

OptIForest: Optimal Isolation Forest for Anomaly Detection

Haolong Xiang¹, Xuyun Zhang^{1*}, Hongsheng Hu², Lianyong Qi³, Wanchun Dou⁴, Mark Dras¹, Amin Beheshti¹ and Xiaolong Xu^{5*}

¹Macquarie University

²CSIRO's Data61

³Qufu Normal University

⁴Nanjing University

⁵Nanjing University of Information Science and Technology

haolong.xiang@hdr.mq.edu.au, {xuyun.zhang, mark.dras, amin.beheshti}@mq.edu.au,
Hongsheng.Hu@data61.csiro.au, lianyongqi@gmail.com, douwc@nju.edu.cn, xlxu@nuist.edu.cn

Abstract

Anomaly detection plays an increasingly important role in various fields for critical tasks such as intrusion detection in cybersecurity, financial risk detection, and human health monitoring. A variety of anomaly detection methods have been proposed, and a category based on the isolation forest mechanism stands out due to its simplicity, effectiveness, and efficiency, e.g., iForest is often employed as a state-of-the-art detector for real deployment. While the majority of isolation forests use the binary structure, a framework LSHiForest has demonstrated that the multi-fork isolation tree structure can lead to better detection performance. However, there is no theoretical work answering the fundamentally and practically important question on the optimal tree structure for an isolation forest with respect to the branching factor. In this paper, we establish a theory on isolation efficiency to answer the question and determine the optimal branching factor for an isolation tree. Based on the theoretical underpinning, we design a practical optimal isolation forest OptIForest incorporating clustering based learning to hash which enables more information to be learned from data for better isolation quality. The rationale of our approach relies on a better bias-variance trade-off achieved by bias reduction in OptIForest. Extensive experiments on a series of benchmarking datasets for comparative and ablation studies demonstrate that our approach can efficiently and robustly achieve better detection performance in general than the state-of-the-arts including the deep learning based methods.

1 Introduction

Detection of anomalies (also known as outliers) is an important machine learning task to capture abnormal patterns or sparse observations in data that collide with the majority showing the expected behaviours [Pang *et al.*, 2021a], and has been deployed in a broad range of fields for criti-

cal applications such as intrusion detection in cybersecurity, financial risk detection, and human or device health monitoring [Ahmed *et al.*, 2016; Chen and Tsourakakis, 2022; Fernando *et al.*, 2021]. While anomalies often take a very small portion of the data or appear infrequently, failing to catch them in a timely manner for further actions can result in severe consequences such as cascading failures in manufacture and deaths in healthcare. Therefore, a variety of unsupervised anomaly detection methods [Ruff *et al.*, 2021], from shallow to deep, have been proposed for different types of anomalies and data types. Note that since collecting labels for data is usually difficult and expensive, especially for anomalies, we herein focus on unsupervised anomaly detection which is more commonly-used in practice. Recently, several deep learning based detection methods, e.g., RDP [Wang *et al.*, 2021] and REPEN [Pang *et al.*, 2018], have shown the advantages of having feature representation learning for anomaly detection [Pang *et al.*, 2021a]. But this does not mean that the traditional (shallow) detection methods will be obsolete because deep neural networks have their own intrinsic limitations such as high computational cost, poor explainability, and difficulty in hyperparameter tuning [Li *et al.*, 2022], or a shallow model can have comparable performance but cost much less, e.g., a very recent work ECOD [Li *et al.*, 2022] just use the statistic analysis of the tail event of a distribution to achieve the good performance over various benchmark datasets. Instead, traditional models can be preferred in specific scenarios, e.g., edge computing where computational resources are limited and medical research where high explainability is in demand. Moreover, these methods can work together with the feature representations learned from deep learning for better performance, e.g., a recent work [Xu *et al.*, 2022] has shown such an example that iForest is successfully used together with deep learning.

Benefiting from ensemble learning [Aggarwal and Sathe, 2015], a category of detection methods based on the isolation forest mechanism [Hariri *et al.*, 2019] stands out of the shallow models due to its excellent simplicity, effectiveness, and efficiency, therefore being really promising for big data anomaly detection. The basic idea is to randomly and recursively partition relatively small samples drawn from a dataset

until all data instances are isolated, which produces a forest of isolation trees. Anomaly scores can be derived from an isolation forest based on an observation that anomalies often have shorter path lengths due to the ease of isolating them from others. As the first instance, iForest [Liu *et al.*, 2008] has been widely recognised in academia and deployed in real applications, e.g., it has been included in *scikit-learn*, a commonly-used machine learning library in Python. While most of isolation forests following iForest use the binary tree structure for data isolation, a framework LSHiForest [Zhang *et al.*, 2017] producing multi-fork isolation trees with the use of similarity hash functions has demonstrated better detection performance. An interesting question arises naturally: What is the optimal branching factor for an isolation tree? However, there is little theoretical work answering this fundamentally important question, and practically this wide gap can hamper the further development of isolation forest based anomaly detection methods.

In this paper, we investigate this interesting problem and establish a theory on the structure optimality of an isolation tree with respect to the branching factor by introducing the notion of isolation efficiency. A constrained optimisation problem is formulated to solve the optimality problem, and an interesting finding is that the optimal branching factor is e (Euler’s number), rather than 2 which is commonly used in existing methods. Based on the theoretical foundation, a practical optimal isolation forest named OptIForest is proposed for efficient and robust anomaly detection. Specifically, we adapt clustering based learning to hash in OptIForest to let the isolation process optimised with more information learned from data. With two key observations on ensemble learning and anomaly distribution in an isolation tree, we design controllable initialisation for agglomerative hierarchical clustering, enabling OptIForest to achieve a better bias-variance trade-off via bias reduction from learning and improve the computation efficiency. Extensive experiments are performed on a suite of benchmarking datasets in our ablation and comparative studies. The results validate our proposed theory on optimal isolation forest, and also show that with respect to detection performance, OptIForest can generally outperform the state-of-the-arts including the deep anomaly detection methods while maintaining a high computation efficiency. The source code is available at <https://github.com/xiagll/OptIForest>.

The main contributions of our work are threefold, summarised as follows: (1) We are the first to formally investigate the optimality problem of isolation tree structure with respect to the branching factor and establish a theory on the optimal isolation forest, offering a theoretical understanding for the effectiveness of the isolation forest mechanism. (2) We innovatively propose a practical optimal isolation forest OptIForest that can enhance both detection performance and computational efficiency by designing a tailored clustering based learning to hash for a good bias-variance trade-off. (3) Results from extensive experiments support our theory well and validate the effectiveness and efficiency of our approach, as well as the advantages over the state-of-the-arts.

