

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

10-2023

Anomaly detection under distribution shift

Tri CAO

Jiawen ZHU

Guansong PANG

Singapore Management University, gspang@smu.edu.sg

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#)

Citation

CAO, Tri; ZHU, Jiawen; and PANG, Guansong. Anomaly detection under distribution shift. (2023). *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), Paris, France, October 2-6*. 6551-6523.

Available at: https://ink.library.smu.edu.sg/sis_research/8415

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylids@smu.edu.sg.

Anomaly Detection under Distribution Shift

Tri Cao, Jiawen Zhu, and Guansong Pang*

School of Computing and Information Systems, Singapore Management University

Abstract

Anomaly detection (AD) is a crucial machine learning task that aims to learn patterns from a set of normal training samples to identify abnormal samples in test data. Most existing AD studies assume that the training and test data are drawn from the same data distribution, but the test data can have large distribution shifts arising in many real-world applications due to different natural variations such as new lighting conditions, object poses, or background appearances, rendering existing AD methods ineffective in such cases. In this paper, we consider the problem of anomaly detection under distribution shift and establish performance benchmarks on four widely-used AD and out-of-distribution (OOD) generalization datasets. We demonstrate that simple adaptation of state-of-the-art OOD generalization methods to AD settings fails to work effectively due to the lack of labeled anomaly data. We further introduce a novel robust AD approach to diverse distribution shifts by minimizing the distribution gap between in-distribution and OOD normal samples in both the training and inference stages in an unsupervised way. Our extensive empirical results on the four datasets show that our approach substantially outperforms state-of-the-art AD methods and OOD generalization methods on data with various distribution shifts, while maintaining the detection accuracy on in-distribution data. Code and data are available at <https://github.com/mala-lab/ADShift>.

1. Introduction

Anomaly Detection (AD) is a crucial task in machine learning that aims to identify rare and unusual patterns in data. It is an important problem in various domains, such as financial domain [1, 4], cybersecurity [69, 72], industrial inspection [5], and medical diagnosis [66, 68]. Due to the difficulty and/or high cost of collecting labeled anomaly data, current AD studies are focused on unsupervised ap-

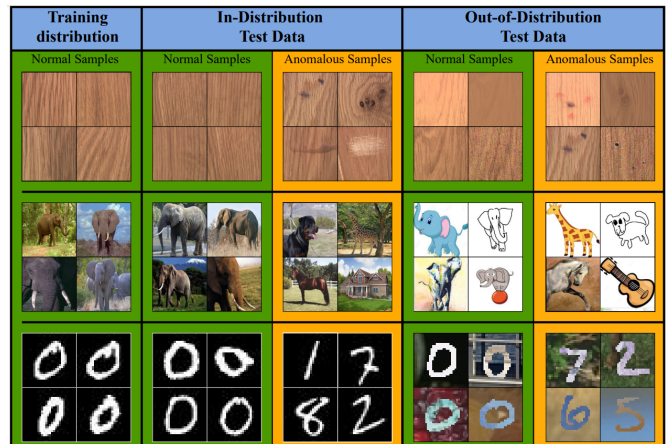


Figure 1: Illustrative samples for anomaly detection under distribution shift. First row: the ‘Wood’ dataset from MVTeC [5]. Second row: the ‘Elephant’ class as normal and the remaining classes as anomaly in PACS [40]. Third row: the ‘0’ class as normal and the remaining classes as anomaly in MNIST [36]/MNIST-M [18]. We aim at distinguishing anomalies from normal data in both in-distribution test data and out-distribution test data

proaches, which aim to learn patterns from a set of normal training samples to identify abnormal samples in test data.

Although existing AD studies have demonstrated promising performance [13, 38, 56, 65], they generally assume that the training and test data are drawn from the same data distribution. However, this assumption is often unrealistic in real-world scenarios as the test data can have large distribution shifts arising in many applications due to different natural variations such as new lighting conditions, object poses, or background appearances, rendering the AD methods ineffective in such cases.

Distribution shift is a ubiquitous problem in different real-world applications, which can significantly degrade the performance of models in various tasks such as image classification, object detection, and segmentation [18, 34, 40, 85]. Many out-of-distribution (OOD) generaliza-

*Corresponding author: G. Pang (panguansong@gmail.com).

tion methods have been introduced to address this problem [8, 19, 24, 29, 37, 43, 48, 55, 86]. These OOD generalization methods rely on large labeled training data from one or multiple relevant domains to learn domain-invariant feature representations. They often require class labels [29, 61, 80], domain labels [9, 74, 82, 88], or the existence of diverse data [24, 86, 89] in the source domain to learn such robust feature representations. However, the training data in the AD task consists of only one class, and the data is monotonous. Consequently, it is difficult to adapt existing OOD generalization techniques to address the AD under distribution shift problem. Trivial adaption of the OOD generalization can fail to learn generalized normality representations, leading to many detection errors, *e.g.*, normal samples with distribution shifts cannot be distinguished from anomalous samples and consequently they are detected as anomaly. As shown by the exemplar data in Fig. 1, normal samples in the in-distribution (ID) test data are very similar to the normal training data, and ID anomalies deviate largely from the normal data; however, due to the distribution shift, the normal samples in the OOD test data are substantially different from the ID normal data in terms of foreground and/or background features, and as a result, these normal samples can be falsely detected as anomaly.

In this paper, we tackle the problem of anomaly detection under distribution shift. It is an *OOD generalization* problem, aiming at learning generalized detection models to accurately detect normal and anomalous samples in test data with distribution shifts, while maintaining the effectiveness on in-distribution test data. This is different from the problem of *OOD detection* [23, 26, 44, 62, 76] that aims to equip supervised learning models with a capability of rejecting OOD/outlier samples as unknown samples for the sake of model deployment safety. This work makes three main contributions in addressing the OOD generalization problem in the AD task:

- We present an extensive study of the distribution shift problem in AD and establish large performance benchmarks under various distribution shifts using four widely-used datasets adapted from AD and OOD generalization tasks. Our empirical results further reveal that existing state-of-the-art (SOTA) AD and OOD generalization methods fail to work effectively in identifying anomalies under distribution shift.
- We then propose a novel robust AD approach to diverse distribution shifts, namely *generalized normality learning* (GNL). GNL minimizes the distribution gap between ID and OOD normal samples in both the training and inference stages in an unsupervised way. To this end, we introduce a normality-preserved loss function to learn distribution-invariant normality representations, which enables GNL to learn generalized

semantics of the normal training data at different feature levels. GNL also utilizes a test time augmentation method to further reduce the the distribution gap during the inference stage.

- Extensive experiments show that our approach GNL substantially outperforms state-of-the-art AD methods and OOD generalization methods by over 10% in AU-CROC on data with various distribution shifts, while maintaining the detection accuracy on the ID test data.

2. Related Work

2.1. Anomaly Detection

One-class Classification. Some early methods for anomaly detection include one-class support vector machine (OC-SVM) [67] and support vector data description (SVDD) [71]. More recently, Deep SVDD [63] uses a deep neural network to identify anomalies with a SVDD objective. A number of methods [12, 21, 64, 78, 81] is then introduced to learn more effective deep one-class description.

