Spot The Odd One Out: Regularized Complete Cycle Consistent Anomaly Detector GAN

Zahra Dehghanian*, Saeed Saravani*, Maryam Amirmazlaghani*, Mohammad Rahmati* *Department of Computer Engineering, Amirkabir University of Technology, Tehran, Iran Emails: {z.dehghanian, s.saravani, mazlaghani, rahmati}@aut.ac.ir

Abstract—This study presents an adversarial method for anomaly detection in real-world applications, leveraging the power of generative adversarial neural networks (GANs) through cycle consistency in reconstruction error. Previous methods suffer from the high variance between class-wise accuracy which leads to not being applicable for all types of anomalies. The proposed method named RCALAD tries to solve this problem by introducing a novel discriminator to the structure, which results in a more efficient training process. Additionally, RCALAD employs a supplementary distribution in the input space to steer reconstructions toward the normal data distribution, effectively separating anomalous samples from their reconstructions and facilitating more accurate anomaly detection. To further enhance the performance of the model, two novel anomaly scores are introduced. The proposed model has been thoroughly evaluated through extensive experiments on six various datasets, yielding results that demonstrate its superiority over existing state-ofthe-art models. The code is readily available to the research community at Github.

Index Terms—Anomaly detection, Generative adversarial network, Cycle consistency, Anomaly score.

I. INTRODUCTION

Discovering dissimilar instances and rare patterns is one of the most essential tasks in real-world data. Anomaly detection is the process of finding such samples, which are known as anomalies [1]. Anomalies are an important aspect of any dataset and play an important role in a wide range of applications. For example, an irregular traffic pattern on a computer network could indicate hacking and data transmission to unauthorized places. Abnormalities in credit card transactions may reveal illicit economic activity [2], or abnormalities in an MRI image may indicate the existence of a malignant tumor [3]. Despite the existence of statistical and machine learning-based methods, designing effective models for detecting anomalies in complex, high-dimensional data spaces is still a major challenge [4].

Generative adversarial networks (GANs) have shown remarkable performance in the field of anomaly detection by overcoming this challenge and modeling the distribution of highdimensional, complex real-world data. In GAN, a generator network is contrasted with a discriminator network; the discriminator attempts to differentiate between the real data and the data produced by the generator network. The generator and discriminator are trained simultaneously; the generator network G records the distribution of the data, and the discriminator D estimates the likelihood whether samples come from real data distribution or are being generated by G. The objective function of the generator G is to maximize the error probability of the network D. This structure leads to a twoplayer game like mini-max games [5].

The ability of adversarial neural networks to represent natural images has previously been demonstrated [6], [7], and their use in processing speech and text [4], as well as medical images [8] is growing. This paper proposes an efficient method for anomaly detection that is based on generative adversarial networks. Similar to many learning-based algorithms, there are two main steps: training and testing. Like other adversarial frameworks, we train the generator and discriminator networks on normal data throughout the training phase so that both are updated in succession. In this case, the joint discriminator is used to make the training of the adversarial structure more stable. The encoder E is trained along with the discriminator and generator network to apply the inverse mapping of the input samples to the latent space.

In order to minimize the distance between input space samples and their reconstruction and also minimize the distance between latent space samples (the inputs that are fed to the generator) and their reconstruction, we proposed a new joint discriminator named D_{xxzz} . In order to find an anomalous sample, the difference between input and its reconstruction is computed and samples with high differences can be determined as anomalies. Finding anomalous samples necessitates an appropriate distance between the input sample and its reconstruction. To achieve this aim, we also train our network on some samples that are produced by the proposed supplementary distribution $\sigma(x)$.In summary, this paper provides the following contributions:

- A new discriminator D_{xxzz} in order to provide complete cycle consistency is proposed.
- A supplementary distribution σ(x) is used in order to bias network towards normal data manifold.
- Two new anomaly scores are introduced.

The structure of this paper is as follows: In the next section, we have a summary of previous works. Preliminaries are provided in Section 3. The proposed model is elaborated in Section 4. The experiments will be discussed in Section 5. We have a conclusion in Section 6, and the final section is devoted to future work.

II. RELATED WORK

Anomaly detection, also known as novelty detection and outlier detection, has been widely studied, as reviewed in [9]– [11]. The previous methods used in this field are generally

A representation learning method learns a mapping for the main characteristics of normal data. One-class support vector machine finds the marginal boundary around the normal data [12]. The isolation forest method is one of the classic machine learning methods. In this method, the tree is built with randomly chosen features, and the anomaly score is the average distance to the root [13]. Deep support vector data description (DSVDD) finds a hypersphere to enclose the representation of normal samples [14]. Liu and Gryllias constructed frequency domain features using cyclic spectral analysis and applied them in the support vector data description (SVDD) framework. This method has been proven robust against outliers and can achieve a high detection rate for detecting anomalies [15]. In [16], researchers presented a new approach to identify imagery anomalies by training the model on normal images altered by geometric transformation. In this model, the classifier calculates the anomaly score using softmax statistics.

Usually, generative models attempt to learn the reconstruction of the data and use this reconstruction to identify anomalous samples [17]. For instance, auto-encoders model the normal data distribution, and the reconstruction error is used as the anomaly score [18], [19]. Deep structured energybased models (DSEBMs) learn an energy-based model and map each sample to an energy score [20]. Deep autoencoding Gaussian mixture model (DAGMM) estimates a mixed Gaussian distribution by using an encoder for normal samples [21]. A recent line of work on anomaly detection has focused on adversarial neural networks. For example, this structure has been used to identify anomalies in medical images [8]. In this work, the inverse mapping into the latent space was performed using the recursive backpropagation mechanism. In [22], a continuation of the prior work, the mapping to the latent space was performed by the encoder in order to reduce computational complexity. In [23], the proposed model was based on bidirectional GAN (BiGAN). The encoder was also responsible for mapping the input data space to the latent space. Unlike the standard GAN structure, where the discriminator takes only the real image and the generated image of the generator network as input, the representation of these images in the latent space was also considered as input to the discriminator network.

In the article [4], anomaly detection was performed by adding a new discriminator in latent space to the adversarial structure to stabilize the training process. In [24] a gan-based model is utilized to detect anomalies in time series. In the context of Mad-gan, the generator is trained on normal time series to learn the underlying patterns and dependencies. It then generates synthetic samples that resemble the normal data distribution. The discriminator is simultaneously trained to differentiate between real and synthetic samples. DCT-GAN proposes a novel approach for detecting anomalies in time series data [25]. The method combines dilated convolutions and transformer architecture within a Generative Adversarial Network (GAN) framework. The dilated convolutions capture long-range dependencies in the time series, while the transformer module learns the temporal relationships between different time steps. The generator network generates synthetic samples, and the discriminator network distinguishes between real and fake samples. By training the GAN on normal time series data, it learns to generate realistic normal samples. Anomalies can be detected by measuring the discrepancy between real and generated samples using a reconstruction loss. Experimental results demonstrate that DCT-GAN outperforms existing methods in terms of anomaly detection accuracy on various benchmark datasets.

