tan et al., 2011), and **HST-Batch (Feedback)** which supports feedback with HiLAD-Batch strategy (with  $T = 50$  and depth=8, a better setting for feedback). For RST, we present the results (**RST-Batch**) with only the original settings ( $T = 30$ , depth=15) (Wu et al., 2014) since it was not competitive with other methods on our datasets. We also compare the HiLAD-Batch variants with the AAD approach (Das et al., 2016) in the batch setting (**AAD-Batch**). Batch data setting is the most optimistic for all algorithms.




Figure 10: Histogram distribution of the angles between score vectors from IFOR and  $\mathbf{w}_{unif}$  for all datasets. The red and green histograms show the angle distributions for anomalies and nominals respectively. Since the red histograms are closer to the left, anomalies are aligned closer to  $\mathbf{w}_{unif}$ .

### 7.1.1 RESULTS FOR HiLAD-BATCH INSTANTIATION

We set the budget  $B$  to 300 for all datasets in the batch setting. The results on the four smaller datasets *Abalone*, *ANN-Thyroid-1v3*, *Cardiotocography*, and *Yeast* are shown in Figure 11. When the algorithm starts from sub-optimal initialization of the weights and with no prior knowledge (HiLAD-Batch (No Prior - Rand)), more number of queries are spent hunting for the first few anomalies, and thereafter detection improves significantly. When the weights are initialized to  $\mathbf{w}_{unif}$ , which is a reliable starting point (HiLAD-Batch (No Prior - Unif) and HiLAD-Batch), fewer queries are required to find the initial anomalies, and typically results in a lower variance in accuracy. Setting  $\mathbf{w}_{unif}$  as prior in addition to informed initialization (HiLAD-Batch) performs better than without the prior (HiLAD-Batch (No Prior - Unif)) on *Abalone*, *ANN-Thyroid-1v3*, and *Yeast*. We believe this is because the prior helps guard against noise.

Figure 12 shows the relative performance (% of anomalies discovered) compared to the **Unsupervised Baseline**. It is clear that HiLAD-Batch discovered up to 250% more anomalies than the baseline algorithm. A key observation to note is that HiLAD-Batch discovers the biggest portion of anomalies in the early phase of the label collection. For *Abalone* dataset

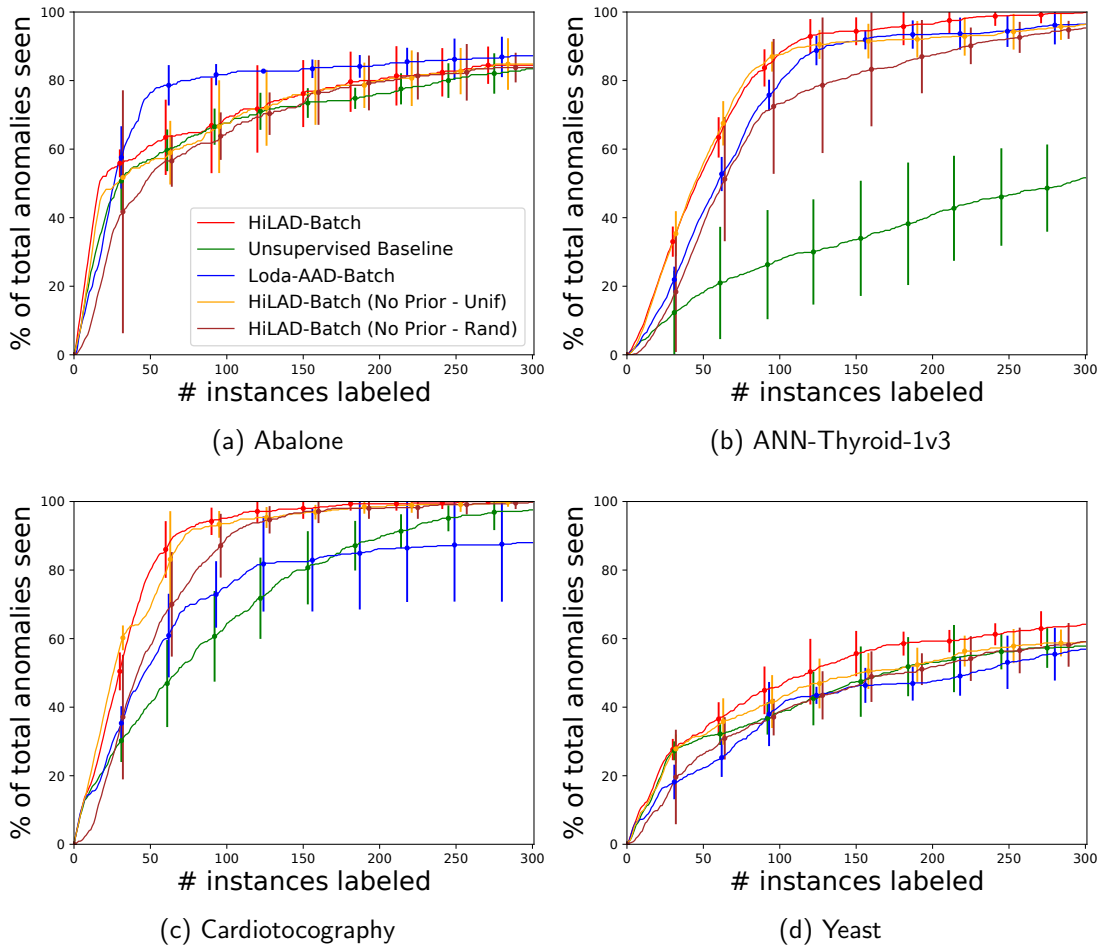


Figure 11: Percentage of total anomalies seen vs. the number of queries for the **smaller** datasets in the **batch** setting.

HiLAD-Batch discovered 40-60% more anomalies in the initial phase whereas Loda-AAD-Batch was under-performing initially and recovered later. The biggest gain we saw was for *Thyroid* dataset where HiLAD-Batch discovered 250% more anomalies compared to the baseline. Other datasets (*Cardiotocography*, *Yeast*) also resulted in higher anomaly discovery using the HiLAD-Batch method.

### 7.1.2 RESULTS FOR DIVERSIFIED QUERY STRATEGY

The diversified querying strategy **Select-Diverse** (Algorithm 2) employs compact descriptions to select instances. Therefore, the evaluation of its effectiveness is presented first. An interactive system can potentially ease the cognitive burden on the analysts by using descriptions to generate a “summary” of the anomalous instances.

We perform a *post hoc* analysis on the datasets with the knowledge of the original classes (Table 2). It is assumed that each class in a dataset represents a different data-generating process. To measure the diversity at any point in our feedback cycle, we compute the



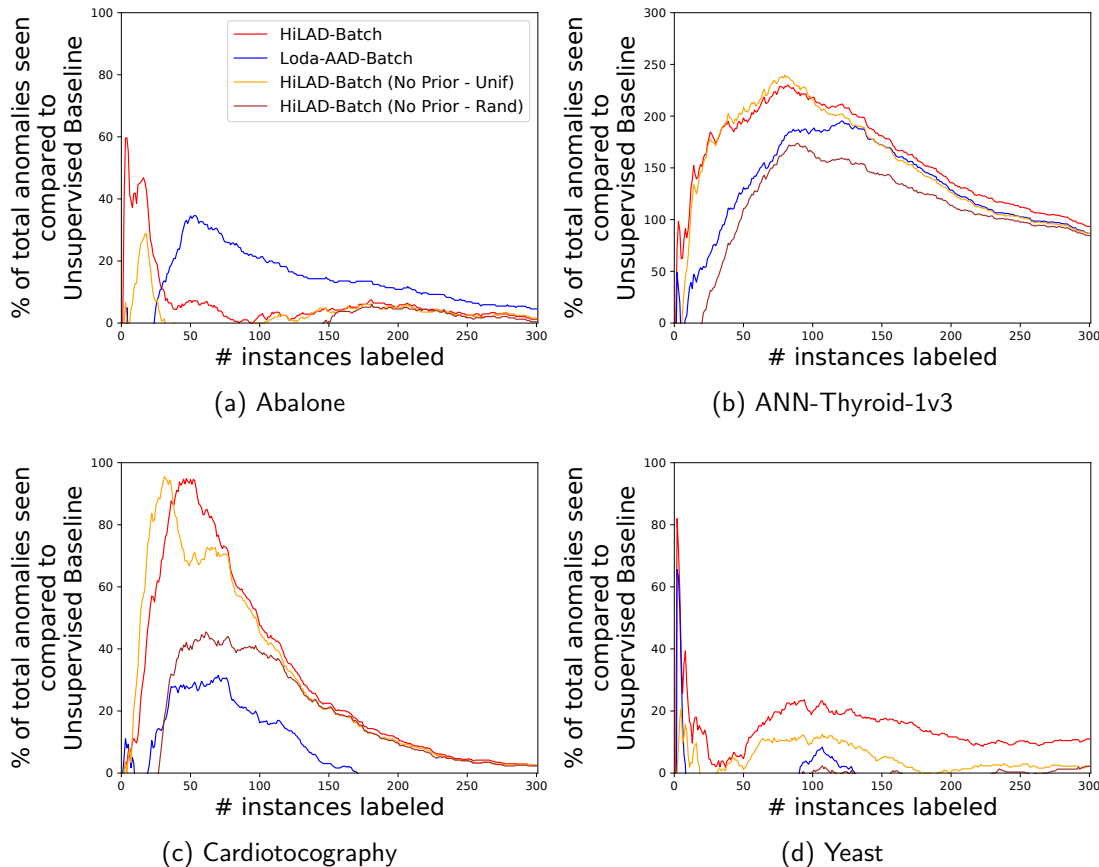


Figure 12: Relative % of anomalies seen compared to the *Unsupervised Baseline* for **smaller** datasets in the **batch** setting.