2 Related Work

Various methods for unsupervised anomaly detection have emerged, including distance-based, density-based, statistical, ensemble-based, and deep learning-based methods [Chandola *et al.*, 2009; Pang *et al.*, 2021a]. This section delves into the crucial methods related to our study.

Deep Learning-based Methods. Recently, deep neural networks are widely explored for anomaly detection, particularly on complex data types [Zong *et al.*, 2018; Zavrtnik *et al.*, 2021]. Techniques like generative adversarial networks (GANs) [Liu *et al.*, 2019], AutoEncoders [Chen *et al.*, 2017], and reinforcement learning have been leveraged to enhance the detection performance [Pang *et al.*, 2021b]. For example, a random distance-based anomaly detection method called REPEN is proposed in [Pang *et al.*, 2018], where learning low-dimensional representation in random subsample is optimised. While these deep methods can have high accuracy with learning feature representation [Zha *et al.*, 2020; Zhang *et al.*, 2021], they often suffer from issues like expensive computations, complex hyperparameter tuning, etc.

Ensemble-based Methods. To achieve robust detection, classical anomaly detection methods are integrated with ensemble learning, e.g., ensemble LOF [Zimek *et al.*, 2013], isolation using Nearest Neighbour Ensemble (iNNE) [Bandaragoda *et al.*, 2014], and average k -NN distance ensemble [Aggarwal and Sathe, 2015]. These methods suffer from heavy computational cost when handling big data. A simpler but more effective method is iForest [Liu *et al.*, 2012] which leverages the principle that anomalies are more likely to be isolated from others. This pioneering work shows the strong ability of the isolation forest mechanism and has been widely adopted in real applications. A sequence of work on isolation forests have been proposed to mitigate the shortcomings in iForest, e.g., SCiForest [Liu *et al.*, 2010] addresses the failure of detecting axis-parallel anomalies and local anomalies by using an optimisation strategy. Another interesting work is LSHiForest which produces multi-fork isolation trees with the use of LSH (locality-sensitive hashing) functions that can hash similar data into the same hash value. The work formally understands the distance metric underlying iForest and enables the isolation forest mechanism widely applicable to any data types where an LSH family can be defined. But why a multi-fork isolation forest performs better has not been tackled in LSHiForest. Deep isolation forest [Xu *et al.*, 2022] has been recently proposed to incorporate isolation forest with random feature representation, extending iForest to complex data types. However, these works still fail to learn information properly from data to reduce the bias of a base detector for a better bias-variance trade-off.

3 Preliminaries and Problem Statement

LSHiForest Framework [Zhang *et al.*, 2017]. Our approach is built on top of the LSHiForest Framework which is an effective and generic anomaly detection framework with the forest isolation mechanism. While the algorithmic procedural and the way of deriving anomaly scores in LSHiForest are quite similar to that in iForest, a key difference is that




Figure 1: Isolating 9 data instances with different tree structures.

LSHiForest makes use of a hash function to determine the branching when isolating data in the recursive tree construction process, i.e., data instances with the same hash value go into the same branch. As a result, the isolation trees can be multi-fork, which is significantly different from the binary case in iForest. Since the hash functions are drawn from an LSH (Locality-Sensitive Hashing) family [Bawa *et al.*, 2005], the data instances falling into the same branch are similar to each other with provably high probability and the isolation process is more natural than binary splitting. Moreover, the LSHiForest framework has high applicability to work with any distance metric with an LSH family. In fact, iForest and its variant have been proved to be two specific instances of the framework with less commonly-used distance metrics. Besides, other instances of the framework with specific distance metrics like Angular distance, Manhattan (ℓ_1) distance, and Euclidean (ℓ_2) distance have been implemented as well, and the instance with Euclidean distance (denoted as L2SH) shows a very efficient, robust, and accurate detection performance. Thus, we are motivated to use LSHiForest and its L2SH instance as the foundation of our approach given their prominent features and excellent performance.

Problem Statement. Although multi-fork isolation trees can be elegantly constructed in LSHiForest and better detection performance can be achieved, the theoretical understanding of this phenomenon is still missing. Accordingly, an interesting question arises naturally: What is the optimal branching factor for an isolation tree? Since the branching factor influences the tree structure which can be regarded as a parameter of the detection model, answering this question is fundamentally important to understand how the branching factor affects the performance of isolation forest based anomaly detection methods, but intrinsically a challenging task. Fig. 1 shows an example of isolating nine data instances with three different tree structures, including 9-fork tree, binary tree, and ternary tree. It is hard to intuitively tell which structure is the best for isolation, so how to address this problem is non-trivial.

Besides the tree structure, another challenge imposed on existing isolation forest based anomaly detection methods is how much information they should learn from the data to facilitate data partition at each internal node. For instance, iForest learns the minimum and maximum of the selected feature to determine the random splitting value, and its variant SCiForest[Liu *et al.*, 2010] learns more information to determine the splitting hyperplane and gains better performance. But it is worth noting that more learning does not mean better detection performance while the bias of a single base detector can be reduced, according to the bias-variance trade-off theory in ensemble learning [Aggarwal and Sathe, 2015]. For example, the recent work [Xu *et al.*, 2022] actually reports that full learning for the neural networks has

poorer detection performance than the random weight initiation. Because LSH hash functions used in LSHiForest are data-independent, no learning is incorporated in LSHiForest for data isolation. Thus, it is conjectured that the detection performance could be further improved if an isolation tree is constructed with the learning to hash technique [Wang *et al.*, 2017] which produces hash values based on learning information from the data. But learning to hash usually incurs higher computational cost than LSH. Therefore, non-trivial effort is required for the design of an anomaly detector based on isolation forest with learning to hash, to target both high computational efficiency and a good bias-variance trade-off.

4 Methodology

4.1 Optimal Isolation Forest

To answer the research question stated above about the optimal branching factor of an isolation tree, this section formulates the problem to derive the solution and discusses its practical implementation. Let T represents an isolation tree, with the branching factor v , the tree depth d , and the number of leaf nodes ψ . For the purpose of theoretical analysis, we can reasonably allow v , d , and ψ to take real values rather than just integers. Actually, v and d can be regarded as average branching factor of the internal nodes and the average height of the leaf nodes, respectively. To study how the branching factor influences the isolation performance of an isolation tree, we assume that T is a perfect tree where all internal nodes have v children and all leaves have the same depth d . Note that the isolation tree structure may be totally different in practice, but the assumption is the average case which is reasonable in the ensemble learning setting. Given T is perfect, we can have the following relationship:

$$\psi = v^d. \quad (1)$$

Definition 1 (Isolation Capacity). *The maximum number of data instances an isolation tree can isolate is defined as the isolation capacity of the tree, denoted as ψ , which is also the number of leaf nodes.*

Given an isolation capacity, it can be interestingly observed that the width (controlled by v) and the depth (controlled by d) of the tree compete with each other, i.e., if the branching factor is smaller, the tree has to be deeper, or vice versa. So, we can have the following definition to capture the overall effect of branching factor and tree depth.