Many existing GAN-based anomaly detection algorithms have weak robustness, in [26] the authors introduce a technique called M3gan, which combines a masking strategy with a mutable filter to diversify the data and improve the robustness of the model. The M3gan approach starts by dividing the multidimensional data into smaller subsets called cells. Each cell is then masked by randomly selecting a subset of its dimensions and replacing their values with noise. This masking strategy helps in preserving the privacy of sensitive information while still allowing for accurate anomaly detection. The deep convolutional autoencoder model is a classic autoencoder model in which the encoder and decoder have a convolutional structure. The anomaly score in this model is the 2-norm of reconstruction errors [27]. In [28] a novel deep learning framework for semi-supervised anomaly detection called deep SAD (semi-supervised anomaly detection) has been proposed. The proposed method, called Deep SAD, combines the benefits of both supervised and unsupervised approaches by leveraging a small amount of labeled normal data and a large amount of unlabeled data. Deep SAD utilizes an autoencoder-based architecture to learn a low-dimensional representation of the input data and employs a novel loss function that encourages the model to assign low reconstruction errors to normal instances while penalizing anomalies.

In this section, we reviewed and categorized the various methods used to clarify abnormal data. In section 3 we will deep into fundamental basis required to deal with proposed model.

III. PRELIMINARIES

In this section, we briefly elaborate the overall idea of a generative adversarial network and then concentrate on the evolution of GAN-based anomaly detection algorithms. The generative adversarial network was first proposed in 2014 by Goodfellow et al [5]. The generator and the discriminator are trained on an M set of $\{\chi^{(i)}\}_{i=1}^{M}$ unlabeled samples. The generator maps selected samples from the latent space z to the input data space. The discriminator attempts to distinguish between the real data $x^{(i)}$ and the data produced by the generator (G). The generator G imitates the input data distribution, whereas the discriminator differentiates between real samples and generator data. In the training phase, the generator G and the discriminator D are alternatively optimized using a stochastic gradient descent approach.

q(x) is the distribution of the input data, and p(z) is the distribution of the generator in the latent space. GAN network training is done by finding a discriminator and a generator that can solve the saddle point problem as $\min_G \max_D V_{GAN}(D, G)$ and the definition of the $V_{GAN}(D, G)$ function is defined as:

$$V_{GAN} = E_{x \sim q(x)}[\log(D(x))] + E_{z \sim p(z)}[\log(1 - D(G(z)))]$$
(1)

Solving this problem concludes that the generator distribution is equal to the true data distribution. It has been proved in [5] that the global optimal discriminator will be obtained if and only if $p_G(x) = q(x)$. By p_G , we mean the distribution learned by the generator. Adversarially learned inference (ALI) [23] attempts to obtain the inverse mapping between input data space and latent space by modeling the joint distribution of encoder as q(x, z) = q(x)e(z|x) and the distribution of generator as p(x, z) = p(z)p(x|z) using the encoder E. Here, e(z|x) is learned by the encoder. The objective function of the ALI model is as follows:

$$\min_{G,E} \max_{D} V_{ALI} = E_{q(x,z)}[\log D(x, E(x))] + E_{p(x,z)}[\log(1 - D(G(z), z))]$$
(2)

where D represents the discriminator, taking x and z as input, and its output value specifies probability of origination of the current inputs from the q(x, z) distribution. Encoder, generator, and discriminator are in their optimal state only if q(x, z) = p(x, z). This has been proved in [23].

Although p and q distributions are apparent, in practice and during model training, they are not necessarily converging to the optimal point. This issue was attributed to the problem of cycle consistency, which was defined as $G(E(x)) \approx \hat{x}$ in [27]. A new framework called ALICE was proposed to solve the above problem by adding the discriminator D_{xx} to the ALI network structure [27]. The objective function of this model is as follows:

$$\min_{E,G} \max_{D_{xz},D_{xx}} V_{ALICE} = V_{ALI} + E_{x \sim q(x)} \left[\log D_{xx}(x,x) + \log(1 - D_{xx}(x,G(E(x)))) \right]$$

This work demonstrates that using a discriminator D_{xx} can achieve the best reconstruction for the input data [29]. In Adversarially Learned Anomaly Detection (ALAD), a conditional distribution was applied to the baseline ALICE model with the inclusion of additional discriminator to stabilize the training process [4]. To deep into the detail, a discriminator D_{zz} is added to the model to ensure the cycle consistency in the latent space, which tries to make the latent space variable and its reconstruction as analogous as possible. By assembling the block proposed in [4] in ALICE framework, the cost function of the ALAD model will finally be as follows:

$$\min_{G,E} \max_{D_{xz}, D_{xx}, D_{zz}} V_{ALAD} = V_{ALICE} + \mathbb{E}_{z \sim p(z)} \left[\log \left(D_{zz}(z, z) \right) + \mathbb{E}_{z \sim p(z)} \left[\log \left(1 - D_{zz}(z, G(E(z))) \right) \right]$$
(4)

In [4], it is claimed that the training model will be stabilized by adding Lipschitz constraints to the discriminators of the GAN model. Moreover, it is shown in practice that with spectral normalization of the weight parameters, the network's performance will be improved [4].

Although the idea of ALAD helps stabilize the cycle, the latent and input space variables are scrutinized in two independent spaces, and the inherent dependence between the variables is ignored. More precisely, the x and its reconstruction are investigated in a separate process from the z and its corresponding reconstruction, while the reconstruction processes of these two pairs of data is along with each other and affects each other directly. To model this reliance, we define a complete cycle and a new discriminator that employs the information contained within the complete cycle.

Another problem with ALAD is that it doesn't take into account the need for weak reconstruction for anomalous samples. In fact, in all of the previous models, it was assumed that if the model were trained with normal data, it would always have a poor reconstruction for anomalous data, even though there is no way to force the model to make weak reconstructions of anomalous samples. The proposed RCALAD model attempts to address this issue by employing the supplementary distribution $\sigma(x)$, which biases all reconstructions toward the normal data distribution, and the model attempts to lower the anomaly score for normal input while increasing it for anomalous input.

IV. PROPOSED MODEL

In this part, we describe our novel adversarial framework for anomaly detection that overcomes the aforementioned drawbacks. The problem of a lack of complete cycle consistency and our corresponding solution will be described first. Then, the necessity of weak reconstruction will be addressed. At the end of this section, two new anomaly scores based on the proposed model are introduced.

A. Complete Cycle Consistency

As mentioned before, in the previous models, the cycle consistency of the input data and latent space variables is examined in two independent procedures. This implies that the reconstruction of the latent space variable \hat{z} and input data \hat{x} are processed separately. In other words, the previous works tried to resemble x and its reconstruction and z and its reconstruction in two separate cycles. The inverse mapping of x was neglected in the z-cycle, and the direct mapping of z was also ignored in the x-cycle. To tackle this problem, we introduce a new complete cycle that simultaneously tries to learn joint distribution, which utilizes neglected parameters in a complete cycle that contain rich information. It should be noted that here the variable z is a sample of the Gaussian distribution given as an input to the generator, and it is not related to the input mapping in the latent space.