Reconstruction-based Methods. One popular AD approach is to use autoencoder (AE) [33]. AE-based anomaly detection learns normal patterns from a dataset to reconstruct new samples, assuming that anomalous samples have higher reconstruction errors due to distribution differences. There are many works following this direction and gaining good performance [20, 25, 57, 60, 79, 83, 84].

Self-supervised Learning Methods. The use of data augmentation techniques is becoming increasingly prevalent in AD. One such strategy involves incorporating synthetic anomalies into datasets that are otherwise free of anomalies [38, 79, 84].

Knowledge Distillation. Another popular line of research is knowledge distillation-based methods. A student-teacher framework with discriminative latent embeddings is introduced in [6]. Many improved versions for AD are then introduced [13, 65, 75]. Anomaly Detection via Reverse Distillation (RD4AD) [13] is the latest one and gains SOTA performances on many datasets.

All these methods are focused on AD with the same distribution in training and test data, which fail to work well on data with distribution shift.

2.2. OOD Generalization

Data Augmentation. One popular approach for OOD generalization is based on data augmentation. Methods in this line involve generating new data samples from existing ones to increase the size and diversity of the training data. The model can then learn more about the underlying data distribution and become more robust to changes in the test data [11, 24, 54, 70, 85, 86].

Unsupervised Learning. By solving pretext tasks, a model can develop general features that are not specific to

the target task. As a result, the model is less likely to be influenced by biases that are unique to a particular domain, which helps to avoid overfitting and increase generalization ability to different unseen data [3, 7, 8, 19, 50, 77].

Although these two types of methods are not designed for AD, they can be easily adapted for AD as they do not require class or domain labels during training. On the other hand, many existing OOD generalization methods, such as domain alignment [27, 31, 41, 42, 49, 53, 87], meta-learning [15–17, 39, 73], and disentangled representation learning [10, 30, 32, 40, 59, 74], require class/domain-related supervision, which are inapplicable for the AD task. A similar issue exists for OOD generalization methods designed for multi-class problems [16, 29]. There are some cross-domain AD methods [14, 46, 47], but they require class labels in the ID data or few training samples from the target domain. By contrast, we focus on unsupervised AD and do not require any OOD data available during training. They focus on video data, while we focus on image data. Additionally, another related research line is on AD in situations involving a ‘near distribution’ scenario [51], where anomalies are semantically similar to the normal distribution. Methods in this line can be more robust to distribution shift than general AD methods, but they do not tackle variations between the distributions of training and testing normal data.

3. Problem Formulation and Challenges

3.1. Problem Formulation

Let \mathcal{X}_s and \mathcal{X}_t denote the source (ID) and target (OOD) distributions, respectively, where \mathcal{X}_s is used for both training and testing phase, while \mathcal{X}_t is only used for inference period. We assume that during training, only normal data from \mathcal{X}_s is available, *i.e.*, $\mathcal{D}_s = \{x \in \mathcal{X}_s \mid y = 0\}$, where $y \in \{0, 1\}$ is the binary label indicating whether x is a normal ($y = 0$) or abnormal ($y = 1$) sample. During testing, data can be normal or abnormal, and can be from either the source or target distribution, *i.e.*, $\mathcal{D}_t = \{x \in \mathcal{X}_s \cup \mathcal{X}_t \mid y = \{0, 1\}\}$. The goal is then to develop an unsupervised anomaly detection model that can effectively handle distribution shift and accurately detect anomalies in \mathcal{D}_t . Specifically, we aim to learn a function $f : \mathcal{X} \rightarrow \mathbb{R}$ that assigns an anomaly score to each sample x in a way such that $\forall x_i, x_j \in \mathcal{D}_t, f(x_i) < f(x_j)$ when $y_i = 0$ and $y_j = 1$.

3.2. The Challenges

The current approaches in AD involve explicit fitting of the normal training data [2, 13, 65, 66]. It can cause the model to learn irrelevant features that are not associated with the appearance of normal data, *e.g.*, the model may mistake domain-specific background information as normal features, resulting in inaccurate anomaly detection when there are distribution shifts presented. OOD generalization




Figure 2: Anomaly scores of RD4AD [13], Mixstyle [89] and our model GNL on PACS [40] when selecting ‘house’ as the normal class and the remaining classes as anomaly classes.

models are also significantly challenged by the studied setting. This is mainly because the training data in the AD task consists of only one class and the data is monotonous, making it difficult to learn and identify patterns that distinguish normal and anomalous instances. Current OOD generalization approaches used in classification, detection, and segmentation need to take into account class labels, domain labels, or the diversity of samples in the training data [16, 29, 73, 74], which often are not applicable to AD tasks. As a result, new methods are required that can effectively address the problem of AD under distribution shift.

Fig. 2 illustrates this issue, where models such as RD4AD [13] (a recent SOTA AD model) and Mixstyle [89] (an OOD generalization method that we use to combine with RD4AD) are seen to struggle with identifying normal samples in the presence of distribution shift, often misclassifying them as anomalous. The overlapping of histograms of the anomaly scores for the normal and abnormal samples indicates that these models have learned features that are not representative of normal data, which can be a major obstacle in detecting anomalies. One of the main reasons is that the background or style features w.r.t. a specific dataset can change due to different natural conditions. As a result, the model may mistake these changed features as anomalies, leading to normal samples in the shifted distribution being classified as anomalous with high anomaly scores. Furthermore, some abnormal samples in the OOD data may possess similar features to background or style features in the training data, leading to them being misclassified with low anomaly scores.

4. Our Approach

To address these challenges, we introduce a novel approach, namely generalized normality learning (GNL). GNL minimizes the distribution gap between ID and OOD normal samples in both the training and inference stages in