Definition 2 (Isolation Area). *The isolation area of an isolation tree, denoted as ϕ , is defined as the product of the branching factor v which controls the width of a tree and the depth d , i.e.,*

$$\phi = v \cdot d. \quad (2)$$

Note that isolation trees of the same isolation area can achieve different isolation capacities. For example, given an isolation area $\phi = 6$, a perfect binary tree (i.e., $v = 2$ and $d = 3$) has the isolation capacity $\psi = 2^3 = 8$, but a perfect ternary tree (i.e., $v = 3$ and $d = 2$) has the isolation capacity $\psi = 3^2 = 9$. The latter case seems more efficient, and we can further define the concept of isolation efficiency as follows.

Definition 3 (Isolation Efficiency). *The isolation efficiency of an isolation tree, denoted as η , is defined as the quotient of the isolation capacity divided by the isolation area, i.e.,*

$$\eta = \frac{\psi}{\phi}. \quad (3)$$

Isolation efficiency fundamentally affects the anomaly detection performance of isolation forests because it is associated with the hardness of distinguishing data instances isolated by a tree. Specifically, higher isolation efficiency leads to stronger distinguishability. Continuing the example above, the isolation efficiency of the binary tree is $\eta = 8/6 \simeq 1.33$ and that of the ternary tree is $\eta = 9/6 = 1.5$. Fig. 1(c) shows that the ternary tree can isolate all the data, while Fig. 1(b) shows that a binary tree with only 3 layers fails to do so.

As the branching factor indeed affects the detection performance in terms of the analysis presented above, it is of both theoretical and practical importance to study the problem of the optimal branching factor. With the concepts introduced above, we can formulate the problem of identifying the optimal branching factor v^* as a constrained optimisation problem as follows.

$$v^* = \underset{v}{\operatorname{argmax}} \eta(v, d) = \frac{v^d}{vd}, \quad (4)$$

s.t. $vd = \Phi$,

where $\eta(v, d)$ is the function of isolation efficiency with respect to branching factor v and tree depth d , and Φ is a constant number representing the fixed isolation area. By solving the optimisation problem, we can have the following theorem.

Theorem 1. *An isolation tree T has the highest isolation efficiency when its branching factor $v = e$, where e is the Euler's number with numerical values around 2.718.*

The theorem can be proved with solving the optimisation problem in (4). Please see Appendix ¹ for mode details, where how the isolation efficiency changes with respect to branching factor with a fixed isolation area is also illustrated. This interesting result reveals that the commonly-used binary tree is not the optimal for the forest isolation mechanism, and explains the better detection accuracy and robustness gained by the multi-fork LSHiForest instances like L2SH.

Definition 4 (Optimal Isolation Tree). *An isolation tree with branching factor $v = e$ is defined as an optimal isolation tree.*

Definition 5 (Optimal Isolation Forest). *A forest consisting of a set of optimal isolation trees is defined as an optimal isolation forest.*

Although the optimal isolation tree is theoretically promising, unfortunately there is no real e -fork branching can be constructed directly in reality. As a practically viable way, we can build an isolation tree with the average branching factor equal to e . To achieve this, we need to generate the branching factors during the tree construction. Let a random variable V denote the branching factor for an isolation tree T , with the sample space $\{v \mid v \in \mathbb{Z} \ \& \ v \geq 2\}$. Let \mathcal{D} be a distribution with the probability of taking value v being: $\Pr(V = v) = p_v$.

¹Please refer to the version of this paper with Appendix in arXiv.

We can generate the branching factors from the distribution \mathcal{D} if it satisfies the following condition:

$$\mathbb{E}(V) = \sum_{v=2}^{+\infty} v \cdot p_v = e. \quad (5)$$

As the branching factor 2 is the only one less than e , there should be a lower bound for p_2 and upper bounds for $p_v, v > 2$, to satisfy the condition mentioned above.

Theorem 2. *To satisfy the condition in Eq.(5), the probability of having a branching factor $V \geq v, v > 2$, should have the following upper bound:*

$$\Pr(V \geq v) \leq \frac{(e-2)}{(v-2)}. \quad (6)$$

The theorem can be proved by letting $p_i = 0$ for $2 < i < v$ and leveraging the inequality $\sum_{i=v}^{+\infty} i p_i \geq v \sum_{i=v}^{+\infty} p_i$. Please see Appendix for more details. The upper bound result can give us an intuition on how the probability decreases when the branching factor increases, e.g., the probability is $e-2 \simeq 0.718$ for $V \geq 3$, $\frac{(e-2)}{2} \simeq 0.359$ for $V \geq 4$, and $\frac{(e-2)}{3} \simeq 0.239$ for $V \geq 5$.

Corollary 1. *To satisfy the condition in Eq.(5), the probability of sampling the branching factor $V = 2$ should follow $p_2 \geq 3 - e \simeq 0.282$.*

This corollary can be simply derived from Theorem 2 when $V \geq 3$, i.e., $p_2 = 1 - \Pr(V \geq 3) \leq 1 - (e-2) = 3 - e$.

There are many possible concrete distributions for \mathcal{D} to be specified, either finite or infinite. For example, by just using binary and ternary branches, we can have a simple finite distribution as follows: $p_2 = 3 - e, p_3 = e - 2$, and $p_i = 0, i \geq 4$. For the infinite case, the probability is often a function of the branching factor. The work in [Russell, 1991] derives a distribution $p_v = \frac{(v-1)}{(v!)}, v \geq 2$, from a stochastic technique. Besides, we can have another infinite distribution $p_v = \frac{(e-1)^2}{2e-1} e^{2-v}, v \geq 2$, to satisfy the condition in Eq.(5), as proved in Appendix. Understanding the distributions of branching factors and the probability bounds can facilitate the decision-makings when one tries to design a more practical optimal isolation forest, as shown in the following section.