difference between the number of unique classes presented to the analyst per query batch averaged across all the past batches. The parameter  $\delta$  for **Select-Diverse** was set to 5 in all experiments. We compare three query strategies in the batch data setup: **HiLAD-Batch-T**, **HiLAD-Batch-D**, and **HiLAD-Batch-R**. **HiLAD-Batch-T** simply presents the top three most anomalous instances per query batch. **HiLAD-Batch-D** employs **Select-Diverse** to present three diverse instances out of the ten most anomalous instances. **HiLAD-Batch-R** presents three instances selected at random from the top ten anomalous instances. Finally, **HiLAD-Batch** greedily presents only the single most anomalous instance for labeling. We find that **HiLAD-Batch-D** presents a more diverse set of instances than both **HiLAD-Batch-T** (solid lines) as well as **HiLAD-Batch-R** (dashed lines) on most datasets. Figure 13b shows that the number of anomalies discovered (on representative datasets) with the diversified querying strategy is similar to the greedy strategy, i.e., *no loss in discovery rate to improve diversity*.

## 7.2 Human-in-the-Loop Learning for Streaming Data Setting

**Experimental Setup.** For **HiLAD-Stream** instantiation, we employ **IFOR** with the number of trees  $T = 100$  and subsample size 256. This is similar to the setup for **HiLAD-Batch**.

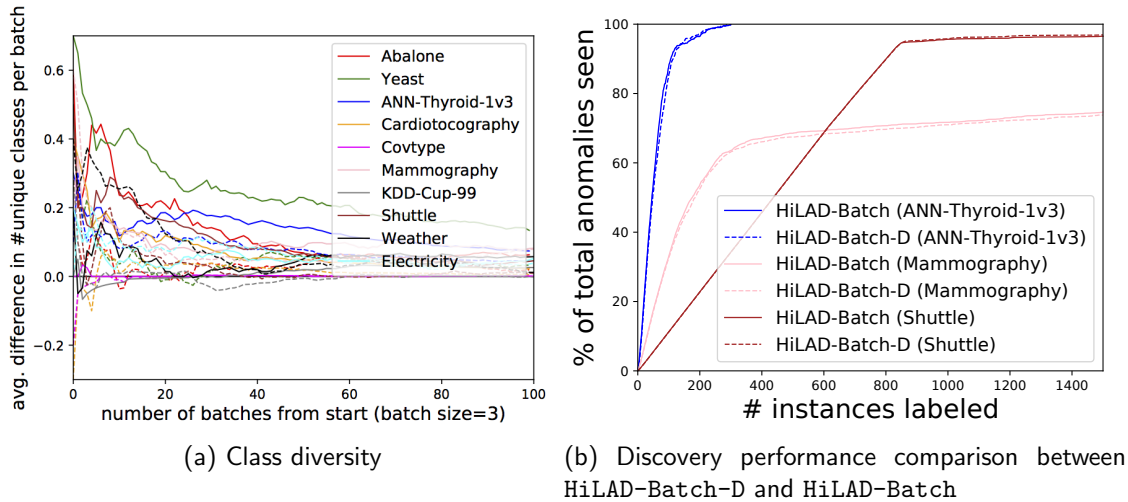


Figure 13: Results comparing diversified querying strategy HiLAD-Batch-D with baseline query strategies HiLAD-Batch-T and HiLAD-Batch-R. The x-axis in (a) shows the number of query batches (of batch size 3). The y-axis shows the difference in the number of unique classes seen averaged across all batches till a particular batch. The solid lines in (a) show the average difference between unique classes seen with HiLAD-Batch-D and HiLAD-Batch-T; the dashed lines show the average difference between HiLAD-Batch-D and HiLAD-Batch-R. (b) presents the discovery performance for HiLAD-Batch-D (dashed line). And the baseline is HiLAD-Batch (solid line)

In all HiLAD-Stream experiments, we set the number of queries per window  $Q = 20$ . The total budget  $B$  and the stream window size  $K$  for the datasets were set respectively as follows: *Covtype* (3000, 4096), *KDD-Cup-99* (3000, 4096), *Mammography* (1500, 4096), *Shuttle* (1500, 4096), *Electricity* (1500, 1024), *Weather* (1000, 1024). These values are reasonable w.r.t the dataset’s size, the number of anomalies, and the rate of concept drift. The maximum number of unlabeled instances residing in memory is  $K$ . When the last window of data arrives, then human-in-the-loop learning is continued with the final set of unlabeled data retained in the memory until the total budget  $B$  is exhausted. The instances are streamed in the same order as they appear in the original public sources. When a new window of data arrives: HiLAD-Stream (KL Adaptive) dynamically determines which trees to replace based on KL-divergence, HiLAD-Stream (Replace 20% Trees) replaces 20% oldest trees, and HiLAD-Stream (No Tree Replace) creates the trees only once with the first window of data and only updates the weights of the fixed leaf nodes using feedback.

### 7.2.1 RESULTS FOR HiLAD-STREAM INSTANTIATION

**Limited memory setting with no concept drift.** The results on the four larger datasets are shown in Figure 14. The performance is similar to what is seen on the smaller datasets. Among the unsupervised algorithms in the batch setting, IFOR (Unsupervised Baseline) and HST (HST-Batch) are competitive, and both are better than RSF (RSF-Batch). With feedback, HiLAD-Batch is consistently the best performer. HST with feedback (HST-Batch (Feedback)) always performs better than HST-Batch. The streaming algorithm with feedback, HiLAD-Stream (KL Adaptive), significantly outperforms Unsupervised Baseline

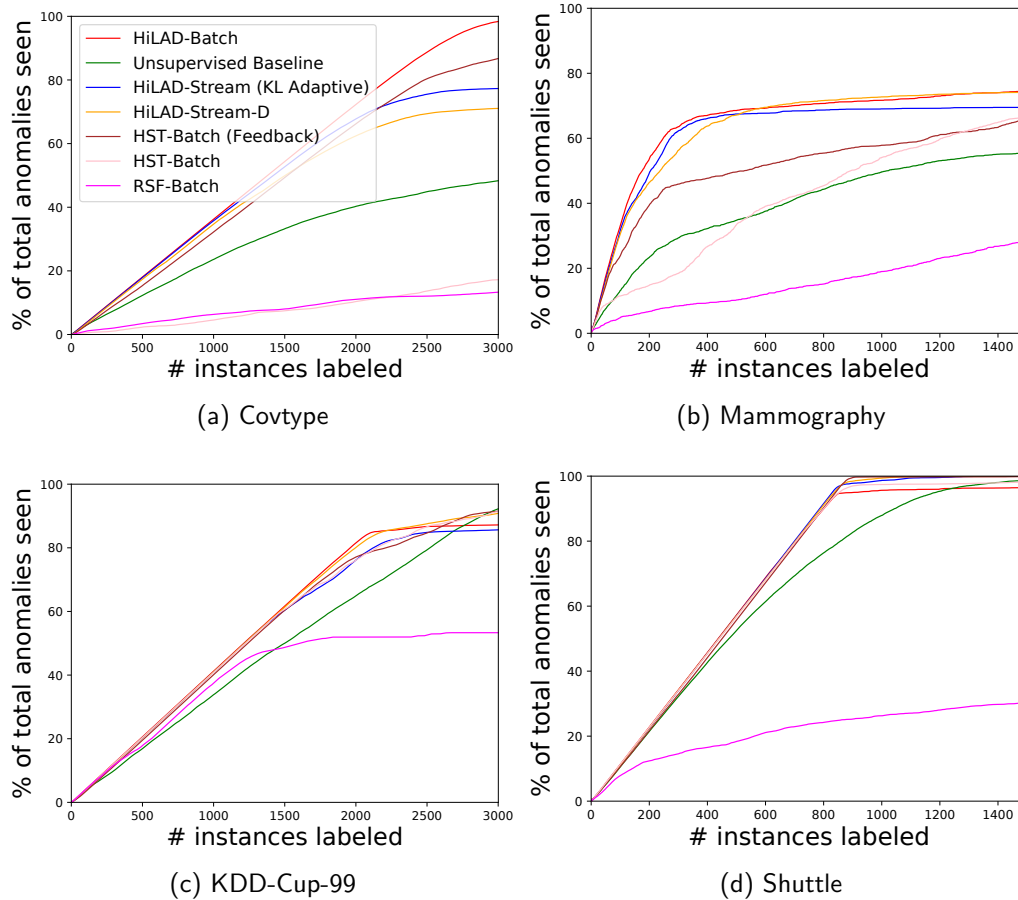


Figure 14: Percentage of total anomalies seen vs. number of queries for the **larger datasets in the limited memory setting**. HiLAD-Stream (KL Adaptive) and HiLAD-Stream-D apply the **Select-Top** and **Select-Diverse** query strategies resp. *Mammography*, *KDD-Cup-99*, and *Shuttle* have no significant drift. *Covtype*, which has a higher drift, is included here for comparison because it is large.

and is competitive with HiLAD-Batch. HiLAD-Stream (KL Adaptive) performs better than HST-Batch (Feedback) as well. HiLAD-Stream-D which presents a more diverse set of instances for labeling performs similar to HiLAD-Stream (KL Adaptive). These results demonstrate that the *feedback-tuned anomaly detectors generalize to unseen data*.