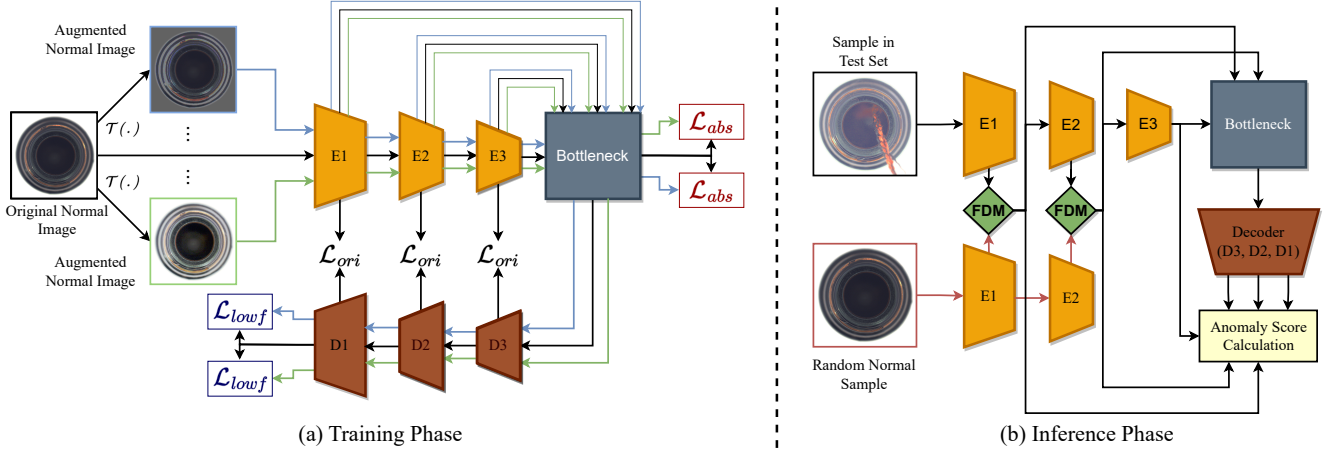


Figure 3: Overview of our approach. (a) Distribution-invariant normality learning in the training phase. (b) Test time augmentation with feature distribution matching in the inference phase.

an unsupervised way. To this end, we introduce a normality-preserved loss function to learn distribution-invariant normality representations, which enables GNL to learn generalized semantics of the normal training data at different feature levels. GNL further utilizes an AD-oriented test time data augmentation method based on feature distribution matching to improve the generalization performance. Fig. 3 describes the two main components of our approach: (a) distribution-invariant normality learning for training, and (b) test time augmentation methods. The two components complement to each other, meaning that the distribution-invariant normality learning process used during training can support the test time augmentation methods used during testing, and vice versa.

4.1. Distribution-invariant Normality Learning

In order to improve the performance of model on OOD datasets while maintaining good performance on ID datasets, we aim to train a student model to be more robust to changes in the distribution of data, while still ensuring that the student overfits on the normal features. Fig. 3 (a) illustrates the training framework.

Our method is built on top of the RD4AD model introduced by Deng et al. [13] that achieves state-of-the-art results on various datasets. The RD4AD framework includes three components: a fixed teacher encoder, a trainable one-class bottleneck embedding module, and a student decoder. When given an input sample, the teacher encoder extracts multi-scale representations, and the student decoder is trained to reconstruct the features from the bottleneck embedding. During testing, the teacher encoder can identify abnormal and OOD features in anomalous samples, but the student decoder fails to reconstruct these features. The model then considers anomalous representations that

have low similarity as highly abnormal.

We propose to incorporate a similarity loss that quantifies the difference between the embedding features of the original samples and those of each transformed normal sample that represents a distinct style from the original data. Specifically, we enforce this loss at both the bottleneck layer and the final block of the decoder. To provide further clarity, we propose the inclusion of a loss term, denoted as \mathcal{L}_{abs} , which is integrated at the bottleneck layer of the encoder. Moreover, we also introduce another loss term, termed as \mathcal{L}_{lowf} , that is added at the final block of the student decoder architecture. Particularly, given a sample $x \in \mathcal{D}_s$, we first apply an augmentation function $\mathcal{T}(\cdot)$ on it, and let $x'_k = \mathcal{T}(x)$ where $k \in [1, N]$ with N is the number of augmented normal samples generated by data augmentation, and ϕ be the mapping that projects the raw image I into the embedding space at the bottleneck layer, then we define \mathcal{L}_{abs} as:

$$\mathcal{L}_{abs} = \sum_{k=1}^N \frac{1}{N} \left\{ \mathcal{L}_{sim}(\phi(x), \phi(x'_k)) \right\}, \quad (1)$$

where $\mathcal{L}_{sim}(\cdot, \cdot)$ is a cosine similarity-based loss function.

Let ω be a reconstruction function from the abstract features to the low-level features at the final block of the decoder, then we further define \mathcal{L}_{lowf} as:

$$\mathcal{L}_{lowf} = \sum_{k=1}^N \frac{1}{N} \left\{ \mathcal{L}_{sim}(\omega(\phi(x)), \omega(\phi(x'_k))) \right\}. \quad (2)$$

We combine these loss functions to introduce the distribution-invariant, normality-preserved loss function:

$$\mathcal{L} = \lambda_{ori} * \mathcal{L}_{ori} + \lambda_{abs} * \mathcal{L}_{abs} + \lambda_{lowf} * \mathcal{L}_{lowf}, \quad (3)$$

where \mathcal{L}_{ori} is the original loss of RD4AD, and λ_{ori} , λ_{abs} , and λ_{lowf} are hyperparameters that determine how much weight should be given to each type of loss function.

We adopt AugMix [24] as the data augmentation method. Still, we remove the augmentation types that have the potential to generate anomalies, *e.g.*, ‘shear_x’, ‘shear_y’, ‘translate_x’, and ‘translate_y’, to ensure that all generated data are normal samples.

Intuitively, the last block of decoder is responsible for reconstructing simple and low-level features, such as edges, corners, and blobs, while the bottleneck layer is responsible for extracting more complex and high-level features. At the bottleneck layer, the abstracted information of the same images from different synthesized methods must be the same, while retaining enough information for reconstruction in the decoder. Therefore, by minimizing the loss function in Eq. 5, GNL learns features from both low-layer CNNs and high-level CNNs respectively to be the same from different distributions generated from a single sample.

4.2. Test Time Augmentation for Anomaly Detection under Distribution Shift

The goal of this component is to address the problem of a mismatch between the distribution of data during testing. To accomplish this, we propose injecting training distribution into the inference samples by using Feature Distribution Matching (FDM) at multi-level layers of the teacher encoder in the inference phase. The proposed testing framework is demonstrated in Fig. 3 (b). Our test time augmentation is applied at the first two residual blocks of the teacher encoder. The inference process from the third residual block onwards, as well as the calculation of the anomaly score, follow the original RD4AD framework without any modifications.

FDM is a group of techniques that aims to reduce the distribution mismatch or discrepancy of data from two different domains. Some previous studies focused on FDM assume that the input features follow a Gaussian distribution [28, 45, 52]. More recently, Zhang et al. [86] introduced a more accurate approach, known as Exact Feature Distribution Matching (EFDM). EFDM precisely matches empirical Cumulative Distribution Functions of image features, resulting in exact feature distribution alignment (as the sample size tends to infinity) and accurate matching of statistical properties like mean, standard deviation, and high-order statistics. Basically, all these FDM techniques are applicable to our proposed framework. Noted that FDM have been used for OOD Generalization, *e.g.*, in Mixstyle [89] and EDFMix [86], but they are used during training with the goal of creating new distribution samples by mixing the subdomain of the samples available in the training set, while we adopt FDM as a component in the inference stage with a different objective.