4.2 OptIForest: Practical Detector Design with Clustering Based Learning to Hash

In this section, we investigate how to implement a practical optimal isolation tree. Since the data-independent LSH hash functions cannot exactly produce a specified number of hash values, we have to leverage learning to hash to achieve this for a specified branching factor. Moreover, learning to hash can capture more information from the data to benefit anomaly detection potentially. There are many types of learning to hash [Wang *et al.*, 2017]. Given the higher accuracy and less quantification loss, the non-parametric hash function $h(\cdot)$ based on nearest vector assignment is adopted herein, i.e.,

$$h(\mathbf{x}) = \underset{k \in \{1, \dots, v\}}{\operatorname{argmin}} \|\mathbf{x} - \mathbf{c}_k\|, \quad (7)$$

where $\{\mathbf{c}_1, \dots, \mathbf{c}_v\}$ is a set of centres of data partitions. Clustering algorithms are employed to compute the centres.

Algorithm 1 Constructing an Optimal Isolation Tree

Input: A dataset (sample) D of size ψ , and cut threshold ϵ .

Output: T_{Opt} -an optimal isolation tree.

```

1: Train LSHiForest to get  $T_{LSH}$ ;           ▷ Pre-training
2: Traverse  $T_{LSH}$  to get  $\Gamma(\epsilon)$ ;
3:  $\mathbb{C}_0 \leftarrow \mathbb{C}_{\Gamma(\epsilon)}$ ;  $\Gamma_0 \leftarrow \Gamma(\epsilon)$            ▷ Initialisation
4: while  $|\mathbb{C}_i| > 1$  do
5:   Generate a branching factor  $v$  from  $\mathcal{D}$ ;
6:   if  $|\mathbb{C}_i| \leq v$  then
7:     return  $T_{Opt}$  with  $(N_{root}, \mathcal{C}') \leftarrow \text{merge}(\mathbb{C}_i)$ ;
8:    $J_{cur} \leftarrow +\infty$ ,  $\mathbb{C}_{cur} \leftarrow NULL$ 
9:   for all  $\mathbb{C}_v = \{\mathcal{C}_{i1}, \dots, \mathcal{C}_{iv}\} \subset \mathbb{C}_i$  do
10:    if  $\text{dist}(\mathbb{C}_v) < J_{cur}$  then
11:       $J_{cur} \leftarrow \text{dist}(\mathbb{C}_v)$ ,  $\mathbb{C}_{cur} \leftarrow \mathbb{C}_v$ 
12:     $(N', \mathcal{C}') \leftarrow \text{merge}(\mathbb{C}_{cur})$ ;           ▷ Optimal identified
13:     $\mathbb{C}_{i+1} \leftarrow \mathbb{C}_i \setminus \mathbb{C}_{cur}$ ,  $\mathbb{C}_{i+1} \leftarrow \mathbb{C}_{i+1} \cup \{\mathcal{C}'\}$ ;
14:     $\Gamma_{i+1} \leftarrow \Gamma_i \setminus \{N_{i1}, \dots, N_{iv}\}$ ,  $\Gamma_{i+1} \leftarrow \Gamma_{i+1} \cup \{N'\}$ ;
    
```

As our goal is to arrange the clusters into a natural hierarchy to form an isolation tree, agglomerative hierarchical clustering is adopted in our approach. The clustering method treats each data instance as a cluster initially and merges similar clusters sequentially until a single cluster is left, forming a hierarchical tree in a bottom-up fashion. While good clustering quality with less data distortion can be obtained, the high computational complexities are often regarded as a downside. With a cubic time complexity, clustering on a small dataset of a fixed size (e.g., 256 in iForest, and 1024 in LSHiForest) still seems too time-consuming. Straightforward adoption of the clustering technique is neither effective nor efficient.

Fortunately, we find that the clustering technique can be optimised based on two key observations in an isolation forest and design a more efficient clustering method. One observation is that anomalies are often located at upper levels of an isolation tree while normal data at lower levels. The quality of clustering at upper levels is more sensitive to anomaly detection than that at lower levels. The other observation is that no learning or full learning in existing methods can lead to poorer detection performance. Therefore, we can let the upper levels of an isolation tree learn more from data to have better clustering quality, while letting the lower levels learn less, aiming to improve detection performance and save computational cost simultaneously.

The basic idea of our approach is to use an isolation tree efficiently produced in LSHiForest to initialise the clusters for agglomerative hierarchical clustering so that the clustering process begins with bigger initial clusters rather than the ones with a single data instance. This can maintain the high clustering quality for upper levels of the hierarchy and save much computational cost because the majority of computation occurs at lower levels. The specific steps for constructing an optimal tree T_{Opt} , are outlined in Algorithm 1, and details are discussed subsequently.

Let T_{LSH} denote an isolation tree in LSHiForest. It can be horizontally partitioned into two disjoint parts by a cut which can be represented as a set of node. Let Γ denote a cut, and $\Gamma = \{N_1, \dots, N_{n_0}\}$, where N_i , $1 \leq i \leq n_0$ represents a

node, and n_0 is the number of the nodes in Γ . The node N_i is either a subtree of T_{LSH} or a leaf (a trivial subtree), and contains a dataset which can be regarded as a cluster denoted as \mathcal{C}_i . Thus, we can obtain a set of clusters $\mathbb{C}_\Gamma = \{\mathcal{C}_1, \dots, \mathcal{C}_{n_0}\}$ from a cut Γ . Then, we can define the concept of ϵ -cut below.

Definition 6 (ϵ -Cut). *Given a cut Γ and its associated clusters $\mathbb{C}_\Gamma = \{\mathcal{C}_1, \dots, \mathcal{C}_{n_0}\}$, for any node $N_i \in \Gamma$, if its cluster size satisfies $|\mathcal{C}_i| \leq \epsilon < |\mathcal{C}_p|$, where ϵ is a threshold and \mathcal{C}_p is the cluster associated with its parent node, Γ is called ϵ -cut and denoted as $\Gamma(\epsilon)$.*

We can construct $\Gamma(\epsilon)$ by simply traversing the tree, and use the associated clusters $\mathbb{C}_{\Gamma(\epsilon)}$ as the initial ones for agglomerative clustering. By tuning $\epsilon \in [1, \psi]$, where ψ is the sample size in isolation forests, we can control the number and size of the clusters to further, and further manipulate the degree of learning. A higher ϵ leads to fewer initial clusters and implies less learning. The initialisation degenerates into the conventional case if $\epsilon = 1$, while if $\epsilon = \psi$, the resultant isolation forest is still T_{LSH} without learning anything.