The Complete Cycle Consistency issue (CCC) declares that for each variable x in the input space if the encoder first estimates the inverse mapping to the latent space, which equals $E(x) = z_x$ and the obtained representation is entered into the generator to generate the network reconstruction from the input variable $G(z_x) = G(E(x)) = \hat{x}$ and the this reconstruction is given to the encoder network in order to calculate the reconstruction in the latent space, that is,



Fig. 1: The information of complete cycle consistency in the proposed model (left side) and using the variables of input data space and latent space in the cycle consistency of the ALAD network (right side).

 $E(\hat{x}) = E(G(z_x)) = \hat{z}_{\hat{x}}$, it is logically expected from any reconstruction-based network that the two variables x and \hat{x} as well as the two variables $\hat{z}_{\hat{x}}$ and z_x have the least possible difference. That is, the CCC issue is defined in such a way that, in any reconstruction-based model, for each input sample and its mapping in the latent space, the network reconstruction for both variables should have a minimum error and maximum similarity.

Without using the CCC, the similarity between the input data and its reconstruction, as well as the similarity between z and its reconstruction, were examined independently and in two separate cycles. It was assumed that they are independent, but we know that these two cycles are entirely dependent on each other, and the assumption of independence is not valid in these two issues. To solve this problem, we proposed to model the dependency by examining the CCC variables in the new discriminator D_{xxzz} and using the information flow in this chain to improve network training for anomaly detection in the best possible way. The difference between the input of D_{xxzz} and the input of D_{zz} used in the ALAD model is represented in Figure 1.

As can be seen in Figure 1, the ALAD model does not use the information of a complete cycle. In order to use the available joint information in a complete cycle, a new variable called $\hat{z}_{\hat{x}}$ is introduced. To calculate this variable, the inverse mapping of the input data x is applied to the generator, and the resulting inverse mapping is calculated again using the decoder. Hence, complete cycle consistency will be provided in this model.

In order to ensure the condition of complete cycle consistency, the new D_{xxzz} discriminator is used with the joint input. It is noteworthy that the effectiveness of the joint discriminators has already been proven once in ALIGAN [23]. Actually, when adding the encoder to the GAN framework, two procedures can be scrutinized. The first one is adding an independent discriminator to train the encoder, and the second one is changing the discriminator input from a single input mode to a joint input mode. It is proven that using a joint discriminator obtains better results. According to the same idea, the input of the joint discriminator D_{xxzz} extracts the most information for model training.

This discriminator uses the quadruple (x, x, z_x, z_x) as the real data and the quadruple of $(x, G(E(x)), z_x, E(G(z_x)))$ as

the fake data. This discriminator attempts to make the input x and network reconstruction, as well as the inverse mapping of the input image in the latent space and its reconstruction by the encoder, as close as possible to each other so that a complete stable loop is provided and the model is trained and stabilized better.

B. Constraint of Weak Reconstruction

In reconstruction-based models, it has always been assumed that if the training and reconstruction of the normal data are properly done, the reconstruction of abnormal data will necessarily be weak and different from the input data. However, experiments indicate that it is not always the case, and sometimes the reconstructed anomalous sample is slightly similar to the input sample. Hence, it won't be easy to recognize it as an abnormal sample. In fact, in none of the previous models, there was not any obligation or control condition to bias the model toward producing poor reconstructions for anomalous samples.

During the training phase, the encoder and generator are only trained on normal samples. As a result, the appropriate z space for normal samples and the reconstruction of them is well modeled. However, because the model has not yet seen the remainder of the space, including abnormal examples, it may map it onto an unknown point in the latent space during the testing phase. In this case, there is no information about mapping anomalous data to its reconstruction. To solve this problem, the supplementary distribution called $\sigma(x)$ is used. Samples from this distribution will cover the input data space, which means that we will produce some extra noisy samples and force the model to generate reconstruction in the normal data manifold. Here we teach the model to produce output in a normal data manifold by feeding it with a wider range of inputs. Through using this distribution, we have further training with unseen data, and as a result, we could have a generalization that will help model produce reconstructions in normal data manifold for a wider range of inputs compare to the absence of the supplementary distribution. Through this extra training, the network learns to reconstruct the normal data class for a relatively more expansive range of inputs. So an appropriate distance is created between the anomalous sample and its reconstruction. This distance is regarded as an appropriate criterion for detecting abnormal samples. Figure 2 illustrates the effect of using $\sigma(x)$ in training routine.

C. RCALAD Model

In this section, we will introduce our proposed model, RCALAD, by integrating both ideas, which means employing the new variable of $\hat{z}_{\hat{x}}$ in the D_{xxzz} discriminator and the $\sigma(x)$ distribution and adding them to the basic model [4]. In this network, we attempted to provide a comprehensive, practical, and compatible framework for all anomaly detection problems by resolving the aforementioned issues of the complete consistency cycle and the necessity of weak reconstruction. Figure 3 depicts a schematic representation of the proposed model. As we can see in Figure 3, an encoder and the generator are trained in the standard structure of the adversarial neural



Fig. 2: Effect of the presence of supplementary distribution in the model training process. These figures indicate the trained models. On the left side, there is no supplementary distribution in the training phase, and on the right side, we had samples from supplementary distribution in training. In this figure, X represents the input data space and Z represents the latent space. Samples are mapped from the latent space to the input data space by the generator G, and the encoder E reverses the mapping. Green circles show normal samples, red crosses represent abnormal samples, and blue stars represent samples generated by the supplementary distribution $\sigma(x)$. The turquoisecolored line shows the value of the abnormality score. As can be seen in Figure 2, if samples from the supplementary distribution $\sigma(x)$ are not present (on the left side of the figure), the abnormality score for the abnormal sample is lower than when these samples(blue stars) are used used (right side). In other words, when there is no supplementary distribution in the training process, after training, the model may reconstruct anomalous samples in a good way, which leads to a small difference between anomalous sample and its reconstruction, and as a result, we will have a low anomaly score, but on the right, the samples from supplementary distribution have biased the model towards the normal manifold. In this case, the model tries to reconstruct an anomalous sample near the normal data manifold, which is far from the anomalous data manifold that leads to a high anomaly score for abnormal data.

network. The inverse mapping from the input data space to the latent space is obtained simply by using the encoder Ein the proposed structure and the direct mapping from the latent space to the input data space is done by the generator G. Here D_{xz} is a new name for a standard discriminator needed for adversarial training of G and E. It is responsible for discriminating between two pairs: (z, G(z) and (x, E(x))). Through this discrimination, the encoder E learns the inverse mapping of input x, and also the generator G learns the direct mapping from the latent space to the input data space. This discriminator determines whether an input variable pair is derived from the input data x distribution and its corresponding point in the latent space or if it is generated by the generator G and sampled from the latent space of z. In order to satisfy the cycle consistency condition in the input and latent space, D_{xx} and D_{zz} discriminators are used, so cycle consistency in both spaces will be modeled independently. Here, a joint discriminator called D_{xxzz} is used to train both generator and encoder networks simultaneously. The D_{xxzz} is introduced to use all the information in a complete cycle. That is, in addition to examining both variables x and z and their reconstruction



Fig. 3: Overall structure of the RCALAD model.