Specifically, given a test sample $p \in \mathcal{D}_t$, we randomly select a training normal sample $q \in \mathcal{D}_s$. These two samples are then fed into the teacher encoder. Let \mathcal{P}^m and \mathcal{Q}^m be the embedded features of p and q at the residual encoding block E^m , respectively, then the testing process is performed as follows:

$$\begin{cases} \mathcal{P}^{m+1} = \text{FDM}(E^{m+1}(\mathcal{P}^m), \mathcal{Q}^{m+1}, \alpha) \\ \mathcal{Q}^{m+1} = E^{m+1}(\mathcal{Q}^m) \\ \mathcal{P}^0 = p, \mathcal{Q}^0 = q, \end{cases} \quad (4)$$

where $m \in \{0, 1\}$ and α is a hyperparameter balancing the severity for mixing the style between the inference sample and the selected normal sample. The processed embedded features \mathcal{P}^1 and \mathcal{P}^2 are then input into the bottleneck layer and participate in the calculation of anomaly scores following the inference process of the original RD4AD.

For the FDM() function above, EFDM [86], which is the SOTA of FDM, is adopted to our method as follow:

$$\text{FDM}(\mathcal{C}, \mathcal{V}, \alpha) : \mathcal{C}_{\tau_i} = (1 - \alpha)\mathcal{C}_{\tau_i} + \alpha\mathcal{V}_{\kappa_i}, \quad (5)$$

where $\{\mathcal{C}_{\tau_i}\}_{i=1}^n$ and $\{\mathcal{V}_{\kappa_i}\}_{i=1}^n$ are sorted values of embedded feature \mathcal{C} and \mathcal{V} in ascending order. Here, n represents the number of elements in vector \mathcal{C} and \mathcal{V} . Note that \mathcal{C} is the embedded feature of the test sample p , which plays the role of carrying the appearance information. \mathcal{V} is the embedded feature of a normal sample q randomly sampled from the training data, carrying the style information.

In essence, the sample q plays a role in conveying distribution information pertaining to the training data. The selection of a random sample is due to the monotonous nature of the data during training, as any sample in the training set is capable of carrying distribution information that represents the training data. Thus, It helps avoid a process for careful sample selection that is often computationally expensive.

By utilizing FDM, our proposed testing process minimizes the disparity between the feature distribution of the inference sample and the feature distribution of normal samples in the training data, in cases where inference samples come from OOD sets. Furthermore, FDM ensures that the feature distribution remains nearly unchanged if inference samples come from ID sets, since the distribution of the test sample is aligned with its own distribution. Therefore, our testing approach can improve performance on OOD data without sacrificing performance on ID data.

5. Experiments

5.1. Datasets

We adapt four datasets from both AD and OOD generalization as the dataset benchmarks for the studied task.

Anomaly Detection. MVTec [5] is a widely-used AD benchmark, which comprises 15 data subsets for industrial defect inspection, including 5 subsets on texture anomalies and 10 subsets on object anomalies. The training dataset consists of 3,629 images in total, all of which are normal images. In contrast, the test dataset contains a total of 1,725 images, comprising both defective and non-defective instances. CIFAR-10 [35] serves as a one-class classification benchmark, featuring 50,000 training and 10,000 test images across 10 equally-sized categories representing diverse natural entities. In order to generate OOD datasets for MVTec and CIFAR-10, we apply 4 types of visual corruptions [22] to MVTec and CIFAR-10: Brightness, Contrast, Defocus Blur, and Gaussian noise. The severity for each type of corruption is set to 3 on MVTec and 5 on CIFAR-10 for obtaining the out-of-distribution data.

OOD Generalization. Two popular OOD benchmarks, MNIST-M [18] and PACS [40], are taken in our experiments. In particular, the primary MNIST [36] is used as the ID data on which the models are trained on, while MNIST-M is used as the OOD set. MNIST and MNIST-M datasets share 10 classes, which correspond to the digits 0 through 9. While MNIST encompasses 70,000 grayscale images of handwritten digits, MNIST-M contains 68,000 OOD images that are synthesized by superimposing random colored patches on the original images from MNIST. PACS is another widely used OOD dataset consisting of 9,991 images, which are shared by seven classes and four domains, namely Art, Cartoon, Photo, and Sketch. We select the images in Photo as the ID data, with the images in Art, Cartoon, and Sketch as the OOD data. The commonly used one-versus-all protocol [58] is used to convert these two datasets into AD datasets with distribution shift, in which samples of one class are used as normal, with the rest of classes as anomaly classes. Furthermore, we perform a multi-class setting on the MNIST/MNIST-M dataset, labeling samples from even-numbered classes as normal, while those from odd-numbered classes are identified as anomalies.

During training, we only use images in the ID dataset, *i.e.*, assuming the OOD data is not available during training. During inference, test sets of both ID and OOD are used.

5.2. Baselines

We conduct a series of experimental evaluations on 4 prominent anomaly detection methods, namely Deep SVDD [63], f-AnoGAN [66], KDAD [65], and RD4AD [13]. These methods stand for popular AD methods and recent state-of-the-art (SOTA) AD models. To evaluate the efficacy of OOD generalization techniques in anomaly detection, we adapt a suite of cutting-edge OOD methods by combining them with the recently proposed RD4AD model, which boosts SOTA performance on multiple datasets. Four different methods are used, includ-

ing three data augmentation-based methods Augmix [24], Mixstyle [89], and EFDM [86], and one self-supervised method Jigsaw [8].

5.3. Implementation Details

Our proposed method GNL is implemented on top of the RD4AD framework. Therefore, we maintain the settings recommended by RD4AD, such as the image size, the optimization method, the way of calculating anomaly score, and other relevant parameters. The details can be found in Appendix. Regarding the specific parameters for our model GNL, we choose $N = 2$ for the number of augmented normal samples generated by data augmentation. We set $\lambda_{ori} = 0.9$, $\lambda_{abs} = 0.05$ and $\lambda_{lowf} = 0.05$ by default for the distribution-invariant, normality-preserved loss function. During the inference phase, we opt for $\alpha = 0.5$ to control the degree of style blending for MVTec, PACS, and CIFAR-10 datasets, while setting $\alpha = 0.9$ for MNIST/MNIST-M, effectively managing the mixing dynamics. We choose EFDM [86] as the FDM technique since it is the latest and shows SOTA performance.

For the AD baselines, we use the official implementation published by the authors of those baselines. However, since the original baselines did not include experiments on the PACS dataset, we use the hyperparameters from MVTec experiments to conduct experiments on the PACS dataset corresponding to each baseline.

For the OOD generalization baselines, we use Augmix with an online augmentation severity of 3. We use all the data augmentation types included in Augmix for MNIST and PACS. However, for the MVTec and CIFAR-10 dataset, we exclude two types of augmentation that overlap with two types of corruptions during testing: Brightness and Contrast. With Mixstyle and EFDM, which are two data augmentation methods at the feature level (rather than at the image level like Augmix), we apply Mixstyle and EFDM to the encoders in the first two network layer according to the settings in RD4AD. As for Jigsaw, we fit the Jigsaw task into the Bottleneck component in RD4AD. All hyperparameters of training are preserved when applying OOD generalization baselines into RD4AD.