Unlike traditional agglomerative hierarchical clustering where a binary tree is generated, we need to achieve multi-fork branches for an optimal isolation tree. The distortion measure is employed to capture the learning loss caused by merging multiple clusters. Let $\mu_{\mathbb{C}}$ denote the merged centre (mean) of a set of clusters $\mathbb{C} = \{\mathcal{C}_1, \dots, \mathcal{C}_v\}$, calculated by

$$\mu_{\mathbb{C}} = \frac{\sum_{i=1}^v \mu_{\mathcal{C}_i} \cdot n_i}{\sum_{i=1}^v n_i}, \quad (8)$$

where $\mu_{\mathcal{C}}$ is the centre of cluster $\mathcal{C} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, calculated by $\mu_{\mathcal{C}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i$. Then, the distortion of merging the clusters in \mathbb{C} can be calculated as

$$\text{dist}(\mathbb{C}) = \sum_{i=1}^v \|\mu_{\mathcal{C}_i} - \mu_{\mathbb{C}}\| \cdot n_i. \quad (9)$$

After generating the branching factor v (Line 5 in Algorithm 1), Line 8-14 shows the details of performing cluster merging. We need to try all the combinations of v clusters and evaluate their distortions. The one with minimum distortion is selected for merging, and a new node is created for the resultant cluster. This process repeats until the number of candidate clusters is not greater than the given branching factor. Finally, all the remaining clusters are merged as the root node of the resultant isolation tree.

It is worth noting the produced tree is not an exact optimal isolation tree if ϵ is not 1, because this practical algorithm takes the trade-off between accuracy and efficiency into account. The most time-consuming part is the evaluation of all the combinations of v clusters, whose time complexity grows exponentially with respect to v . Thus, to be efficient, our approach only leverages binary and ternary branches for the tree construction, i.e., we adopt a finite distribution, which is technically reasonable because the probability drops considerably when the branching factor increases. But note that other distributions can be used if other clustering techniques like K-means are used.

Once a forest of optimal isolation trees are built as an optimal isolation forest for OptIForest, it can perform anomaly

Datasets	AUC-ROC (%)							AUC-PR (%)						
	iForest	LSHiForest	ECOD	REPEN	RDP	DiForest	OptIForest	iForest	LSHiForest	ECOD	REPEN	RDP	DiForest	OptIForest
AD	69.3 ± 1.9	77.9 ± 0.4	69.8 ± 0	70.0 ± 2.2	88.7 ± 0.3	76.8 ± 0.7	77.4 ± 0.7	40.1 ± 4.9	46.2 ± 0.6	48.3 ± 0	37.7 ± 4.0	72.6 ± 0.7	51.8 ± 2.7	43.8 ± 3.1
campaign	70.9 ± 1.0	67.7 ± 0.6	77.5 ± 0	61.1 ± 4.4	76.3 ± 0.8	69.3 ± 0.9	74.7 ± 0.4	28.5 ± 1.3	24.7 ± 1.0	35.6 ± 0	17.8 ± 3.9	37.2 ± 0.9	27.5 ± 1.3	32.0 ± 0.5
Arrhythmia	79.7 ± 1.0	77.5 ± 0.5	82.3 ± 0	73.7 ± 3.6	75.5 ± 0.5	76.3 ± 1.1	79.6 ± 0.8	47.5 ± 1.4	38.6 ± 0.7	49.4 ± 0	37.4 ± 3.5	32.0 ± 0.6	38.2 ± 1.2	45.1 ± 1.2
cardio	93.0 ± 0.7	90.4 ± 0.6	95.0 ± 0	91.5 ± 2.9	88.1 ± 0.6	93.0 ± 0.5	92.8 ± 1.3	57.0 ± 2.8	48.9 ± 1.0	67.6 ± 0	53.3 ± 10.6	53.9 ± 1.5	58.7 ± 1.9	58.9 ± 3.7
backdoor	72.7 ± 2.9	89.2 ± 0.9	84.9 ± 0	86.8 ± 1.6	91.0 ± 2.1	92.0 ± 0.5	92.7 ± 0.5	4.5 ± 0.7	27.3 ± 2.6	9.6 ± 0	12.5 ± 1.7	3.5 ± 0.8	39.4 ± 3.3	51.7 ± 8.7
KDDCup99	97.0 ± 0.6	96.4 ± 0.2	91.1 ± 0	95.9 ± 0.6	41.0 ± 3.1	88.5 ± 0.7	97.4 ± 0.1	48.6 ± 7.0	32.6 ± 1.1	48.5 ± 0	44.1 ± 1.9	15.4 ± 0.9	16.7 ± 0.5	43.0 ± 0.1
Celeba	69.4 ± 2.4	72.5 ± 0.7	72.3 ± 0	84.3 ± 2.2	86.0 ± 0.6	67.9 ± 1.7	79.2 ± 1.9	6.3 ± 0.9	6.8 ± 0.3	8.5 ± 0	10.7 ± 1.8	10.4 ± 0.6	5.5 ± 0.6	8.1 ± 1.0
mnist	80.2 ± 1.8	85.3 ± 0.6	83.8 ± 0	67.6 ± 10.6	85.1 ± 1.6	83.7 ± 1.3	85.5 ± 1.0	27.7 ± 3.2	38.3 ± 1.0	30.5 ± 0	20.4 ± 10.2	36.7 ± 2.4	33.0 ± 2.1	40.7 ± 1.9
Census	60.1 ± 1.8	62.6 ± 0.4	66.8 ± 0	62.7 ± 1.5	65.3 ± 0.4	59.4 ± 1.1	67.8 ± 1.0	7.1 ± 0.3	7.5 ± 0.1	8.6 ± 0	7.7 ± 0.2	8.6 ± 0.1	6.9 ± 0.2	8.8 ± 0.2
Donors	76.6 ± 1.0	74.5 ± 0.7	74.0 ± 0	83.2 ± 1.7	96.2 ± 1.1	67.8 ± 1.2	77.1 ± 3.2	11.9 ± 0.8	10.5 ± 0.6	13.6 ± 0	15.5 ± 1.3	43.2 ± 6.1	8.0 ± 0.3	11.3 ± 1.6
Cover	88.0 ± 2.1	93.6 ± 0.5	93.3 ± 0	86.6 ± 5.7	51.2 ± 1.3	76.4 ± 4.0	90.9 ± 0.8	6.4 ± 0.9	9.0 ± 0.8	11.6 ± 0	5.3 ± 2.0	2.0 ± 1.1	4.0 ± 0.8	6.5 ± 0.7
http	99.9 ± 0	93.3 ± 0	97.9 ± 0	99.4 ± 0.1	99.3 ± 0.1	99.3 ± 0.1	99.4 ± 0.1	90.2 ± 7.9	34.2 ± 0.6	14.5 ± 0	39.5 ± 2.4	36.2 ± 0.8	35.1 ± 0.4	35.4 ± 1.5
smtp	90.5 ± 0.8	86.9 ± 0.9	88.0 ± 0	90.9 ± 1.3	69.8 ± 1.1	84.6 ± 0.5	92.4 ± 0.6	0.4 ± 0	56.9 ± 3.1	50.7 ± 0	26.9 ± 10.8	21.7 ± 3.9	55.2 ± 7.0	36.9 ± 13.7
Ionosphere	84.5 ± 0.5	91.2 ± 0.2	76.8 ± 0	86.5 ± 2.3	82.7 ± 0.6	94.3 ± 0.5	93.4 ± 0.3	80.8 ± 0.5	89.6 ± 0.4	66.3 ± 0	81.0 ± 3.6	80.7 ± 0.6	93.0 ± 0.6	92.3 ± 0.5
Satellite	70.3 ± 1.8	76.7 ± 0.5	74.6 ± 0	74.0 ± 3.0	60.9 ± 0.6	69.2 ± 1.3	78.6 ± 0.7	65.0 ± 2.1	63.5 ± 0.6	66.1 ± 0	70.3 ± 3.9	61.3 ± 0.7	47.5 ± 1.3	71.5 ± 0.9
Shuttle	99.7 ± 0.1	97.2 ± 0.8	99.7 ± 0	99.4 ± 0.1	99.6 ± 0.9	96.3 ± 1.2	98.0 ± 0.4	97.6 ± 0.5	40.1 ± 2.7	95.0 ± 0	91.9 ± 0.6	89.7 ± 1.1	56.5 ± 6.5	64.0 ± 4.3
Spam	61.7 ± 2.9	70.7 ± 0.2	65.6 ± 0	73.4 ± 0.3	74.7 ± 0.5	65.0 ± 0.9	71.1 ± 0.2	46.8 ± 3.0	59.1 ± 0.5	51.8 ± 0	62.7 ± 0.5	63.0 ± 0.2	59.6 ± 0.5	58.1 ± 0.5
Vowel	75.3 ± 1.2	90.8 ± 0.7	40.8 ± 0	77.9 ± 5.5	67.2 ± 1.2	94.4 ± 1.4	90.0 ± 1.2	12.6 ± 1.7	29.5 ± 2.1	2.8 ± 0	14.5 ± 6.5	7.6 ± 3.2	41.1 ± 9.0	32.4 ± 4.3
Waveform	69.6 ± 2.3	70.7 ± 0.8	71.5 ± 0	63.2 ± 9.5	66.9 ± 0.9	63.2 ± 3.4	74.4 ± 1.3	9.7 ± 0.9	9.8 ± 0.4	9.4 ± 0	7.6 ± 2.4	8.7 ± 1.5	11.2 ± 3.3	11.3 ± 1.0
Wine	63.9 ± 0.7	67.0 ± 0.3	62.5 ± 0	64.6 ± 2.2	61.7 ± 0.4	65.9 ± 1.2	66.4 ± 0.5	7.8 ± 0.4	8.9 ± 0.2	8.1 ± 0	8.1 ± 0.7	8.1 ± 0.2	8.5 ± 0.5	8.1 ± 0.2
Average	74.3 ± 1.4	81.9 ± 0.5	78.8 ± 0	79.6 ± 3.1	81.2 ± 0.8	79.2 ± 1.2	83.9 ± 0.9	34.8 ± 2.5	34.1 ± 1.0	34.8 ± 0	33.3 ± 3.6	37.4 ± 1.4	34.9 ± 2.2	38.0 ± 3.7