 \hat{x} and \hat{z} in the corresponding space, their joint distribution is used in D_{xxzz} during the process of detecting anomalous samples. By using D_{xxzz} more information is available to determine whether the input data is anomalous or not. This network is responsible for determining between quadruple samples of (x, x, z_x, z_x) and $(x, G(E(x)), z_x, E(G(z_x)))$. In fact, the discriminator D_{xxzz} tries to maximize the similarity between x and its reconstruction G(E(x)), and it also attempts to make the mapping of the input image in the latent space z_x and its reconstruction $E(G(z_x))$ as similar as feasible. To cover much more of the latent space, the $\sigma(x)$ block is added to this model. Using this distribution, new samples are made from the input data space and then mapped into the latent space of the normal data. Finally, the objective function of the proposed model is as follows:

$$\min_{G,E} \max_{D_{xxzz}, D_{xz}, D_{xx}, D_{zz}} V_{RCALAD} (D_{xxzz}, D_{xz}, D_{xx}, D_{zz}, E, G) = V_{ALAD} + E_{x \sim \sigma(x)} [\log (1 - D_{xz}(x, E(x)))] + E_{x \sim q(x)} [\log D_{xxzz}(x, x, E(x), E(x))] + E_{\boldsymbol{x} \sim q(x)} [1 - \log D_{xxzz}(x, G(E(x)), E(x), E(G(E(x))))]$$
(5)

D. Anomaly Scores

The main goal of this proposed model is to detect anomalies based on the accurate reconstruction of normal input data, whereas abnormal samples are reconstructed in a weak manner. One of the key elements in anomaly detection is the definition of the anomaly score for calculating the distance between the input sample and the reconstruction provided by the network [4]. Some of the anomaly scores that were used in previous models are as follows:

$$A_{L_1}(x) = \|x - \hat{x}\|_1,$$

$$A_{L_2}(x) = \|x - \hat{x}\|_2,$$

$$A_{\text{Logits}}(x) = \log \left(D_{xx}(x, \hat{x}) \right),$$

$$A_{\text{Features}}(x) = \|f_{xx}(x, x) - f_{xx}(x, \hat{x})\|_1$$
(6)

Here, logit means the raw output of the discriminators, while feature means the output of the layer preceding logit. In our proposed model, since the new discriminator D_{xxzz} offers the ability of extracting additional information, it is crucial to specify the new anomaly scores in order to make use of this information. Therefore, two new anomaly scores are introduced.

The first anomaly score presented in this paper is called $A_{fm}(x)$. This score uses the D_{xxzz} discriminator feature space to calculate distance between samples and their reconstruction. For this purpose, the output of the second-to-last layer is used as features. Our anomaly score is defined as follows:

$$A_{fm}(x) = \|f_{xxzz}(x, x, z_x, z_x) - f_{xxzz}(x, \hat{x}, z_x, \hat{z}_{\hat{x}})\|_1$$
(7)

In this equation, f(.) represents the activation function of the next to last layer in the D_{xxzz} discriminator structure. The concept used in the definition of this score is using the confidence level of the discriminator on the quality of the reconstructions provided by the network. In other words, if reconstruction is performed well, the sample belongs to the trained normal data of the network. Thus, the higher value of this criterion means the greater difference in reconstructions and so higher possibility of input data's abnormality. The second point in this article is presented with the aim of maximizing the use of information in the model for anomaly detection. In this section, the A_{all} criterion is defined. The score is the sum of the outputs of all discriminators, including D_{xx} , D_{zz} and D_{xxzz} .

In fact, since all the discriminators in the proposed model are trained only based on the normal samples and the reconstruction for all the input data space is biased towards the normal data space, it is expected that the input image data and its reconstruction will look different, and the discriminators can easily identify these anomalous inputs. The mathematical expression of this criterion is given in the following equation:

$$A_{all}(x) = \frac{1}{3} (D_{xxzz} (x, \hat{x}, z_x, \hat{z}_{\hat{x}}) + D_{xx} (x, \hat{x}) + D_{zz} (z_x, \hat{z}_{\hat{x}}))$$
(8)

The criterion A_{all} tries to utilize all discriminators' information. During the training phase, the discriminators learn to pay attention to the difference between the pairs of (x, x) and (x, \hat{x}) as well as the pairs of (z_x, z_x) and $(z_x, \hat{z}_{\hat{x}})$. It means, the farther \hat{x} from x or $\hat{z}_{\hat{x}}$ from z_x , it will be easier for the discriminators to recognize the data origin. In the proposed model, by adding the distribution of $\sigma(x)$ and biasing all the reconstruction towards the normal data distribution, the reconstruction error for abnormal data is increased and the discriminators' output can be considered as a reliable criterion for abnormality detection. Finally, the recommended anomaly scores can be viewed according to the algorithm 1.

V. EXPERIMENTS

This section compares the proposed RCALAD model with prominent anomaly detection models. To test the models on a fair basis, the reported outcomes for all the implemented models are based on tabular data obtained from the average of ten runs, and for each class of image data, they are based on the average of three runs. The anomaly score used in tabular data is A_{all} score and, for image data, it is A_{fm} score. The reason Algorithm 1 Process of calculating anomaly scores in Regularized Complete Adverarially Learned Anomaly Detection

Input: $x \sim p_{x_{\text{Test}}}(x)$, E, G, D_{xx} , D_{zz} , D_{xxzz} , f_{xxzz} % f_{xxzz} is the feature layer of D_{xxzz} **Output:** $A_{all}(x)$, $A_{fm}(x)$ % A is the anomaly score procedure INFERENCE $z_x \leftarrow E(x)$ %Encode samples, Construct latent Embedding $\hat{x} \leftarrow G(z_x)$ %Reconstruct samples $\hat{z}_{\hat{x}} \leftarrow E(\hat{x})$ %Reconstruct latent Embedding $A_{fm}(x) \leftarrow ||f_{xxzz}(x, x, z_x, z_x) - f_{xx}(x, \hat{x}, z_x, \hat{z}_{\hat{x}})||_1$ $A_{all}(x) \leftarrow \frac{1}{3}(D_{xxzz}(x, \hat{x}, z_x, \hat{z}_{\hat{x}}) + D_{xx}(x, \hat{x}) + D_{zz}(z_x, \hat{z}_{\hat{x}}))$ **return** $A_{all}(x)$, $A_{fm}(x)$

end procedure

for choosing these scores based on data type will be discussed later. Moreover, the ALAD model is implemented and the results of the best anomaly score $A_{Features}$ are reported. For other models, the available results are adopted from [4], [17]. In testing phase, anomalous samples and normal samples are fed to the trained model. We expect that the model produce a high anomaly scores for anomalous samples and low anomaly scores for normal samples. We know that for KDDCup99, arrhythmia, thyroid, and musk datasets 20, 15, 2.5 and 3.2 percent of the data are anomalous samples, respectively, so we consider the samples with high anomaly scores as anomalous samples. For the CIFAR-10 and SVHN 90 percent of the data are anomalous.