Following previous studies [13, 14, 63, 65, 66], we evaluate the performance of our anomaly detection methods using a metric called the Area Under the ROC Curve (AU-ROC). This metric is commonly used to assess how well a given method is able to distinguish between normal and anomalous data points. The results are averaged over three independent runs.

5.4. Comparison Results

The performance of our model GNL and the baselines on MVTec, CIFAR-10, MNIST, and PACS are shown in Tables 1, 2, 3 and 4, respectively. Note that due to space

limitations, the performances in all four tables are the average results of the classes per dataset. Detailed results are presented in Appendix. Overall, GNL can significantly outperforms SOTA AD models and OOD generalization methods in detecting anomalies on the OOD test data, while at the same time maintaining the detection accuracy on the ID data. Below we discuss the results in detail.

5.4.1 Performance of AD Methods

In general, we observe a significant drop in the AUC scores of all AD methods, Deep SVDD, f-AnoGAN, KDAD and RD4AD, on the OOD data across all four datasets used. This indicates that their performance is severely affected by the distribution shift. In particular, the performance of all AD models is promising on the MNIST set. However, this performance is reduced by about 30-40% when the models are tested on the MNIST-M set, which contains variations that are not present in the original MNIST set. Similar trends are observed in the PACS dataset, where the models' performance is also significantly affected by the distribution shifts in the OOD data. The models perform well on the Photo data, which is the ID data, but their performance drops significantly on the three OOD datasets, Art, Cartoon and Sketch. On MVTec and CIFAR-10, the performance still drops but is less severe than on the other two sets.

Method	ID	OOD			
	MVTec	Brightness	Contrast	Blur	Noise
Deep SVDD	69.98	55.18	50.07	68.82	59.11
f-AnoGAN	75.65	48.36	49.29	37.98	39.10
KDAD	85.50	83.81	64.03	84.17	82.04
RD4AD	98.64	96.50	94.12	98.9	90.14
Augmix	96.29	95.10	94.51	95.39	90.99
Mixstyle	98.58	96.60	94.45	98.27	88.92
EFDM	98.64	96.78	94.77	98.25	89.29
Augmix+Mixstyle	96.78	96.86	94.57	98.73	90.12
Augmix+EFDM	97.04	96.83	95.21	98.11	90.18
Jigsaw	73.97	73.36	67.88	73.88	72.60
GNL (Ours)	97.99	97.43	97.46	97.77	94.10

Table 1: AUROC (%) results on MVTec and its four corruptions. The best performance is **boldfaced**.

5.4.2 Performance of Combined OOD Generalization and AD Methods

Our results in Tables 1, 2, 3 and 4 indicate that the detection performance cannot be significantly improved by combining different OOD generalization techniques with the recent SOTA AD model RD4AD on the four datasets. This lack of improvement can be attributed to the fact that these OOD methods attempt to increase the diversity of data by enriching the available data based on its own distribution. However, because the training data in AD is typically

Method	ID	OOD			
	CIFAR	Brightness	Contrast	Blur	Noise
Deep SVDD	64.62	59.13	55.94	62.13	54.46
f-AnoGAN	70.25	54.62	57.23	60.74	51.76
KDAD	84.21	75.91	64.37	63.49	56.87
RD4AD	84.62	75.89	65.34	66.67	58.82
Augmix	82.83	74.15	62.48	66.92	57.36
Mixstyle	83.68	76.07	63.87	65.74	57.74
EFDM	83.92	76.19	63.92	64.81	57.63
Augmix+Mixstyle	83.87	76.02	65.55	63.89	58.04
Augmix+EFDM	82.96	75.73	64.39	63.83	57.14
Jigsaw	71.29	66.86	61.45	60.12	55.29
Ours	82.29	77.94	66.13	64.04	61.51

Table 2: AUROC (%) results on CIFAR-10 and its four corruptions.

Method	One-vs-All		Multi-class	
	ID	OOD	ID	OOD
Deep SVDD	97.73	49.92	86.94	51.19
f-AnoGAN	97.52	52.72	88.45	51.85
KDAD	98.87	54.87	90.43	52.84
RD4AD	98.89	58.09	88.70	51.74
Augmix	98.26	59.61	88.76	52.19
Mixstyle	98.84	57.22	87.36	52.13
EFDM	98.62	57.23	87.78	52.36
Augmix+Mixstyle	98.12	58.89	89.23	52.45
Augmix+EFDM	98.24	58.91	90.04	52.64
Jigsaw	98.90	58.51	87.29	52.87
GNL (Ours)	96.91	70.87	88.59	58.50

Table 3: AUROC results (%) on in-distribution (MNIST) and out-of-distribution (MNIST-M) datasets for one-vs-all and multi-class settings.

Method	ID	OOD		
	Photo	Art	Cartoon	Sketch
Deep SVDD	40.87	53.42	41.23	39.48
f-AnoGAN	61.34	50.15	52.42	63.77
KDAD	88.17	62.86	62.64	51.40
RD4AD	81.49	61.07	60.34	55.06
Augmix	76.35	60.50	58.96	57.86
Mixstyle	78.23	60.93	60.93	54.89
EFDM	78.47	60.55	62.15	55.63
Augmix+Mixstyle	76.12	60.16	61.29	55.76
Augmix+EFDM	77.28	60.93	63.18	56.67
Jigsaw	62.19	52.55	53.83	62.15
GNL (Ours)	87.67	65.62	67.96	62.39

Table 4: AUROC results (%) on in-distribution (Photo) and out-of-distribution (Art, Cartoon, Sketch) datasets.

monotonous and unimodal, these OOD methods often fail to generate data samples that significantly deviate from the original data distribution. As a result, the added diversity of generated data is not sufficient to significantly improve the performance of AD models.

Moreover, these OOD techniques also have a tendency

to generate undesired anomaly data, which is akin to injecting noise into the training data, thereby reducing the performance of AD models on the in-distribution dataset.

5.4.3 Performance of Our Method GNL

On the MVTEC AD dataset in Table 1, our method shows remarkable improvement in performance on the OOD dataset, while maintaining the performance on the ID data. In fact, our method achieves a highly comparable AUROC score of 0.9799 on the original MVTEC ID data, while also obtaining an impressive AUROC score of 0.9743 on the Brightness, 0.9746 on Contrast, 0.9777 on the Defocus_blur dataset, and 0.9410 on Gaussian Noise, which are significant improvements over the other methods. These results demonstrate the robustness and effectiveness of our GNL model to diverse distribution shifts.

Our experimental findings on the CIFAR-10 dataset exhibit a close resemblance to the outcomes observed on the MVTEC dataset, as depicted in Table 2. Our approach attains a competitive AUROC score of 0.8229 on the native MVTEC ID data. Notably, our method yields enhanced AUROC scores of 0.7794 for Brightness, 0.6613 for Contrast, 0.6404 for Defocus.blur, and 0.6151 for Gaussian Noise datasets, showcasing often large enhancements compared to other method, especially on the Contrast and Noise cases.