Figure 16 shows drift detection on datasets which do not have any drift. The streaming window size for each dataset was set commensurate to its total size: *Abalone*(512), *ANN-Thyroid-1v3*(512), *Cardiotocography*(512), *Covtype*(4096), *Electricity*(1024), *KDD-Cup99*(4096), *Mammography*(4096), *Shuttle*(4096), *Weather*(1024), and *Yeast*(512).

**Streaming setting with concept drift.** Figure 17 shows the results after integrating drift detection and label feedback with HiLAD-Stream for the datasets which are expected to have significant drift. Both *Covtype* and *Electricity* show more drift in the data than *Weather* (top row in Figure 17). The rate of drift determines how fast the model should be updated. If we update the model too fast, then we loose valuable knowledge gained

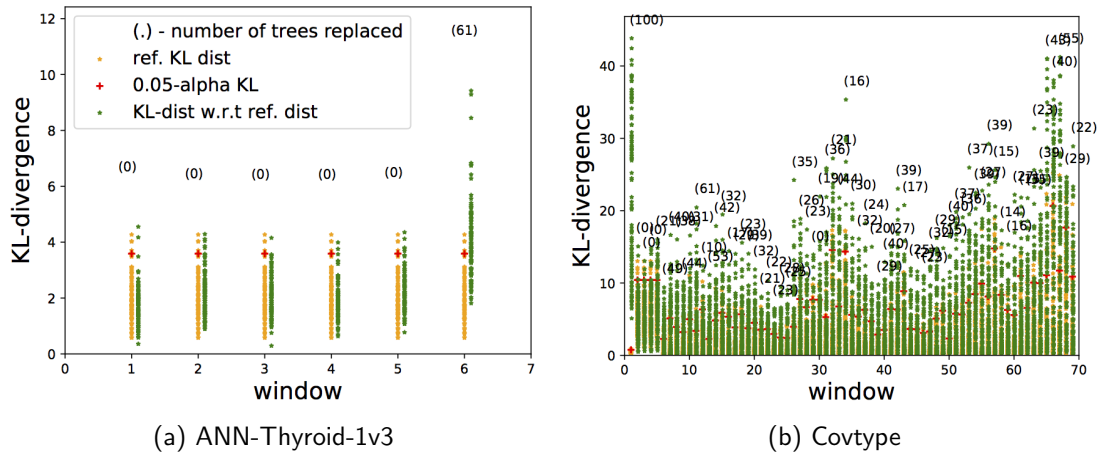


Figure 15: Results for drift detection across windows. (a) When there is no drift, such as in *ANN-Thyroid-1v3*, then no trees are replaced for most of the windows, and the older model is retained. (b) If there is drift, such as in *Covtype*, then the trees are more likely to be replaced.

from feedback which is still valid. On the other hand, if we update the model too slowly, then the model continues to focus on the stale subspaces based on the past feedback. It is hard to find a common rate of update that works well across all datasets (such as replacing 20% trees with each new window of data). Figure 17 (bottom row) shows that the adaptive strategy (**HiLAD-Stream (KL Adaptive)**) which replaces obsolete trees using KL-divergence, as illustrated in Figure 17 (top row), is robust and competitive with the best possible configuration, i.e., **HiLAD-Batch**.

### 7.2.2 PARAMETER SENSITIVITY ANALYSIS FOR **HiLAD-Stream**

**HiLAD-Stream** depends on drift detection measure. It has several approaches to replace ensemble members: i) KL-adaptive, ii) a fixed amount of tree replacement, and iii) no replacement. *KL-adaptive* case, we focused on selection of 95 percentile threshold to discard less important ensemble members(trees). We changed the decision threshold to observe **HiLAD-Stream**'s sensitivity. Similar exploration was performed for fixed tree replacement parameters from a wide range of parameters 5%, 10%, 15%, 20% 25%. Table 3 presents our algorithms detection performance with the variation of parameters. KL based threshold selection (95%) is the most stable one. In a real-world scenario, determining the fixed replacement factors is infeasible, and KL-divergence based threshold approach will be the recommended choice. We present support for this claim in Tables 7, 6, 8 for various number of feedback. Our recommended choice seems to discover comparable amount of anomalies compared to fixed settings. We highlighted the performance of our algorithm bold and underlined the best result achieved with fixed drift amount. In a practical scenario, the amount of drift is unknown hence setting the right configuration is always a challenging decision. Whereas, our principled approach **HiLAD-Stream** determines such thresholds dynamically and can reduce the burden from human experts on setting the fixed parameter.

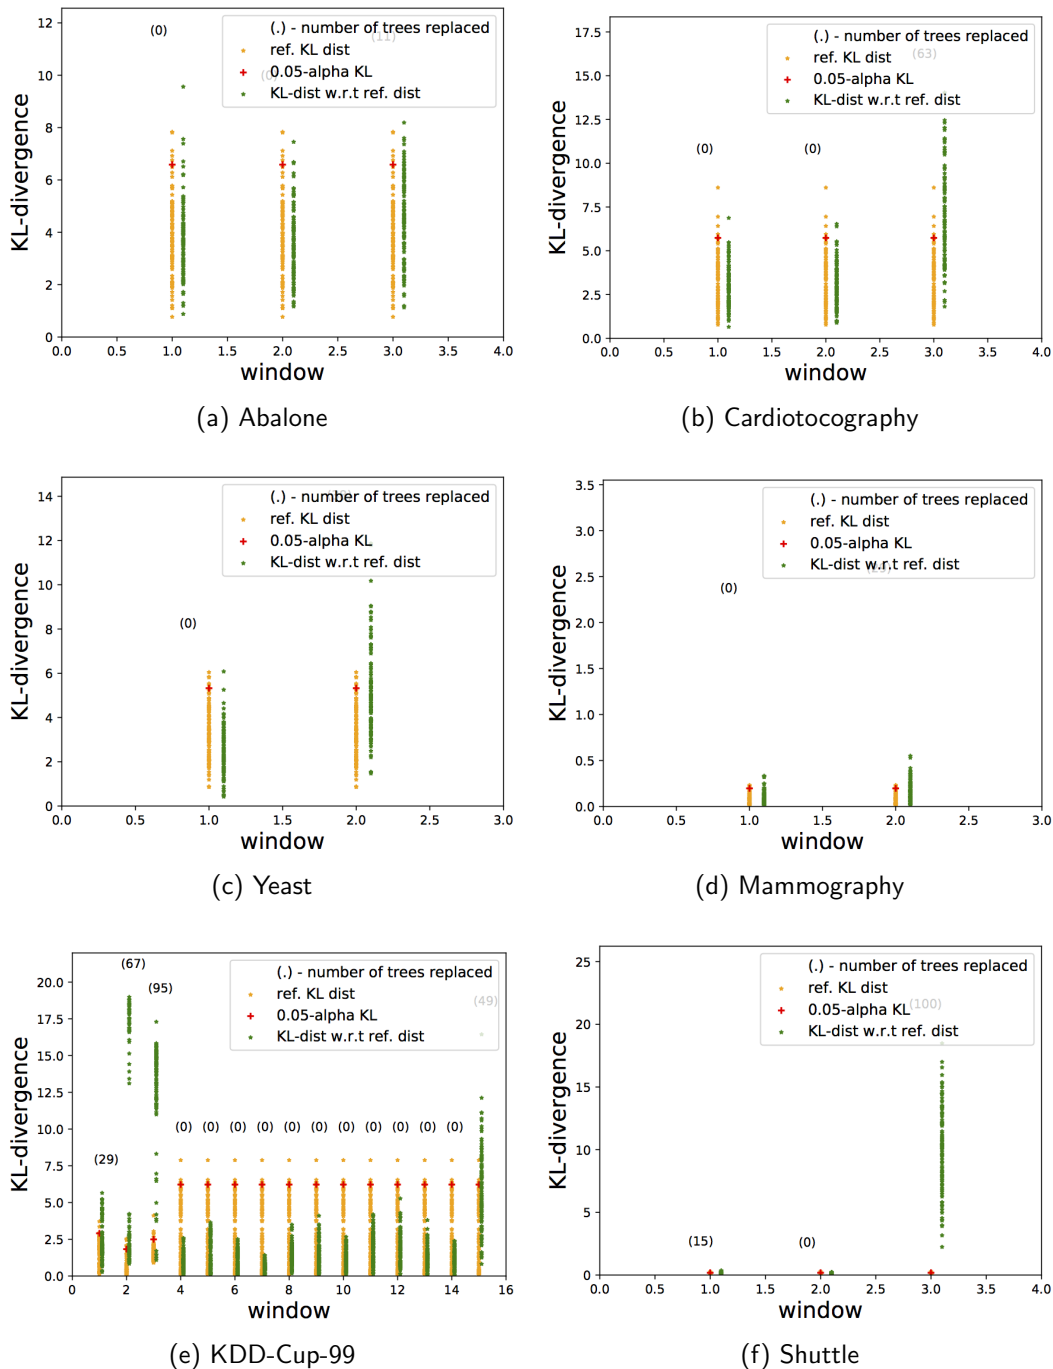


Figure 16: The last data window in each dataset usually has much fewer instances and therefore, its distribution is very different from the previous window despite there being no data drift. **Therefore, we ignore the drift in the last window.** We did not expect *Abalone*, *Cardiocotography*, *KDDCup99*, *Mammography*, *Shuttle*, *Yeast*, and *ANN-Thyroid-1v3* (Figure 15a) to have much drift in the data. This can also be seen in the plots where most of the windows in the middle of streaming did not result in too many trees being replaced (the numbers in the parenthesis are mostly zero).

