



Enhancing interpretability and generalizability in extended isolation forests

Alessio Arcudi^{a,*}, Davide Frizzo^a, Chiara Masiero^b, Gian Antonio Susto^a

^a University of Padova, Italy

^b Statwolf Data Science Srl, Italy

ARTICLE INFO

Keywords:

Anomaly detection
Explainable artificial intelligence
Industry 5.0
Interpretable machine learning
Outlier detection
Unsupervised learning

ABSTRACT

Anomaly Detection (AD) focuses on identifying unusual patterns in complex datasets and systems. While Machine Learning and Decision Support Systems (DSS) are effective for this, simply detecting anomalies often falls short in real-world scenarios, especially in engineering contexts where diagnostics and maintenance are essential. Users need clear explanations behind anomaly predictions to understand the root causes and trust the model. The unsupervised nature of AD complicates the development of interpretable tools. To address this, we propose the Extended Isolation Forest Feature Importance (ExIFFI), a new approach that explains the predictions of the Extended Isolation Forest (EIF), applicable to all Isolation Forest models that split using hyperplanes. ExIFFI provides both global and local explanations by analyzing feature importance.

Additionally, we introduce Enhanced Extended Isolation Forest (EIF⁺), an improved version of EIF, designed to better detect unseen anomalies by modifying the splitting strategy of hyperplanes. We compare various unsupervised AD methods across five synthetic and eleven real-world datasets using the Average Precision metric. EIF⁺ consistently outperforms EIF in all scenarios, demonstrating superior generalization. To validate the interpretability, we propose a new metric — AUC_{FS} (Area Under the Curve of Feature Selection) — which uses feature selection as a performance indicator. ExIFFI proves more effective than other unsupervised interpretation methods, excelling in 8 out of 11 real-world datasets and correctly identifying anomalous features in synthetic datasets. Finally, we provide open-source code to encourage further research and reproducibility.

1. Introduction

Machine Learning (ML) and Artificial Intelligence (AI) play a central role in the ongoing socio-economic change, revolutionizing various sectors such as manufacturing (Krafft et al., 2020; Xu et al., 2018; Leng et al., 2024; Arunthavanathan et al., 2021), medicine (Houssein et al., 2021), smart farming (Kaur and Bhattacharya, 2023; Jin and Xu, 2024) and the Internet of Things (Klaib et al., 2021; Kaur and Bhattacharya, 2024). With the increasing deployment of ML across various industries, new challenges have arisen, particularly due to the complexity and lack of transparency in these systems.

Moreover, the diversity of end users, many of whom may not have expertise in data-driven methods, further complicates the effective use of these technologies. To address this, it is crucial to develop algorithms that can explain the structure and predictions of ML models, making them accessible and understandable to a broader audience.

Many scientific works identify explainability¹ as a key factor to enable the successful adoption of ML-based systems (Confalonieri et al., 2021; Doshi-Velez and Kim, 2017; Linardatos et al., 2020).

For tabular data, several methods are available to explain a model's predictions (Molnar, 2020). One common approach is to calculate the importance of each feature in the model's predictions. This involves evaluating the influence of each feature on both individual predictions (referred to as “Local Importance”) and the overall dataset (“Global Importance”). By evaluating feature importance, users gain a deeper understanding of how the model utilizes input data to generate predictions.

Despite the remarkable recent advancements in eXplainable Artificial Intelligence (XAI), most approaches are designed for supervised

* Corresponding author.

E-mail addresses: alessio.arcudi@studenti.unipd.it (A. Arcudi), davide.frizzo.1@phd.unipd.it (D. Frizzo), chiara.masiero@statwolf.com (C. Masiero), gianantonio.susto@unipd.it (G.A. Susto).

¹ While authors in the literature use the term ‘interpretability’ and ‘explainability’ associated with slightly different concepts when associated to Machine Learning/Artificial Intelligence as presented in Gilpin et al. (2018), in this work we will use both terms interchangeably.

tasks, leaving unsupervised tasks, like Anomaly Detection (AD), rarely discussed in the literature.

AD, also referred to as Outlier Detection,² is a field of ML that focuses on identifying elements that live outside the standard “normal” behavior observed in the majority of the dataset (Hawkins, 1980).

Explainability is crucial for Anomaly Detection (AD) approaches, particularly in industrial settings. For instance, AD methods used to monitor industrial machinery can significantly benefit from clear interpretations of detected anomalies. Understanding the root cause of an anomaly can help reduce machine failures, minimize energy loss, conserve resources, and lower production costs. Additionally, the need for interpretable algorithms extends beyond technical benefits, lack of explanations can undermine trust in the model’s output, especially in high-stakes decision-making scenarios. In such cases, explainability is not just a preference but a legal requirement (Commission, 2020).

Isolation Forest (IF) (Liu et al., 2008) is a popular Anomaly Detection (AD) method due to its high accuracy, low computational cost, and relatively simple internal mechanism. It operates by recursively splitting the feature space using randomly chosen axis-aligned hyperplanes, allowing it to isolate anomalies with minimal partitions. To enhance the interpretability of IF, specialized approaches like Depth-based Isolation Forest Feature Importance (DIFFI) (Carletti et al., 2023) have been developed. DIFFI takes advantage of IF’s internal structure to provide both global and local explanations of the model’s decisions.

However, the one-dimensional partitioning process used by IF can introduce artifacts that degrade anomaly detection and reduce the quality of feature explanations. To address this limitation, the Extended Isolation Forest (EIF) has been proposed (Hariri et al., 2021). EIF improves upon IF by utilizing oblique partitions, which help eliminate these artifacts and enhance both anomaly detection and feature explanation.

According to the literature, the EIF is among the top-performing unsupervised anomaly detection (AD) methods (Bouman et al., 2023). However, EIF lacks built-in interpretability for features and struggles with robustness when encountering novel anomalies. To address these limitations and encourage the adoption of isolation-based unsupervised AD approaches, this paper makes three key contributions:

1. We propose the Extended Isolation Forest Feature Importance (ExIFFI), the first (to the best of our knowledge) model-specific approach which generalize the explanations to both IF models and the extended version EIF.
2. We present EIF⁺, a refined version of the EIF designed to optimally model the space around the training data distribution, enhancing the algorithm’s performance on unseen data. This innovation aims to bolster the model’s generalization ability, reliably identifying anomalies even without prior knowledge of their potential locations or without being present in the training dataset. ExIFFI applies to EIF⁺, too.
3. We benchmark the novel EIF⁺ and ExIFFI against state-of-the-art isolation based AD and explainability approaches on 16 public datasets and showcase their effectiveness and computational efficiency. To facilitate investigation and reproducibility, we introduce a novel functionally-grounded quantitative evaluation, named AUC_{FS} score, that measure the effectiveness of the interpretation exploiting Feature Selection as a proxy task.

These contributions address critical gaps in the interpretability and robustness of the EIF, which are essential to bridge for its broader acceptance and adoption over the traditional IF, particularly given EIF’s superior performance, enhanced robustness, and reduced bias. Moreover, this work paves the way for the wider adoption and application

² In this paper, we will refer to ‘Outlier Detection’ and ‘Anomaly Detection’ alternatively, always referring to the same unsupervised task of finding anomalous data points.

of feature importance methods across various domains by establishing a clear and systematic framework for comparing feature importance explainability approaches.

The paper is organized as follows. Section 2 surveys relevant research. Section 3.1 provides an introduction to the IF algorithm, followed by a discussion of the EIF algorithm in Section 3.2. The newly proposed EIF⁺ model is introduced in Section 3.3. Sections 4.1 and 4.2 present the DIFFI and ExIFFI interpretation algorithms, along with the graphical tools used to illustrate their results.

Section 5 describes the experimental setup for evaluating EIF⁺ and ExIFFI analyzing anomaly detection and interpretability performances. Section 6 details the experiments, including graphical evaluations of three datasets in Sections 6.1, 6.2, while Section 6.3 presents the numerical AD performance results for 16 datasets. Section 6.4 covers the interpretation results, Section 6.5 analyzes the correlation between interpretation and AD models results, and Section 6.6 provides an ablation study on scalability for large datasets.

Finally, limitations, conclusions, and future research directions are discussed in Section 7.

The Appendix includes an extensive description of datasets Appendix A.1, an analysis comparing EIF and IF Appendix A.2, an analysis providing a deep explanation on why ExIFFI is not depth-based as DIFFI Appendix A.3, an ablation study on the EIF⁺ hyperparameter η Appendix A.4, and a comparison of different interpretations through the Normalized Discounted Cumulative Gain (NDCG) metric Appendix A.5.

2. Related work

Root Cause Analysis (RCA) is essential for identifying the underlying causes of anomalies in model behavior, enhancing our understanding of system outputs. Papageorgiou et al. (2022) categorize RCA methods into two main types: probabilistic and deterministic.

Probabilistic RCA methods, such as Bayesian Networks (Kitson et al., 2023), model relationships between variables to infer causal connections (Amin et al., 2021a). These approaches are particularly effective in fields like medical diagnostics and risk assessment, where they can trace the origins of specific events by representing probabilistic dependencies (Yu and Zhao, 2020; Amin et al., 2019; Khakzad et al., 2011; Amin et al., 2021b). However, they often require substantial domain expertise, are computationally intensive, and may struggle with scalability in large datasets with complex interdependencies.

In contrast, Deterministic RCA methods focus on evaluating the importance of individual features in contributing to specific outcomes without modeling the entire system. These methods are divided into ad-hoc and post-hoc approaches (Molnar, 2020). For example, ad-hoc methods like the feature importance measure in Random Forests (Kursa and Rudnicki, 2010) are straightforward to implement and interpret but may lack generalizability due to their dependence on specific models. Post-hoc methods, such as SHAP (Lundberg and Lee, 2017a) and permutation importance (Altmann et al., 2010), offer model-agnostic interpretability, though they often come at the cost of higher computational demands and do not inherently provide causal insights.

Thus, the choice between probabilistic and deterministic RCA methods hinges on the specific needs of the problem at hand. Probabilistic methods are suited for uncovering causal relationships but are more complex and resource-intensive, while deterministic methods offer greater flexibility and ease of use, especially when a model-agnostic solution is required, though they may provide less depth in causal inference.

Explainability in Unsupervised Anomaly Detection (AD) remains a significant challenge, where Explainable Anomaly Detection (XAD) focuses on extracting relevant knowledge from AD models regarding relationships either present in the data or learned by the model. Recent surveys highlight the importance of interpretability in AD, particularly

as it becomes a regulatory requirement in various applications (Li et al., 2023b).

Unsupervised AD methods can be categorized into parametric and non-parametric approaches, each presenting unique advantages and challenges, especially in terms of explainability. Parametric models express relationships through a finite set of parameters, facilitating straightforward predictions. However, these models often lack clear interpretability, necessitating additional tools for explanation. For instance, the widely-used Anomaly Autoencoder (AE) model (Aggarwal, 2015) faces interpretability challenges, leading to the development of methods like RXP (Residual eXplainer) (Oliveira et al., 2022), which provides explanations by analyzing deviations in reconstructed features.

On the other hand, Non-parametric AD approaches do not rely on predefined parametric functions and instead detect anomalies by leveraging the inherent structure and patterns in the data. This often results in intrinsic interpretability. For example, ECOD (Empirical-Cumulative-Distribution-based Outlier Detection) (Li et al., 2023a) derives explanations directly from data distributions, and IF is enhanced by the DIFFI algorithm (Carletti et al., 2023), which uses partitioning processes to explain outliers. This intrinsic interpretability makes non-parametric methods particularly valuable when understanding the decision-making process is crucial, with DIFFI showing effectiveness in industrial applications (Carletti et al., 2020; Brito et al., 2022) by providing both local and global explanations.

In conclusion, the selection between parametric and non-parametric unsupervised AD methods should be guided by application needs. Parametric models offer simpler predictions but often require supplementary tools for interpretability, whereas non-parametric models provide greater flexibility and inherent interpretability by utilizing data patterns directly, making them ideal for scenarios where understanding the decision-making process is critical.

This paper addresses a significant gap in the literature by introducing a novel interpretability algorithm specifically tailored for models based on the IF framework, including the Extended Isolation Forest (EIF) model. While existing methods like DIFFI offer some interpretability of IF model, they do not fully accommodate the unique structure of EIF. Our proposed algorithm enhances the interpretability of EIF and similar models, offering both local and global insights into anomaly detection processes. This contribution is crucial for advancing Explainable Anomaly Detection (XAD) by providing more accessible and effective interpretability solutions for IF-based models, particularly in complex, unsupervised settings.

3. Isolation-based approaches for anomaly detection

Next, we provide some notions about isolation-based approaches for AD. This family of methodologies, stemming from the Isolation Forest (Liu et al., 2008), identifies outliers as samples that can be easily separated from the others, i.e., through a reduced number of splitting hyperplanes.

3.1. Isolation forest

Isolation Forest is a widely used ML model for AD (Liu et al., 2008). It generates a set of N random trees, called isolation trees, that are able to identify anomalous elements in a dataset based on their position in the tree structure. The idea behind this approach is that anomalies, on average, are located at the beginning of the trees because they are easier to separate from the rest of the dataset.

Assume that we have n training data $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, where $\mathbf{x}_i \in \mathbb{R}^d$. The IF algorithm chooses at random one dimension $q \in \{1, \dots, d\}$ and a split value $p \in [\min_{i \in \{1, \dots, n\}} x_{i,q}, \max_{i \in \{1, \dots, n\}} x_{i,q}]$. The dataset is then divided into two subsets, the left one $L = \{\mathbf{x}_i | x_{i,q} \leq p\}$ and the right one $R = \{\mathbf{x}_i | x_{i,q} > p\}$. This procedure is calculated iteratively until the whole forest is built. Suppose the size of the dataset is excessive,

meaning that the number of samples makes the construction of the trees too slow; in that case, it is demonstrated by Liu et al. (2008) that it is better to build the forest using only a random subsample \tilde{X} with \tilde{n} elements for each tree. Not only does this keep the computational cost low, it also improves IF's ability to identify anomalies clearly. Once the model has built an isolation forest, to determine which data points live outside the dataset distribution, the algorithm computes an anomaly score for each of them. This value is based on the average depth among trees where each data point is isolated.

The depth of a point x in a tree t , denoted by $h_t(x)$, is the cardinality of the set of nodes \mathcal{P}_x^t that it has to pass to reach the leaf node, i.e.:

$$h_t(x) = |\mathcal{P}_x^t| \quad \text{where} \quad \mathcal{P}_x^t = \{k_{s_1}, k_{s_2}, \dots, k_{s_h}\}, \quad (1)$$

Let $E(h(x))$ be the mean value of depths reached among all trees for a single data point x . Then, according to Preiss (2000), the anomaly score is defined by the function:

$$s(x, \tilde{n}) = 2^{-\frac{E(h(x))}{c(\tilde{n})}}, \quad (2)$$

where $c(N)$ is the normalizing factor defined as the average depth of an unsuccessful search in a Binary Search Tree (Hariri et al., 2021), i.e.:

$$c(N) = 2H(N-1) - \frac{2(N-1)}{N}. \quad (3)$$

and $H(i)$ is the harmonic number that can be estimated by:

$$H(i) = \ln(i) + 0.5772156649 \text{ (Euler constant)}. \quad (4)$$

IF's fast execution with low memory requirement is a direct result of building partial models and requiring only a significantly small sample size as compared to the given training set. This capability is due to the fact IF has the main goal to quickly isolate anomalies more than modeling the normal distribution, contrarily to what other detection methods do (Ruff et al., 2021).

3.2. Extended isolation forest

Although IF is one of the most popular and effective AD algorithms, it has some drawbacks due to its partition strategy. One of them is related to considering only hyper-planes that are orthogonal with respect to the directions of the feature space, as Hariri et al. show in Hariri et al. (2021). In some cases, this bias leads to the formation of some artifacts where points are associated with low anomaly scores even if they are clearly anomalous. These areas are in the intersection of the hyperplanes orthogonal to the dimensions associated with the detection of inliers (Fig. A.9), creating misleading score maps and, in some cases, wrong predictions.

As a consequence, Hariri et al. in Hariri et al. (2021) tried to improve the AD algorithm by correcting this bias with a different and more general algorithm called Extended IF. Instead of selecting a hyper-plane orthogonal to a single input dimension chosen at random, they suggested picking a random point \mathbf{p} and a random vector \mathbf{v} that are consequently used to build a hyperplane that will split the space at each node of the binary tree into two smaller subspaces, as the IF does. Therefore, they proposed a more generic and extended paradigm, while maintaining the fast execution and low memory requirements of the original IF.

In the following, we will briefly describe the working principle of the EIF model. Let us consider a set X of n elements $\mathbf{x}_i \in \mathbb{R}^d$ with $i \in \{1, \dots, n\}$. In order to evaluate the anomaly score of the samples, the EIF model (Hariri et al., 2021) generates a forest of N binary trees:

$$\mathcal{T} = \{t_0, t_1, \dots, t_N\}. \quad (5)$$

Every tree $t \in \mathcal{T}$ is built from a bootstrap sample $X_t \subset X$. At each node k of t , a random hyperplane H_k^t is selected by picking a random point \mathbf{p} inside the distribution space limits, and a normalized random vector \mathbf{v} defined as

$$\mathbf{v} \sim \frac{\mathbf{Z}}{\|\mathbf{Z}\|_2}, \quad \text{where } \mathbf{Z} = (Z_1, \dots, Z_d)^\top \text{ and} \\ Z_i \sim \mathcal{N}(0, 1) \forall i \in \{1, \dots, d\}. \quad (6)$$

Lesouple et al. in Lesouple et al. (2021) pointed out that this way of drawing a random hyperplane may lead to the creation of empty branches, that are merely artifacts of the algorithm. To avoid this problem, they propose a particular way of selecting random hyperplanes, that we also use in our implementation of the EIF model.

As done in Lesouple et al. (2021) and Hariri et al. (2021), we first select a random hyperplane by drawing a unit vector \mathbf{v} defined as in (6). which will serve as the normal vector of the hyperplane. Then, as done in Lesouple et al. (2021), the point where the hyperplane passes is obtained by a scalar α drawn uniformly between the minimum value and the maximum value of the points of the dataset X in the direction of the random vector previously drawn. Therefore,

$$\alpha \in [\min \{\mathbf{x} \cdot \mathbf{v} \forall \mathbf{x} \in X\}, \max \{\mathbf{x} \cdot \mathbf{v} \forall \mathbf{x} \in X\}], \quad \mathbf{p} = \alpha \mathbf{v}. \quad (7)$$

The algorithm starts by generating a random hyperplane $H_{k_0}^t$ defined by the intercept point \mathbf{p} and \mathbf{v} as described in Eqs. (6) and (7).