A. Datasets

In order to evaluate the performance of the proposed model and scrutinize its efficiency from different viewpoints, various datasets with diverse characteristics are used. The proposed method is tested on the available image and tabular datasets. For tabular datasets, four datasets, including KDDCup99 [30], arrhythmia [31], thyroid [32], and musk [33], are used. The KDDCup99 dataset contains nearly 5000000 samples from 41 dimensions, 34 of which are continuous and 7 of which are categorical. We then utilized one-hot representation to encode categorical characteristics, yielding a total of 121 features following this encoding. Arrhythmia is a medical collection related to cardiac arrhythmia with 16 classes. This database contains 274 attributes and it has 490 samples. We applied our methods on raw data. Also, thyroid is a three-class dataset related to thyroid disease. The "hyperfunction" class, consisting of 2.5% of the dataset set, is treated as anomaly. It contains 3772 samples and each sample has 6 continuous features. The Musk Anomaly Detection dataset was created to classify six classes of molecular musk. This dataset has 3062 samples with 166 features. These four datasets, 20, 15, 2.5 and 3.2 percent of the data are anomalous samples, respectively. Hence, in the test phase, after calculating the anomaly score, the aforementioned proportion of the data that has the highest anomaly score is classified as an anomaly. In order to assess the proposed model on these datasets, F1, recall, and precision criteria are used. Two datasets, CIFAR-10 [34] and SVHN [35], are considered for the image datasets. CIFAR-10 consists of 60,000 32x32 color images containing one of 10 object classes, with 6000 images per class. SVHN dataset includes nearly 100000 32x32 labeled real world images of house numbers taken from Google Street View. Pixels of these two datasets were scaled to be in range [-1,1]. Both of these datasets have ten classes, and, like the previous works, one class is considered as the normal class and the other nine classes as the abnormal class. The model is trained only on normal data manifold. As a result, we should train the model specifically for each class. The criterion used to evaluate the model on the image dataset is the area under the receiver operating curve (AUROC). For all the datasets that are used, 80% of the data are used for training, and 20% are used for testing. Validation data is chosen from 25% of the training data.

B. Experiments on the Tabular Datasets

Evaluation results of the proposed RCALAD model and other state-of-the-art models on tabular datasets (KDDCup99, arrhythmia, thyroid, and musk) are summarized in Table I. The structures used in the generator, discriminator, and encoder networks are all fully connected layers with nonlinear activation functions. It should be noted that, in this step, N(0, I) distribution is used as $\sigma(x)$. In comparison to existing models, the proposed model has a successful performance on the arrhythmia and musk datasets, as shown in Table I. Our model is also the best according to F1 criteria on the KDD dataset, but it takes second place on the thyroid dataset due to the exceptional performance of the IF model. The reason for this phenomenon can be attributed to the nature of the data in this dataset. Since there are various features in this dataset, only a few of them are informative; therefore, the results of classic models such as IF, which are based on feature selection, are better. An idea to improve the proposed model results on the thyroid dataset is to use models such as IF in the preprocessing step to select more informative features for training the model.

C. Experiments on the image Datasets

In this section, the performance of the proposed model on CIFAR-10 and SVHN image data is scrutinized in two separate tables. All the experiment are done in one vs. all setting. It means that we consider one class as normal and the nine others as anomalies, we train our model only on the normal class. For each dataset, we train a specific model for each normal class. We report the results for each class separately, as well as the average performance for whole classes.

As in Tables II and III, the proposed model has significantly improved the results on the CIFAR-10 dataset. The results show that the model performs best in the half of classes and gets competitive results on others. In addition to being superior in seven classes on SVHN dataset, the proposed model also performs the best in the average of all classes. Furthermore our model decreases class-wise variance on image datasets that leads to more reliable results. We provide some example of input images with their reconstruction in 4 and 5. As you can see our model produces reconstructions in normal data manifold even for anomalous samples. For example the proposed model can reconstruct a dog image as a car. In order to validate the superior performance of the RCALAD model, we also conducted the Wilcoxon signed-rank test. In Appendix I, the P-values and statistics are reported.



Fig. 4: Reconstruction of normal and abnormal inputs on CIFAR-10 dataset. The first row is normal inputs and the second row is their reconstruction. The third row are anomalies and the fourth row is corresponding reconstruction for anomalies.



Fig. 5: Reconstruction of normal and abnormal inputs on SVHN dataset. The first row is normal inputs and the second row is their reconstruction. The third row are anomalies and the fourth row is corresponding reconstruction for anomalies.

D. Ablation Studies

In this section, we examine the effectiveness of each component added to the basic model on both kinds of datasets. In these experiments, the average results of the model are repeated in the presence and absence of the discriminator D_{xxzz} and the supplementary distribution $\sigma(\mathbf{x})$. Through these experiments, we want to find out how each proposed part affects the final results.

According to Tables IV and V, the proposed RCALAD model achieves the best result in the presence of both parts. In scrutinizing the role of the D_{xxzz} discriminator, this discriminator has improved the accuracy on the CIFAR-10 dataset to an optimal level but has not made significant improvement on the SVHN dataset. In terms of the role of the $\sigma(x)$ distribution, it performed well on the CIFAR-10 dataset and enhanced the AUROC criterion. However, when applied to the SVHN dataset, it reduced the AUROC criteria by a tiny amount when compared to the base model. Still, its inclusion in the final model resulted in the extraction of new information and a

TABLE I: Output results of the proposed model in comparison with the basic models on the tabular data set.

Model		KDDCu	р		Arhythmi	a		Thyroid	l		Musk	
	Prec.	Recall	F1 score	Prec.	Recall	F1 score	Prec.	Recall	F1 score	Prec.	Recall	F1 score
IF [36]	92.1	93.7	92.9	51.4	54.6	53.0	70.1	71.4	70.2	47.9	47.7	47.5
OC-SVM [12]	74.5	85.2	79.5	53.9	40.8	45.1	36.3	42.3	38.8	_	_	_
DSEBMr [20]	85.1	64.7	73.2	15.1	15.1	15.1	4.0	4.03	4.0	_	_	_
DSEBMe [20]	86.1	64.4	73.9	46.6	45.6	46.0	13.1	13.1	13.1	_	_	_
AnoGAN [8]	87.8	82.9	88.6	41.1	43.7	42.4	44.1	46.8	45.4	3.0	3.1	3.1
DAGMM [29]	92.9	94.2	93.6	49.0	50.7	49.8	47.6	48.3	47.8	_	_	_
ALAD [4]	94.2	95.7	95.0	50.0	53.1	51.5	22.9	21.5	22.2	58.1	59.0	58.3
DSVDD [13]	89.8	94.9	92.1	35.3	34.3	34.7	22.2	23.6	23.2	_	_	_
RCALAD	95.3	95.6	95.4	58.8	62.5	60.6	53.7	51.5	52.6	62.9	63.3	63.1

TABLE II: Output results of the proposed model compared to the basic models on the CIFAR-10 dataset.