Table 1: AUC-ROC and AUC-PR performance (mean ± standard deviation) of all methods. Our OptIForest method outperforms others.

detection with the same way to derive anomaly scores as LSHiForest. Although a practical isolation tree can consists of two parts with the upper layers resulting from clustering and the lower layers from LSHiForest, the hash function interface makes the two implementations (LSH functions and the non-parametric learning to hash function in Eq.(7)) no difference as to anomaly score computation.

5 Experiments

5.1 Experiment Setting

Baselines. Our method OptIForest is compared with six state-of-the-art anomaly detection methods: iForest, LSHiForest, ECOD, REPEN, RDP, and DiForest. We evaluate all methods on 20 widely-used benchmark datasets [Pang *et al.*, 2019; Han *et al.*, 2022; Li *et al.*, 2022]. We refer the reader to Appendix for more details of the baselines and datasets.

Metrics. We conventionally use the Area Under Receiver Operating Characteristic Curve (AUC-ROC) and Area Under the Precision-Recall Curve (AUC-PR) as the performance metrics [Kurt *et al.*, 2020; Wang *et al.*, 2021]. To ensure a fair comparison, we follow the optimal parameter settings of the baseline methods. Please see Appendix for more details about the metrics and the parameter settings. All experiments are run 15 times and averaged results are reported.

5.2 Comparison Study Results and Discussion

AUC-ROC Results. Table 1 illustrates that OptIForest is the most robust method that achieves the best performance on most datasets. Specifically, OptIForest has the highest average AUC-ROC score among all the compared methods, surpassing the second best method by 2%. Additionally, RDP and LSHiForest perform well and are stable on most datasets. ECOD produces exceptional results on certain datasets, but performs poorly on others like “Vowel” and “CD” due to its requirement on the data distribution. Note that ECOD has zero standard deviation because it is a deterministic method.

The rest methods iForest, REPEN, and DiForest only perform well on a few datasets, indicating their lack of robustness across all datasets. Thus, it can be seen that no learning in iForest or DiForest fail to achieve a robust performance due to the lack of knowledge learned from data. LSHiForest and OptIForest with better tree structures can achieve bias reduction, and learning from data further makes OptIForest perform better and even outperform the deep detector.

AUC-PR Results. The results in Table 1 show that OptIForest consistently performs best across various datasets with the highest average AUC-PR score, showing the superiority of our method. Specifically, our method outperforms the second best one by about 0.6%. The top two results of each dataset are highlighted in bold. Although iForest and ECOD have exceptional results on a few datasets, they perform poorly on others. REPEN has the lowest average AUC-PR score and performs poorly on most datasets. The AUC-PR results show the same trend as the AUC-ROC results, supporting that our OptIForest performs the best performance.

Execution Time. We use execution time as the efficiency metric to compare OptIForest with other methods, and the results are reported in Table 3 in Appendix. OptIForest has much shorter execution time than the deep learning methods of REPEN and RDP for most datasets. As expected, OptIForest reasonably has longer execution time than other isolation forests where no learning is conducted. So, it can be seen that our approach strikes a good balance between execution efficiency and detection performance. It is appealing that OptIForest can achieve better performance than deep learning methods but takes much less execution cost.