Thus the hyperplane splits the dataset X_t into two subsets, $L_{k_0}^t = \{\mathbf{x} | \mathbf{x} \in X_{k_0}^t, \mathbf{v}_{k_0}^t \cdot \mathbf{x} > \mathbf{v}_{k_0}^t \cdot \mathbf{p}_{k_0}^t\}$, $R_{k_0}^t = X_t \setminus L_{k_0}^t$. (8)

The root node of the tree, k_0 , is the space splitting made by the hyperplane $H_{k_0}^t$. Then, this method is applied recursively to the subsets $L_{k_0}^t$ and $R_{k_0}^t$ until the max number of splits is reached, which corresponds to the preset max depth of the tree or when the set to split has only one element.

Thanks to this hierarchical tree structure, to evaluate if an element is an anomaly, the model extracts the path of the point $x \in X_t$ from the root to the leaf nodes down the tree. Then, as IF does (Ruff et al., 2021), the EIF algorithm uses the average depth of the point in each tree to evaluate the anomaly score, according to the paradigm that the anomalies can be isolated with few partitions. The average depth of the point in the trees will be translated to an anomaly score according to Eq. (2), as in IF.

3.3. EIF⁺ a novel enhancement of EIF algorithm

Lesouple et al. (2021) showed that the EIF algorithm presented by Hariri et al. in Hariri et al. (2021) can create misleading empty branches. On the other hand, we observed that the solution proposed in Lesouple et al. (2021) hinders the ability to generalize well in the space around the distribution. Actually, generalization ability is very important in the context of AD, since an anomaly is a point that is outside the normal distribution of the data.

Therefore, we propose EIF⁺, a novel approach that enhances the EIF methodology, based on the modification introduced by Lesouple et al. (2021). EIF⁺ aims to better adapt to the space surrounding the training data distribution, thus improving the performances of the EIF model when the model is fitted without anomalies in the training set. This goal is achieved by choosing splitting hyperplanes with an *ad hoc* procedure, in Fig. 1 it is possible to observe the difference.

Let $\mathcal{A} = \{\mathbf{x} \cdot \mathbf{v} \forall \mathbf{x} \in X\}$ be the set of the point projections along the hyperplane's orthogonal direction \mathbf{v} . As in Lesouple et al. (2021), the point \mathbf{p} is defined as $\mathbf{p} = \alpha \mathbf{v}$. However, instead of drawing α using the interval defined by the minimum and maximum values of \mathcal{A} as in Fig. 1(a), EIF⁺ draws it from a normal distribution $\mathcal{N}(\mathbb{E}[\mathcal{A}], \eta\sigma(\mathcal{A}))$ as shown in Fig. 1(b). Even if EIF⁺ can generate empty branches, they contribute to the formation of a model that exhibits enhanced generalizability, as will be shown in Section 6.

4. Interpretability of isolation-based anomaly detection models

In the context of anomaly detection models, interpretability is crucial because it allows us to understand and trust the model's decisions, especially in critical applications such as fraud detection, network security, and medical diagnostics. An interpretable model enables practitioners to verify whether the model is identifying meaningful patterns or if it is influenced by spurious correlations, ultimately ensuring that the anomalies detected are actionable and reliable.

With this in mind, we move on to the task of interpreting the predictions made by the isolation-based models introduced in Section 3. Interpretation algorithms aim to explain the latent patterns identified by the models, thereby enriching our comprehension of their outputs. Drawing inspiration from the DIFFI algorithm developed by Carletti et al. (2023) for the IF, we introduce ExIFFI, a novel model-specific algorithm designed to interpret the results generated by the IF, EIF and its variant, the EIF⁺.

4.1. Depth-based isolation forest feature importance

Depth-based Isolation Forest Feature Importance (DIFFI) was the first unsupervised model-specific method addressing the need to interpret the IF model (Carletti et al., 2023). It exploits the structure of the trees in the IF algorithm to understand which features are the most relevant to discriminate whether a point is an outlier or not. In particular, a meaningful feature should isolate the anomalies as soon as possible, and create a high imbalance when isolating anomalous points (the opposite being true for inliers).

We briefly explain how DIFFI works and we introduce the related notation, as we will leverage it when introducing the novel ExIFFI approach to explain EIF and EIF⁺ predictions. From Carletti et al. (2023), we define:

Induced Imbalance Coefficients λ : given an internal node k of an isolation tree, as defined in Section 3.1, let $n(k)$ be the number of points that the node divides, being $n_l(k)$ and $n_r(k)$ the number of points on the left and the right child, respectively. The coefficient measuring the induced imbalance of the node v is:

$$\lambda(v) = \begin{cases} 0 & \text{if } n_l(k) = 0 \text{ or } n_r(k) = 0 \\ \tilde{\lambda}(k) & \text{otherwise} \end{cases} \quad (9)$$


where

$$\tilde{\lambda}(k) = g\left(\frac{\max(n_l(k), n_r(k))}{n(k)}\right) \quad \text{and} \quad g(a) = \frac{a - \lambda_{\min}(n)}{2(\lambda_{\max}(n) - \lambda_{\min}(n))} + 0.5. \quad (10)$$

In the previous equation, λ_{\min} and λ_{\max} denote the minimum and maximum scores that can be obtained a priori given the number of data $n(k)$.

Cumulative feature importances \mathbf{I} : it is a vector of dimension d (i.e., the number of features) where the j th component is the feature importance of the j th feature. In Carletti et al. (2023), authors distinguish between \mathbf{I}_I , created based on predicted inliers, and \mathbf{I}_O , based on the outliers. Concerning \mathbf{I}_I , the procedure starts with the initialization $\mathbf{I}_I = \mathbf{0}_d$. Considering the subset of predicted inliers for the tree t , $\mathcal{P}_{I,t}$, for each predicted inlier $x_I \in \mathcal{P}_{I,t}$, DIFFI iterates over the internal nodes in its path $Path(x_I, t)$. If the splitting feature associated with the generic internal node v is f_j , then we update the j th component of \mathbf{I}_I by adding the quantity:

$$\Delta = \frac{1}{h_t(x_I)} \lambda_I(v) \quad (11)$$



(a) Hyperplane selection EIF utilized by the algorithm developed by Lesouple et al. [39]. This algorithm selects the point p by uniformly sampling it from an interval depicted in light blue within the figure.

(b) Hyperplane selection EIF⁺, a novel approach for drawing the point p . This new method involves sampling the point p from a normal distribution centered around the mean point of the dataset projection onto the direction described by vector v (represented as red dots in the figure). The distribution is visually depicted as a blue bell shape.

Fig. 1. Comparison of hyperplane selection methods: (a) EIF method by Lesouple et al. and (b) the novel EIF⁺ method.

The same procedure applies for $x_O \in \mathcal{P}_{O,I}$ and \mathbf{I}_O . Intuitively, at each point and each step, the vector accumulates the imbalance produced by each feature. The imbalance is measured by the previously defined λ and weighted by the depth of the node. This means that, according to DIFFI, features that allowed to isolate the points sooner, are considered to be more useful.

Features counter \mathbf{V} : it is used to compensate for uneven random feature sampling that might bias the calculated cumulative feature importance. At each passage of a point through a node, the entry of the counter corresponding to the splitting feature is incremented by one. As for the cumulative feature importance, two feature counters are calculated, the one for predicted inliers \mathbf{V}_I and the one for outliers \mathbf{V}_O .

Based on the above-introduced quantities, DIFFI computes the Global Feature Importance by looking at the weighted ratio between outliers and inliers cumulative feature importance:

$$GFI = \frac{\mathbf{I}_O/\mathbf{V}_O}{\mathbf{I}_I/\mathbf{V}_I}$$

4.2. Extended Isolation Forest Feature Importance (ExIFFI)

Drawing inspiration from DIFFI, we introduce ExIFFI, the Extended Isolation Forest Feature Importance, a generalization of the DIFFI algorithm that is able to rank the importance of the features in deciding whether a sample is an anomaly or not for any of the IF models where the space is split by an hyperplane.

As seen in Section 3.2, a node $k \in \mathcal{K}_t$ in an EIF tree $t \in \mathcal{T}$ corresponds to an hyperplane \mathcal{H}_k^t , that splits the subset $X_k^t \subseteq X^t \subseteq X$. Using the Lesouple et al. (2021) correction of the EIF, \mathcal{H}_k^t is completely defined by means of a vector orthogonal to its direction \mathbf{v}_k^t , defined as in (6), and a point \mathbf{p}_k^t that belongs to it, defined as in (7).

The hyperplane \mathcal{H}_k^t separates the points in a set of elements on the left side of the hyperplane and a set of elements on the right side of the hyperplane.

$$\begin{aligned} L_k^t &= \{\mathbf{x} | \mathbf{x} \in X_k^t, \mathbf{v}_k^t \cdot \mathbf{x} > \mathbf{v}_k^t \cdot \mathbf{p}_k^t\}, \\ R_k^t &= X_k^t \setminus L_k^t. \end{aligned} \quad (12)$$

ExIFFI computes a vector of feature importances for each node of the tree, based on two intuitions:

- The importance of the node k for a point x is higher when k creates a greater inequality between the number of elements on each side of the hyperplane, and x is in the smaller subset. Indeed, greater inequality means a higher grade of isolation for the points in the smaller set.
- For node k , the relative importance of the j th feature is determined by the projection of the normal vector of the splitting hyperplane onto that feature. If the split occurs along a single feature, that feature will receive the entire importance score. If the splitting hyperplane is oblique, the importance scores of multiple features will be calculated based on their respective projections onto the normal vector of the hyperplane.

Thus, we assign an importance value function to every node of the trees that is the projection on the normal vector of the splitting hyperplane of the quotient between the cardinality of the sample before a particular node and after it, following the path of a sample \mathbf{x} . Thus, knowing that the splitting hyperplane \mathcal{H}_k^t of that node is determined by the pair $\{\mathbf{v}_k^t, \mathbf{p}_k^t\}$ ³:

$$\lambda_k^t(\mathbf{x}) = \begin{cases} \left(\frac{|X_k^t|}{|L_k^t|} \right) \text{abs}(\mathbf{v}_k^t), & \text{if } \mathbf{v}_k^t \cdot \mathbf{x} > \mathbf{v}_k^t \cdot \mathbf{p}_k^t \\ \left(\frac{|X_k^t|}{|R_k^t|} \right) \text{abs}(\mathbf{v}_k^t), & \text{otherwise} \end{cases} \quad (13)$$

For example, consider a node k with a hyperplane defined by $\mathbf{v}_k^t = [0.5, 0.7, 0.2]$ and $\mathbf{p}_k^t = [2.0, 2.5, 1.0]$. Suppose a point $\mathbf{x} = [2.5, 3.0, 1.5]$. We first calculate the dot products:

$$\mathbf{v}_k^t \cdot \mathbf{x} = 0.5 \times 2.5 + 0.7 \times 3.0 + 0.2 \times 1.5 = 1.25 + 2.1 + 0.3 = 3.65$$

$$\mathbf{v}_k^t \cdot \mathbf{p}_k^t = 0.5 \times 2.0 + 0.7 \times 2.5 + 0.2 \times 1.0 = 1.0 + 1.75 + 0.2 = 2.95$$

Since $\mathbf{v}_k^t \cdot \mathbf{x} = 3.65$ is greater than $\mathbf{v}_k^t \cdot \mathbf{p}_k^t = 2.95$, the point \mathbf{x} lies in R_k^t . If the subsets contain $|L_k^t| = 90$ and $|R_k^t| = 10$ points out of $|X_k^t| = 100$, the importance vector $\lambda_k^t(\mathbf{x})$ would be:

$$\lambda_k^t(\mathbf{x}) = \left(\frac{100}{10} \right) \times \text{abs}([0.5, 0.7, 0.2]) = [10 \times 0.5, 10 \times 0.7, 10 \times 0.2] = [5, 7, 2]$$

The vector of importances evaluated in the tree t for a point x is the sum of all the importance vectors of all the nodes that the element x passed through on its path to the leaf node in the tree t defined in Eq. (1):

$$\mathbf{I}_t(x) = \sum_{k \in \mathcal{P}_x^t} \lambda_{t,k}(x). \quad (14)$$

³ With $\text{abs}(\mathbf{v})$ we refer to the positive part of every element of the vector \mathbf{v} .

We then calculate the sum of the importance vector of the point x for all the trees in \mathcal{T} :

$$\mathbf{I}(x) = \sum_{t \in \mathcal{T}} \mathbf{I}_t(x) \quad (15)$$

We define $\mathbf{V}(x)$ as the sum of the vectors orthogonal to the hyperplanes of the nodes that an element x passes in a tree, then we calculate the sum of the values in all the trees:

$$\mathbf{V}(x) = \sum_{t \in \mathcal{T}} \sum_{k \in \mathcal{P}_x^t} \mathbf{v}'_k \quad (16)$$

4.2.1. ExIFFI: Global feature importance

To globally evaluate the importance of the features, we divide X into the subset of predicted inliers $\mathcal{Q}_I = \{\mathbf{x}_i \in X | \hat{y}_i = 0\}$ and the one of predicted outliers $\mathcal{Q}_O = \{\mathbf{x}_i \in X | \hat{y}_i = 1\}$ where $\hat{y}_i \in \{0, 1\}$ is the binary label produced by the thresholding operation indicating whether the corresponding data point \mathbf{x}_i is anomalous ($\hat{y}_i = 1$) or not ($\hat{y}_i = 0$).

We define the global importance vectors for the inliers and for the outliers by summing out the importance values introduced in Eq. (15):

$$\mathbf{I}_I = \sum_{\mathbf{x} \in \mathcal{Q}_I} \mathbf{I}(\mathbf{x}), \quad \mathbf{I}_O = \sum_{\mathbf{x} \in \mathcal{Q}_O} \mathbf{I}(\mathbf{x}). \quad (17)$$

Likewise the sum of the orthogonal vectors:

$$\mathbf{V}_I = \sum_{\mathbf{x} \in \mathcal{Q}_I} \mathbf{V}(\mathbf{x}), \quad \mathbf{V}_O = \sum_{\mathbf{x} \in \mathcal{Q}_O} \mathbf{V}(\mathbf{x}). \quad (18)$$

Due to stochastic sampling of hyperplanes, in order to avoid the bias generated by the fact that it is possible that some dimensions are sampled more often than others, the vectors of importance have to be divided by the sum of the orthogonal vectors.

$$\hat{\mathbf{I}}_I = \frac{\mathbf{I}_I}{\mathbf{V}_I}, \quad \hat{\mathbf{I}}_O = \frac{\mathbf{I}_O}{\mathbf{V}_O}. \quad (19)$$

To evaluate which are the most important features to discriminate a data as an outlier we divide the importance vector of the outliers by the one of the inliers. Eq. (19), and we obtain the global feature importance vector in the same vein as in the DIFFI algorithm (Carletti et al., 2023):

$$\mathbf{GFI} = \frac{\hat{\mathbf{I}}_O}{\hat{\mathbf{I}}_I}. \quad (20)$$

For clarity, the pseudocode of the ExIFFI algorithm for the computation of the Global Feature Importance is reported in 1.

4.2.2. ExIFFI: Local feature importance

The Local Feature Importance assumes significance primarily within the context of anomalous data points, especially from the point of view of applications. Indeed, providing explanations of samples deemed anomalous eases decision-making by domain experts, who can subsequently tailor their responses based on the salient features driving the anomaly of a single point. Let us take into account an element x , Eq. (15) gives a vector of importances $\mathbf{I}(x)$ of the sample x for each feature. Then the vector $\mathbf{V}(x)$ is the normalization factor of the feature importance. Thus, the Local Feature Importance (LFI) of an element x is the quotient:

$$\mathbf{LFI}(x) = \frac{\mathbf{I}(x)}{\mathbf{V}(x)}. \quad (21)$$

The pseudocode for the LFI computation is analogous to the one reported in 1 with the difference that Eq. (21) is used in place of Eq. (20) in the last step.

Algorithm 1: ExIFFI Algorithm

Input : Isolation based AD model F , Input dataset D
Output: Global Feature Importance vector $GFI \in \mathbb{R}^p$
 $\mathbf{I}(\mathbf{x}) \leftarrow (0)_p \quad \forall \mathbf{x} \in D$;
// Initialize all importance vectors to zero
foreach x in D **do**
 foreach Isolation Tree t in F **do**
 $\mathcal{P}_x^t = \{k \in t \mid x \text{ traverses node } k \text{ in tree } t\}$;
 $\mathbf{I}_t(x) \leftarrow \sum_{k \in \mathcal{P}_x^t} \lambda_{k,t}(x)$ // Aggregate the importance scores of all nodes in the path from the root to the leaf node containing x
 $\mathbf{V}_t(x) \leftarrow \sum_{k \in \mathcal{P}_x^t} v_{k,t}(x)$;
 // Aggregate the vectors orthogonal to hyperplanes of the nodes in the path from the root to the leaf node containing x
 $\mathbf{I}(\mathbf{x}) \leftarrow \sum_{t \in F} \mathbf{I}_t(\mathbf{x})$;
 // Aggregate the importance scores for x over all the trees in the forest F
 $\mathbf{V}(\mathbf{x}) \leftarrow \sum_{t \in F} \sum_{k \in \mathcal{P}_x^t} \mathbf{v}'_k$;
 // Aggregate the vectors of the direction of the cut for x over all the trees in the forest F
 end
end
 $\mathcal{P}_I, \mathcal{P}_O \leftarrow \text{Predict}(D, F)$;
// Get the predicted inliers and outliers
 $\mathbf{I}_I \leftarrow \sum_{\mathbf{x} \in \mathcal{P}_I} \mathbf{I}(\mathbf{x})$;
 $\mathbf{I}_O \leftarrow \sum_{\mathbf{x} \in \mathcal{P}_O} \mathbf{I}(\mathbf{x})$;
// Sum together the importance vectors of all inliers and outliers
 $\mathbf{V}_I \leftarrow \sum_{\mathbf{x} \in \mathcal{P}_I} \mathbf{V}(\mathbf{x})$;
 $\mathbf{V}_O \leftarrow \sum_{\mathbf{x} \in \mathcal{P}_O} \mathbf{V}(\mathbf{x})$;
// An analogous operation is done for the orthogonal vectors
 $\hat{\mathbf{I}}_I \leftarrow \frac{\mathbf{I}_I}{\mathbf{V}_I}$;
 $\hat{\mathbf{I}}_O \leftarrow \frac{\mathbf{I}_O}{\mathbf{V}_O}$;
// Compute the final normalized importance vectors separately for inliers and outliers
 $GFI \leftarrow \frac{\hat{\mathbf{I}}_O}{\hat{\mathbf{I}}_I}$;
return GFI ;

4.2.3. Visualizing explanations

Miller (2019) defines interpretability in AI models as the extent to which a human can comprehend the rationale behind a decision. To be effective, interpretability should deliver a clear and comprehensible representation of how inputs influence outputs, even for individuals who are not experts in the field.

To bolster users' trust in the model, relying solely on a series of numerical Scores is insufficient. Providing a series of summary scores and comprehensible graphical representations of these scores may help the evaluation of the model outputs and bolster its interpretational efficacy.