Normal	DCAE [37]	DSEBM [20]	DAGMM [29]	IF [36]	AnoGAN [8]	ALAD [4]	RCALAD
Airplane	59.1	41.4	56.0	60.1	67.1	64.7	68.4
car	57.4	57.1	56.0	50.8	54.7	45.7	57.2
Bird	48.9	61.9	53.8	49.2	52.9	67.0	69.6
Cat	58.4	50.1	51.2	55.1	54.5	59.2	67.2
Deer	54.0	73.2	52.2	49.8	65.1	72.7	71.9
Dog	62.2	60.5	49.3	58.5	60.3	52.8	65.1
Frog	51.2	68.4	64.9	42.9	58.5	69.5	70.3
Horse	58.6	53.3	55.3	55.1	62.5	44.8	59.6
Ship	76.8	73.9	51.9	74.2	75.8	73.4	70.5
Truck	67.3	63.6	54.2	58.9	66.5	43.2	57.6
Mean	59.4	60.3	54.4	55.5	61.8	59.3	65.7

TABLE III: Output results of the proposed model compared to the basic models on the SVHN dataset.

Normal	OCSVM [12]	DSEBMr [20]	DSEBMe [20]	IF [36]	AnoGAN [8]	ALAD [4]	RCALAD
0	52.0	56.1	53.4	53.0	57.3	58.7	60.4
1	48.6	52.3	52.1	51.2	57.0	62.8	59.2
2	49.7	51.9	51.8	52.3	53.1	55.2	54.9
3	50.9	51.8	51.7	52.2	52.6	53.8	55.8
4	48.4	52.5	52.4	49.1	53.9	58.0	58.5
5	51.1	52.4	52.3	52.4	52.8	56.1	56.2
6	50.1	52.1	52.2	51.8	53.2	57.4	59.4
7	49.6	53.4	55.3	52.0	55.0	58.8	58.0
8	45.0	51.9	52.5	52.3	52.2	55.2	56.1
9	52.5	55.8	52.7	53.7	53.1	57.3	58.3
Mean	50.2	52.9	52.4	51.6	54.0	57.3	57.7

more comprehensive view. To sumup, when both D_{xxzz} and $\sigma(\mathbf{x})$ are present in training phase, we can achieve the best results.

E. Evaluating the Sufficiency of the D_{xxzz}

By adding the D_{xxzz} discriminator, is there a need for D_{xx} and D_{zz} discriminators or not? To answer this question correctly, we performed some experiments, whose results are summarized in Table VI. In fact, in this section, in addition to the above-mentioned question, the result of adding D_{xxzz} discriminator in basic models such as ALI and ALICE are investigated.

In Table VI we examine the effects of all discriminators on different models. According to this table and as expected from the theoretical results, adding the D_{xxzz} discriminator to the general frameworks had the highest efficiency. As a result, eliminating D_{xx} has less of an impact on the model because some of the information it collects is covered by the D_{xxzz} discriminator. However, it is apparent that removing D_{zz} , which assesses the similarity of z and its reconstruction in an independent cycle, decreases the accuracy. As can be seen, employing these three discriminators is seen to be the most effective, as the D_{xxzz} discriminator alone does not cover all aspects. It is worth noting that this remark applies to both image and tabular datasets.

F. Scores Evaluation

In this section, the proposed anomaly scores are evaluated and compared with the anomaly scores presented in previous work [4]. As shown in Table VII, on tabular data, the raw output of the D_{xxzz} discriminator A_{all} outperforms other anomaly scores. The performance of the feature-based score A_{fm} on image data is remarkable, as can be seen in Table VIII. This difference in the performance of the introduced scores can be attributed to the difference in the number of features in these two types of datasets. Given the fact that the number of features on tabular data is less than those on image data, the outputs of all discriminators are enough to detect anomalous data. However, in the image datasets, the output of the next-to-last layer contains more information to distinguish

TABLE IV: Effects of the various proposed sections in improving the results of the tabular datasets.

Model	Precision	Recall	F1 score		
KDDCup	99				
Baseline(ALAD)	94.4	95.7	95.0		
Baseline + D_{xxzz} (CALAD)	95.9	95.7	95.8		
Baseline + $\sigma(x)$ (RALAD)	94.3	95.5	94.9		
Baseline + D_{xxzz} + $\sigma(\mathbf{x})(\mathbf{RCALAD})$	95.3	95.6	95.4		
Arrhy	thmia				
Baseline(ALAD)	50.0	53.1	51.5		
Baseline + D_{xxzz} (CALAD)	57.4	60.5	57.5		
Baseline + $\sigma(x)$ (RALAD)	54.6	56.5	55.5		
Baseline + D_{xxzz} + $\sigma(\mathbf{x})$ (RCALAD)	58.8	62.5	60.6		
Thyroid	l				
Baseline(ALAD)	22.9	21.5	22.2		
Baseline + D_{xxzz} (CALAD)	52.9	51.8	52.3		
Baseline + $\sigma(x)$ (RALAD)	43.1	45.7	44.3		
Baseline + D_{xxzz} + $\sigma(\mathbf{x})$ (RCALAD)	53.7	51.5	52.6		
Musk					
Baseline(ALAD)	50.0	53.1	51.5		
Baseline + D_{xxzz} (CALAD)	57.4	60.5	57.5		
Baseline + $\sigma(x)$ (RALAD)	54.6	56.5	55.5		
Baseline + D_{xxzz} + $\sigma(\mathbf{x})$ (RCALAD)	62.9	63.3	63.1		

TABLE V: Effects of the various proposed sections in improving the results of the image datasets.

Model	AUROC
CIFAR-10	
Baseline(ALAD)	59.3
Baseline + D_{xxzz} (CALAD)	63.4
Baseline + $\sigma(x)$ (RALAD)	64.2
Baseline + D_{xxzz} + $\sigma(\mathbf{x})$ (RCALAD)	65.7
SVHN	
Baseline(ALAD)	57.3
Baseline + D_{xxzz} (CALAD)	57.6
Baseline + $\sigma(x)$ (RALAD)	56.8
Baseline + D_{xxzz} + $\sigma(x)$ (RCALAD)	57.7

TABLE VI: Assessing the performance of the model in the presence or absence of each of the discriminators.

model	D_{zz}	D_{xx}	D_{xxzz}	Precision	Recall	F1 score		
KDDCup99								
ALAD	yes	yes	no	94.2	95.7	95.0		
$ALI + D_{xxzz}$	no	no	yes	93.8	95.1	94.4		
$ALI + D_{zz} + D_{xxzz}$	yes	no	yes	94.6	95.5	95.0		
ALICE + D_{xxzz}	no	yes	yes	94.1	94.5	94.7		
CALAD	yes	yes	yes	95.9	95.7	95.8		
RCALAD	yes	yes	yes	95.3	95.6	95.4		
		Arr	hythmia					
ALAD	yes	yes	no	50.0	53.1	51.5		
$ALI + D_{xxzz}$	no	no	yes	52.2	52.9	52.5		
ALI+ D_{zz} + D_{xxzz}	yes	no	yes	57.1	58.2	57.6		
ALICE + D_{xxzz}	no	yes	yes	54.3	56.1	55.1		
CALAD	yes	yes	yes	57.4	60.5	57.5		
RCALAD	yes	yes	yes	58.8	62.5	60.6		

between normal and abnormal data. In this way, the A_{fm} score excelled on the image datasets, while A_{all} performed well on tabular datasets.