On the MNIST/MNIST-M dataset in Table 3, GNL consistently and significantly outperforms all other methods on the OOD data MNIST-M, increasing by at least 10 AUROC scores. Compared to the best performer – RD4AD – on the ID dataset that obtains an AUROC score of 0.9889, GNL exhibits a small decline and obtains an AUROC score of 0.9691. However, a significant improvement in performance is observed on the MNIST-M dataset, with an AUROC score of 0.7087 compared to 0.5809 for RD4AD. Regarding the multi-class setting, the results indicate its increased challenge compared to one-vs-all setting. Our model still maintains superior performance on OOD data while also excelling on ID data.

Similarly, GNL achieves consistently more superior AUROC performance on all four OOD datasets of PACS in Table 4. In particular, GNL obtains AUROC scores of 0.6562, 0.6796, and 0.6239 for the Art, Cartoon, and Sketch datasets, respectively, increasing by at least 5% on the Art and Cartoon datasets over the competing models. Compared to RD4AD, our method not only largely improves the OOD performance, but also enhances its performance on the ID data, the Photo data. This is because the Photo data contains multiple sub-domains, and RD4AD can be susceptible to overfitting on a specific sub-domain in the training data. By contrast, our method helps to mitigate this issue by learning more generalized normality representations, which improves performance across all sub-domains

within the Photo data. The performance of GNL on the ID data is also highly comparable to the best performer KDAD, 0.8767 vs. 0.8817, whereas GNL outperforms KDAD on the three OOD datasets by about 3%-10% in AUROC.

5.5. Robustness to Various Distribution Shift Levels

Fig. 4 presents the results of the robustness of GNL to varying levels of distribution shift, using the best competing methods RD4AD, Augmix and EFDM as baselines. The experiments are done on MVTEC with increasing levels of ‘Contrast’ corruption. Notably, the performance of the baselines exhibits a significant decline as the severity of corruption amplifies. The reason behind this phenomenon is intuitive as increased corruption severity introduces more substantial distribution variance, making it arduous for the models to discern between anomalous and normal samples. Our proposed method, on the other hand, demonstrates remarkable stability in performance across multiple levels of distribution shift. Our method maintains stable performance when the severity is between 1 and 3, and reduces to an AUROC of about 0.90 when the severity is 4 and 5, decreasing about 5% AUROC vs. about 30%-35% decrease in the competing methods. These results indicate strong robustness of GNL to heavy distribution shifts.




Figure 4: AUROC results on MVTEC with varying severity of the ‘Contrast’ corruption.

5.6. Ablation Study

We examine the importance of two main components: Distribution-invariant Normality Learning (DINL) using \mathcal{L}_{abs} and \mathcal{L}_{ori} individually or simultaneously (in addition to \mathcal{L}_{ori}), and AD-oriented Test Time Augmentation (ATTA) on the PACS dataset, with RD4AD as the baseline. The results are reported in Table 5. The experiment results show that \mathcal{L}_{abs} and \mathcal{L}_{lowf} positively contribute to the superior performance of DINL from low-level and high-level features respectively; and they can complement each other

when combining them in DINL. Looking more broadly, two main components, DINL and ATTA, also positively contribute to the superior performance of GNL. In particular, the experimental results show that if only the test time augmentation is applied, we gain about 2% AUROC improvement over the baseline on the OOD datasets, but it leads to a slight performance decrease on the ID data. When DINL is applied, it results in substantial improvement across both ID and OOD datasets, having 4%-7% AUROC improvement. When both are applied, we obtain the best performance, resulting in further substantial AUROC improvement. This indicates that both components, one reducing the distribution gap during training and another reducing the gap during inference, can well complement each other.

Method	ID		OOD	
	Photo	Art	Cartoon	Sketch
Baseline	81.49	61.07	60.34	55.06
\mathcal{L}_{abs} only	82.02	60.59	63.93	56.81
\mathcal{L}_{lowf} only	82.90	61.27	62.25	55.52
DINL	85.71	62.34	65.63	57.12
ATTA	81.05	64.36	62.04	57.04
DINL+ATTA	87.67	65.62	67.96	62.39

Table 5: AUROC results (%) of ablation study.




Figure 5: AUROC results using varying α . The smaller the α value, the lower the severity of style transfer.

5.7. Hyperparameter Analysis

Fig. 5 depicts how the performance of our model GNL changes with varying α , which is a hyperparameter in ATTA. The results suggest that the effectiveness of our model remains consistent across different α values on the Photo and Cartoon datasets. In contrast, the model’s ability to detect anomalies on the Art data appears to improve as α increases. However, for the Sketch data, the model’s performance reaches its maximum at $\alpha = 0.4$ and slightly

decreases as α increases further. Overall, a medium value, e.g., $\alpha = 0.5$, is generally recommended in practice.

We evaluate the hyperparameter sensitivity of our key component DINL using four settings of the three λ hyperparameters: λ_{ori} , λ_{abs} , and λ_{lowf} , with their sum set to one to ease the analysis. $\lambda_{abs} = \lambda_{lowf}$ is used as the features learned by them are considered equally important for the task. The results on PACS are shown in Table 6. DINL shows good robustness across different hyperparameter ratios in the three losses.

$\lambda_{ori}; \lambda_{abs}; \lambda_{lowf}$	Photo (ID)	Art	Cartoon	Sketch
0.95, 0.025, 0.025	85.90	62.71	64.28	59.97
0.90, 0.050, 0.050	85.71	62.34	65.63	57.12
0.85, 0.075, 0.075	85.60	63.23	65.03	58.57
0.80, 0.100, 0.100	84.89	61.45	65.12	56.84

Table 6: AUROC using various λ settings.

5.8. Time and Space Efficiency

For space complexity, our method improves the training objective and the inference of RD4AD without altering its architecture, thereby avoiding any increase in the number of parameters.

As shown in Table 7, in terms of time efficiency, our method’s training duration is slightly longer than RD4AD, but this additional time yields substantial performance improvements. As for the inference, our approach remains reasonably responsive.

Model	Training (per epoch)	Inference (per image)
RD4AD	2.5726	0.0282
Ours	7.2557	0.0356

Table 7: Runtime (s) on the ‘Dog’ dataset of PACS using one RTX 3090 24GB GPU.

6. Conclusion

In this work we propose a novel approach, namely GNL, to addressing the problem of anomaly detection in the presence of distribution shifts. GNL improves the generalization of the detection model by reducing the distribution gap between ID and OOD normal data in both training and inference stages. We also present comprehensive performance benchmarks and reveal that combined AD and OOD generalization methods do not work well for this task. Our approach is specifically designed for the OOD generalization in the AD task and shows significant improvement over the competing baselines. As shown in our results, our approach GNL is also robust to heavy distribution shifts. Overall, our approach represents an important contribution to unsupervised anomaly detection, as it addresses a more realistic problem that has not been adequately studied before.