To achieve these goals, two distinct graphical representations are proposed

GFI Score Plot The Score Plot is the average importance score across runs of the algorithm. For clarity, we introduce the Score Plot, a horizontal bar plot displaying average Global Importance Scores across training executions (Fig. 2(a)).

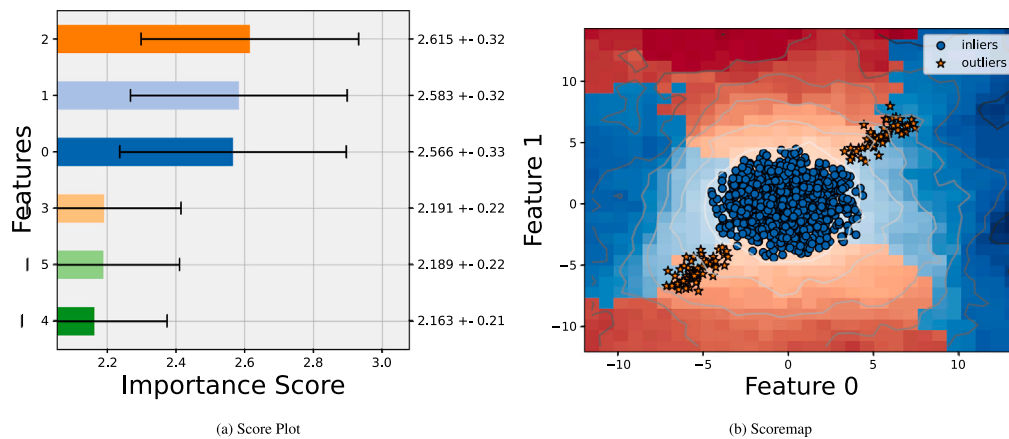


Fig. 2. Score Plot 2(a), and Scoremap 2(b) of the dataset Bisect3d.

LFI Scoremap The Scoremap provides local interpretability by computing Local Importance Scores for feature pairs. Each grid point represents the feature with the highest score, shown in red or blue for the first and second features respectively. Darker shades indicate higher Feature Importance Scores. Scatter plots overlay data points, distinguishing inliers (blue dots) and outliers (red stars). Contour plots offer contextual insights of the anomaly score regions. Focusing on the most important feature pair is common practice, facilitating anomaly dispersion analysis. For instance, in the Bisect3D dataset (Fig. 2(b)), this strategy was employed.

The GFI Score Plot represents the outcome of multiple model runs, 40 in this paper (Section 6), to account for the stochastic nature of IF-based models like EIF and EIF⁺. This approach recognizes the inherent randomness of these models and leverages multiple runs to capture variability, enhancing the reliability of interpretability results and providing deeper insight into the influence of individual features on model output.

To illustrate the interpretation of the graphical representation, we use the Bisect3D dataset introduced in Appendix A.1. In this dataset, anomalies are placed along the bisector of Features 0, 1, and 2, while inliers are centered around the origin, following a normal distribution across six dimensions. Consequently, Features 0, 1, and 2 play a key role in detecting outliers, and due to model stochasticity, any of these three features can be top-ranked in different executions.

The Score Plot in Fig. 2(a) shows nearly identical Importance Scores for these three features, with the rest having consistently lower scores. The feature ranking may vary with additional model runs.

The scatter plot in Fig. 2(b) highlights that outliers, represented by stars, deviate from the ball of inliers. The color distribution in the Scoremap indicates how the Local Importance Score varies across feature space. Lighter shades dominate the center, while the further from the distribution of inliers, the darker the color are, we observe that the anomalies lay in between the change of the color of the most important feature, underlying how the outliers are distributed in the intersection of the two features.

To sum up, ExIFFI serves as the feature importance mechanism for the EIF, extending the ideas from DIFFI, which focuses on the IF. While DIFFI evaluates feature importance by assessing the imbalance created by single-feature splits in IF trees, ExIFFI adapts this approach to the more complex hyperplane splits in EIF. By leveraging the tree structure, ExIFFI determines how each feature contributes to creating imbalanced splits, ranking features based on their importance in detecting anomalies within the EIF framework. This structural method allows for understanding feature importance both globally, across the entire dataset, and locally, for individual data points.

5. Experimental setup

To evaluate the performance of the EIF⁺ model and the interpretability of ExIFFI, we conducted experiments using various benchmark datasets, detailed in Appendix A.1. The implementation for EIF⁺, ExIFFI, and the experimental setup are publicly available at <https://github.com/alessioarcudi/ExIFFI>. In Section 5.1, we present in detail how we will conduct the experiments about the model's ability to detect anomalies, especially under low contamination conditions, and provide a comparative analysis against other AD models. Additionally, in Section 5.2, we present instead the proxy tasks used in order to evaluate the accuracy of ExIFFI's interpretability and comparing it with the ad-hoc method DIFFI and ECOD and various post-hoc interpretation techniques.

Following the evaluation of the datasets, we analyze the soundness and computational efficiency of these algorithms. In Section 6.5 we compute the correlation between the local importance scores provided by different interpretation algorithms with the anomaly scores provided by the explained AD model with the aim of identifying a strong linear relationship between these two quantities, which serves as a proof of the correctness of the explanations. Finally, in Section 6.6 we compare the algorithms considering the time needed to fit the model, predict the results and computing feature importances. This examination highlights the efficiency of the proposed models.

5.1. Performance evaluation

We assess the performance of the EIF⁺ model by benchmarking it against several anomaly detection (AD) models. This evaluation includes both well-established models from the literature and a state-of-the-art unsupervised AD model. The models compared are:

- Traditional Isolation-based approaches such as Isolation Forest (IF) and Extended Isolation Forest (EIF).
- The novel Deep Isolation Forest (DIF), an advanced Isolation-based model integrating deep learning techniques for enhanced AD (Xu et al., 2023).
- The Anomaly AutoEncoder (Aggarwal, 2015), a deep learning-based AD model, recognized for its robust performance in unsupervised AD scenarios.
- The Empirical Cumulative distribution based Outlier Detection (ECOD) (Li et al., 2023a), a non-parametric AD method which uses the Empirical Cumulative Distribution function (ECDF) of each single feature to compute the anomaly score, following the intuition that anomalous points should be placed on the tails of distributions.

In our experiments, we will evaluate the performance of anomaly detection (AD) models across different levels of contamination within the training dataset. This analysis is vital, as highlighted by Kim et al. (2023), who stress the importance of examining the robustness of AD models when the “normality assumption” (that all training data are inliers) does not hold.

It is crucial to assess how AD models respond to varying levels of contamination in the training data, where normal and anomalous samples are mixed without labels. This evaluation not only reveals the models’ effectiveness in real-world scenarios but also highlights the impact of anomalies concentrated in specific regions of the data space. When anomalies are misclassified as normal, the model may adapt to these outliers, distorting its representation of the overall data distribution. As a result, the model’s ability to correctly identify and isolate true anomalies can be significantly diminished.

Specifically, our analysis studies anomaly contamination’s effect on average precision, focusing on EIF versus EIF⁺ performance differences. This thorough assessment benchmarks EIF⁺ against traditional methods and showcases its generalization ability.

With reference to the analysis of the performances varying the contamination, we consider two distinct experimental scenarios chosen to mirror the varied conditions under which AD models might be deployed, offering a comprehensive understanding of their adaptability and effectiveness:

- Scenario I: We fit and evaluate the models using the entire dataset.
- Scenario II: We fit the models exclusively on the inliers within the dataset, that correspond to the scenario of 0 contamination of outliers. Subsequently, we assess the Average Precision scores of these models when applied to the entirety of the dataset.

Finally in Section 6.3 the performances of the AD models are analyzed in terms of average precision, precision and AUC-ROC score. Results are detailed in the Table 2.

5.2. Interpretability evaluation

Following the performance assessment, we shift our focus to the interpretability of AD models, evaluating the ExIFFI algorithm, an ad-hoc interpretation algorithm for the isolation based anomaly detections such as IF, EIF and EIF⁺.

Our evaluation is carried out in two stages. Initially, we qualitatively examine the ExIFFI algorithm’s interpretive method through the plots detailed in Section 4.2.3, observing how the resulting visualizations provide insights on the nature of the anomalies.

We then proceed to assess our model interpretive effectiveness by comparing it with alternative interpretation algorithms.

- Ad-hoc algorithm named DIFFI, introduced in Section 4.1.
- Leveraging the ad-hoc interpretation of ECOD model (Li et al., 2023a) we developed a local and global interpretation on its feature outlier scores to create a common environment of comparison. The model assigns an outlier score s_i to each feature i given a sample x , and the paper analyzes feature outliers by comparing each feature’s score $s_i(x)$ to the 99th percentile P_{99} , saying that the closer s_i is to P_{99} , the more outlier-like the feature. Then we defined the LFI as:

$$I_{\text{local}}(x) = \frac{1}{1 + (s_i(x) - P_{99})^2}$$

The GFI, similar to ExIFFI, is the mean importance of the outliers P_O features divided by the mean importance of the inliers P_I :

$$I_I = \frac{\sum_{x \in P_O} I(x)}{|P_O|} \frac{|P_I|}{\sum_{x \in P_I} I(x)}$$

- Post-hoc method that leverages Random Forest as a surrogate model. This approach involves using an inherently interpretable model to estimate the predictions of the initial Isolation Forest models (IF, EIF, or EIF⁺), with interpretations derived from the feature importance scores that the surrogate model provides.

In the absence of a one-size-fits-all metric or fully comprehensible prior knowledge on the features, assessing the performances of such algorithms is challenging due to the dependency on various factors, including input data complexity, model intricacies, and end-user interpretability, then the evaluation becomes a delicate balancing act between quantifiable measures and qualitative insights.

To quantitatively evaluate model interpretability, a proxy task is often utilized as an indirect means to assess model performance. In this paper the Feature Selection acts as such a proxy, where Importance Scores from interpretability methods prioritize input features. Then we assess the interpretation algorithm performance based on how it prioritizes key features versus less important ones offering insight into its effectiveness using the Feature Selection. Additionally, we add to these results those from a casual feature selection to gauge the overall performance trends more broadly.

We then introduce a metric to quantify the goodness of a Feature Selection, the Area Under the Curve of Feature Selection (AUC_{FS}), taking inspiration from the methodological evaluation used by Turbe et al. in Turbé et al. (2023). AUC_{FS} is calculated as the difference between the area under the curve when shrinking the dimensions of the dataset on the most important features (AUC_M) and the area under the curve when considering the least important features (AUC_L), $AUC_{FS} = AUC_M - AUC_L$. This metric reflects the distinction in selecting significant features or dropping them, with a larger gap indicating more effective prioritization, as on one hand losing vital information leads to a noticeable performance decline when important features are excluded early on, on the other hand maintaining the most important information will result in maintaining the performances or even increasing them. The pseudocode 2 presents in details how the AUC_{FS} score is calculated.

Algorithm 2: Feature Selection Proxy Task

```

Input : Isolation based AD model  $F$ ,  $GFI$  ranking  $\mathcal{R}$ , dataset
           $X$ 
Output: Average Precision vector  $AP$ 
 $AP_{\text{least}} \leftarrow [0, \dots, 0]$ ;
 $AP_{\text{most}} \leftarrow [0, \dots, 0]$ ;
 $X_{\text{least}} \leftarrow X$ ;
 $X_{\text{most}} \leftarrow X$ ;
for  $i \leftarrow 1$  to  $p$  do
     $AP_{\text{least}}[i] \leftarrow \text{average\_precision}(F, X_{\text{least}})$ ;
     $AP_{\text{most}}[i] \leftarrow \text{average\_precision}(F, X_{\text{most}})$ ;
     $X_{\text{least}} \leftarrow X_{\text{least}} \setminus X_{\text{least}}[:, \mathcal{R}[i]]$ ;
    // Inverse Approach: Remove the least
    // important feature
     $X_{\text{most}} \leftarrow X_{\text{most}} \setminus X_{\text{most}}[:, \mathcal{R}[i]]$ ;
    // Direct Approach: Remove the most important
    // feature
end
return  $(AP_{\text{least}}, AP_{\text{most}})$ ;

```

There exist various post-hoc interpretation algorithms like SHAP (Lundberg and Lee, 2017b) that could be integrated as a comparison in the proxy task. However, due to their computational burden, these algorithms are often less suitable for tree-based models, which are widely employed in industrial settings due to their speed and low memory requirements. Our experiments confirmed the unsuitability of SHAP in our evaluation 6.6, even considering the faster variants of SHAP (Lundberg et al., 2018), the required computational time necessary to apply it to the benchmark datasets is excessive, making it not viable in practice.

Table 1

Experimental datasets overview: The first column lists the dataset names; the second specifies the total number of instances; the third details the count of outliers; the fourth shows contamination rates; the fifth indicates the total features; the final column reveals the dataset's dimensionality.

	Data n	Anomalies	Contam %	Features d	Size	Dataset type
Xaxis	1 100	100	9.09	6	(Low)	Synthetic
Bisect	1 100	100	9.09	6	(Low)	Synthetic
Bisect3D	1 100	100	9.09	6	(Low)	Synthetic
Bisect3D_Skewed	1 100	100	9.09	6	(Low)	Synthetic
Bisect6D	1 100	100	9.09	6	(Low)	Synthetic
Anthyroid	7 200	534	7.56	6	(Low)	Real
Breastw	683	239	52.56	9	(Middle)	Real
Cardio	1 831	176	9.60	21	(High)	Real
Glass	214	29	13.55	9	(Middle)	Real
Ionosphere	351	126	35.71	33	(High)	Real
Pendigits	6 870	156	2.27	16	(Middle)	Real
Pima	768	268	34.89	8	(Middle)	Real
Shuttle	49 097	3 511	7.15	9	(Middle)	Real
Wine	129	10	7.75	13	(Middle)	Real
Diabetes	85 916	8 298	9.65	4	(Low)	Real
Moodify	276 260	42 188	15.27	11	(Middle)	Real

Our interpretability evaluation bypasses human experiments due to their high cost in terms of required time and effort, focusing instead on metric-based assessments. Like other studies (Turbé et al., 2023; Wojtas and Chen, 2020), we adopt, according to the Doshi-Velez's et al. taxonomy (Doshi-Velez and Kim, 2017), a Functionally-grounded evaluation by using a proxy task (Feature Selection). This method is neither specific nor expensive, making it ideal for our purposes. It provides a practical framework for assessing interpretability without the significant resources required for human-based evaluations, aligning with our research constraints and objectives.

6. Experimental evaluation of EIF⁺ and ExIFFI

This section evaluates EIF⁺ and ExIFFI, demonstrating their improvements over existing models. We focus on performance, particularly interpretability and anomaly detection. The analysis also considers their practicality in various settings by examining computational time.

We use a benchmark of 16 datasets with labeled anomalies to evaluate the performances of the models. This benchmark includes 5 synthetic datasets, which were designed to highlight the differences in model performance and to provide a ground truth about anomalies and model interpretation, as well as 11 open source datasets based on real applications. In the rest of the paper, the datasets will be indicated using a typewriter font. Table 1 summarizes the key characteristics of the datasets that were examined. A detailed description can be found in Appendix A.1.

In the following sections, we present the qualitative results for two synthetic datasets: Xaxis and Bisect3D_Skewed (Section 6.1), along with the real-world dataset glass (Section 6.2). Next, we provide the quantitative results of the AD performance and its interpretation in Sections 6.3 and 6.4. Additional graphical results can be found in the repository at <https://github.com/alessioarcudi/ExIFFI>.

6.1. Synthetic datasets

We evaluate the performance of EIF⁺ and ExIFFI in detecting anomalies and identifying the main features for the detection of anomalous points through the graphical representations illustrated in Section 4.2.3.

Synthetic datasets are designed to showcase the functionality and behavior of the models and explanation algorithms, both in situations where a single feature drives the detection of an anomaly and in cases where multiple features contribute to anomaly identification. By constructing these datasets, we can replicate specific conditions and assess how effectively the model interprets anomalies when the underlying anomaly-generating mechanism is fully understood. This method enables a precise evaluation of the model's interpretative abilities.

The construction of each synthetic dataset is inspired by from the research paper authored by Carletti et al. (2023). The inliers are points that lie within a p -dimensional sphere centered at the origin with a radius r . These inliers are sampled from a uniform distribution in p dimensions, denoted as $\mathcal{U}_p([-r, r])$, and are included in the dataset if their L2 norm is less than r . Mathematically, this set of inliers is represented as $\mathcal{D}_I = \{x \sim \mathcal{U}_p([-r, r]) \mid \|x\|^2 \leq r\}$. This ensures that all inlier points are uniformly distributed within the defined boundary, providing a well-defined structure for the normal data.

Outliers, on the other hand, are generated using a different approach to simulate anomalies. These outliers are uniformly distributed along one or more axes, characterized by k anomalous features. The remaining $p - k$ features act as noise, sampled from a Gaussian distribution $\mathcal{N}(0, 1)$. The anomalous points are determined by the following parameters:

- A unit vector $\mathbf{u} \in \mathbb{S}^{n-1}$, where \mathbb{S}^{n-1} represents the unit sphere in \mathbb{R}^n , and $u_i \neq 0$ if i is an anomalous feature,
- A scalar distance d from the origin,
- A random variable x , sampled from a uniform distribution $x \sim \mathcal{U}([\min, \max])$,
- A noise vector ϵ , where each anomalous feature is perturbed by Gaussian noise: $\epsilon_i \sim \mathcal{N}(0, 1)$ if i is anomalous.

Thus, the outlier points are generally defined as:

$$\mathcal{D}_O = \left\{ d \cdot \mathbf{u} + x \cdot \mathbf{u} + \epsilon \mid d \in \mathbb{R}, \mathbf{u} \in \mathbb{S}^{n-1}, x \sim \mathcal{U}([\min, \max]), \right. \\ \left. \epsilon_i \sim \mathcal{N}(0, 1) \text{ for each anomalous feature } i \right\}$$

This equation captures the generation of outliers as a combination of a scaled unit vector, a random displacement, and added noise, reflecting the anomalous behavior in the dataset.

In this section, we examine anomalies in datasets with $p = 6$ dimensions, varying the number of anomalous features, k . Two datasets are constructed to represent different scenarios: Xaxis, where $k = 1$, $r = 5$, $d = 5$, $\min = 0$, $\max = 5$ and anomalies occur along a single dimension thus $u = [1, 0, 0, 0, 0, 0]$, and Bisect3D_Skewed, where $k = 3$, $r = 5$, $d = 5$, $\min = 0$, $\max = 5$ and anomalies are distributed along a skewed vector ($u = [4, 3, 2, 0, 0, 0] / \|[4, 3, 2, 0, 0, 0]\|$). These datasets are used to illustrate how feature importance scores change as anomalies along one feature or along multiple feature with an order of how much each feature is important to determine the anomalies.