G. Choosing appropriate supplementary distribution

In this section, we examine different probability distributions as $\sigma(\mathbf{x})$ in order to find the best option and its effect on the performance of the proposed model. We consider N(0, I), N(0, 2I), and U(-1, +1) as supplementary distributions. As we can see in table IX N(0, I) serves the best results.

H. Statistical Test

The Wilcoxon Rank Test [?], also known as the Mann Whitney U Test, is a non-parametric statistical test used to

compare two independent samples. It is often described as the non-parametric version of the two-sample t-test. The test only makes the first two assumptions of independence and equal variance and does not assume that the data have a known distribution. The test is used to determine if there is a significant difference between the two groups being compared. The research hypothesis is that the proposed method would have a higher accuracy than the previous methods. The null hypothesis is that there would be no difference in accuracy between the proposed method and the previous methods. As can be seen in Table 1-3, the highest accuracy after the proposed model belongs to the ALAD model, and for this reason, this statistical test has been performed with the results of this model. The result of this test is reported in Table X.

Score	Precision	Recall	F1 score					
	KDDCup	599						
A_{L_1}	90.8	91.0	90.9					
A_{L_2}	90.1	90.0	90.0					
$\tilde{A_{Logits}}$	91.6	91.6	91.6					
$A_{Features}$	91.2	91.7	91.1					
A_{fm}	93.2	93.7	93.0					
A_{all}	92.3	90.0	92.1					
	Arrhythr	nia						
A_{L_1}	35.2	37.5	36.4					
A_{L_2}	35.2	37.5	36.4					
A_{Logits}	55.8	59.3	57.6					
$A_{Features}$	23.2	25.0	24.2					
A_{fm}	44.1	46.8	45.4					
A_{all}	61.7	65.6	63.7					
	Thyroid							
A_{L_1}	49.8	49.0	49.9					
A_{L_2}	50.1	50.0	50.0					
$\tilde{A_{Logits}}$	49.6	49.7	49.7					
$A_{features}$	51.2	51.7	51.5					
A_{fm}	52.2	51.2	51.7					
A_{all}	53.7	51.5	52.6					
	Musk							
A_{L_1}	59.7	59.3	59.5					
A_{L_2}	60.0	60.1	60.1					
$\bar{A_{Logits}}$	58.6	58.9	58.8					
$A_{Features}$	58.2	58.8	58.8					
A_{fm}	61.1	61.8	61.4					
A_{all}	62.9	63.3	63.1					

TABLE VII: Comparing the performance of the proposed anomaly scores with other scores on tabular data

TABLE VIII: Comparing the performance of the proposed anomaly scores with other scores on image datasets.

Anomaly Score	AUROC
CIFAR-1	0
A_{L_1}	63.4
A_{L_2}	63.2
A_{Logits}	62.9
$A_{Features}$	63.1
A_{fm}	65.7
A_{all}	64.7
SVHN	
A_{L_1}	57.7
A_{L_2}	56.3
A_{Logits}	53.6
$A_{Features}$	57.6
A_{fm}	57.7
$\dot{A_{all}}$	57.6

TABLE IX: Comparing the performance of different supplementary distribution on the RCALAD model.

		KDDCUP		Arrhythmia		
$\sigma(\mathbf{x})$	Prec.	Recall	F_1	Prec.	Recall	F_1
$\mathcal{N}(0, \boldsymbol{I})$	0.629	0.633	0.631	0.588	0.625	0.606
$\mathcal{N}(0, 2\mathbf{I})$	0.626	0.633	0.629	0.580	0.629	0.603
$\mathcal{U}(-1,1)$	0.608	0.604	0.606	0.584	0.633	0.607

TABLE X: Result of Wilcoxom Rank Test.

Dataset	#run	statistics	P-value
Arrhythmia	10	0.0	0.00195
KDD99	10	7.0	0.03710
SVHN	30	157.0	0.12413
CIFAR10	30	56.0	0.00011

The results of the Wilcoxon signed-rank test indicate that the proposed method had a significantly better performance than the previous methods in three over four datasets. As mentioned in section 5.1, Dataset SVHN pertains to real images of the house number. This dataset proved to be particularly challenging due to the presence of noise. To further investigate the performance of the proposed method on the SVHN dataset, we conduct a follow-up experiment consisting of 10 tests per class, as opposed to 3 tests described in Section 5. The results of this extended experiment yield a test statistic value of 1898.0 and a corresponding p-value of 0.03109, providing evidence of the superior performance of the proposed method on the challenging SVHN dataset.

VI. CONCLUSION

This research presents a unique and novel solution for high variance rate problem in anomaly detection tasks through the use of generative adversarial neural networks (GANs). The proposed framework utilizes a generator and an encoder to learn the mapping between the input data space and latent space and incorporates two discriminators, D_{xx} and D_{zz} , to improve the stability of the training process and ensure a consistent cycle. To further enhance the information utilization of the cycle, a new discriminator, D_{xxzz} , is introduced. Experiments reveal that considering cycle consistency for input data for the latent space variable z as one complete joint cycle provides additional information for the model that improves the results. The proposed model also incorporates a supplementary distribution $\sigma(x)$ to influence the network output to align with the normal data distribution, leading to even more precise anomaly detection. The supplementary distribution helps the model map unseen data in the normal data manifold, and as a result, an appropriate distance between the anomalous sample and its reconstruction will be achieved. The proposed model outperforms existing models for both tabular and image datasets in terms of anomaly detection and reducing class-wise variance in image datasets.

VII. FUTURE WORK

The RCALAD model has yielded promising results in the field of anomaly detection. However, like other GAN-based models, it is susceptible to robustness challenges. To address these issues, existing techniques, as outlined in [38], [39] can be employed.

In the present study, the anomaly ratio parameter was an input during test time. A potential avenue for future research would be to explore methods that can effectively determine the optimal value for this parameter without requiring it to be supplied as input.

Further, another potential area for future direction could be learning importance coefficients for each discriminator based on the type of data being analyzed (e.g., tabular or imagery data) and incorporating these coefficients into the cost function for calculating anomaly scores.

REFERENCES

- D. Yao, X. Shu, L. Cheng, S. J. Stolfo, Anomaly detection as a service: challenges, advances, and opportunities, Synthesis Lectures on Information Security, Privacy, and Trust 9 (3) (2017) 1–173.
- [2] C. Jiang, J. Song, G. Liu, L. Zheng, W. Luan, Credit card fraud detection: A novel approach using aggregation strategy and feedback mechanism, IEEE Internet of Things Journal 5 (5) (2018) 3637–3647.
- [3] X. Dai, M. Bikdash, Distance-based outliers method for detecting disease outbreaks using social media, in: SoutheastCon 2016, IEEE, 2016, pp. 1–8.