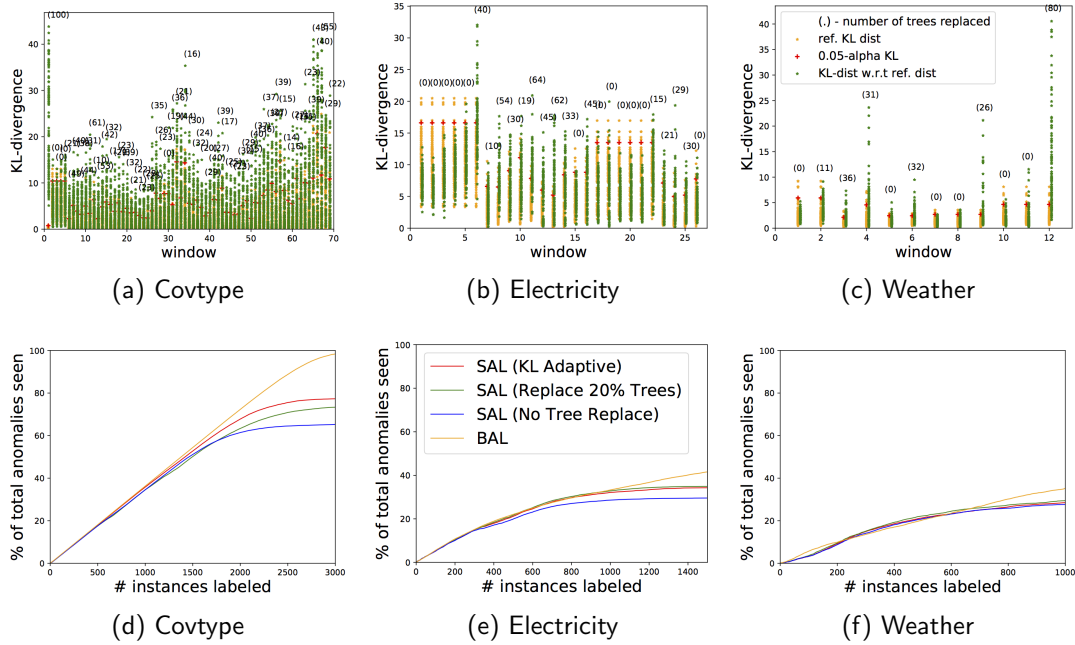


Figure 17: Results for integrated drift detection and label feedback with HiLAD-Stream algorithm. The top row shows the number of trees replaced per window when a drift in the data was detected relative to previous window(s). The bottom row shows the percentage of total anomalies seen vs. number of queries for the **streaming datasets with significant concept drift**.

Table 3: The fraction of total anomalies seen for all three datasets for streaming experiment. We explored our algorithms tree replacement fraction using both (fixed replacement and KL-threshold) approaches. It seems that KL-Divergence based technique with 95% showed stable performance over fixed tree replacement (marked as **bold**). We also underlined the best results obtained from all datasets.

Tree replacement technique	Amount	Datasets (# of feedbacks)		
		Electricity (1500)	Weather (1000)	Covtype (3000)
Fixed	5%	33.32	26.28	72.26
	10%	34.29	27.91	72.59
	15%	35.49	27.80	72.90
	20%	35.98	<u>29.13</u>	72.64
	25%	35.96	28.25	75.03
	30%	35.89	27.90	77.86
	50%	38.28	25.76	76.21
	80%	38.64	20.63	63.49
KL threshold (in percentile)	92.5%	32.70	25.08	<u>80.58</u>
	95%	<b>35.33</b>	<b>28.08</b>	<b>80.44</b>
	97.5%	35.14	28.32	79.76

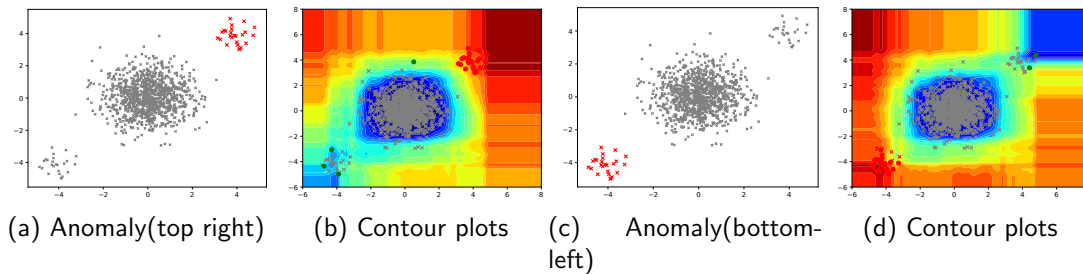


Figure 18: Class-wise feedback experiment with synthetic datasets. Figure (a) and (c) are presenting two types of anomalies located two corners. And (b), (d) are showing how the algorithm rank data points across the data space.

### 7.2.3 ANOMALY CHARACTERIZATION FROM LABEL FEEDBACK USING HiLAD

HiLAD can identify anomalies from different subspaces based on feedback. It takes into account the users’ preferences (data point type and spatial location) to re-weight different subspaces. In contrast, classical anomaly detection approaches always propose candidates from the same ranked list. We demonstrate the performance of HiLAD on three cases (one on synthetic, and two others from representative real-world datasets (*Yeast* and *Abalone*)).

1. **Synthetic dataset** with two groups of anomalies created to provide controlled feedback. It’s clear from the score contour plots (Figure 18) that our algorithm pays attention where the domain expert is interested.
2. **Abalone** contains two anomaly classes (*Class 3* and *Class 21*). In our experiments, we provided feedback from one class as anomalies and found that HiLAD ranked anomalies from that class higher as shown in Figure 19.
3. **Yeast** contains three anomaly classes (*POX*, *VAC*, and *ERL*). The results of the experiment where we provide feedback on only one class of anomalies is shown in Figure 20.

We see similar behavior of the algorithm for classes *VAC*(Figure 20c) and *ERL*(Figure 20d) as we did for both classes of the *Abalone* dataset. However, for class *POX* we see an initial higher gain with HiLAD-Batch than the unsupervised baseline before the unsupervised baseline overtakes HiLAD-Batch. We took a closer look into the cause for this using t-SNE plots (Figure 21). Here we see that there is one clustered group of anomalies (lower left) that gets discovered fast initially by HiLAD-Batch and thus explains its initial steeper rise. But then, rest of the anomalies belonging to *POX* are mixed within the nominal data and hence harder to discover. Since these might be spread out among nominals, positive feedback on one such anomaly does not get propagated to other anomalies efficiently. Moreover, a negative feedback on a nominal data instance (a very likely scenario for this situation) lowers the priority of an anomaly that is present in its proximity.




Figure 19: Abalone dataset contains anomalies from two different classes. They are known as Class 3 and Class 21. The above figures are (a) feedback with HiLAD without class preferences (b) domain expert is interested in examples from *Class 21*. (c) *Class 3* is the point of interest. In all cases, we notice that getting feedback on a particular class can help us to recognize more anomalies from any particular class. (Note that for Figure (b) and (c) Y-axis represents what fraction of anomalies were discovered from that class.)




Figure 20: Yeast contains anomalies from three classes known as *VAC*, *POX*, *ERL*. These plots are presenting HiLAD-Batch is able to identify anomalous examples from a specific type of examples. (Note that for Y-axis represents what fraction of anomalies were discovered from that class.)

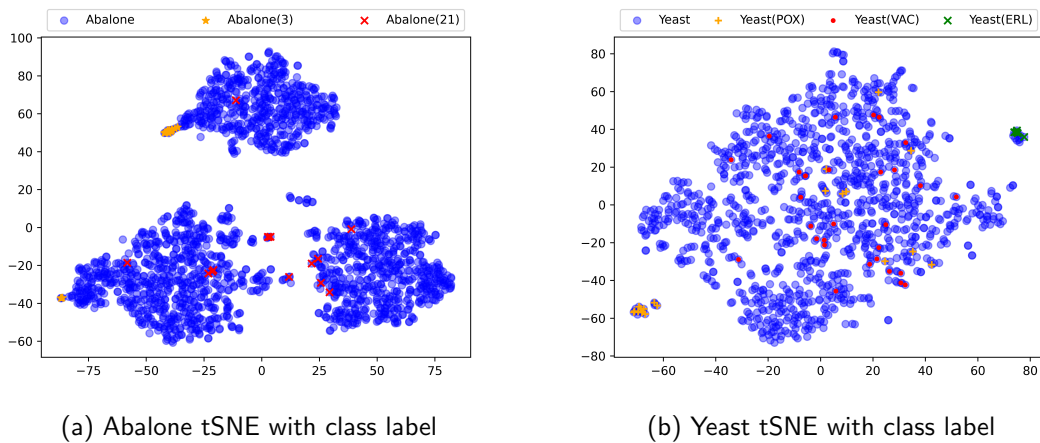


Figure 21: tSNE plot for Abalone and Yeast. Samples for each class are marked and color codes. They are widespread on the data space



### 7.3 Comparison with Recent Semi-Supervised Approaches

We propose effective instantiations of the generic human-in-the-loop learning framework HiLAD for anomaly discovery. Prior works including (Siddiqui et al., 2018), (Das et al., 2016), (Das et al., 2017) are notable methods in this line of work. All these prior methods used ensemble based approaches as a black-box system (Pevný, 2016). However, we used tree-based ensembles / Isolation Forest (Liu et al., 2008) and exploited the inherent structure of the trees to effectively utilize human feedback. Our experimental evaluation considers all these works as baselines. Moreover, we did a fine-grained experimental comparison with (Siddiqui et al., 2018). Figure 22 shows the comparison of HiLAD-Batch with feedback-guided anomaly detection via online optimization (**Feedback-guided Online**) (Siddiqui et al., 2018). The results for *KDDCup99* and *Covtype* could not be included for **Feedback-guided Online** because their code<sup>3</sup> resulted in error (*Segmentation Fault*) when run with 3000 feedback iterations (a reasonable budget for the large datasets).