For a deeper explanation on the synthetic datasets generation process please refer to the pseudocode presented in Appendix 3. The other dataset used in the quantitative evaluation are presented in the Appendix A.1 and are named Bisect, Bisect3D and Bisect6D

Table 2

Precision report: Scenario 1 (S1) and Scenario 2 (S2) performance metrics for various anomaly detection models across different datasets.

Scenario	Model	Measure	Synthetic															Real World	Mean
			Xaxis	Bisec	Bisec3D	Bisec6D	Bisec3D_Skewed	Anthyroid	Breastw	Cardio	Diabetes	Glass	Ionosphere	Moodify	Pendigits	Plma	Shuttle		
S1	IF	Avg Prec	0.09	0.75	0.95	0.99	0.6	0.30	0.96	0.53	0.46	0.36	0.79	0.35	0.28	0.51	0.98	0.26	0.57
		Prec	0.11	0.74	0.85	0.98	0.51	0.32	0.90	0.44	0.46	0.34	0.65	0.30	0.31	0.52	0.96	0.20	0.54
		ROC AUC	0.51	0.85	0.91	0.98	0.63	0.63	0.90	0.69	0.70	0.62	0.73	0.59	0.64	0.63	0.97	0.56	0.72
	EIF	Avg Prec	0.97	0.97	0.96	0.96	0.77	0.19	0.91	0.55	0.47	0.35	0.82	0.33	0.23	0.49	0.90	0.21	0.63
		Prec	0.89	0.91	0.89	0.92	0.71	0.24	0.86	0.59	0.48	0.27	0.69	0.37	0.24	0.49	0.87	0.20	0.60
		ROC AUC	0.93	0.95	0.93	0.95	0.84	0.59	0.86	0.77	0.71	0.58	0.76	0.63	0.61	0.61	0.93	0.56	0.76
	EIF+	Avg Prec	0.91	0.94	0.93	0.93	0.75	0.20	0.88	0.50	0.44	0.31	0.84	0.28	0.21	0.49	0.70	0.17	0.59
		Prec	0.89	0.83	0.71	0.81	0.64	0.28	0.85	0.52	0.49	0.24	0.72	0.27	0.41	0.50	0.71	0.00	0.55
		ROC AUC	0.93	0.90	0.84	0.89	0.8	0.61	0.84	0.73	0.72	0.56	0.78	0.57	0.69	0.61	0.84	0.45	0.73
	DIF	Avg Prec	0.06	0.16	0.30	0.19	0.25	0.21	0.44	0.58	0.10	0.30	0.87	0.26	0.34	0.41	0.54	0.07	0.31
		Prec	0.06	0.20	0.28	0.16	0.32	0.26	0.49	0.52	0.90	0.20	0.76	0.26	0.39	0.41	0.60	0.00	0.36
		ROC AUC	0.48	0.56	0.60	0.53	0.62	0.60	0.46	0.73	0.49	0.54	0.81	0.56	0.68	0.54	0.78	0.45	0.59
	AE	Avg Prec	0.52	0.89	0.97	0.97	0.83	0.21	0.61	0.62	0.45	0.41	0.76	0.45	0.22	0.43	0.90	0.14	0.59
		Prec	0.46	0.82	0.92	0.95	0.78	0.25	0.44	0.62	0.47	0.44	0.60	0.53	0.33	0.44	0.95	0.10	0.57
		ROC AUC	0.70	0.90	0.95	0.97	0.87	0.59	0.41	0.79	0.70	0.68	0.68	0.72	0.65	0.57	0.97	0.51	0.73
	ECOD	Avg Prec	0.1	0.99	0.97	1.00	0.82	0.26	0.85	0.65	0.29	0.38	0.56	0.5	0.26	0.46	0.91	0.19	0.57
		Prec	0.1	0.99	0.9	1.00	0.75	0.3	0.75	0.53	0.35	0.41	0.54	0.48	0.35	0.45	0.87	0.1	0.55
		ROC AUC	0.5	0.99	0.94	1.00	0.86	0.62	0.74	0.74	0.64	0.66	0.64	0.69	0.67	0.58	0.93	0.51	0.73
S2	IF	Avg Prec	0.10	0.99	1.00	1.00	0.98	0.45	0.99	0.68	0.26	0.65	0.89	0.72	0.37	0.58	0.98	0.61	0.70
		Prec	0.00	0.95	1.00	1.00	0.96	0.43	0.94	0.60	0.31	0.68	0.79	0.73	0.42	0.55	0.97	0.50	0.68
		ROC AUC	0.45	0.97	1.00	1.00	0.97	0.69	0.94	0.78	0.62	0.82	0.83	0.84	0.70	0.66	0.98	0.72	0.81
	EIF	Avg Prec	0.99	0.99	0.99	0.99	0.97	0.45	0.98	0.74	0.55	0.57	0.90	0.65	0.27	0.54	0.97	0.57	0.76
		Prec	1.00	1.00	1.00	0.99	0.98	0.41	0.93	0.72	0.61	0.68	0.82	0.64	0.36	0.57	0.95	0.60	0.77
		ROC AUC	1.00	1.00	1.00	0.99	0.98	0.68	0.93	0.84	0.78	0.82	0.86	0.79	0.67	0.67	0.97	0.78	0.86
	EIF+	Avg Prec	0.99	1.00	0.99	0.99	0.97	0.46	0.99	0.76	0.61	0.69	0.96	0.66	0.36	0.58	0.91	0.79	0.80
		Prec	1.00	1.00	0.99	1.00	0.95	0.46	0.95	0.73	0.56	0.79	0.87	0.66	0.45	0.57	0.96	0.70	0.79
		ROC AUC	1.00	1.00	0.99	1.00	0.97	0.70	0.95	0.85	0.76	0.88	0.90	0.80	0.72	0.67	0.97	0.83	0.87
	DIF	Avg Prec	0.97	0.85	0.79	0.36	0.6	0.39	0.60	0.79	0.13	0.85	0.96	0.72	0.46	0.49	0.99	0.72	0.67
		Prec	0.90	0.78	0.72	0.32	0.51	0.39	0.66	0.70	0.17	0.82	0.89	0.73	0.53	0.46	0.97	0.70	0.64
		ROC AUC	0.94	0.87	0.84	0.62	0.73	0.67	0.64	0.83	0.54	0.90	0.90	0.84	0.76	0.59	0.98	0.83	0.78
	AE	Avg Prec	1.00	1.00	1.00	1.00	0.95	0.44	0.98	0.83	0.50	0.53	0.87	0.65	0.24	0.50	0.90	0.45	0.74
		Prec	1.00	1.00	1.00	1.00	0.9	0.42	0.94	0.70	0.52	0.62	0.75	0.68	0.33	0.54	0.94	0.50	0.74
		ROC AUC	1.00	1.00	1.00	1.00	0.94	0.69	0.93	0.76	0.73	0.78	0.80	0.81	0.66	0.64	0.97	0.72	0.84
	ECOD	Avg Prec	0.14	0.97	0.99	1.00	0.92	0.32	0.97	0.64	0.4	0.45	0.7	0.58	0.28	0.5	0.95	0.25	0.63
		Prec	0.15	0.9	0.99	1.00	0.86	0.32	0.91	0.7	0.44	0.55	0.56	0.55	0.35	0.51	0.93	0.2	0.62
		ROC AUC	0.53	0.94	0.99	1.00	0.92	0.63	0.91	0.67	0.69	0.74	0.66	0.73	0.67	0.62	0.96	0.56	0.76

Performance

In the Xaxis dataset anomalies are restricted to specific intervals along a single feature (feature 0), while the rest of the data points are uniformly spread within a 6-dimensional space. Detecting these anomalies poses a challenge for algorithms like ECOD and Isolation Forest, as they evaluate features individually. In particular in high-dimensional spaces, anomalies confined to one dimension appear normal across others, making them harder to detect. This challenge is particularly evident in the well-known bias of Isolation Forest, pointed out by Hariri et al. (2021), where there is a lower anomaly score in the bundles orthogonal to all the axis but one around the data distribution. Since anomalies in this case are aligned with one axis, they are inside the bundle so they have a lower anomaly scores than anomalies with the same distance from the data distribution but in a combination of features.

This occurs because Isolation Forest as well as ECOD isolates each feature independently, leading to a skewed perception when other features behave normally. This bias motivated the development of the EIF algorithm, which enhances anomaly detection by partitioning the data space based on linear combinations of multiple features, rather than treating them individually. As a result, EIF achieves substantially better performance in detecting anomalies in datasets where the anomalies are spread across a single dimension, such as the Xaxis dataset, which Isolation Forest and ECOD struggle to identify.

In Bisect3D_Skewed the anomalies are aligned along the direction $u = [4, 3, 2, 0, 0, 0]$, then they are correlated to multiple features.

This improves the detection performance of both ECOD and Isolation Forest, as the anomalies exhibit more distinct behavior in the full feature space.

In the two synthetic datasets presented, and in the results of the synthetic datasets in the Table 2, the DIF perform poorly. This is due to the fact that the anomalies are distributed in a specific manner that does not leverage the strengths of DIF. DIF excels in detecting anomalies when there are complex, non-linear relationships in the data. However, when anomalies are confined to a specific feature (like in the feature 0 case) or are linearly aligned (along the direction $[4, 3, 2, 0, 0, 0]$), the model may not effectively distinguish them from the normal data because these scenarios might not exhibit the complex interactions between features that DIF is designed to exploit. This leads to the model's inability to properly isolate and detect these anomalies.

Finally, for what concerns the EIF and EIF+ models they showcase excellent performances with Average Precision values close to the perfect score of 1, as detailed in Table 2

Interpretation

The Fig. 4 shows the importance scores of ExIFFI applied to EIF+ in Scenario II. The bar charts in the top row in Figs. 4(a) and 4(b) illustrate the importance of each feature of the two datasets: Xaxis and Bisect3D_Skewed. The importance scores not only demonstrate the model's ability to accurately identify the key features contributing to the anomalies, assigning the highest scores to the most relevant

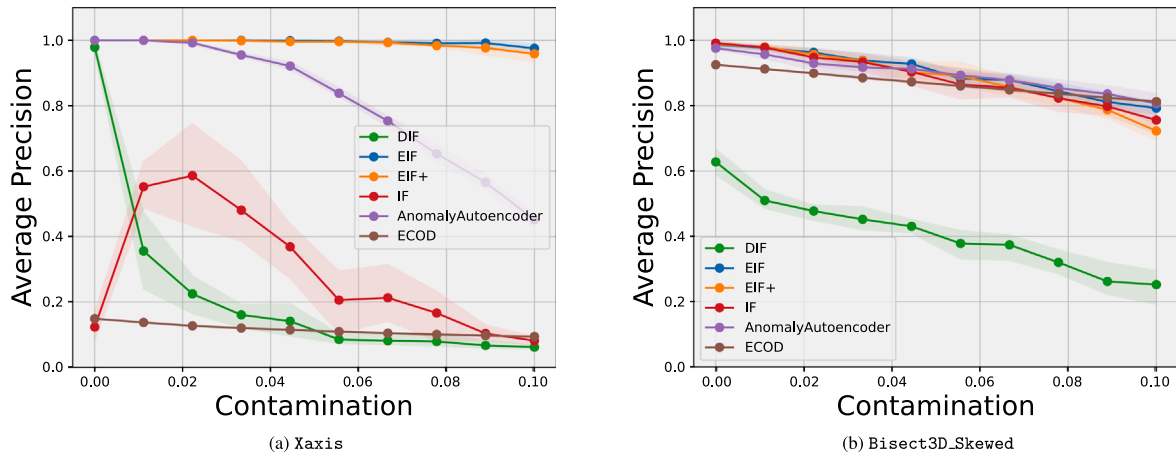


Fig. 3. Figs. 3(a), 3(b) showcase the behavior of the average precision metric as the contamination level in the training set increases for different AD models. In particular Fig. 3(a) is dedicated to Xaxis while 3(b) refers to Bisect3D_Skewed.

features, but also highlight its capability to distinguish between features with varying scores when anomalies are unevenly distributed among them. This nuanced differentiation underscores the model's effectiveness in capturing the relative significance of each feature based on the specific distribution of anomalies.

In the bottom row, the local importance score maps depict the relative importance of Feature 0 versus Feature 1 across different points in a 2D space. The color gradient indicates which feature is more important at each location (red for Feature 1 and blue for Feature 0). The maps demonstrate that the model consistently identifies the correct dominant feature across different regions, with clear transitions between blue and red areas, indicating precise differentiation between the influence of Feature 0 and Feature 1 in different parts of the data space. This ability to distinguish local importance highlights ExIFFI's robustness and effectiveness in feature selection under Scenario II.

The Table 3 presents an evaluation of feature selection methods applied to the two Scenarios presented in Section 5.1, highlighting key insights into their performance. The IF-based feature selection (IF_ExIFFI) yields poor interpretation results in the case of Xaxis dataset, which is consistent with the model's overall bad performance, indicating a direct link between the quality of the model and the reliability of the feature importance it generates. Similarly, the Random Forest-based method (IF_RF) also shows low interpretation results, reinforcing the notion that the accuracy of a model's performance is crucial for effective feature importance evaluation.

On the other hand, the other feature selection methods, including those based on EIF and EIF+, demonstrate very similar and strong performance across both scenarios. This consistency suggests that these methods successfully identified the most important features, contributing to their reliable performance across different conditions. Notably, the EIF+_{ExIFFI} and EIF_{RF} methods performed particularly well, especially in Scenario II, highlighting their superior generalization ability to unseen anomalies.

6.2. Real-world datasets

Real-world data presents a multifaceted benchmarking challenge, mainly due to the intricate distributions and the subtle distinction between inliers and outliers. These complexities often complicate the task of reliably identifying anomalies, which in turn can affect the interpretability of results. Our experiments aimed to shed light on the efficacy of our models in discerning the features that affect the most the anomalies.

We analyzed in details the results of the dataset named Glass due to its unique characteristics and the insights they provide into anomaly detection and feature interpretation. The Glass dataset offers clear

definitions of features and anomalous classes, supported by literature that helps validate the correctness of interpretation results. Moreover, this dataset's particular distribution, where some inliers are far from the normal distribution and some outliers are close to inliers, highlights the significant impact of contamination levels on model performance.

For a detailed view of the quantitative results of the datasets examined, refer to the Tables 2 and 3.

Performance

Fig. 5 shows that EIF+ significantly outperforms EIF when contamination levels are low, highlighting its superior generalization in scenarios where the model is fitted primarily on normal data. This improvement is evident in Fig. 6(b), where some inliers (blue dots) are far from the distribution and some outliers (marked with orange star) are inside the normal distribution, making it difficult for the model to distinguish between normal and anomalous data when fitted to the entire dataset. By focusing on normal data, EIF+ better discerns the two, leading to its marked advantage over EIF in these conditions.

A similar pattern can be observed for DIF, which also experiences substantial performance gains when contamination levels are low. While DIF struggles in high contamination scenarios, much like other models, its deep learning architecture allows it to better map data into informative subspaces when contamination decreases. This enables DIF to capture the underlying structure of the normal data more effectively, thereby isolating the remaining anomalies more accurately. However, as noted in another table, DIF lacks robustness; despite achieving high scores with some datasets in particular in Scenario II as it is possible to see in Table 2, its performance drastically drops in other datasets such as the synthetic ones, breastw, diabetes and anthyroid. Additionally, DIF's fit and prediction times analyzed in Section 6.6 are significantly longer compared to simpler models like IF and EIF, making it computationally expensive. In contrast, EIF+ not only improves upon EIF's accuracy but also maintains relatively robust results compared to the benchmark models and very low fit and prediction times, offering a balanced trade-off between performance and computational efficiency.

Interpretation

In this study, we chose to present the results from Scenario II due to the significant overlapping of anomalies with inliers in Scenario I, which led to poor performance. The strong overlap in Scenario I made it challenging for the models to accurately distinguish between normal and anomalous instances, resulting in less reliable feature importance interpretations. As a result, the findings from Scenario II, where the

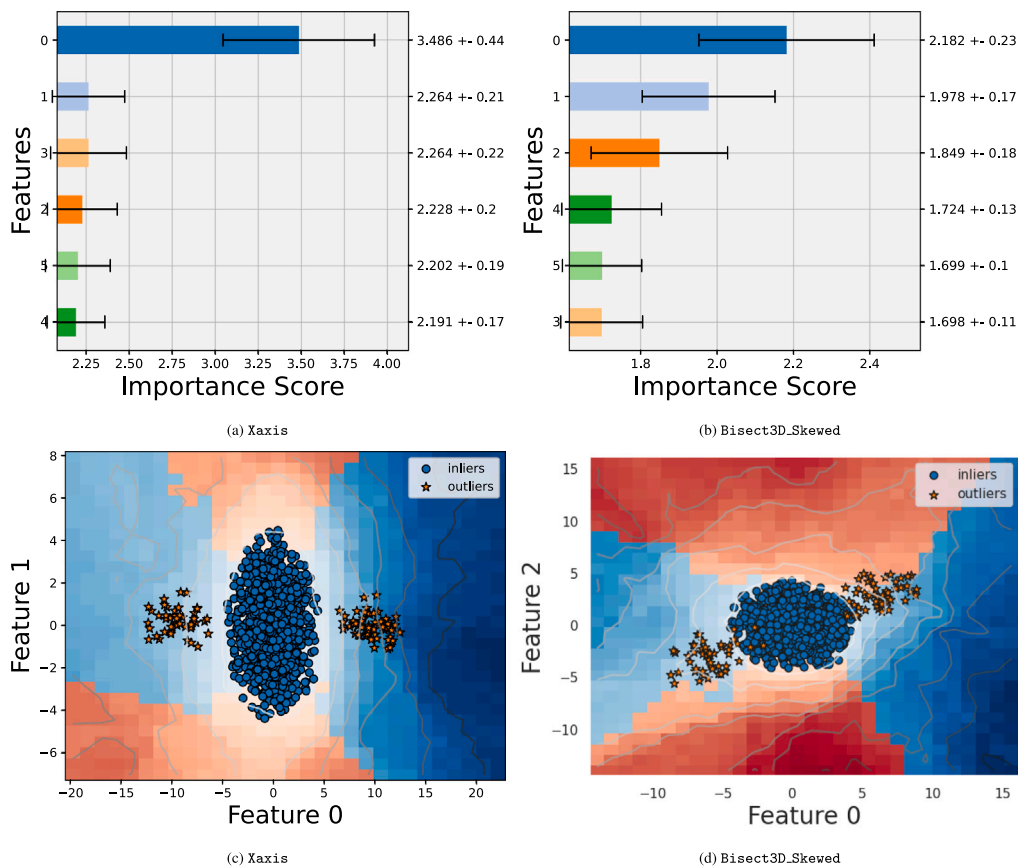


Fig. 4. Figs. 4(a), 4(b) display the Score Plots for the Xaxis and Bisect3D_Skewed datasets, respectively. Figs. 4(c) and 4(d) outline, instead, the Local Scoremaps of these two datasets. The feature importances used to generate these plots were obtained using the ExIFFI algorithm applied to the EIF⁺AD model.