- [4] H. Zenati, M. Romain, C.-S. Foo, B. Lecouat, V. Chandrasekhar, Adversarially learned anomaly detection, in: 2018 IEEE International conference on data mining (ICDM), IEEE, 2018, pp. 727–736.
- [5] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, Advances in neural information processing systems 27 (2014).
- [6] A. Radford, L. Metz, S. Chintala, Unsupervised representation learning with deep convolutional generative adversarial networks, arXiv preprint arXiv:1511.06434 (2015).
- [7] A. Creswell, T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, A. A. Bharath, Generative adversarial networks: An overview, IEEE signal processing magazine 35 (1) (2018) 53–65.
- [8] T. Schlegl, P. Seeböck, S. M. Waldstein, U. Schmidt-Erfurth, G. Langs, Unsupervised anomaly detection with generative adversarial networks to guide marker discovery, in: International conference on information processing in medical imaging, Springer, 2017, pp. 146–157.
- [9] R. Kaur, S. Singh, A survey of data mining and social network analysis based anomaly detection techniques, Egyptian informatics journal 17 (2) (2016) 199–216.
- [10] A. Zimek, E. Schubert, H.-P. Kriegel, A survey on unsupervised outlier detection in high-dimensional numerical data, Statistical Analysis and Data Mining: The ASA Data Science Journal 5 (5) (2012) 363–387.
- [11] M. A. Pimentel, D. A. Clifton, L. Clifton, L. Tarassenko, A review of novelty detection, Signal processing 99 (2014) 215–249.
- [12] B. Schölkopf, R. C. Williamson, A. Smola, J. Shawe-Taylor, J. Platt, Support vector method for novelty detection, Advances in neural information processing systems 12 (1999).
- [13] L. Ruff, R. Vandermeulen, N. Goernitz, L. Deecke, S. A. Siddiqui, A. Binder, E. Müller, M. Kloft, Deep one-class classification, in: International conference on machine learning, PMLR, 2018, pp. 4393– 4402.
- [14] S. Liang, Y. Li, R. Srikant, Enhancing the reliability of outof-distribution image detection in neural networks, arXiv preprint arXiv:1706.02690 (2017).
- [15] C. Liu, K. Gryllias, A semi-supervised support vector data descriptionbased fault detection method for rolling element bearings based on cyclic spectral analysis, Mechanical Systems and Signal Processing 140 (2020) 106682.
- [16] I. Golan, R. El-Yaniv, Deep anomaly detection using geometric transformations, Advances in neural information processing systems 31 (2018).
- [17] Z. Yang, I. S. Bozchalooi, E. Darve, Regularized cycle consistent generative adversarial network for anomaly detection, arXiv preprint arXiv:2001.06591 (2020).
- [18] D. T. Nguyen, Z. Lou, M. Klar, T. Brox, Anomaly detection with multiple-hypotheses predictions, in: International Conference on Machine Learning, PMLR, 2019, pp. 4800–4809.
- [19] S. Pidhorskyi, R. Almohsen, G. Doretto, Generative probabilistic novelty detection with adversarial autoencoders, Advances in neural information processing systems 31 (2018).
- [20] S. Zhai, Y. Cheng, W. Lu, Z. Zhang, Deep structured energy based models for anomaly detection, in: International conference on machine learning, PMLR, 2016, pp. 1100–1109.
- [21] T. Schlegl, P. Seeböck, S. M. Waldstein, G. Langs, U. Schmidt-Erfurth, fanogan: Fast unsupervised anomaly detection with generative adversarial networks, Medical image analysis 54 (2019) 30–44.
- [22] H. Zenati, C. S. Foo, B. Lecouat, G. Manek, V. R. Chandrasekhar, Efficient gan-based anomaly detection, arXiv preprint arXiv:1802.06222 (2018).
- [23] V. Dumoulin, I. Belghazi, B. Poole, O. Mastropietro, A. Lamb, M. Arjovsky, A. Courville, Adversarially learned inference, arXiv preprint arXiv:1606.00704 (2016).
- [24] D. Li, D. Chen, B. Jin, L. Shi, J. Goh, S.-K. Ng, Mad-gan: Multivariate anomaly detection for time series data with generative adversarial networks, in: Artificial Neural Networks and Machine Learning–ICANN 2019: Text and Time Series: 28th International Conference on Artificial Neural Networks, Munich, Germany, September 17–19, 2019, Proceedings, Part IV, Springer, 2019, pp. 703–716.
- [25] Y. Li, X. Peng, J. Zhang, Z. Li, M. Wen, Dct-gan: Dilated convolutional transformer-based gan for time series anomaly detection, IEEE Transactions on Knowledge and Data Engineering (2021).
- [26] Y. Li, X. Peng, Z. Wu, F. Yang, X. He, Z. Li, M3gan: A masking strategy with a mutable filter for multidimensional anomaly detection, Knowledge-Based Systems 271 (2023) 110585.
- [27] C. Li, H. Liu, C. Chen, Y. Pu, L. Chen, R. Henao, L. Carin, Alice: Towards understanding adversarial learning for joint distribution matching, Advances in neural information processing systems 30 (2017).

- [28] L. Ruff, R. A. Vandermeulen, N. Görnitz, A. Binder, E. Müller, K.-R. Müller, M. Kloft, Deep semi-supervised anomaly detection, arXiv preprint arXiv:1906.02694 (2019).
- [29] B. Zong, Q. Song, M. R. Min, W. Cheng, C. Lumezanu, D. Cho, H. Chen, Deep autoencoding gaussian mixture model for unsupervised anomaly detection, in: International conference on learning representations, 2018.
- [30] S. Stolfo, W. Fan, W. Lee, A. Prodromidis, P. Chan, KDD Cup 1999 Data, UCI Machine Learning Repository, DOI: https://doi.org/10.24432/C51C7N (1999).
- [31] H. Guvenir, B. Acar, H. Muderrisoglu, R. Quinlan, Arrhythmia, UCI Machine Learning Repository, DOI: https://doi.org/10.24432/C5BS32 (1998).
- [32] R. Quinlan, Thyroid Disease, UCI Machine Learning Repository, DOI: https://doi.org/10.24432/C5D010 (1987).
- [33] D. Chapman, A. Jain, Musk (Version 1), UCI Machine Learning Repository, DOI: https://doi.org/10.24432/C5ZK5B (1994).
- [34] A. Krizhevsky, G. Hinton, et al., Learning multiple layers of features from tiny images (2009).
- [35] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, A. Y. Ng, Reading digits in natural images with unsupervised feature learning (2011).
- [36] F. T. Liu, K. M. Ting, Z.-H. Zhou, Isolation forest, in: 2008 eighth ieee international conference on data mining, IEEE, 2008, pp. 413–422.
- [37] A. Makhzani, B. J. Frey, Winner-take-all autoencoders, Advances in neural information processing systems 28 (2015).
- [38] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, X. Chen, Improved techniques for training gans, Advances in neural information processing systems 29 (2016).
- [39] R. Chalapathy, A. K. Menon, S. Chawla, Robust, deep and inductive anomaly detection, in: Joint European Conference on Machine Learning and Knowledge Discovery in Databases, Springer, 2017, pp. 36–51.