#### 7.3.1 COMPARISON WITH LIMITED LABELED SET

In this section, we present a comparison with SOEL using the same experimental setup followed in (Li et al., 2023). SOEL is a semi-supervised algorithm which, as per its original design, does not expect a continuous user feedback cycle. Instead, SOEL assumes that a limited amount of labeled data is already available. Under the experimental setup, the datasets are first split into two parts (*train*, *test*). Next, a fixed number of labels  $B$  are collected by following one of the nine query strategies from previous works: (1) *Mar* (Görnitz et al., 2013), (2) *Hybr1* (Görnitz et al., 2013), (3) *Pos1* (Pimentel et al., 2020), (4) *Pos2* (Barnabé-Lortie, Bellinger, & Japkowicz, 2015), (5) *Rand1* (Ruff et al., 2018), (6) *Rand2* (Trittenbach, Englhardt, & Böhm, 2021), (7) *Hybr2*, (8) *Hybr3* (Ning, Chen, Zhou, & Wen, 2022), and (9) *SOEL* (Li et al., 2023). Finally, the F1-score for each query strategy is reported averaged over five runs. The train-test splits and initialization were random for each experiment run. *Hybr2* is attributed to a strategy discussed in [https://github.com/shubhomoydas/ad\\_examples](https://github.com/shubhomoydas/ad_examples) but was implemented incorrectly in (Li et al., 2023) with minimum euclidean distance between queried instances. The correct implementation *maximizes* the euclidean distance to improve diversity.

Our first set of comparisons with SOEL (Table 4) are based on the four tabular datasets from (Li et al., 2023). In these datasets, more than 30% samples are considered anomalies, especially in the *test* set. **HiLAD** performed best on two out of four datasets. **SOEL** performed best on *Satellite*, and **Hybr1** was best on *Ionosphere*. Arguably, this setup has a much higher proportion of anomalies than in real-world scenarios.

Our second set of comparisons with SOEL are based on the seven datasets from Table 2 where % of anomalies present in the test set are more realistic ( $< 10\%$ ). We follow the same training methodology as in the previous experiment and report the results in Table 5. Here, **HiLAD** achieved the best F1-score on six out of seven datasets. A key takeaway from both sets of results (Table 4 and Table 5) is that HiLAD’s performance is consistent across different datasets.

Our third set of experiments are again based on datasets from Table 2. Here, we compare the performance of algorithms as the amount of labeled data  $B$  during training time is

3. <https://github.com/siddiqmd/FeedbackIsolationForest>(retrieved on 10-Oct-2018)

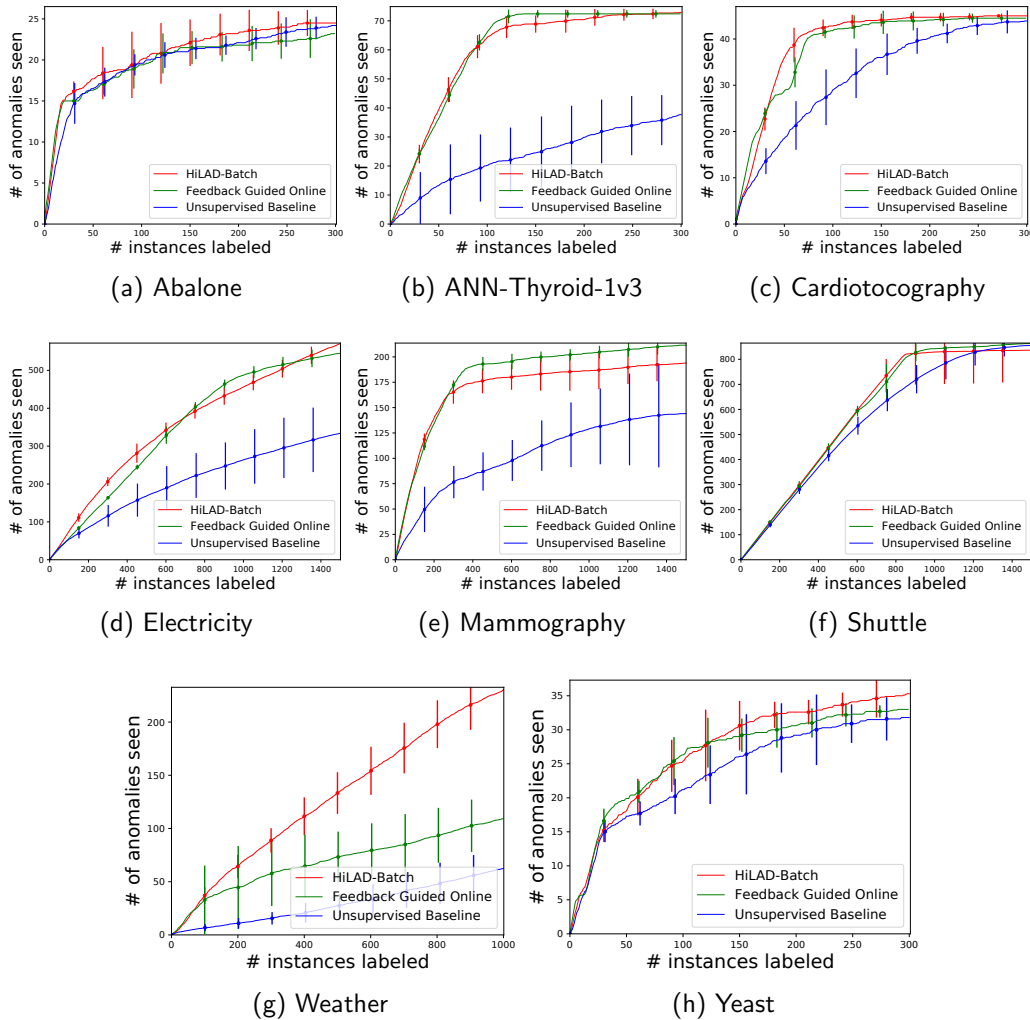


Figure 22: Results comparing HiLAD-Batch with feedback-guided anomaly detection via online optimization (Siddiqui et al., 2018). HiLAD-Batch is the tree-based model implemented in our codebase and employs the AAD loss (anomalies score higher than  $\tau$ -th quantile score and nominals lower). Feedback Guided Online employs the linear loss in Siddiqui et al., 2018. Unsupervised Baseline is the unsupervised Isolation Forest baseline. Both approaches perform similar on most datasets. While HiLAD-Batch has slightly poor accuracy on *Mammography* than Feedback Guided Online, HiLAD-Batch performs much better on *Weather*.

increased in the following order: 10, 20, 30, 40, 50. The results are reported in Figure 23. For some datasets, we didn't have 50 anomalous samples present. One interesting observation is that with 10 labeled samples, HiLAD was not the best performer for *Cardiotocography*. However, as more labeled samples were added, HiLAD's performance improved, unlike other methods whose performance didn't.

Table 4: F1-score (%) with standard deviation for anomaly detection on tabular data when the query budget  $|B| = 10$ . HiLAD performs the best on two of the four datasets whereas SOEL and Hybr1 performed the best on the rest two.

Dataset name	Mar	Hybr1	Pos1	Pos2	Rand1	Rand2	Hybr2	Hybr3	SOEL	HiLAD
BreastW	81.6±0.7	83.3±2.0	58.6±7.7	81.3±0.8	87.1±1.0	82.9±1.1	55.0±6.0	79.6±4.9	93.9±0.5	<b>0.958±0.00</b>
Ionosphere	91.9±0.3	<b>92.3±0.5</b>	56.1±6.2	91.1±0.8	91.1±0.3	91.9±0.6	64.0±4.6	88.2±0.9	91.8±1.1	0.843±0.02
Pima	50.1±1.3	49.2±1.9	48.5±0.4	52.4±0.8	53.6±1.1	51.9±2.0	53.8±4.0	48.4±0.7	55.5±1.2	<b>0.635±0.02</b>
Satellite	64.2±1.2	66.2±1.7	57.0±3.0	56.7±3.2	67.7±1.2	66.6±0.8	48.6±6.9	56.9±7.0	<b>71.1±1.7</b>	0.674±0.00

Table 5: F1-score (%) with standard deviation for anomaly detection on tabular data when the query budget  $|B| = 10$ . HiLAD performs the best on six of the seven datasets.