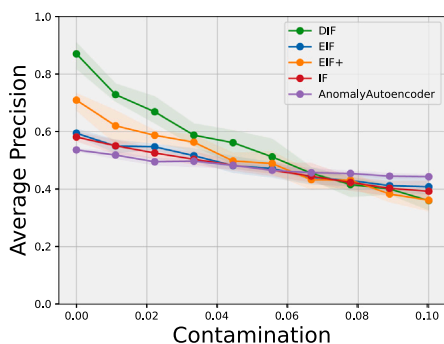


Fig. 5. Figure showcases the behavior of the average precision metric as the contamination level in the training set increases for different AD models trained on the glass dataset.

model was trained solely on inliers, provide a clearer and more robust analysis. For completeness.

As observable in Fig. 6(a), Barium (Ba) is emerged as the most important feature, followed by Potassium (K). This aligns with domain knowledge that we have on the dataset since the anomalous points represent the chemical composition of the headlamp glass, confirming the correctness of the approach, as Barium is known to be the key differentiator between headlamp and window glass (Mrabet et al., 2024).

The scoremap reveals that deviations in Potassium levels are not strongly correlated with the anomaly distribution. In fact, as shown in Fig. 6(b), the anomalous points (represented by orange stars) do not

deviate from the normal distribution along the Potassium feature. However, some inlier points clearly deviate in their Potassium composition and influence the interpretation output.

6.3. Performance report

Table 2 examines the performance achieved by different AD models across various datasets. The analysis focuses on Average Precision (Avg Prec), Precision (Prec), and ROC-AUC score metrics, under the two different scenarios; with outliers in the training dataset (S1) and without outliers (S2).

In S1 even if for each dataset the most performing model vary, the EIF showcases superior Avg Prec, Prec, and ROC AUC scores on average. This improvement underscores EIF’s enhanced anomaly detection capability, particularly in handling complex data structures, it is a model that robustly have high performances across all the possible datasets. However, the EIF⁺ model does not outperform EIF within this scenario, indicating its modifications do not significantly impact performance under standard test conditions, it remains the second most performing model, together with AE. S2’s findings mark a distinct improvement in EIF⁺’s performance over EIF, notably in real-world datasets. In fact EIF⁺ achieves parity or surpasses EIF in Avg Prec and Prec metrics in all the dataset but Shuttle, demonstrating its refined ability to detect anomalies. This enhanced performance is attributable to EIF⁺’s algorithmic adjustments, optimized to grasp the unseen outliers within complex data distributions more effectively than EIF. Moreover EIF⁺ demonstrate enhanced performances by figuring as the most performing model in 6 out of 11 real world datasets and reaching very high performances in synthetic ones. Thus it results as the most performing model on average across the datasets.

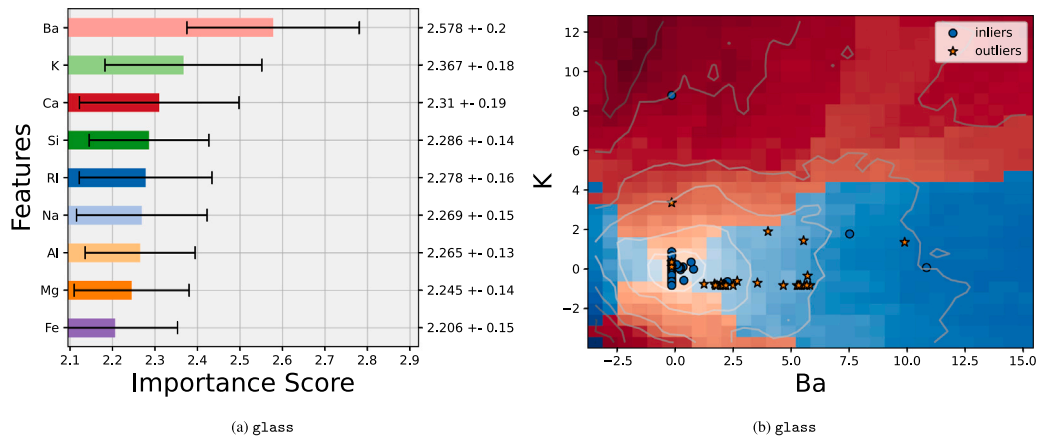


Fig. 6. Figs. 6(a), 6(b) display the Score Plot and Local Scoremaps for the glass dataset. The feature importances used to generate these plots were obtained using the ExIFFI algorithm applied to the EIF+ AD model.




Fig. 7. Average precision values for EIF and EIF+ across multiple datasets in two scenarios and the mean precision over all datasets. Results for both Scenario I and Scenario II are displayed for each dataset.

In Fig. 7 we observe how in Scenario I the results indicate that EIF performs slightly better than EIF+ across most datasets. This is reflected by marginally higher average precision values for EIF in the majority of the datasets. Although the difference is not substantial, it suggests that EIF has a small performance edge under these specific conditions. Both algorithms appear to be competitive in this scenario, with only slight variations in their performance.

In contrast, Scenario II demonstrates a significant improvement in the performance of EIF+, which outperforms EIF across almost all datasets except for the shuttle dataset, where EIF maintains a slight advantage. The mean results summarize this trend, showing higher average precision values for EIF+ in Scenario II. This suggests that EIF+ exhibits greater generalization capability when the anomalies are not in the training set, as reflected by its overall better performance across the majority of datasets in this scenario.

6.4. Importance score

Table 3 presents the performance of the newly introduced AUC_{FS} metric on 16 benchmark datasets in the two different Scenarios and the interpretations are evaluated with the EIF model or the EIF+. As detailed in 5, ExIFFI is benchmarked against DIFFI, an ad-hoc interpretability algorithm of the IF, post-hoc interpretability approach utilizing feature importance from a Random Forest surrogate model and the intrinsically interpretable model ECOD.

For synthetic datasets, the AUC_{FS} scores are generally high across all models showing that all the models found the right anomalous features, with the notable exception of the Xaxis dataset in the cases of IF_ExIFFI, IF_RF and ECOD present wrong results due to the lack of precision of the underlying models.

In the case of real-world datasets, the unique definitions of inliers and outliers complicate the interpretability performance assessment of the models. This issue is especially evident in the Breastw, Pima, and Moodify datasets, where the AUC_{FS} values are similar and predominantly negative.

For the other eight datasets, the data supports the effectiveness of ExIFFI, which consistently records the highest AUC_{FS} values in nearly all datasets and scenarios, except for the glass dataset in Scenario I. In ExIFFI applied to EIF or EIF+ have the highest average score in any context.

The analysis also reveals that the most effective interpretations by ExIFFI are generally those associated to the most performing isolation forest model, indicating that the AUC_{FS} metric values are strongly correlated with the efficacy of a specific prediction in a given dataset and scenario (i.e., Scenario I or Scenario II).

In conclusion, this comprehensive analysis confirms the effectiveness of the combined use of EIF+ and ExIFFI, achieving optimal results in both anomaly detection and interpretability.

Table 3

AUC_{FS} metric values for different interpretation algorithms with average precision evaluated with EIF and EIF⁺ in Scenario I (S1) and Scenario II (S2) across different datasets. The highest values are highlighted in bold.

Scenario	Model	Measure	Synthetic					Real World											Mean
			Xaxis	Bisec	Bisec3D	Bisec6D	Bisec3D_Skewed	Anthyroid	Breastw	Cardio	Diabetes	Glass	Ionosphere	Moodify	Pendigits	Pima	Shuttle	Wine	
S1	EIF	DIFFI	4.73	4.12	3.42	0.02	2.92	0.96	0.11	4.65	-0.36	1.59	2.36	-1.48	0.11	-0.48	-0.51	0.17	1.40
		IF_ExIFFI	-4.72	4.13	3.45	0.18	3.62	1.12	0.25	1.93	0.84	0.30	3.29	-1.68	2.19	-0.31	2.76	0.68	1.12
		EIF_ExIFFI	4.71	4.15	3.39	0.08	3.67	0.98	0.04	5.41	0.96	0.46	1.82	-1.78	1.58	-0.31	-1.43	0.45	1.51
		EIF ⁺ _ExIFFI	4.71	4.07	3.42	0.16	3.63	0.98	0.07	4.42	1.05	0.40	0.61	-1.51	1.86	-0.30	0.45	0.12	1.51
		IF_RF	-2.82	4.01	3.41	0.19	3.39	0.45	-0.14	1.48	-0.30	-1.73	-1.91	-0.80	1.48	-0.36	0.44	-3.09	0.23
		EIF_RF	4.72	4.05	3.44	0.16	3.59	-0.68	-0.18	1.35	0.49	-1.71	-1.97	-0.84	-0.70	-0.44	1.69	-2.41	0.66
		EIF ⁺ _RF	4.72	4.02	3.39	0.16	3.58	-1.19	-0.16	0.32	0.46	-1.46	-1.87	-1.88	-1.40	-0.28	1.90	-2.81	0.46
	ECOD	-4.7	3.35	3.58	0.21	3.63	-0.10	0.09	3.09	-1.01	-0.42	-0.61	2.50	2.56	-0.07	2.20	-0.03	0.89	
	EIF ⁺	DIFFI	4.54	4.16	3.32	-0.02	2.75	0.91	0.08	4.39	-0.34	1.55	-1.13	-1.11	0.23	-0.49	-0.69	0.43	1.16
		IF_ExIFFI	-4.49	4.12	3.46	0.18	3.11	1.10	0.28	2.40	0.81	0.48	2.68	-1.16	2.17	-0.20	2.78	0.62	1.15
		EIF_ExIFFI	4.55	4.04	3.38	0.05	3.23	0.93	0.02	5.12	0.82	0.56	1.61	-1.37	1.76	-0.32	1.22	0.45	1.63
		EIF ⁺ _ExIFFI	4.52	4.05	3.48	.004	3.16	0.97	0.08	4.25	0.74	0.52	0.45	-0.99	1.92	-0.27	0.73	0.18	1.49
		IF_RF	-2.68	3.93	3.44	0.07	3.00	0.42	-0.14	1.28	-0.30	-1.31	1.76	-0.72	1.37	-0.30	0.36	-2.33	0.49
		EIF_RF	4.56	3.95	3.43	0.05	3.11	-0.67	-0.16	1.35	0.38	-1.51	-1.85	-0.71	-0.36	-0.41	1.73	-1.57	0.70
EIF ⁺ _RF		4.57	4.01	3.40	0.08	3.16	-1.16	-0.14	0.36	0.40	-1.30	-1.56	-1.43	-0.97	-0.28	1.62	-2.05	0.54	
ECOD	-4.47	3.31	3.42	-0.04	3.19	-0.12	-0.01	2.82	-0.71	-0.11	-0.34	2.07	2.71	-0.12	2.72	0.20	0.91		
S2	EIF	DIFFI	2.83	2.92	2.79	-0.04	1.50	2.04	0.06	4.76	-1.45	2.11	-1.33	-3.66	-0.38	0.25	-1.48	-3.79	0.45
		IF_ExIFFI	-4.73	3.84	3.01	0.06	4.07	2.25	0.08	3.40	-1.01	3.45	0.89	4.39	2.68	0.37	2.88	8.07	2.10
		EIF_ExIFFI	4.72	3.84	3.01	0.11	4.05	2.03	0.06	8.28	0.35	2.55	4.31	-0.35	1.88	-0.21	0.76	7.12	2.65
		EIF ⁺ _ExIFFI	4.74	3.82	3.02	0.07	3.99	2.18	0.02	8.60	0.82	1.76	2.26	-0.23	1.22	-0.03	0.01	9.00	2.58
		IF_RF	-2.85	3.82	3.05	0.10	4.04	2.05	0.09	7.69	-0.05	2.18	-1.23	2.07	1.28	0.19	0.90	4.41	1.73
		EIF_RF	4.73	3.83	3.03	0.06	4.05	2.03	0.10	7.78	0.31	-1.63	1.33	1.63	-0.84	0.12	2.74	3.94	2.06
		EIF ⁺ _RF	4.74	3.83	3.04	0.04	4.06	2.09	0.10	8.24	0.33	1.94	3.58	1.47	-0.24	0.06	2.13	4.87	2.52
	ECOD	-4.73	3.02	3.05	0.12	4.07	-0.25	0.07	6.26	-0.69	-0.71	-2.54	4.29	2.82	0.01	2.84	-1.59	1.00	
	EIF ⁺	DIFFI	2.728	2.92	2.83	-0.12	1.52	2.00	0.06	4.33	-1.48	2.44	1.92	-3.36	-0.76	0.30	-1.45	-3.49	0.65
		IF_ExIFFI	-4.74	3.82	3.01	0.09	4.03	2.23	0.08	3.40	-0.93	3.41	0.26	4.23	3.18	0.36	3.07	9.40	2.18
		EIF_ExIFFI	4.74	3.83	3.02	0.13	4.04	1.92	0.08	8.38	0.31	2.43	2.84	-0.20	2.23	-0.20	0.31	8.12	2.62
		EIF ⁺ _ExIFFI	4.74	3.83	3.03	0.08	4.03	2.17	0.06	8.73	1.17	2.26	1.41	-0.04	1.32	-0.02	-0.13	9.25	2.62
		IF_RF	-2.85	3.83	3.03	0.11	4.04	1.99	0.06	7.88	-0.12	2.28	-1.08	2.48	1.38	0.21	0.64	4.53	1.78
		EIF_RF	4.74	3.84	3.01	0.06	4.02	1.97	0.09	7.85	0.35	-1.64	0.69	2.17	-1.28	0.17	2.11	4.36	2.03
EIF ⁺ _RF		4.74	3.83	3.01	0.07	4.01	2.02	0.07	8.30	0.29	1.63	2.17	2.08	-0.74	0.14	2.11	5.64	2.46	
ECOD	-4.74	3.03	3.03	0.07	4.04	-0.16	0.07	6.04	-0.37	0.17	-2.06	4.57	2.75	0.05	2.90	-0.60	1.17		

6.5. Correlation table

In this section we want to verify the soundness of the results produced by the interpretation algorithms analyzed in this study. This is achieved by computing the correlation between the Local Importance Scores produced by an interpretation algorithm and the Anomaly Scores assigned by the explained AD models to different samples. Intuitively if the interpretation model is able to correctly detect anomalous points then, given an input data point x , high LFI scores should correspond to an high value on the Anomaly Score (i.e. x is an outlier). Contrarily low importance values will be paired with a reduced Anomaly Score (i.e. x is an inlier). In order to compute the correlation between the LFI score and the Anomaly Scores the former were aggregated through the sum in order to obtain a single value.

In Table 4 most of the correlation values exceed 0.9 proving the effectiveness of the proposed interpretation methods which are able to provide sufficiently high relevance scores to features describing anomalous samples. An exception is represented by the ECOD and IF_ExIFFI algorithms which perform poorly in correlation terms (i.e. lower than 0.75) on the Xaxis dataset. The reason under these poor performances can be found on the limitations of these two AD models which fail to detect outliers on Xaxis as already noted in Section 6.1.

6.6. Time scaling experiments

In this section, we compare the computational efficiency of different algorithms, evaluating the time required for model fitting, making predictions, and interpreting results. Our examination, depicted in Fig. 8, presents how these durations adjust with dataset size increases—both by sample number and feature count, on a logarithmic scale. This investigation is crucial for understanding each algorithm's practical use

in varying scenarios, such as industrial-sized datasets, where the sample size and the feature space are generally larger than the ones encountered in the benchmark datasets used for evaluation in 5. Figs. 8(a), 8(d) details fitting times varying the sample size and number of features respectively, Figs. 8(b), 8(e) outline prediction times, and Figs. 8(c), 8(f) the time to determine importance scores, providing insight into the scalability and efficiency of these computational processes.

Analyzing computational efficiency, our observations reveal significant differences in the fitting, prediction, and interpretation speeds of various algorithms, as depicted in Figs. 8(a), 8(d), 8(b), 8(e), and 8(c), 8(f). Specifically, the fitting time, as shown in Fig. 8(a), demonstrates exponential increases for both the DIF and Anomaly Autoencoder models. The Anomaly Autoencoder exhibits a steeper rise due to weight optimization requirements, contrasting with DIF, which utilizes deep neural networks with random weights for nonlinear input data transformation, hence avoiding lengthy NN training periods.

Additionally, ECOD scales exponentially as the sample size of the dataset increases. Although ECOD is faster than IF-based models when the sample size is small, it becomes slower when the dataset contains more than 30k points. Interestingly, the number of features does not significantly affect the efficiency of ECOD, making it less sensitive to feature-dimensionality compared to other models.

For the experiment in which time effectiveness is tested by varying the number of features in the dataset, the Anomaly Autoencoder remains the most time-consuming method, while the trend of DIF is closer to that of the Isolation-based methods (i.e., EIF, EIF⁺, IF), as shown in 8(d).

The prediction time analysis, referenced in Fig. 8(b), highlights DIF as the most time-intensive model. This is attributed to DIF's necessity to aggregate predictions across an ensemble of neural networks and representation spaces, significantly extending the prediction process.

Table 4

Correlation values between local importance scores and anomaly scores for different interpretation algorithms in Scenario I (S1) and Scenario II (S2) across different datasets. The highest values are highlighted in bold.