References

- [1] Mohiuddin Ahmed, Abdun Naser Mahmood, and Md Rafiqul Islam. A survey of anomaly detection techniques in financial domain. *Future Generation Computer Systems*, 55:278–288, 2016.
- [2] Samet Akcay, Amir Atapour-Abarghouei, and Toby P Breckon. Ganomaly: Semi-supervised anomaly detection via adversarial training. In *Computer Vision–ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part III 14*, pages 622–637. Springer, 2019.
- [3] Isabela Albuquerque, Nikhil Naik, Junnan Li, Nitish Keskar, and Richard Socher. Improving out-of-distribution generalization via multi-task self-supervised pretraining. *arXiv preprint arXiv:2003.13525*, 2020.
- [4] Archana Anandakrishnan, Senthil Kumar, Alexander Shtatnikov, Tanveer Faruque, and Di Xu. Anomaly detection in finance: editors’ introduction. In *KDD 2017 Workshop on Anomaly Detection in Finance*, pages 1–7. PMLR, 2018.
- [5] Paul Bergmann, MiFchael Fauser, David Sattlegger, and Carsten Steger. Mytec ad—a comprehensive real-world dataset for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9592–9600, 2019.
- [6] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4183–4192, 2020.
- [7] Silvia Bucci, Antonio D’Innocente, Yujun Liao, Fabio M Carlucci, Barbara Caputo, and Tatiana Tommasi. Self-supervised learning across domains. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):5516–5528, 2021.
- [8] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2229–2238, 2019.
- [9] Fabio Maria Carlucci, Paolo Russo, Tatiana Tommasi, and Barbara Caputo. Hallucinating agnostic images to generalize across domains. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 3227–3234. IEEE, 2019.
- [10] Prithvijit Chattopadhyay, Yogesh Balaji, and Judy Hoffman. Learning to balance specificity and invariance for in and out of domain generalization. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX 16*, pages 301–318. Springer, 2020.
- [11] Chen Chen, Wenjia Bai, Rhodri H Davies, Anish N Bhuva, Charlotte H Manisty, Joao B Augusto, James C Moon, Nay Aung, Aaron M Lee, Mihir M Sanghvi, et al. Improving the generalizability of convolutional neural network-based segmentation on cmr images. *Frontiers in cardiovascular medicine*, 7:105, 2020.
- [12] Yuanhong Chen, Yu Tian, Guansong Pang, and Gustavo Carneiro. Deep one-class classification via interpolated gaussian descriptor. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 383–392, 2022.
- [13] Hanqiu Deng and Xingyu Li. Anomaly detection via reverse distillation from one-class embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9737–9746, 2022.
- [14] Choubo Ding, Guansong Pang, and Chunhua Shen. Catching both gray and black swans: Open-set supervised anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7388–7398, 2022.
- [15] Qi Dou, Daniel Coelho de Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. *Advances in Neural Information Processing Systems*, 32, 2019.
- [16] Yingjun Du, Jun Xu, Huan Xiong, Qiang Qiu, Xiantong Zhen, Cees GM Snoek, and Ling Shao. Learning to learn with variational information bottleneck for domain generalization. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X 16*, pages 200–216. Springer, 2020.
- [17] Yingjun Du, Xiantong Zhen, Ling Shao, and Cees GM Snoek. Metanorm: Learning to normalize few-shot batches across domains. In *International Conference on Learning Representations*, 2021.
- [18] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015.
- [19] Muhammad Ghifary, W Bastiaan Kleijn, Mengjie Zhang, and David Balduzzi. Domain generalization for object recognition with multi-task autoencoders. In *Proceedings of the IEEE international conference on computer vision*, pages 2551–2559, 2015.
- [20] Dong Gong, Lingqiao Liu, Vuong Le, Budhaditya Saha, Moussa Reda Mansour, Svetha Venkatesh, and Anton van den Hengel. Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1705–1714, 2019.
- [21] Sachin Goyal, Aditi Raghunathan, Moksh Jain, Harsha Vardhan Simhadri, and Prateek Jain. Drocc: Deep robust one-class classification. In *International conference on machine learning*, pages 3711–3721. PMLR, 2020.
- [22] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *Proceedings of the International Conference on Learning Representations*, 2019.
- [23] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. In *ICLR*, 2017.
- [24] Dan Hendrycks, Norman Mu, Ekin D Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. *arXiv preprint arXiv:1912.02781*, 2019.

- [25] Jinlei Hou, Yingying Zhang, Qiaoyong Zhong, Di Xie, Shiliang Pu, and Hong Zhou. Divide-and-assemble: Learning block-wise memory for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8791–8800, 2021.
- [26] Yen-Chang Hsu, Yilin Shen, Hongxia Jin, and Zsolt Kira. Generalized odin: Detecting out-of-distribution image without learning from out-of-distribution data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10951–10960, 2020.
- [27] Shoubu Hu, Kun Zhang, Zhitang Chen, and Laiwan Chan. Domain generalization via multidomain discriminant analysis. In *Uncertainty in Artificial Intelligence*, pages 292–302. PMLR, 2020.
- [28] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, pages 1501–1510, 2017.
- [29] Zeyi Huang, Haohan Wang, Eric P Xing, and Dong Huang. Self-challenging improves cross-domain generalization. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 124–140. Springer, 2020.
- [30] Maximilian Ilse, Jakub M Tomczak, Christos Louizos, and Max Welling. Diva: Domain invariant variational autoencoders. In *Medical Imaging with Deep Learning*, pages 322–348. PMLR, 2020.
- [31] Xin Jin, Cuiling Lan, Wenjun Zeng, and Zhibo Chen. Feature alignment and restoration for domain generalization and adaptation. *arXiv preprint arXiv:2006.12009*, 2020.
- [32] Aditya Khosla, Tinghui Zhou, Tomasz Malisiewicz, Alexei A Efros, and Antonio Torralba. Undoing the damage of dataset bias. In *Computer Vision—ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part I 12*, pages 158–171. Springer, 2012.
- [33] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [34] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, pages 5637–5664. PMLR, 2021.
- [35] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research). 2009. URL <http://www.cs.toronto.edu/kriz/cifar.html>, 5, 2009.
- [36] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [37] Alexander Lehner, Stefano Gasperini, Alvaro Marcos-Ramiro, Michael Schmidt, Mohammad-Ali Nikouei Mahani, Nassir Navab, Benjamin Busam, and Federico Tombari. 3d-vfield: Adversarial augmentation of point clouds for domain generalization in 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17295–17304, 2022.
- [38] Chun-Liang Li, Kihyuk Sohn, Jinsung Yoon, and Tomas Pfister. Cutpaste: Self-supervised learning for anomaly detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9664–9674, 2021.
- [39] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Sequential learning for domain generalization. In *Computer Vision—ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part I*, pages 603–619. Springer, 2021.
- [40] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, pages 5542–5550, 2017.
- [41] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5400–5409, 2018.
- [42] Haoliang Li, YuFei Wang, Renjie Wan, Shiqi Wang, Tie-Qiang Li, and Alex Kot. Domain generalization for medical imaging classification with linear-dependency regularization. *Advances in Neural Information Processing Systems*, 33:3118–3129, 2020.
- [43] Quande Liu, Cheng Chen, Jing Qin, Qi Dou, and Pheng-Ann Heng. Feddg: Federated domain generalization on medical image segmentation via episodic learning in continuous frequency space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1013–1023, 2021.
- [44] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. *Advances in neural information processing systems*, 33:21464–21475, 2020.
- [45] Ming Lu, Hao Zhao, Anbang Yao, Yurong Chen, Feng Xu, and Li Zhang. A closed-form solution to universal style transfer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5952–5961, 2019.
- [46] Yiwei Lu, Frank Yu, Mahesh Kumar Krishna Reddy, and Yang Wang. Few-shot scene-adaptive anomaly detection. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pages 125–141. Springer, 2020.
- [47] Hui Lv, Chen Chen, Zhen Cui, Chunyan Xu, Yong Li, and Jian Yang. Learning normal dynamics in videos with meta prototype network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15425–15434, 2021.
- [48] Qixiang Ma, Longyu Jiang, Wenxue Yu, Rui Jin, Zhixiang Wu, and Fangjin Xu. Training with noise adversarial network: A generalization method for object detection on sonar image. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 729–738, 2020.
- [49] Divyat Mahajan, Shruti Tople, and Amit Sharma. Domain generalization using causal matching. In *International Conference on Machine Learning*, pages 7313–7324. PMLR, 2021.