5.3 Ablation Study Results and Discussion

To understand how the branching factor, the cut threshold, and the sampling size influence the performance of OptIForest, we performed detailed ablation studies on eight datasets with a range of data types and sizes.

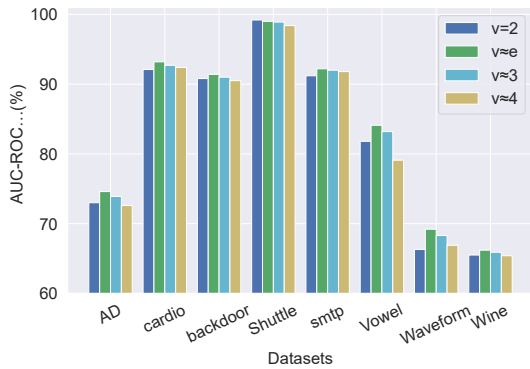


Figure 2: Detection performance changes w.r.t. branching factor v . A branching factor closer to e leads to a better performance.

Branching Factor. To eliminate the effects of feature learning, we use data-independent baseline (i.e., $\epsilon = \psi$) to analyze the impact of the branching factor. Besides, it is difficult to implement the average branch of e in practical experiment because many of the underlying branches just produce 2 branch forks. However, we can still analyze the AUC-ROC results for branches that are either near or far from e . The AUC-ROC results for different branching factors are presented in Fig. 2. When the branching factor is close to e , the AUC-ROC results are the best on almost all datasets. This result means that the best anomaly detection accuracy can be achieved if the branching factor satisfies the condition: $v \approx e$. These experimental results validate Theorem 1 in Section 4.1.

Cut Threshold ϵ . The cut threshold is studied by raising the branching factor to the power of e , as small changes of the threshold do not greatly affect detection accuracy. We also study the boundary condition of the cut threshold ($\epsilon = 512$). Fig. 3 displays AUC-ROC results and standard deviations for different ϵ . It can be seen that the curve increases as the threshold increases in four large datasets with data sizes larger than 10,000 or dimension sizes larger than 1,000, indicating little learning is required for large datasets (with $\epsilon = e^6$ as a reference). Conversely, more learning is necessary in four small datasets, and $\epsilon = e^4$ serves as a good reference point to balance accuracy and time efficiency. A comparison between the results of $\epsilon = 512$ with others illustrates that appropriate learning results in better outcomes than not learning on most datasets. But it is exceptional in the “AD” and “vowel” datasets, where not learning yields better results. This could be attributed to the fact that the isolation forest without learning exactly has the average branch of e , which achieves a favorable bias-variance trade-off.

Sampling Size. To facilitate the observation of the results, the sampling size is set from 2^6 to 2^{11} with exponential increase, as the results in this range can display significant change. The cut threshold of the same sampling size will impact the AUC-ROC results, so we determine the best results for each sampling group as the final outcomes. It can be concluded from Fig. 4 that the AUC-ROC results improve with the increase of the sampling size on most datasets. The AUC-

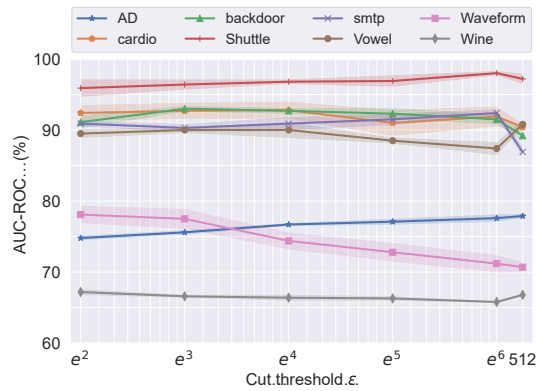


Figure 3: Detection performance changes w.r.t. cut threshold ϵ . In general, a certain level of learning outperforms the case of no learning ($\epsilon = 512$ means no learning).

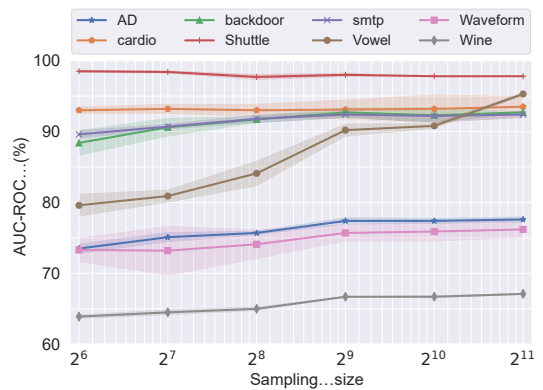


Figure 4: Detection performance changes w.r.t. sampling size ψ . In general, the performance is saturated when ψ reaches a threshold ($2^9 = 512$ herein).

ROC results remain stable once the sampling size exceed 2^9 , except for the “cardio” and “Shuttle” datasets, which keep stable on any sampling size.

6 Conclusion and Future Work

Isolation Forest is an appealing anomaly detection method given its salient characteristics, but it lacks a theoretical foundation about the structure optimality of an isolation tree. In this paper, we have introduced the concept of isolation efficiency and formulated a constrained optimisation problem to derive the optimal branching factor for an isolation tree. We have shown that an optimal isolation tree or forest is theoretically with the branching factor e . Furthermore, we have developed a practical optimal isolation forest OptIForest which can achieve both high computational efficiency and a good bias-variance trade-off by designing a novel clustering based learning to hash for data isolation. We have conducted extensive experiments on a variety of benchmarking datasets for both ablation and comparative studies, and the results have confirmed the effectiveness and efficiency of OptIForest. In the future, we plan to design a version in the context of federated learning where data are scattered across multiple clients.

Acknowledgments

This work is supported in part by the ARC DECRA Fund of the Australian Government under Grant No. DE210101458, and the National Natural Science Foundation of China under Grant No. 92267104 and No. 42050102.