Scenario	Model	Synthetic					Real World										
		Xaxis	Bisec	Bisec3D	Bisec6D	Bisec3D_Skewed	Anthyroid	Breastw	Cardio	Diabetes	Glass	Ionosphere	Moodify	Pendigits	Pima	Shuttle	Wine
S1	DIFFI	0.83	0.91	0.92	0.92	0.91	0.89	0.84	0.91	0.88	0.85	0.9	0.88	0.92	0.92	0.91	0.93
	IF_ExIFFI	0.76	0.9	0.91	0.92	0.87	0.91	0.91	0.92	0.85	0.92	0.97	0.87	0.91	0.91	0.94	0.9
	EIF_ExIFFI	0.95	0.95	0.95	0.95	0.94	0.91	0.95	0.93	0.91	0.92	0.96	0.75	0.95	0.89	0.88	0.93
	EIF ⁺ _ExIFFI	0.93	0.93	0.94	0.93	0.92	0.92	0.92	0.94	0.95	0.94	0.96	0.93	0.92	0.94	0.79	0.96
	ECOD	0.55	0.85	0.85	0.94	0.78	0.88	0.98	0.93	0.86	0.91	0.98	0.87	0.93	0.86	0.88	0.91
S2	DIFFI	0.84	0.95	0.95	0.96	0.93	0.9	0.92	0.93	0.91	0.9	0.91	0.92	0.92	0.92	0.91	0.94
	IF_ExIFFI	0.74	0.93	0.92	0.93	0.89	0.9	0.94	0.92	0.88	0.94	0.88	0.91	0.9	0.91	0.92	0.93
	EIF_ExIFFI	0.95	0.97	0.96	0.96	0.95	0.92	0.97	0.95	0.93	0.94	0.94	0.76	0.95	0.88	0.91	0.94
	EIF ⁺ _ExIFFI	0.92	0.95	0.95	0.94	0.93	0.9	0.97	0.93	0.89	0.94	0.92	0.92	0.92	0.94	0.89	0.95
	ECOD	0.58	0.84	0.84	0.94	0.77	0.88	0.9	0.93	0.88	0.9	0.99	0.86	0.93	0.85	0.88	0.89

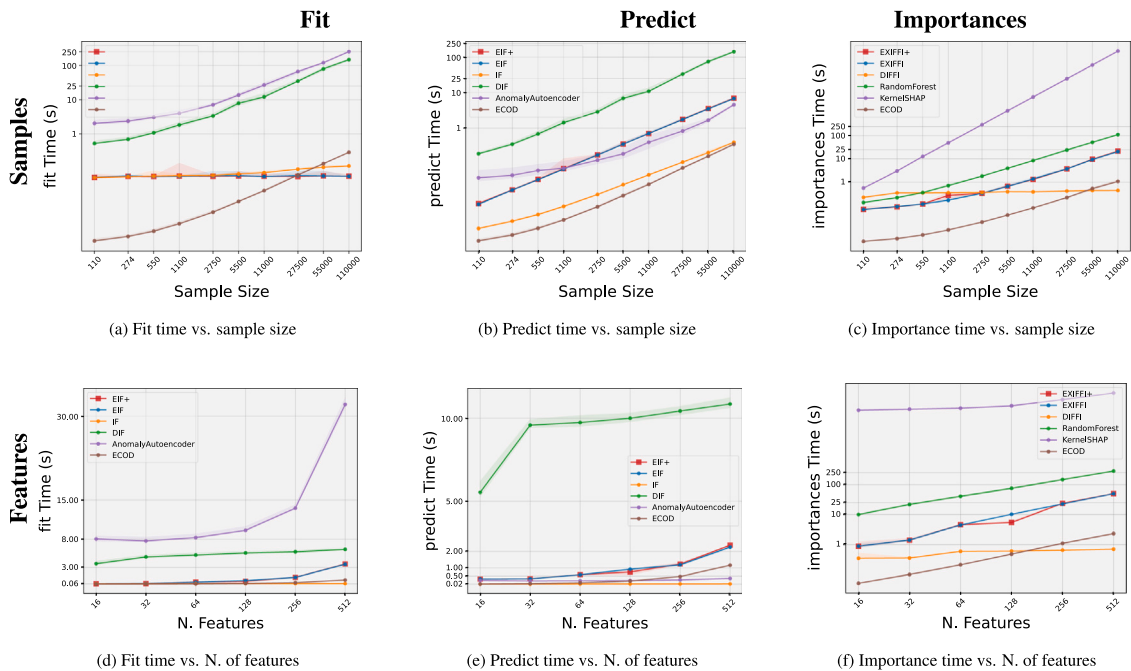


Fig. 8. Time scaling experiments display the execution times trend of various AD and XAD algorithms varying the sample size (fixing the number of features to 6) and the cardinality of the feature space (fixing the sample size to 5000).

In 8(e), instead, while DIF is still the slowest model, the Anomaly Autoencoder, ECOD and IF models prediction time can be considered independent on the number of features but only dependent on the sample size. Finally, EIF and EIF⁺ depict an increasing trend with, however, much smaller values than the ones of DIF.

In the case of interpretation, Fig. 8(c) identifies DIFFI as the fastest algorithm for calculating importance scores. EIF and EIF⁺ experience

a marked increase in computation time for this task beyond 11,000 samples. Nonetheless, their performance remains better compared to the RandomForest post-hoc interpretability method, which suffers from exponential growth in execution time with larger sample and feature sizes. KernelSHAP, on the other hand, produces out-of-scale time results, making it unsuitable as an interpretation algorithm for large datasets due to its inefficiency.

ECOD, in contrast, is in line with the time complexity of ExIFFI, with a constant factor of difference since ECOD does not need to reconstruct the importance scores because they are intrinsically embedded in its output. This analysis illustrates the efficiency and adaptability of EIF⁺ and ECOD in providing rapid interpretation across varying dataset scales.

7. Conclusions

This paper delves into the important area of unsupervised AD, a critical task for identifying unusual patterns or behaviors in data. While detecting anomalies is fundamental, the paper emphasizes that this is often not enough for practical use cases. Users not only need to know that an anomaly has occurred, but also understand why the model made a certain prediction. This understanding is key for root cause analysis and for building trust in the model's decisions.

The primary contribution of this work is the introduction of ExIFFI, an interpretability method inspired by DIFFI that aims to provide both local and global interpretability, particularly for Isolation Forest (IF)-based models. The qualitative results demonstrate that the approach produces comprehensible results, as shown through its application on three datasets: *Xaxis*, *Bisect3D_skewed*, and *glass*. Additionally, the work presents EIF⁺, a variant of the Extended Isolation Forest (EIF) model, which is designed to improve generalization performance. A thorough comparative analysis of Isolation-based and deep learning-based anomaly detection approaches is provided, which represents one of the most comprehensive studies in this domain.

The experimental results support the utility of EIF⁺ and ExIFFI, showing that EIF⁺ yields robust performance across multiple datasets, outperforming benchmark models in certain scenarios. Specifically, EIF⁺ was the top-performing model in 6 out of 11 real-world datasets with the highest average score across the datasets in Scenario II and the second-best model after EIF in Scenario I. The ExIFFI method applied to IF-based models consistently showed superior performance across all scenarios, as evaluated using the AUC_{FS} metric, which is based on the Unsupervised Feature Selection proxy task.

Furthermore, the applicability of EIF⁺ and ExIFFI to large-scale datasets, such as those encountered in industrial settings, was assessed through time-scaling experiments, with results presented in the paper's corresponding section.

Looking ahead, future research could explore innovative ways to utilize the information encoded in the splitting nodes of decision trees, particularly when multiple variables compete for relevance in the model. An additional research direction could investigate the development of a more general interpretability algorithm that could be applied to all types of Isolation Forest models, including those with different splitting functions, such as Hypersphere IF (Choudhury et al., 2021).

Finally, despite the promising results of ExIFFI, a key limitation of the current study is the lack of evaluation in a real-world industrial scenario with field-specific expertise. While ExIFFI provides interpretability and facilitates understanding of anomaly detection models, its true impact on human comprehension and decision-making processes has not yet been fully explored. The study did not include experiments involving human feedback, where experts in the field could interact with the model and assess how well the interpretability methods aid in understanding the data. Without such evaluations, the extent to which ExIFFI improves human understanding of model predictions and supports more informed decisions remains uncertain. Future work should focus on assessing the practical value of ExIFFI by conducting experiments that involve domain experts and gathering insights on how interpretability impacts human interaction with data-driven systems.

CRedit authorship contribution statement

Alessio Arcudi: Writing – original draft, Writing – review & editing, Visualization, Software, Project administration, Methodology, Formal analysis, Conceptualization. **Davide Frizzo:** Writing – original draft, Writing – review & editing, Visualization, Software, Formal analysis. **Chiara Masiero:** Writing – original draft, Writing – review & editing, Conceptualization, Methodology, Funding. **Gian Antonio Susto:** Writing – original draft, Writing – review & editing, Conceptualization, Methodology, Funding.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

This work was partially carried out within the MICS (Made in Italy - Circular and Sustainable) Extended Partnership and received funding from Next-GenerationEU (Italian PNRR - M4C2, Invest 1.3 - D.D. 1551.11-10-2022, PE00000004). Moreover this study was also partially carried out within the PNRR research activities of the consortium iNEST (Interconnected North-East Innovation Ecosystem) funded by the European Union Next-GenerationEU (Piano Nazionale di Ripresa e Resilienza (PNRR) - Missione 4 Componente 2, Investimento 1.5 - D.D. 1058 23/06/2022, ECS00000043). This work was also cofunded by the European Union in the context of the Horizon Europe project 'AIMS5.0 - Artificial Intelligence in Manufacturing leading to Sustainability and Industry5.0' Grant agreement ID: 101112089. Moreover, we acknowledge financial support under the National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.1, "Progetti di ricerca di Rilevante Interesse Nazionale - PRIN" Call for tender No. 104 published on 2.2.2022 by the Italian Ministry of University and Research (MUR), funded by the European Union - NextGenerationEU - Project Title "Smart Monitoring and Production Optimization for Zero-waste Semiconductor Manufacturing: SMART-IC" - CUP C53D23003670006.

Appendix

A.1. Datasets

A.1.1. Synthetic datasets

We generated three additional datasets, following the same procedure for the inliers and outliers definition as the one described in Section 6: *Bisect*, *Bisect3D* and *Bisect6D*, with the aim to explore how performance is affected when anomalies are defined by multiple features. These new datasets position the anomalies along bisectors in higher-dimensional subspaces (i.e. 2,3 and 6 dimensions respectively), incorporating multiple features to define the outliers. By doing so, we are able to assess the robustness of the algorithms in detecting anomalies that span multiple dimensions.

The mathematical construction of these synthetic datasets ensures a controlled and systematic variation in the number of anomalous features, providing valuable insights into the behavior of the algorithms in complex, high-dimensional settings.

- In the *Bisect* dataset, the inliers are uniformly distributed within a six-dimensional sphere. The outliers are positioned along the bisector of the subspace formed by Feature 0 and Feature 1, using the direction vector $v = [1, 1, 0, 0, 0, 0]$. Noise is added to the remaining four dimensions to introduce variability, and the outliers are scaled and perturbed accordingly. The graphical results of the ExIFFI explanations for this dataset are accessible at <https://github.com/alessioarcudi/ExIFFI> and they showcase how the most important features, with similar scores, are Feature 0 and 1 as expected from the generation process.

- In the Bisect3D dataset, the inliers follow the same distribution as in Bisect. The outliers are aligned along the bisector of the subspace spanned by Feature 0, Feature 1, and Feature 2, with the direction vector $v = [1, 1, 1, 0, 0, 0]$. Noise is added to the remaining three dimensions to simulate variability. This dataset is very similar to Bisect3D_Skewed, presented in 6, since also in this case anomalies are aligned along three features. The main difference lays in the vector v used for the generation. In Bisect3D v has the same amplitude across all the three anomalous directions, making outliers deviate from normal points by the same amount along Feature 0,1 and 2. Contrarily in Bisect3D_Skewed the outlyingness of the different variables is different. As a result in the Score Plot of Bisect3D, depicted in Fig. 2(a), the importance scores of the top 3 features are almost identical while there is a clear distinction in the Score Plot of Bisect3D_Skewed, outlined in Fig. 4(b).
- Finally in Bisect6D anomalies are distributed along all the six features composing the input space while inliers are distributed as usual with the addition of some white noise to induce diversity. Consequently the direction vector is defined as $v = [1, 1, 1, 1, 1, 1]$. This dataset was produced to observe the behavior of ExIFFI in an extreme scenario where all dimensions are affected by anomalous points. The outcome is that all six features are assigned more or less the same importance score and thus the final feature ranking is decided by the stochasticity of the algorithm. For a more detailed description of the results we refer the interested reader to <https://github.com/alessioarcudi/ExIFFI>.

For clarity the pseudocode for the generation of the synthetic datasets employed in this paper is provided in 3.

A.1.2. Real-world datasets

Most of the Real-World datasets used in our experiments come from the widely used Outlier Detection DataSets (ODDS) library introduced by Rayana (2016). Unlike synthetic datasets, these reflect the complexity of real-world scenarios. However, many of these datasets were originally designed for different purposes, such as multi-class classification, and later adapted for anomaly detection (AD) by undersampling the least-represented class to create outliers, while merging the remaining classes as normal data. This transformation can present challenges, as the undersampled minority class may not truly represent anomalies, complicating the evaluation of AD methods. Moreover, AD models designed to detect isolated anomalies may not always align with the provided labels, further complicating interpretability, as discussed in Section 6.2.

Many real-world datasets also lack detailed information on features and labels, making it difficult to evaluate model performance and interpret Feature Importance scores. An exception is the Glass dataset, detailed by Carletti et al. (2019), which allows for deeper analysis of feature influence on anomalies (see Section 6.2). The absence of such detailed information poses significant challenges to assessing models qualitatively and understanding feature impact on anomalies.

To address these limitations, two new datasets, Diabetes and Moodify, were added. These datasets provide detailed information on the semantics of their features and labels, allowing for more thorough evaluation of AD and interpretation results in alignment with domain knowledge.

Here we introduce the real-world datasets with details about the context of the features and labels

- **Anthyroid:** The Anthyroid dataset (Quinlan, 1987) is part of the UCI Machine Learning Repository. It is a three-class Classification dataset. The aim of the classification task associated with this dataset is to detect whether a patient is hypothyroid (i.e. the patient has an underactive thyroid) or not. For this scope, the three classes used refer to the functioning conditions

Algorithm 3: Synthetic Dataset Generation with Direction Vector u for Outliers

```

Input : Number of samples to generate  $n$ , Number of inliers
          $n_I$ , Number of outliers  $n_O$ , Number of features  $p$ ,
         Radius of the inliers ball  $r$ , Radius of the outlier
         distance  $d$ , Number of anomalous features  $k$ , Bounds
         for anomalies [min,max], Direction vector  $v \in \mathbb{R}^k$ 

Output: Generated synthetic dataset  $D \in \mathbb{R}^{n \times p}$ 
 $D_I \leftarrow$  empty list;
 $D_O \leftarrow$  empty list;
// Initialize empty lists for inliers and
// outliers
// --- Inliers Generation ---
while  $\text{len}(D_I) < n_I$  do
    // Sample a random  $p$ -dimensional vector from a
    // Uniform distribution in  $[-r, r]$ 
    inlier_point  $\sim U_p([-r, r])$ ;
    // Check if the L2 norm of the point is less
    // than or equal to  $r$ 
    if  $(\| \text{inlier\_point} \|^2) \leq r$  then
        |  $D_I \leftarrow D_I \cup \text{inlier\_point}$ ;
    end
end
// --- Normalize the direction vector  $v$  ---
 $u \leftarrow v / \|v\|$ ;
// Normalize  $v$  to obtain the unit direction
// vector  $u$ 
// --- Outliers Generation ---
for  $i \leftarrow 1$  to  $n_O$  do
    // Sample a scalar  $x$  from a uniform
    // distribution in [min,max]
     $x \sim U([\text{min}, \text{max}])$ ;
    // Create an empty vector for the outlier
    // point
    outlier_point  $\leftarrow \text{zeros}(p)$ ;
    // Calculate the values of the first  $k$ 
    // anomalous features based on direction  $u$  and
    // added noise
    for  $j \leftarrow 1$  to  $k$  do
        | outlier_point[ $j$ ]  $\leftarrow d * u[j] + x * v[j] + \mathcal{N}(0, 1)$ ;
    end
    // Generate remaining  $p - k$  noisy features
    // from a Gaussian distribution
    for  $j \leftarrow k + 1$  to  $p$  do
        | outlier_point[ $j$ ]  $\leftarrow \mathcal{N}(0, 1)$ ;
    end
    // Add the generated outlier point to the
    // dataset
     $D_O \leftarrow D_O \cup \text{outlier\_point}$ ;
end
// Combine inliers and outliers to form the
// final dataset
 $D \leftarrow D_I \cup D_O$ ;
return  $D$ ;

```

of the thyroid that can be *normal* (not hypothyroid), *hyperfunctioning* (overactive thyroid) and *subnormal*. The dataset originally contained 15 categorical and 6 numerical attributes. In order to adapt it to an Anomaly Detection task only the numerical attributes were considered and the *hyperfunctioning* and *subnormal* function classes were considered as part of the outliers while the *normal* functioning samples are used to build the inlier class.

The six numerical features represent, respectively, the following quantities: TSH, T3, TT4, T4U, and FTI.

- **TSH (Thyroid-Stimulating Hormone):** TSH is an hormone produced by the pituitary gland. An underactive thyroid (hypothyroidism) is associated with high levels of TSH while low levels of TSH are associated with an overactive thyroid (hyperthyroidism).
 - **T3 (Triiodothyronine):** T3 is one of the thyroid hormones. It plays a crucial role in regulating the metabolism. It is important to measure the levels of this hormone, specifically in cases of hyperthyroidism.
 - **TT4 (Total Thyroxine):** This quantity represents the total amount of the T4 thyroid hormone in the blood. T4 levels are associated to the overall thyroid hormone production.
 - **T4U (Thyroxine-Binding Globulin):** T4U measures the level of thyroxine-binding globulin, a protein binding to thyroid hormones in the blood. T4U levels are connected to the thyroid hormone availability in the body.
 - **FTI (Free Thyroxine Index):** Taking into account the T4 and T4U levels, the FTI provides an estimate of the amount of free thyroxine (T4) in the blood. It is used to assess the free, active thyroid hormone levels in the body.
- **Breast:** The Breast dataset (Wolberg, 1992) is a Binary Classification dataset where the target is the presence of breast cancer or not in a patient. The peculiarity of the samples contained in this dataset is that they are formed by categorical features. Normally, the Anomaly Detection models described in this paper are not built to deal with categorical features but in this particular case the high number of levels (e.g. The age variable has nine levels: 10–19, 20–29, ...; the tumor-size feature has levels equal to 0–4,5–9,10–14, ...) characterizing the Breast dataset's attributes makes it possible to consider them as numerical features. The dataset is composed of samples coming from the clinical cases of Dr. Wolberg collected in a time span going from January 1989 to November 1991. The dataset is composed by 9 features representing the following quantities: Clump Thickness, Uniformity of Cell Size, Uniformity of Cell Shape, Marginal Adhesion, Single Epithelial Cell Size, Bare Nuclei, Bland Chromatin, Normal Nucleoli and Mitoses.
 - **Cardio:** The Cardio (Campos and Bernardes, 2010) dataset is part of the UCI Machine Learning Repository and its complete name is Cardiotocography since it contains measurements of fetal heart rate (FHR) and uterine contraction (UC) on cardiotocograms.⁴ This dataset contains 3 classes that were assigned by expert obstetricians: *normal*, *supsect*, and *pathologic*. In order to use this dataset for Anomaly Detection the *pathologic* class was discarded and the *pathologic* class was downsampled to 176 points to maintain unbalance with respect to the *pathologic* class. There are 21 features.⁵
 - **Glass:** The Glass dataset (German, 1987) is originally used for multi-class classification to distinguish 7 different types of glasses. There are 9 features: the first one is the Refractive Index while the others measure the concentration of Magnesium (Mg), Silicon (Si), Calcium (Ca), Iron (Fe), Sodium (Na), Aluminum (Al), Potassium (K) and Barium (Ba) in the glass samples. The seven original glass groups were divided into two: Groups 1,2,3 and 4 represent Window Glasses while the remaining ones are non-window glasses: containers glass, tableware glass and headlamp glass. Among the non-window glasses, headlamp glasses were considered as outliers while Window Glasses are labeled as inliers in order to convert this dataset to be used for Anomaly Detection. According to some prior knowledge on the subject the Barium (Ba) and Potassium (K) concentration should be decisive in distinguishing between headlamp and window glasses. In fact Barium is considered the crucial element to perform the distinction between headlamp and window glasses since it is usually added to the category of headlamp glasses in order to improve their optical properties. Another potential key attribute is Potassium (K) that is frequently exploited to enhance strength and thermal resistance of window glasses.
 - **Ionosphere:** The Ionosphere dataset (Sigillito et al., 1989) contains measurements collected by a radar in Goose Bay, Labrador. The measurement system consisted of 16 high-frequency antennas with a transmitted power in the order of 6.4 kW. The targets to measure were free electrons in the Ionosphere. Originally this is built as a Binary Classification dataset where the two classes are *good* if the captured electrons show some kind of structure in the ionosphere and *bad* otherwise. The signal is processed with an autocorrelation function depending on the time of a pulse and the pulse number. There are 17 pulse numbers and each of them is identified by a pair of features for a total of 34 features.
 - **Pima:** The Pima dataset⁶ comes from the National Institute of Diabetes, Digestive and Kidney Diseases. It is a Binary Classification dataset whose aim is to predict whether a patient has diabetes or not. This dataset is the result of the application of some constraints to a larger dataset. In Pima, in fact, the data considered are obtained from female patients of at least 21 years of age coming from the Pima Indian heritage. The dataset contains 8 features indicating some typical diagnostic measurements: Number of pregnancies, Plasma Glucose Concentration, Diastolic Blood Pressure, Triceps skin thickness, 2-h serum if insulin, BMI (Body Mass Index), Diabetes Pedigree function, and Age.
 - **Pendigits:** The Pendigits dataset collects data regarding handwritten digits produced by Forty-four human writers. The task associated with this dataset is the one of recognizing the correct written digit. The dataset is composed of 16 features and the number of objects per class (i.e. per digit) was reduced by a factor of 10% to increase the unbalance between classes in order to adapt the dataset as a benchmark for the evaluation of Anomaly Detection models.
 - **Shuttle:** The Shuttle dataset (Newman et al., 1993) describes radiator positions in a NASA space shuttle. The samples are characterized by 9 attributes. Originally, the dataset is used for Multi-Class Classification with the target variable containing seven possible classes, which are: Radiator Flow, Fpv Close, Fpv Open, High, Bypass, Bpv Close, and Bpv Open. Besides the normal Radiator Flow class about 20% of the data points describe anomalous situations. To reduce the amount of anomalies the Radiator Flow class is used to form the inlier class while a stratified sampling procedure was applied to classes 2,3,5,6 and 7. Finally, data coming from class 4 were discarded.
 - **Wine:** The Wine dataset (Aeberhard and Forina, 1991) is part of the UCI Machine Learning repository and it was originally created as a 3-class classification dataset. In fact, it contains data resulting from a chemical analysis of wines grown in the same region of Italy but obtained from 3 different cultivars. So the aim of a Classification model applied to this dataset would be to correctly predict the original culture of wine given its chemical properties.