Dataset	Rand1	Rand2	Mar	Pos1	Pos2	Hybr1	Hybr2	Hybr3	SOEL	HiLAD
Abalone	0.28±0.04	0.31±0.15	0.3±0.18	0.14±0.14	0.23±0.2	0.32±0.21	0.25±0.05	0.31±0.1	0.28±0.09	<b>0.47±0.13</b>
ANN-Thyroid-1v3	0.1±0.11	0.15±0.12	0.1±0.11	0.05±0.02	0.11±0.1	0.13±0.13	0.19±0.08	0.09±0.09	0.06±0.03	<b>0.44±0.11</b>
Cardiotocography	0.18±0.13	0.23±0.09	0.29±0.07	0.09±0.12	0.24±0.1	0.27±0.08	<b>0.59±0.1</b>	0.41±0.15	0.46±0.11	0.57±0.08
KDD-Cup-99	0.24±0.12	0.3±0.16	0.15±0.06	0.32±0.15	0.23±0.14	0.29±0.13	0.4±0.1	0.34±0.05	0.3±0.15	<b>0.87±0.01</b>
Mammography	0.02±0.02	0.01±0.01	0±0	0.01±0.02	0±0.01	0.02±0.02	0.01±0.01	0.01±0.01	0.01±0.01	<b>0.34±0.06</b>
Shuttle	0.31±0.14	0.36±0.19	0.42±0.23	0.49±0.4	0.52±0.16	0.46±0.17	0.66±0.29	0.55±0.18	0.63±0.1	<b>0.83±0.04</b>
Yeast	0.24±0.07	0.21±0.09	0.24±0.02	0.17±0.12	0.24±0.05	0.28±0.02	0.21±0.11	0.25±0.09	0.29±0.02	<b>0.33±0.05</b>

## 7.4 Broader Applicability of Insights and Algorithmic Ideas

We employed the tree-based approach as a device to demonstrate how our generic mechanism for incorporating labeled anomalies can be applied to ensembles of anomaly detectors. This technique can be applied in situations where an ensemble of detectors can be trained inexpensively (e.g., using feature bagging (Lazarevic & Kumar, 2005), LODA (Pevný, 2016), IFOR (Liu et al., 2008) etc.) or with more sophisticated methods (Ted et al., 2013). We have focused on tree-based methods in this paper due to their popularity and frequent state-of-the-art performance. Our algorithm makes an assumption that the data is *i.i.d.* distributed in the form of tabular data with feature values and hence, we did not apply it *directly* to raw data such as text and images. However, this is not a serious limitation. Deep-learning-based anomaly detection methods that use deep learning to extract features (Pang, Shen, Cao, & Hengel, 2021) can employ our proposed algorithm on the representation of the data in the latent space where it is assumed that the data is *i.i.d.* distributed.

The proposed technique treats the subspaces at the tree leaves to be ensemble components. GLAD (Islam, Das, Doppa, & Natarajan, 2020) uses the insights from this algorithm and extends it to the extreme case where each instance is treated as a subspace; and then, instead of trees, uses neural networks to incorporate feedback using a conceptually similar loss function as proposed this paper. The strong results of GLAD algorithm demonstrate that the key ideas behind the proposed algorithm are versatile.

## 7.5 Summary of Experimental Findings

We briefly summarize the main findings of our empirical evaluation of HiLAD framework.

- Uniform prior over the weights of the scoring function with tree-based ensembles to rank data instances is very effective for human-in-the-loop anomaly detection. The histogram distribution of the angles between score vectors from Isolation Forest (IFOR) and  $\mathbf{w}_{unif}$  show that anomalies are aligned closer to  $\mathbf{w}_{unif}$ .

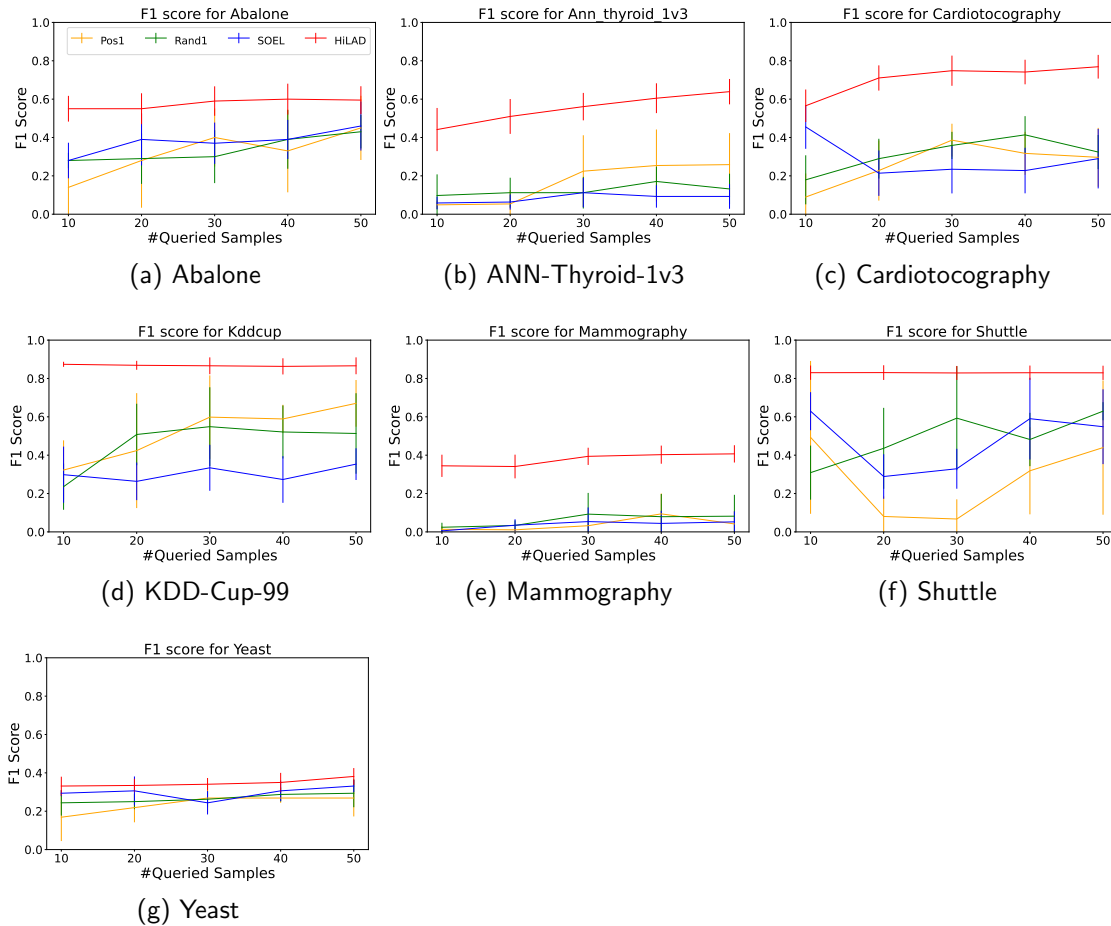


Figure 23: Comparison of F1 scores on the seven datasets with different query budgets (10, 20, 30, 40, and 50). HiLAD performs better as it gets more labeled samples.

- The diversified query selection strategy (**Select-Diverse**) based on compact description improves diversity over greedy query selection strategy with no loss in anomaly discovery rate.
- The KL-Divergence based drift detection algorithm is very robust in terms of detecting and quantifying the amount of drift. In the case of limited memory settings with no concept drift, feedback tuned anomaly detectors generalize to unseen data. In the streaming setting with concept drift, our HiLAD-Stream algorithm is robust and competitive with the best possible configuration, namely, HiLAD-Batch.
- HiLAD can discover higher anomalies with limited labeled feedback. With more labeled samples, it discovers more anomalies compared to baseline methods.

## 8. Summary and Future Work

This paper studied a human-in-the-loop learning framework using tree-based anomaly detector ensembles for discovering and interpreting anomalous data instances. We first explained the reason behind the empirical success of tree-based anomaly detector ensembles and called attention to an under-appreciated property that makes them uniquely suitable for label-efficient human-in-the-loop anomaly detection. We demonstrated the practical utility of this property by designing efficient learning algorithms to support this framework. We also showed that the tree-based ensembles can be used to compactly describe groups of anomalous instances to discover diverse anomalies and to improve interpretability. To handle streaming data settings, we developed a novel algorithm to detect the data drift and associated algorithms to take corrective actions. This algorithm is not only robust, but can also be employed broadly with any ensemble anomaly detector whose members can compute sample distributions such as tree-based and projection-based detectors.

Our immediate future work includes deploying human-in-the-loop anomaly detection algorithms in real-world systems to measure their accuracy and usability (e.g., qualitative assessment of interpretability and explanations). Developing algorithms for interpretability and explainability of anomaly detection systems is a very important future direction.

## Acknowledgments

This research was supported in part by the AgAID AI Institute for Agriculture Decision Support, supported by the National Science Foundation and United States Department of Agriculture - National Institute of Food and Agriculture award #2021-67021-35344, and in part by Contract W911NF15-1-0461 with the US Defense Advanced Research Projects Agency (DARPA) Communicating with Computers Program and the Army Research Office.

## Appendix A. Parameter Sensitivity for Streaming Algorithms

We provide comprehensive experimental results to evaluate the performance of our active anomaly discovery algorithm for streaming data setting on the dataset described in section 7.2.2. Table 6, 7, 8 show the results for different experimental configurations. These results demonstrate that irrespective of the label feedback budget and data drift our proposed KL-based threshold selection approach remains stable.