References

- [Aggarwal and Sathe, 2015] Charu C Aggarwal and Saket Sathe. Theoretical foundations and algorithms for outlier ensembles. *Acm sigkdd explorations newsletter*, 17(1):24–47, 2015.
- [Ahmed *et al.*, 2016] Mohiuddin Ahmed, Abdun Naser Mahmood, and Jiankun Hu. A survey of network anomaly detection techniques. *Journal of Network and Computer Applications*, 60:19–31, 2016.
- [Bandaragoda *et al.*, 2014] Tharindu R Bandaragoda, Kai Ming Ting, David Albrecht, Fei Tony Liu, and Jonathan R Wells. Efficient anomaly detection by isolation using nearest neighbour ensemble. In *IEEE International conference on data mining workshop (ICDM workshop)*, pages 698–705. IEEE, 2014.
- [Bawa *et al.*, 2005] Mayank Bawa, Tyson Condie, and Prasanna Ganesan. Lsh forest: self-tuning indexes for similarity search. In *Proceedings of the 14th international conference on World Wide Web (WWW)*, pages 651–660, 2005.
- [Chandola *et al.*, 2009] Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM computing surveys*, 41(3):1–58, 2009.
- [Chen and Tsourakakis, 2022] Tianyi Chen and Charalampos Tsourakakis. Antibenford subgraphs: Unsupervised anomaly detection in financial networks. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (SIGKDD)*, pages 2762–2770, 2022.
- [Chen *et al.*, 2017] Jinghui Chen, Saket Sathe, Charu Aggarwal, and Deepak Turaga. Outlier detection with autoencoder ensembles. In *Proceedings of the 2017 SIAM international conference on data mining (SDM)*, pages 90–98. SIAM, 2017.
- [Fernando *et al.*, 2021] Tharindu Fernando, Harshala Gammulle, Simon Denman, Sridha Sridharan, and Clinton Fookes. Deep learning for medical anomaly detection—a survey. *ACM Computing Surveys*, 54(7):1–37, 2021.
- [Han *et al.*, 2022] Songqiao Han, Xiyang Hu, Hailiang Huang, Minqi Jiang, and Yue Zhao. Adbench: Anomaly detection benchmark. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [Hariri *et al.*, 2019] Sahand Hariri, Matias Carrasco Kind, and Robert J Brunner. Extended isolation forest. *IEEE Transactions on Knowledge and Data Engineering*, 33(4):1479–1489, 2019.
- [Kurt *et al.*, 2020] Mehmet Necip Kurt, Yasin Yilmaz, and Xiaodong Wang. Real-time nonparametric anomaly detection in high-dimensional settings. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [Li *et al.*, 2022] Zheng Li, Yue Zhao, Xiyang Hu, Nicola Botta, Cezar Ionescu, and George Chen. Ecod: Unsupervised outlier detection using empirical cumulative distribution functions. *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [Liu *et al.*, 2008] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *IEEE international conference on data mining (ICDM)*, pages 413–422, 2008.
- [Liu *et al.*, 2010] Fei Tony Liu, Kai Ming Ting, and Zhi Hua Zhou. On detecting clustered anomalies using sciforest. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases (ECML-PKDD)*, pages 274–290. Springer, 2010.
- [Liu *et al.*, 2012] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation-based anomaly detection. *ACM Transactions on Knowledge Discovery from Data*, 6(1):1–39, 2012.
- [Liu *et al.*, 2019] Yezheng Liu, Zhe Li, Chong Zhou, Yuanchun Jiang, Jianshan Sun, Meng Wang, and Xiangnan He. Generative adversarial active learning for unsupervised outlier detection. *IEEE Transactions on Knowledge and Data Engineering*, 32(8):1517–1528, 2019.
- [Pang *et al.*, 2018] Guansong Pang, Longbing Cao, Ling Chen, and Huan Liu. Learning representations of ultrahigh-dimensional data for random distance-based outlier detection. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining (SIGKDD)*, pages 2041–2050, 2018.
- [Pang *et al.*, 2019] Guansong Pang, Chunhua Shen, and Anton van den Hengel. Deep anomaly detection with deviation networks. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining (SIGKDD)*, pages 353–362, 2019.
- [Pang *et al.*, 2021a] Guansong Pang, Chunhua Shen, Longbing Cao, and Anton Van Den Hengel. Deep learning for anomaly detection: A review. *ACM Computing Surveys*, 54(2):1–38, 2021.
- [Pang *et al.*, 2021b] Guansong Pang, Anton van den Hengel, Chunhua Shen, and Longbing Cao. Toward deep supervised anomaly detection: Reinforcement learning from partially labeled anomaly data. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining (SIGKDD)*, pages 1298–1308, 2021.
- [Ruff *et al.*, 2021] Lukas Ruff, Jacob R Kauffmann, Robert A Vandermeulen, Grégoire Montavon, Wojciech Samek, Marius Kloft, Thomas G Dietterich, and Klaus-Robert Müller. A unifying review of deep and shallow anomaly detection. *Proceedings of the IEEE*, 109(5):756–795, 2021.
- [Russell, 1991] KG Russell. Estimating the value of e by simulation. *The American Statistician*, 45(1):66–68, 1991.
- [Wang *et al.*, 2017] Jingdong Wang, Ting Zhang, Nicu Sebe, Heng Tao Shen, et al. A survey on learning to hash. *IEEE*

transactions on pattern analysis and machine intelligence, 40(4):769–790, 2017.

- [Wang *et al.*, 2021] Hu Wang, Guansong Pang, Chunhua Shen, and Congbo Ma. Unsupervised representation learning by predicting random distances. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2950–2956, 2021.
- [Xu *et al.*, 2022] Hongzuo Xu, Guansong Pang, Yijie Wang, and Yongjun Wang. Deep isolation forest for anomaly detection. *arXiv preprint arXiv:2206.06602*, 2022.
- [Zavrtanik *et al.*, 2021] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Reconstruction by inpainting for visual anomaly detection. *Pattern Recognition*, 112:107706, 2021.
- [Zha *et al.*, 2020] Daochen Zha, Kwei-Herng Lai, Mingyang Wan, and Xia Hu. Meta-aad: Active anomaly detection with deep reinforcement learning. In *IEEE International conference on data mining (ICDM)*, pages 771–780. IEEE, 2020.
- [Zhang *et al.*, 2017] Xuyun Zhang, Wanchun Dou, Qiang He, Rui Zhou, Christopher Leckie, Ramamohanarao Kotagiri, and Zoran Salcic. Lshiforest: A generic framework for fast tree isolation based ensemble anomaly analysis. In *IEEE international conference on data engineering (ICDE)*, pages 983–994, 2017.
- [Zhang *et al.*, 2021] Huayi Zhang, Lei Cao, Peter VanNostrand, Samuel Madden, and Elke A Rundensteiner. Elite: Robust deep anomaly detection with meta gradient. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (SIGKDD)*, pages 2174–2182, 2021.
- [Zimek *et al.*, 2013] Arthur Zimek, Matthew Gaudet, Ricardo JGB Campello, and Jörg Sander. Subsampling for efficient and effective unsupervised outlier detection ensembles. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining (SIGKDD)*, pages 428–436, 2013.
- [Zong *et al.*, 2018] Bo Zong, Qi Song, Martin Renqiang Min, Wei Cheng, Cristian Lumezanu, Daeki Cho, and Haifeng Chen. Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In *International conference on learning representations (ICLR)*, 2018.