⁴ A cardiotocogram is a medical test monitoring fetal hearth rate and uterine contractions during pregnancy. It is used to assess the health status of the fetus and the progress of labor during pregnancy and childbirth.

⁵ The complete list of features names can be found at: <https://archive.ics.uci.edu/dataset/193/cardiotocography>.

⁶ <https://www.kaggle.com/datasets/uciml/pima-indians-diabetes-database>.

As it is usually done in these cases the dataset was adapted to test Anomaly Detection models considering the data from two cultures as inliers and the ones beholding to the last culture as the outliers. The dataset is composed by 13 features representing the following quantities: Alcohol, Malic Acid, Ash, Alkalinity of ash, Magnesium, Total phenols, Flavanoids, Nonflavanoid phenols, Proanthocyanins, Color Intensity, Hue, OD280/OD315⁷ of diluted wines, and Proline.

- **Diabetes:** The Diabetes dataset⁸ is a Binary Classification dataset with medical data about patients used to predict whether they have diabetes or not. There are 4 categorical and 4 numerical variables. In order to use the ExIFFI model only the numerical variables were considered:
 - Age: The age ranges between 0 and 80 and it can be an important factor since Diabetes is usually diagnosed in older adults.
 - BMI (Body Mass Index): Measure of body fat based on weight and height. High BMI values are linked to a higher risk of diabetes.
 - HbA1c-Level: The Hemoglobin A1c Level it is a measure of a person's average blood sugar level over the past 2–3 months. A high level of HbA1c is usually associated with high diabetes risk.
 - blood-glucose-level: The amount of glucose in the blood-stream at a given time. High glucose levels are a key factor to detect Diabetes.

Finally, the target variable is a binary variable indicating the presence or absence of diabetes in the patient. Following the usual protocol for Real World Datasets, the inlier group will be represented by healthy patients while the ones affected by diabetes will be placed in the outlier group.

- **Moodify:** Moodify is a recommendation app⁹ that classifies Spotify songs according to the emotions transmitted to the users. The system is trained on the Moodify dataset¹⁰ which contains features regarding the main characteristics of about 278.000 songs and the target is a 4 levels categorical variable with the following coding:

1. *Sad*
2. *Happy*
3. *Energetic*
4. *Calm*

The label with the lowest frequency (15%) in the dataset is *Calm* so that was used to create the outliers group while the inlier group is formed by the songs that belong to the other three classes.

The dataset is composed of 11 numerical features describing different musical attributes of a song. Except for the Loudness variable, which expresses the overall loudness of the track in decibels (dB), all the features have values contained in the [0, 1] interval. The variable names are the following: Acousticness, Danceability, Energy, Instrumentalness, Liveness, Loudness, Speechiness, Valence, and Tempo. Other variables are Duration (ms) and spec-rate.

⁷ These codes refers to two measurements of the optical density of a wine sample performed at two different wavelengths: 280 nm and 315 nm. These measurements are typically exploited to assess the protein content and the stability of wine.

⁸ <https://www.kaggle.com/datasets/iammustafatz/diabetes-prediction-dataset>.

⁹ More details on the Moodify project can be found at: <https://github.com/orzanai/Moodify>.

¹⁰ <https://www.kaggle.com/datasets/abdullahorzan/moodify-dataset>.

A.2. Draw-backs of isolation forest

Next, we delve deeper into the artifacts of the IF and how the EIF is able to avoid them. The primary distinction among the mentioned models lies in their approach to determining splitting hyperplanes. In the case of the IF, it operates under the constraint of splitting the feature space along a single dimension. Consequently, the constructed splitting hyperplane is invariably orthogonal to one dimension while remaining parallel to the others.

The EIF model instead relaxes this constraint, it allows multiple splits in the feature space along different dimensions, thus the splitting hyperplane can be oriented differently in each isolation tree. This relaxation helps to provide a more expressive way to capture complex data distributions while maintaining performances. Moreover, as Hariri et al. pointed out in Hariri et al. (2021), the EIF avoids the creation of regions where the anomaly score is lower only due to the imposed constraints. In Fig. A.9 we can observe that the effect of the constraint is twofold: (i) it generates bundles in the hyperplanes orthogonal to the main directions, and (ii) it creates “artifacts”, i.e. low anomaly score zones in their intersections, as we can observe in the Scoremap A.9(a).

As we progressively relax the IF constraint on splitting directions, these artifacts tend to disappear. We simulated the evolution of the anomaly score surface by allowing the hyperplane to split along a new feature in each step. Figure vividly demonstrates the gradual elimination of artifact-prone regions.

A.3. Non depth-based importance

In general, it is accurate to assert that the amount of feature importance attributed by a particular node k to a given sample x , provided that the node effectively separates the samples, should be greater if the node is closer to the root of the tree, and the importance score should decrease as the node becomes deeper within the tree. This principle is derived from the same concept used to determine anomaly scores, which are higher for samples that on average are in leafs closer to the root of the tree.

Two distinct approaches exist for assigning importance scores to nodes: the one used in the ExIFFI algorithm and the one proposed by the DIFFI algorithm. The DIFFI algorithm importance scores are assigned to nodes based on a weighting scheme that inversely relates the score to the depth of the node. This approach is parameterized to ensure that scores decrease as nodes become deeper within the tree.

The ExIFFI algorithm is an alternative approach for assigning importance scores to nodes in a decision tree, which differs from the DIFFI algorithm by not using the inverse of the depth to decrease the score as the node becomes deeper in the tree. This stems from the fact that the maximum acceptable score for a node is directly linked to the number of elements it must partition. As the depth of the node increases, this number inevitably decreases. Therefore, we assert that it is redundant and potentially misleading to incorporate an additional parameter to weigh the importance score.

The underlying idea is that as the node becomes deeper in the tree, it is responsible for splitting fewer and fewer elements, and therefore, it should not be penalized by reducing its importance score because the structure of the algorithm itself decreases the importance on average after every step deeper in the tree. In contrast, a node that splits many elements closer to the root of the tree may be more important because it has a greater impact on the separation of the data, but its importance should not be guaranteed by an external factor.

The ExIFFI algorithm adjusts the maximum score acceptable for a node based on the number of elements it needs to split, which allows for a fair comparison of the importance of nodes across different depths in the tree. This approach can help to avoid overestimating the importance of nodes that are close to the root and underestimating the importance of nodes that are deeper in the tree, a risk that the DIFFI algorithm does not take into account.




Fig. A.9. Scoremap showing the differences between IF and EIF in the two dimensional space dataset Bimodal described in Appendix A.1.1. Source: Figure inspired from Hariri et al. (2021).

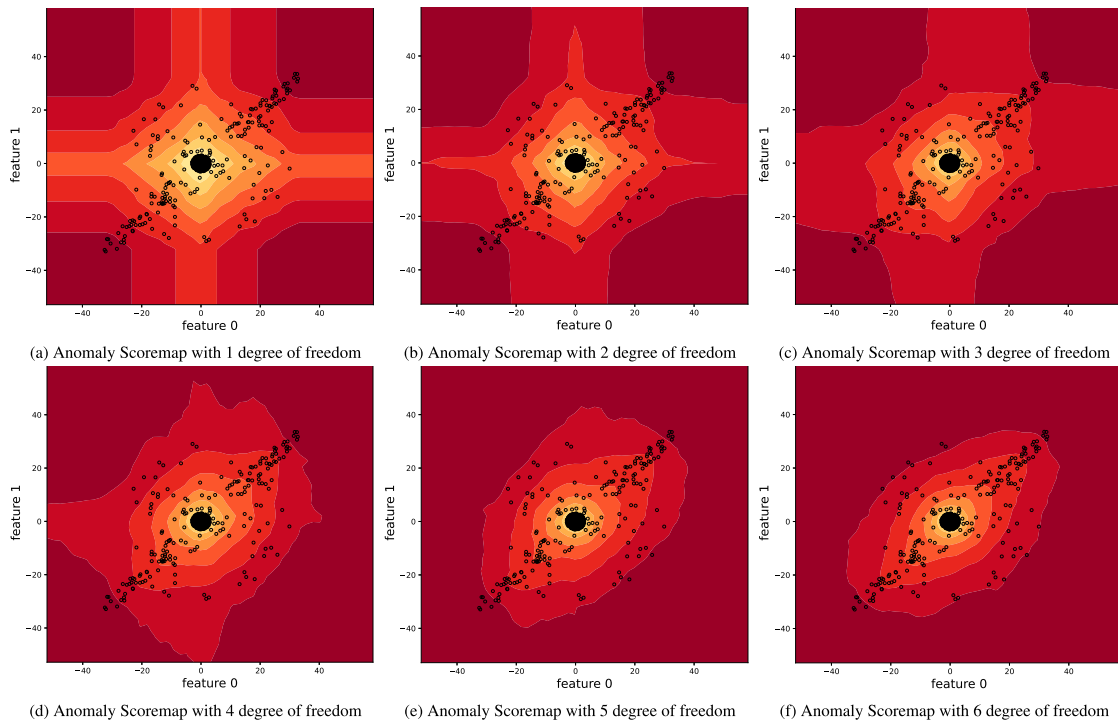


Fig. A.10. Scoremap resulting from adding one degree of freedom per time while searching for splitting hyperplanes. In this way we move from the classic IF A.10(a) to the EIF A.10(f). The experiment is performed on the six dimensional space dataset Bisect described in Appendix A.1.1.

We conducted an analysis of the ExIFFI algorithm with and without the depth parameter, in order to identify differences in the evolution of the score at different depths of the trees. To achieve this, we plotted the average score of the nodes at various levels of depth in the forests, as shown in Fig. A.11.

The results of this analysis indicate that the score generally decreases without the use of a depth parameter. Although the ExIFFI algorithm without the depth parameter produced higher scores, the overall shape of the graph is still descending as nodes are evaluated deeper in the tree. This suggests that the ExIFFI algorithm is effective in assigning importance scores to nodes, even when the depth parameter is not utilized.

In summary, our analysis suggests that the ExIFFI algorithm is a robust method for assigning importance scores to nodes in decision trees, with or without the use of a depth parameter.

A.4. EIF⁺: the effects of η

In this section the results of the ablation study on the η hyperparameter of the EIF⁺ model for the datasets covered in this Appendix are reported.

The ablation study consists in evaluating the performance of the EIF⁺ model, through the Average Precision metric, for 25 linearly spaced values between 0.5 and 5 of the parameter η , whose interpretation is given in 3.3.

After a careful observation of the plots collected in Fig. A.12 it is possible to conclude that for the Bisect, Breastw, Ionosphere and Shuttle the Average Precision of EIF⁺ is not affected by the value of the parameter η . In fact, as it can be noticed comparing to the blue line inserted in the plot to represent the Average Precision value achieved by the EIF model, in these sets of data the improvement achieved by EIF⁺ on its counterpart EIF is minimal since both models

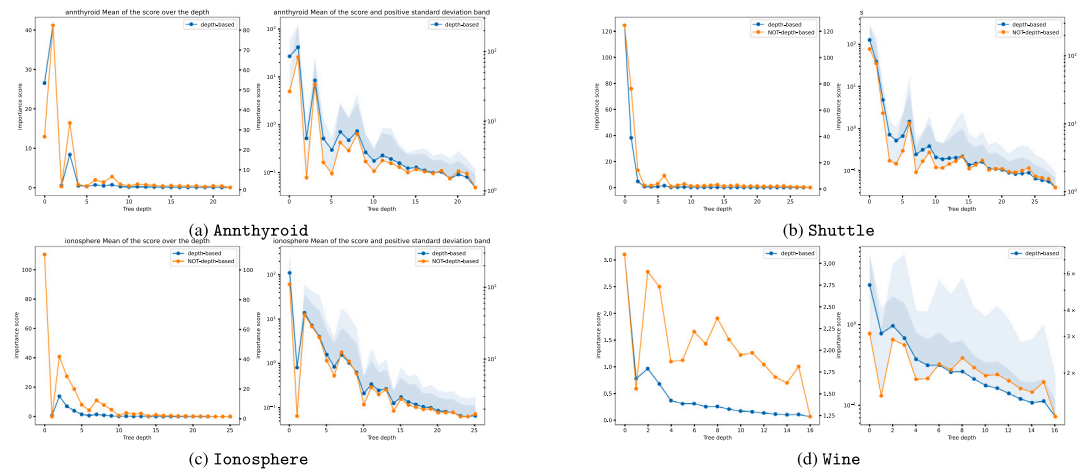


Fig. A.11. Analysis of the average scores of the nodes related to their depth.

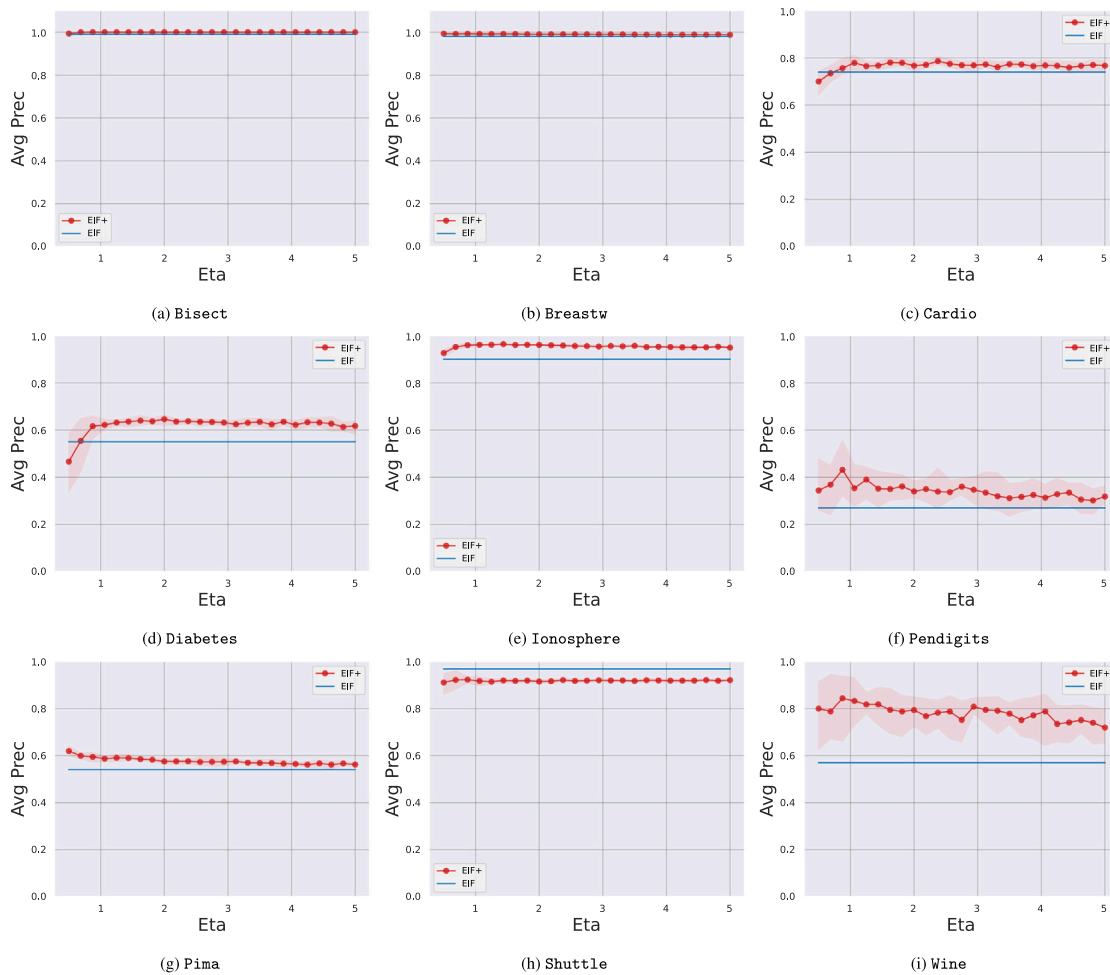


Fig. A.12. Exploration of how average precision changes with variations in the η parameter.

already achieve almost perfect precision scores. As a consequence the model remain unchanged independently on the η value since anomalies are detectable utilizing the traditional oblique cuts employed by EIF.