## Appendix B. Comparison with Streaming Anomaly Detectors

We performed a comparison with recent streaming based anomaly detection methods including RRCF (Guha et al., 2016) and the classical baseline (Liu et al., 2008). Table 9, 10, and 11 demonstrate the better anomaly discovery performance of our HiLAD-Stream approach for all three datasets, namely, *Electricity*, *Covtype*, and *Weather* respectively.

Table 6: Fraction of total anomalies discovered for *Electricity* dataset. We varied the fraction of replaced trees for both fixed and KL-Divergence based approaches. KL-Divergence based approach with 95% threshold shows stable results compared with other parameter settings (we marked this row **bold**). We also underlined the best results obtained for different label feedback budget. Same observation applies to the other two datasets as well. Therefore, we recommend this setting (KL-Divergence with 95% threshold) can be used as the default configuration for all datasets.

Tree replacement technique	Amount	Feedback received						
		0	250	500	750	1000	1250	1500
Fixed	5%	0.07	12.49	20.98	28.20	31.52	32.93	33.32
	10%	0.07	12.60	21.40	28.56	32.21	33.75	34.29
	15%	0.07	<u>12.60</u>	<u>21.66</u>	29.15	33.01	34.77	35.49
	20%	0.07	12.48	21.65	<u>29.53</u>	33.34	35.42	35.98
	25%	0.07	12.27	21.10	29.01	33.18	35.32	35.96
	30%	0.07	12.38	21.16	28.78	33.02	35.14	35.89
	50%	0.07	12.10	20.55	28.50	33.81	36.76	38.28
	80%	0.07	11.84	20.08	27.98	<u>33.83</u>	<u>37.14</u>	<u>38.64</u>
KL	92.5%	0.07	12.32	20.89	28.16	31.04	32.32	32.70
threshold	<b>95.0%</b>	<b>0.07</b>	<b>12.39</b>	<b>21.33</b>	<b>28.88</b>	<b>32.83</b>	<b>34.58</b>	<b>35.33</b>
(in percentile)	97.5%	0.07	12.41	21.60	29.43	32.85	34.56	35.14

Table 7: Fraction of total anomalies discovered for *Covtype* dataset. We varied the fraction of replaced trees for both fixed and KL-Divergence based approaches. KL-Divergence based approach with 95% threshold shows stable results compared with other parameter settings (we marked this row **bold**). We also underlined the best results obtained for different label feedback budget.

Tree Replacement		Feedback received						
		0	500	1000	1500	2000	2500	3000
Fixed	5%	0.04	17.80	35.60	52.55	64.66	70.77	72.26
	10%	0.04	17.72	35.10	51.58	65.71	71.47	72.59
	15%	0.04	17.64	34.67	50.75	64.99	71.28	72.90
	20%	0.04	17.60	34.59	50.29	65.27	71.18	72.64
	25%	0.04	17.57	34.59	50.37	66.06	73.16	75.03
	30%	0.04	17.55	34.67	50.36	66.72	75.77	77.86
	50%	0.04	17.38	33.92	48.78	64.63	74.19	76.21
	80%	0.04	17.22	33.70	48.45	59.87	63.13	63.49
KL	92.5%	0.04	17.91	35.84	<u>53.04</u>	68.58	78.03	<u>80.58</u>
threshold	<b>95.0%</b>	<b>0.04</b>	<b>17.91</b>	<b>35.85</b>	<b>53.01</b>	<b>69.17</b>	<b>78.42</b>	<b>80.44</b>
(in percentile)	97.5%	0.04	17.88	35.80	52.97	68.76	77.85	79.76

Table 8: Fraction of total anomalies discovered for *Weather* dataset. We varied the fraction of replaced trees for both fixed and KL-Divergence based approaches. KL-Divergence based approach with 95% threshold shows stable results compared with other parameter settings (we marked this row **bold**). We also underlined the best results obtained for different label feedback budget.

Tree Replacement technique	Amount	Feedback received				
		0	250	500	750	1000
Fixed	5%	0.15	11.22	20.12	24.16	26.28
	10%	0.15	11.33	20.43	25.44	27.91
	15%	0.15	11.45	20.78	25.46	27.80
	20%	0.15	11.60	<u>21.23</u>	26.05	<u>29.13</u>
	25%	0.15	11.40	20.95	<u>26.22</u>	28.25
	30%	0.15	11.75	20.82	25.44	27.90
	50%	0.15	11.20	19.71	23.60	25.76
	80%	0.15	11.11	17.39	20.14	20.63
KL	92.5%	0.15	11.01	19.05	22.88	25.08
	<b>95.0%</b>	<b>0.15</b>	<b>11.42</b>	<b>20.30</b>	<b>25.38</b>	<b>28.08</b>
	97.5%	0.15	<u>11.81</u>	21.13	25.81	28.32

Table 9: Comparison of fraction of total anomalies discovered for the *Electricity* dataset.

Algorithms	Feedback received						
	0	250	500	750	1000	1250	1500
IFor	0.07	7.51	15.15	20.59	24.59	27.51	28.99
RRCF	0.00	3.35	6.12	8.6	9.98	12.68	15.30
HiLAD-Stream(KL-95%)	<b>0.07</b>	<b>12.39</b>	<b>21.33</b>	<b>28.88</b>	<b>32.83</b>	<b>34.58</b>	<b>35.33</b>

Table 10: Comparison of fraction of total anomalies discovered for the *Covtype* dataset.

Algorithms	Feedback received						
	0	500	1000	1500	2000	2500	3000
IFor	0.04	12.11	17.88	23.99	31.43	36.48	40.90
RRCF	0.00	0.33	0.76	1.20	1.67	2.03	2.62
HiLAD-Stream(KL-95%)	<b>0.04</b>	<b>17.91</b>	<b>35.85</b>	<b>53.01</b>	<b>69.17</b>	<b>78.42</b>	<b>80.44</b>

Table 11: Comparison of fraction of total anomalies discovered for the *Weather* dataset.

Algorithms	Feedback received				
	0	250	500	750	1000
IFor	0.15	1.33	4.10	9.10	14.31
RRCF	0.15	2.4	6.10	8.54	11.59
HiLAD-Stream(KL-95%)	<b>0.15</b>	<b>11.42</b>	<b>20.30</b>	<b>25.38</b>	<b>28.08</b>

## References

- Abe, N., Zadrozny, B., & Langford, J. (2006). Outlier detection by active learning. In *Proceedings of the Twelfth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pp. 504–509.
- Aggarwal, C. C., & Sathe, S. (2017). *Outlier Ensembles: An Introduction*. Springer Publishing Company, Incorporated.
- Almgren, M., & Jonsson, E. (2004). Using active learning in intrusion detection. In *17th IEEE Computer Security Foundations Workshop, (CSFW)*, p. 88.
- Angiulli, F., & Fassetti, F. (2007). Detecting distance-based outliers in streams of data. In *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, pp. 811–820.
- Balcan, M., Broder, A. Z., & Zhang, T. (2007). Margin based active learning. In *Proceedings of 20th Annual Conference on Learning Theory (COLT)*, pp. 35–50.
- Balcan, M., & Feldman, V. (2015). Statistical active learning algorithms for noise tolerance and differential privacy. *Algorithmica*, 72(1), 282–315.
- Barnabé-Lortie, V., Bellinger, C., & Japkowicz, N. (2015). Active learning for one-class classification. In *2015 IEEE 14th international conference on machine learning and applications (ICMLA)*, pp. 390–395. IEEE.
- Breunig, M. M., Kriegel, H., Ng, R. T., & Sander, J. (2000). LOF: identifying density-based local outliers. In *Proceedings of the 2000 ACM International Conference on Management of Data (SIGMOD)*, pp. 93–104.
- Campos, G. O., Zimek, A., & Meira, W. (2018). An unsupervised boosting strategy for outlier detection ensembles. In *Advances in Knowledge Discovery and Data Mining: 22nd Pacific-Asia Conference, PAKDD 2018, Melbourne, VIC, Australia, June 3-6, 2018, Proceedings, Part I 22*, pp. 564–576. Springer.
- Cao, X., Liu, W., & Tsang, I. W. (2023). Data-efficient learning via minimizing hyperspherical energy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(11), 13422–13437.
- Cao, X., & Tsang, I. W. (2022). Shattering distribution for active learning. *IEEE Transactions on Neural Networks and Learning Systems*, 33(1), 215–228.
- Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM computing surveys (CSUR)*, 41(3), 15:1–15:58.
- Cheng, L., Wang, Y., Liu, X., & Li, B. (2020). Outlier detection ensemble with embedded feature selection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34, pp. 3503–3512.
- Chiang, A., & Yeh, Y. (2015). Anomaly detection ensembles: In defense of the average. In *IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology, Volume III*, pp. 207–210.
- Cohn, D. A., Atlas, L. E., & Ladner, R. E. (1994). Improving generalization with active learning. *Machine Learning*, 15(2), 201–221.