For other datasets, where the data distribution is intricated and EIF-like partitions are not as effective as the novel partition approach introduced with EIF+, the effect of η is more visible in the sense that there are some contained oscillations on the Average Precision values as the hyperparameter is varied.

Concluding, comparing the diverse plots produced it is possible to infer that the optimal value for the parameter analyzed in this section is highly dependent on the specific dataset structure and thus the suggested approach is to perform an hyperparameter optimization through a validation set in order to find the optimal value of η .

A.5. Explanations comparison through normalized discounted cumulative gain

In Section 6.4 different interpretation algorithm for Anomaly Detection were compared through the newly introduced AUC_{FS} metric. In this section we consider the employment of another quantitative metric for the evaluation of the effectiveness of different XAD methods: the Normalized Discounted Cumulative Gain (NDCG). The NDCG measure was originally introduced for the evaluation of the rank quality in the field of Information Retrieval (Järvelin and Kekäläinen, 2002). In this experiment NDCG is re adapted to perform the comparison between the feature ranking produced by the computation of the Global Feature Importance scores with a dataset's ground truth feature ranking. For this reason this evaluation was conducted solely on the synthetic datasets, the only ones equipped with ground truth anomalies and feature rankings as explained in 6.

In order to obtain the NDCG score to evaluate how close a feature ranking produced by an interpretation algorithm is to the ideal one, the Discounted Cumulative Gain (DCG) of the raw GFI scores will be compared with a relevance vector.

The DCG is computed considering the GFI values of the different features together with a discount factor that takes into account the position of a feature in the list. Considering a dataset \mathcal{X} composed by p features and a GFI feature vector $GFI = [I_1, \dots, I_p]$, the DCG score for dataset \mathcal{X} can be defined as follows:

$$DCG = \sum_{i=0}^p \frac{I_i}{\log_2(i+1)}$$

Finally to obtain the NDCG score, DCG is normalized by the Ideal DCG (IDCG) (i.e. the DCG score of the relevance vector representing the ideal feature ranking):

$$NDCG = \frac{DCG}{IDCG}$$

The relevance vector has to be crafted in order to represent the ground truth importance score ranking, which may not be unique. In this specific study, in fact, there are multiple correct rankings for each different synthetic datasets. Indeed, since normal features are sampled as noise their order in the feature ranking is not relevant. The aim of the evaluation through the NDCG score is to assess the ability of a XAD model to correctly rank the crafted anomalous features in the top positions.

For example, in the *Xaxis* dataset an explanation can be considered effective if it is able to place Feature 0 in the first position of the ranking, while in *Bisect3D_Skewed* a good interpretation is the one that places Feature 0,1 and 2 in the top 3 spots, in this precise order. For this reason the relevance vectors for these two datasets can be constructed, respectively, as follows: $v = [1, 0, 0, 0, 0, 0]$, $v = [4, 3, 2, 0, 0, 0]$.

As it can be noticed from Table A.5 most of the interpretation models considered are able to correctly rank the abnormal features, leading to a perfect NDCG score of 1. There are a few exceptions represented by the DIFFI and ECOD interpretability algorithms on the *Xaxis* dataset which obtain lower NDCG scores. This results comes without a surprise, in fact the struggle of the IF and ECOD model on the *Xaxis* dataset are already deeply documented in Sections Section 6 and in Tables 2 3 4

Data availability

I have shared the link to my data/code in the paper.

Table A.5

Normalized discounted cumulative gain score for different interpretation algorithms in Scenario I (S1) and Scenario II (S2) across the synthetic datasets introduced in this study. Scores for the proposed algorithm (i.e. ExIFFI) are enclosed in gray cells.

Scenario	Model	Synthetic				
		Xaxis	Bisec	Bisec3D	Bisec6D	Bisec3D_Skewed
S1	DIFFI	1.0	1.0	1.0	0.9	1.0
	IF_ExIFFI	0.35	1.0	1.0	0.99	1.0
	EIF_ExIFFI	1.0	1.0	1.0	0.99	1.0
	EIF ⁺ _ExIFFI	1.0	1.0	1.0	0.99	1.0
	IF_RF	1.0	1.0	1.0	0.99	1.0
	EIF_RF	1.0	1.0	1.0	0.99	1.0
	EIF ⁺ _RF	1.0	1.0	1.0	0.99	1.0
	ECOD	0.35	0.69	1.0	0.99	1.0
	S2	DIFFI	0.63	0.91	0.96	0.67
IF_ExIFFI		0.35	1.0	1.0	0.99	1.0
EIF_ExIFFI		1.0	1.0	1.0	0.99	1.0
EIF ⁺ _ExIFFI		1.0	1.0	1.0	0.99	1.0
IF_RF		1.0	1.0	1.0	0.99	1.0
EIF_RF		1.0	1.0	1.0	0.99	1.0
EIF ⁺ _RF		1.0	1.0	1.0	0.99	1.0
ECOD		0.35	0.69	1.0	0.99	1.0

References

- Aeberhard, S., Forina, M., 1991. Wine. <http://dx.doi.org/10.24432/C5PC7J>, UCI Machine Learning Repository.
- Aggarwal, C.C., 2015. Outlier analysis. pp. 98–101.
- Altmann, A., Tološi, L., Sander, O., Lengauer, T., 2010. Permutation importance: a corrected feature importance measure. *Bioinformatics* 26 (10), 1340–1347.
- Amin, M.T., Khan, F., Ahmed, S., Imtiaz, S., 2021a. A data-driven Bayesian network learning method for process fault diagnosis. *Process Saf. Environ. Prot.* 150, 110–122.
- Amin, M.T., Khan, F., Ahmed, S., Imtiaz, S., 2021b. A data-driven Bayesian network learning method for process fault diagnosis. *Process Saf. Environ. Prot.* 150, 110–122. <http://dx.doi.org/10.1016/j.psep.2021.04.004>, URL <https://www.sciencedirect.com/science/article/pii/S0957582021001804>.
- Amin, M.T., Khan, F., Imtiaz, S., 2019. Fault detection and pathway analysis using a dynamic Bayesian network. *Chem. Eng. Sci.* 195, 777–790. <http://dx.doi.org/10.1016/j.ces.2018.10.024>, URL <https://www.sciencedirect.com/science/article/pii/S0009250918307371>.
- Arunthavanathan, R., Khan, F., Ahmed, S., Imtiaz, S., 2021. An analysis of process fault diagnosis methods from safety perspectives. *Comput. Chem. Eng.* 145, 107197. <http://dx.doi.org/10.1016/j.compchemeng.2020.107197>, URL <https://www.sciencedirect.com/science/article/pii/S0098135420312400>.
- Bouman, R., Bukhsh, Z., Heskes, T., 2023. Unsupervised anomaly detection algorithms on real-world data: how many do we need?. arXiv preprint [arXiv:2305.00735](https://arxiv.org/abs/2305.00735).
- Brito, L.C., Susto, G.A., Brito, J.N., Duarte, M.A., 2022. An explainable artificial intelligence approach for unsupervised fault detection and diagnosis in rotating machinery. *Mech. Syst. Signal Process.* 163, 108105. <http://dx.doi.org/10.1016/j.ymssp.2021.108105>, URL <https://www.sciencedirect.com/science/article/pii/S0888327021004891>.
- Campos, D., Bernardes, J., 2010. Cardiotocography <http://dx.doi.org/10.24432/C51S4N>, UCI Machine Learning Repository.
- Carletti, M., Maggipinto, M., Beghi, A., Antonio Susto, G., Gentner, N., Yang, Y., Kyek, A., 2020. Interpretable anomaly detection for knowledge discovery in semiconductor manufacturing In: 2020 Winter Simulation Conference. WSC, pp. 1875–1885. <http://dx.doi.org/10.1109/WSC48552.2020.9384026>.
- Carletti, M., Masiero, C., Beghi, A., Susto, G.A., 2019. Explainable machine learning in industry 4.0: Evaluating feature importance in anomaly detection to enable root cause analysis In: 2019 IEEE International Conference on Systems, Man and Cybernetics. SMC, IEEE, pp. 21–26.
- Carletti, M., Terzi, M., Susto, G.A., 2023. Interpretable anomaly detection with DIFFI: Depth-based isolation forest feature importance
- Choudhury, J., Ky, P., Ren, Y., Shi, C., 2021. Hypersphere for branching node for the family of isolation forest algorithms In: 2021 IEEE International Conference on Smart Computing. SMARTCOMP, pp. 418–423. <http://dx.doi.org/10.1109/SMARTCOMP52413.2021.00090>.

- Commission, E., 2020. On artificial intelligence—a European approach to excellence and trust URL https://ec.europa.eu/info/publications/white-paper-artificial-intelligence-european-approach-excellence-and-trust_en.
- Confalonieri, R., Coba, L., Wagner, B., Besold, T.R., 2021. A historical perspective of explainable artificial intelligence Wiley Interdiscip. Rev.: Data Min. Knowl. Discov. 11 (1), e1391.
- Doshi-Velez, F., Kim, B., 2017. Towards a rigorous science of interpretable machine learning arXiv preprint arXiv:1702.08608.
- German, B., 1987. Glass identification <http://dx.doi.org/10.24432/C5WW2P>, UCI Machine Learning Repository.
- Gilpin, L.H., Bau, D., Yuan, B.Z., Bajwa, A., Specter, M., Kagal, L., 2018. Explaining explanations: An overview of interpretability of machine learning In: 2018 IEEE 5th International Conference on Data Science and Advanced Analytics. DSAA, IEEE, pp. 80–89.
- Hariri, S., Kind, M.C., Brunner, R.J., 2021. Extended isolation forest IEEE Trans. Knowl. Data Eng. 33 (4), 1479–1489. <http://dx.doi.org/10.1109/TKDE.2019.2947676>.
- Hawkins, D.M., 1980. Identification of Outliers. Vol. 11, Springer.
- Houssein, E.H., Emam, M.M., Ali, A.A., Suganthan, P.N., 2021. Deep and machine learning techniques for medical imaging-based breast cancer: A comprehensive review Expert Syst. Appl. 167, 114161.
- Järvelin, K., Kekäläinen, J., 2002. Cumulated gain-based evaluation of IR techniques ACM Trans. Inf. Syst. 20, 422–446. <http://dx.doi.org/10.1145/582415.582418>.
- Jin, B., Xu, X., 2024. Forecasting wholesale prices of yellow corn through the Gaussian process regression Neural Comput. Appl. 36, 8693–8710. <http://dx.doi.org/10.1007/s00521-024-09531-2>.
- Kaur, G., Bhattacharya, M., 2023. Intelligent fault diagnosis for AIT-based smart farming applications IEEE Sens. J. 23 (22), 28261–28269. <http://dx.doi.org/10.1109/JSEN.2023.3321417>.
- Kaur, G., Bhattacharya, M., 2024. Green fault tolerant aIoT-enabled mobile sink data collection scheme in sensor networks IEEE Trans. Veh. Technol. 1–10. <http://dx.doi.org/10.1109/TVT.2024.3400880>.
- Khakzad, N., Khan, F., Amyotte, P., 2011. Safety analysis in process facilities: Comparison of fault tree and Bayesian network approaches Reliab. Eng. Syst. Saf. 96 (8), 925–932. <http://dx.doi.org/10.1016/j.res.2011.03.012>, URL <https://www.sciencedirect.com/science/article/pii/S0951832011000408>.
- Kim, M., Yu, J., Kim, J., Oh, T.-H., Choi, J.K., 2023. An iterative method for unsupervised robust anomaly detection under data contamination IEEE Trans. Neural Netw. Learn. Syst.
- Kitson, N.K., Constantinou, A.C., Guo, Z., Liu, Y., Chobtham, K., 2023. A survey of Bayesian network structure learning Artif. Intell. Rev. 56 (8), 8721–8814.
- Klaib, A.F., Alsrhein, N.O., Melhem, W.Y., Bashtawi, H.O., Magableh, A.A., 2021. Eye tracking algorithms, techniques, tools, and applications with an emphasis on machine learning and internet of things technologies Expert Syst. Appl. 166, 114037.
- Krafft, M., Sajtos, L., Haenlein, M., 2020. Challenges and opportunities for marketing scholars in times of the fourth industrial revolution J. Interact. Mark. 51 (1), 1–8.
- Kursa, M.B., Rudnicki, W.R., 2010. Feature selection with the boruta package J. Stat. Softw. 36 (11), 1–13. <http://dx.doi.org/10.18637/jss.v036.i11>, URL <https://www.jstatsoft.org/index.php/jss/article/view/v036i11>.
- Leng, J., Zhu, X., Huang, Z., Li, X., Zheng, P., Zhou, X., Mourtzis, D., Wang, B., Qi, Q., Shao, H., Wan, J., Chen, X., Wang, L., Liu, Q., 2024. Unlocking the power of industrial artificial intelligence towards industry 5.0: Insights, pathways, and challenges J. Manuf. Syst. 73, 349–363. <http://dx.doi.org/10.1016/j.jmsy.2024.02.010>.
- Lesouple, J., Baudoin, C., Spigai, M., Tourneret, J.-Y., 2021. Generalized isolation forest for anomaly detection Pattern Recognit. Lett. 149, 109–119. <http://dx.doi.org/10.1016/j.patrec.2021.05.022>.
- Li, Z., Zhao, Y., Hu, X., Botta, N., Ionescu, C., Chen, G.H., 2023a. ECOD: Unsupervised outlier detection using empirical cumulative distribution functions IEEE Trans. Knowl. Data Eng. 35 (12), 12181–12193. <http://dx.doi.org/10.1109/tkde.2022.3159580>.
- Li, Z., Zhu, Y., Van Leeuwen, M., 2023b. A survey on explainable anomaly detection ACM Trans. Knowl. Discov. Data 18 (1), <http://dx.doi.org/10.1145/3609333>.
- Linardatos, P., Papastefanopoulos, V., Kotsiantis, S., 2020. Explainable ai: A review of machine learning interpretability methods Entropy 23 (1), 18.
- Liu, F.T., Ting, K.M., Zhou, Z.-H., 2008. Isolation forest In: 2008 Eighth IEEE International Conference on Data Mining. pp. 413–422. <http://dx.doi.org/10.1109/ICDM.2008.17>.
- Lundberg, S.M., Erion, G.G., Lee, S.-I., 2018. Consistent individualized feature attribution for tree ensembles arXiv preprint arXiv:1802.03888.
- Lundberg, S.M., Lee, S.-I., 2017a. A unified approach to interpreting model predictions In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (Eds.), Advances in Neural Information Processing Systems 30. Curran Associates, Inc., pp. 4765–4774.
- Lundberg, S.M., Lee, S.-I., 2017b. A unified approach to interpreting model predictions In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (Eds.), Advances in Neural Information Processing Systems. Vol. 30, Curran Associates, Inc., URL https://proceedings.neurips.cc/paper_files/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf.
- Miller, T., 2019. Explanation in artificial intelligence: Insights from the social sciences Artificial Intelligence 267, 1–38. <http://dx.doi.org/10.1016/j.artint.2018.07.007>, URL <https://www.sciencedirect.com/science/article/pii/S0004370218305988>.
- Molnar, C., 2020. Interpretable Machine Learning. Lulu. com.
- Mrabet, H., Khattech, I., Bouzidi, S., Kechiche, L., Jbeli, A., Al Harbi, N., Bouzidi, C., Muñoz, F., Balda, R., 2024. Influence of barium substitution on the physical, thermal, optical and luminescence properties of Sm³⁺-doped metaphosphate glasses for reddish orange light applications RSC Adv. 14, 2070–2079. <http://dx.doi.org/10.1039/D3RA08015C>.
- Newman, C., D.J., B., Merz, C., 1993. Statlog (Shuttle) UCI Machine Learning Repository. <https://doi.org/10.24432/C5WS31>.
- Oliveira, D.F.N., Vismari, L.F., Nascimento, A.M., de Almeida, J.R., Cugnasca, P.S., Camargo, J.B., Almeida, L., Gripp, R., Neves, M., 2022. A new interpretable unsupervised anomaly detection method based on residual explanation IEEE Access 10, 1401–1409. <http://dx.doi.org/10.1109/ACCESS.2021.3137633>.
- Papageorgiou, K., Theodosiou, T., Rapti, A., Papageorgiou, E.I., Dimitriou, N., Tzavaras, D., Margetis, G., 2022. A systematic review on machine learning methods for root cause analysis towards zero-defect manufacturing Front. Manuf. Technol. 2, 972712.
- Preiss, B., 2000. Data Structures and Algorithms with Object-Oriented Design Patterns in Java. Worldwide series in computer science, Wiley, URL <https://books.google.it/books?id=ywpRAAAAMAAJ>.
- Quinlan, R., 1987. Thyroid disease <http://dx.doi.org/10.24432/C5D010>, UCI Machine Learning Repository.
- Rayana, S., 2016. ODDS library URL <https://odds.cs.stonybrook.edu>.
- Ruff, L., Kauffmann, J.R., Vandermeulen, R.A., Montavon, G., Samek, W., Kloft, M., Dietterich, T.G., Müller, K.-R., 2021. A unifying review of deep and shallow anomaly detection Proc. IEEE 109 (5), 756–795. <http://dx.doi.org/10.1109/JPROC.2021.3052449>.
- Sigillito, V., Wing, S., Hutton, L., Baker, K., 1989. Ionosphere <http://dx.doi.org/10.24432/C5W01B>, UCI Machine Learning Repository.
- Turbé, H., Bjelogrić, M., Lovis, C., Mengaldo, G., 2023. Evaluation of post-hoc interpretability methods in time-series classification Nat. Mach. Intell. 5 (3), 250–260.
- Wojtas, M., Chen, K., 2020. Feature importance ranking for deep learning Adv. Neural Inf. Process. Syst. 33, 5105–5114.
- Wolberg, W., 1992. Breast cancer wisconsin (original) <http://dx.doi.org/10.24432/C5HP4Z>, UCI Machine Learning Repository.
- Xu, M., David, J.M., Kim, S.H., et al., 2018. The fourth industrial revolution: Opportunities and challenges Int. J. Financ. Res. 9 (2), 90–95.
- Xu, H., Pang, G., Wang, Y., Wang, Y., 2023. Deep isolation forest for anomaly detection IEEE Trans. Knowl. Data Eng..
- Yu, W., Zhao, C., 2020. Robust monitoring and fault isolation of nonlinear industrial processes using denoising autoencoder and elastic net IEEE Trans. Control Syst. Technol. 28 (3), 1083–1091. <http://dx.doi.org/10.1109/TCST.2019.2897946>, URL <https://www.scopus.com/inward/record.uri?eid=s2-2.0-85084107882&doi=10.1109%2FTCST.2019.2897946&partnerID=40&md5=b1de4b374ad9eddefe8a05c524e4cbea>. Cited by: 165.