- Das, S., Wong, W., Dietterich, T. G., Fern, A., & Emmott, A. (2016). Incorporating expert feedback into active anomaly discovery. In *IEEE 16th International Conference on Data Mining (ICDM)*, pp. 853–858.
- Das, S., Wong, W.-K., Fern, A., Dietterich, T. G., & Siddiqui, M. A. (2017). Incorporating expert feedback into tree-based anomaly detection. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD), IDEA Workshop*.
- Dasgupta, S., Kalai, A. T., & Monteleoni, C. (2009). Analysis of perceptron-based active learning. *Journal of Machine Learning Research (JMLR)*, 10, 281–299.
- Dheeru, D., & Karra Taniskidou, E. (2017). UCI machine learning repository..
- Ditzler, G., & Polikar, R. (2013). Incremental learning of concept drift from streaming imbalanced data. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 25, 2283–2301.
- Domingues, R., Filippone, M., Michiardi, P., & Zouaoui, J. (2018). A comparative evaluation of outlier detection algorithms: Experiments and analyses. *Pattern Recognition*, 74, 406–421.
- Emmott, A., Das, S., Dietterich, T. G., Fern, A., & Wong, W. (2015). Systematic construction of anomaly detection benchmarks from real data. *CoRR*, abs/1503.01158.
- Freund, Y., Seung, H. S., Shamir, E., & Tishby, N. (1997). Selective sampling using the query by committee algorithm. *Machine learning*, 28(2-3), 133–168.
- Görnitz, N., Kloft, M., Rieck, K., & Brefeld, U. (2013). Toward supervised anomaly detection. *Journal of Artificial Intelligence Research (JAIR)*, 46, 235–262.
- Guha, S., Mishra, N., Roy, G., & Schrijvers, O. (2016). Robust random cut forest based anomaly detection on streams. In *Proceedings of the 33rd International Conference on Machine Learning, (ICML)*, pp. 2712–2721.
- Gupta, M., Gao, J., Aggarwal, C., & Han, J. (2014). Outlier detection for temporal data: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 26, 2250–2267.
- Harries, M., & of New South Wales., U. (1999). *Splice-2 comparative evaluation [electronic resource] : electricity pricing*. University of New South Wales, School of Computer Science and Engineering [Sydney].
- He, J., & Carbonell, J. G. (2007). Nearest-neighbor-based active learning for rare category detection. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 633–640.
- Ikeda, Y., Ishibashi, K., Nakano, Y., Watanabe, K., Tajiri, K., & Kawahara, R. (2018). Human-assisted online anomaly detection with normal outlier retraining. In *Proc. ACM SIGKDD Workshop*.
- Islam, M. R., Das, S., Doppa, J. R., & Natarajan, S. (2020). Glad: Glocalized anomaly detection via active feature space suppression. *International Conference on Machine Learning Workshop on Human in the Loop Learning*, 152.
- Kalai, A. T., Klivans, A. R., Mansour, Y., & Servedio, R. A. (2008). Agnostically learning halfspaces. *SIAM Journal on Computing*, 37(6), 1777–1805.

- Kearns, M. J. (1998). Efficient noise-tolerant learning from statistical queries. *Journal of the ACM (JACM)*, 45(6), 983–1006.
- Kim, M., Kim, J., Yu, J., & Choi, J. K. (2023). Active anomaly detection based on deep one-class classification. *Pattern Recognition Letters*, 167, 18–24.
- Kirner, E., Schubert, E., & Zimek, A. (2017). Good and bad neighborhood approximations for outlier detection ensembles. In *International Conference on Similarity Search and Applications*, pp. 173–187. Springer.
- Kontaki, M., Gounaris, A., Papadopoulos, A., Tsihlias, K., & Manolopoulos, Y. (2011). Continuous monitoring of distance-based outliers over data streams. In *2011 IEEE 27th International Conference on Data Engineering*, pp. 135–146.
- Lazarevic, A., & Kumar, V. (2005). Feature bagging for outlier detection. In *Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pp. 157–166.
- Li, A., Qiu, C., Kloft, M., Smyth, P., Mandt, S., & Rudolph, M. (2023). Deep anomaly detection under labeling budget constraints. In *International Conference on Machine Learning*, pp. 19882–19910. PMLR.
- Liu, F. T., Ting, K. M., & Zhou, Z. (2008). Isolation forest. In *Proceedings of the IEEE International Conference on Data Mining (ICDM)*, pp. 413–422.
- Micenková, B., McWilliams, B., & Assent, I. (2014). Learning outlier ensembles: The best of both worlds—supervised and unsupervised. In *Proceedings of the ACM SIGKDD 2014 Workshop on Outlier Detection and Description under Data Diversity (ODD2)*. New York, NY, USA, pp. 51–54. Citeseer.
- Monteleoni, C. (2006). Efficient algorithms for general active learning. In *Proceedings of 19th International Conference on Learning Theory (COLT)*, pp. 650–652.
- Na, G. S., Kim, D., & Yu, H. (2018). Dilof: Effective and memory efficient local outlier detection in data streams. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- Ning, J., Chen, L., Zhou, C., & Wen, Y. (2022). Deep active autoencoders for outlier detection.. pp. 1–13. Springer.
- Nissim, N., Cohen, A., Moskovitch, R., Shabtai, A., Edry, M., Bar-Ad, O., & Elovici, Y. (2014). ALPD: active learning framework for enhancing the detection of malicious PDF files. In *Proceedings of IEEE Joint Intelligence and Security Informatics Conference (JISIC)*, pp. 91–98.
- Pang, G., Shen, C., Cao, L., & Hengel, A. V. D. (2021). Deep learning for anomaly detection: A review. *ACM Computing Surveys (CSUR)*, 54(2), 1–38.
- Pang, G., Ting, K. M., Albrecht, D. W., & Jin, H. (2016). Zero++: Harnessing the power of zero appearances to detect anomalies in large-scale data sets. *Journal of Artificial Intelligence Research (JAIR)*, 57, 593–620.
- Pevný, T. (2016). LODA: Lightweight online detector of anomalies. *Machine Learning*, 102(2), 275–304.

- Pichara, K., & Soto, A. (2011). Active learning and subspace clustering for anomaly detection. *Intelligent Data Analysis*, *15*, 151–171.
- Pimentel, T., Monteiro, M., Veloso, A., & Ziviani, N. (2020). Deep active learning for anomaly detection. In *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE.
- Ruff, L., Vandermeulen, R., Göernitz, N., Deecke, L., Siddiqui, S. A., Binder, A., Müller, E., & Kloft, M. (2018). Deep one-class classification. In *International conference on machine learning*, pp. 4393–4402. PMLR.
- Ruff, L., Vandermeulen, R. A., Görnitz, N., Binder, A., Müller, E., Müller, K.-R., & Kloft, M. (2019). Deep semi-supervised anomaly detection. In *International Conference on Learning Representations*.
- Salehi, M., Leckie, C., Bezdek, J., Vaithianathan, T., & Zhang, X. (2016). Fast memory efficient local outlier detection in data streams. *IEEE Transactions on Knowledge and Data Engineering*, *28*, 3246–3260.
- Settles, B. (2012). *Active Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.
- Siddiqui, M. A., Fern, A., Dietterich, T. G., Wright, R., Theriault, A., & Archer, D. W. (2018). Feedback-guided anomaly discovery via online optimization. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Stokes, J. W., Platt, J. C., Kravis, J., & Shilman, M. (2008). ALADIN: Active learning of anomalies to detect intrusions. *Technique Report, Microsoft Network Security Redmond, WA, 98052*.
- Tan, S. C., Ting, K. M., & Liu, F. T. (2011). Fast anomaly detection for streaming data. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 1511–1516.
- Ted, E., Goldberg, H. G., Memory, A., Young, W. T., Rees, B., Pierce, R., Huang, D., Reardon, M., Bader, D. A., Chow, E., et al. (2013). Detecting insider threats in a real corporate database of computer usage activity. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pp. 1393–1401.
- Trittenbach, H., & Böhm, K. (2019). One-class active learning for outlier detection with multiple subspaces. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management(CIKM)*, pp. 811–820. ACM.
- Trittenbach, H., Englhardt, A., & Böhm, K. (2021). An overview and a benchmark of active learning for outlier detection with one-class classifiers. *Expert Systems with Applications*, *168*, 114372.
- Veeramachaneni, K., Arnaldo, I., Korrapati, V., Bassias, C., & Li, K. (2016). AI<sup>2</sup> : Training a big data machine to defend. In *Second IEEE International Conference on Big Data Security on Cloud*, pp. 49–54.

- Wang, T., Rudin, C., Doshi-Velez, F., Liu, Y., Klampfl, E., & MacNeille, P. (2016). Bayesian rule sets for interpretable classification. In *Proceedings of 16th IEEE International Conference on Data Mining (ICDM)*, pp. 1269–1274.
- Woods, K. S., Doss, C. C., Bowyer, K. W., Solka, J. L., Priebe, C. E., & Kegelmeyer, W. P. (1993). Comparative evaluation of pattern recognition techniques for detection of microcalcifications in mammography. *International Journal of Pattern Recognition and Artificial Intelligence*, 07(06), 1417–1436.
- Wu, K., Zhang, K., Fan, W., Edwards, A., & Yu, P. S. (2014). Rs-forest: A rapid density estimator for streaming anomaly detection. In *Proceedings of IEEE International Conference on Data Mining (ICDM)*, pp. 600–609.
- Yan, S., & Zhang, C. (2017). Revisiting perceptron: Efficient and label-optimal learning of halfspaces. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 1056–1066.
- Yang, D., Rundensteiner, E. A., & Ward, M. O. (2009). Neighbor-based pattern detection for windows over streaming data. In *Proceedings of the 12th international conference on extending database technology: advances in database technology*, pp. 529–540.
- Zhao, Y., & Hryniewicki, M. K. (2019). Dcso: dynamic combination of detector scores for outlier ensembles..