- [50] Udit Maniyar, Aniket Anand Deshmukh, Urun Dogan, Vineeth N Balasubramanian, et al. Zero shot domain generalization. *arXiv preprint arXiv:2008.07443*, 2020.
- [51] Hossein Mirzaei, Mohammadreza Salehi, Sajjad Shahabi, Efstratios Gavves, Cees GM Snoek, Mohammad Sabokrou, and Mohammad Hossein Rohban. Fake it until you make it: Towards accurate near-distribution novelty detection. In *The Eleventh International Conference on Learning Representations*, 2022.
- [52] Youssef Mroueh. Wasserstein style transfer. In *International Conference on Artificial Intelligence and Statistics*, pages 842–852. PMLR, 2020.
- [53] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International conference on machine learning*, pages 10–18. PMLR, 2013.
- [54] Sebastian Otálora, Manfredo Atzori, Vincent Andrearczyk, Amjad Khan, and Henning Müller. Staining invariant features for improving generalization of deep convolutional neural networks in computational pathology. *Frontiers in bioengineering and biotechnology*, 7:198, 2019.
- [55] Cheng Ouyang, Chen Chen, Surui Li, Zeju Li, Chen Qin, Wenjia Bai, and Daniel Rueckert. Causality-inspired single-source domain generalization for medical image segmentation. *IEEE Transactions on Medical Imaging*, 2022.
- [56] Guansong Pang, Chunhua Shen, Longbing Cao, and Anton Van Den Hengel. Deep learning for anomaly detection: A review. *ACM computing surveys (CSUR)*, 54(2):1–38, 2021.
- [57] Hyunjong Park, Jongyoun Noh, and Bumsub Ham. Learning memory-guided normality for anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14372–14381, 2020.
- [58] Pramuditha Perera, Ramesh Nallapati, and Bing Xiang. Ogan: One-class novelty detection using gans with constrained latent representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2898–2906, 2019.
- [59] Vihari Piratla, Praneeth Netrapalli, and Sunita Sarawagi. Efficient domain generalization via common-specific low-rank decomposition. In *International Conference on Machine Learning*, pages 7728–7738. PMLR, 2020.
- [60] Masoud Pourreza, Bahram Mohammadi, Mostafa Khaki, Samir Bouindour, Hichem Snoussi, and Mohammad Sabokrou. G2d: Generate to detect anomaly. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2003–2012, 2021.
- [61] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12556–12565, 2020.
- [62] Jie Ren, Peter J Liu, Emily Fertig, Jasper Snoek, Ryan Poplin, Mark Depristo, Joshua Dillon, and Balaji Lakshminarayanan. Likelihood ratios for out-of-distribution detection. *Advances in neural information processing systems*, 32, 2019.
- [63] Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft. Deep one-class classification. In *International conference on machine learning*, pages 4393–4402. PMLR, 2018.
- [64] Mohammad Sabokrou, Mahmood Fathy, Guoying Zhao, and Ehsan Adeli. Deep end-to-end one-class classifier. *IEEE transactions on neural networks and learning systems*, 32(2):675–684, 2020.
- [65] Mohammadreza Salehi, Niousha Sadjadi, Soroosh Baselizadeh, Mohammad H Rohban, and Hamid R Rabiee. Multiresolution knowledge distillation for anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14902–14912, 2021.
- [66] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Georg Langs, and Ursula Schmidt-Erfurth. f-anogan: Fast unsupervised anomaly detection with generative adversarial networks. *Medical image analysis*, 54:30–44, 2019.
- [67] Bernhard Schölkopf, John C Platt, John Shawe-Taylor, Alex J Smola, and Robert C Williamson. Estimating the support of a high-dimensional distribution. *Neural computation*, 13(7):1443–1471, 2001.
- [68] Nina Shvetsova, Bart Bakker, Irina Fedulova, Heinrich Schulz, and Dmitry V Dylov. Anomaly detection in medical imaging with deep perceptual autoencoders. *IEEE Access*, 9:118571–118583, 2021.
- [69] Md Amran Siddiqui, Jack W Stokes, Christian Seifert, Evan Argyle, Robert McCann, Joshua Neil, and Justin Carroll. Detecting cyber attacks using anomaly detection with explanations and expert feedback. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2872–2876. IEEE, 2019.
- [70] Aman Sinha, Hongseok Namkoong, Riccardo Volpi, and John Duchi. Certifying some distributional robustness with principled adversarial training. *arXiv preprint arXiv:1710.10571*, 2017.
- [71] David MJ Tax and Robert PW Duin. Support vector data description. *Machine learning*, 54:45–66, 2004.
- [72] Chee-Wooi Ten, Junho Hong, and Chen-Ching Liu. Anomaly detection for cybersecurity of the substations. *IEEE Transactions on Smart Grid*, 2(4):865–873, 2011.
- [73] Bailin Wang, Mirella Lapata, and Ivan Titov. Meta-learning for domain generalization in semantic parsing. *arXiv preprint arXiv:2010.11988*, 2020.
- [74] Guoqing Wang, Hu Han, Shiguang Shan, and Xilin Chen. Cross-domain face presentation attack detection via multi-domain disentangled representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6678–6687, 2020.
- [75] Guodong Wang, Shumin Han, Errui Ding, and Di Huang. Student-teacher feature pyramid matching for anomaly detection. *arXiv preprint arXiv:2103.04257*, 2021.
- [76] Haotao Wang, Aston Zhang, Yi Zhu, Shuai Zheng, Mu Li, Alex J Smola, and Zhangyang Wang. Partial and asymmetric contrastive learning for out-of-distribution detection in long-tailed recognition. In *International Conference on Machine Learning*, pages 23446–23458. PMLR, 2022.

- [77] Shujun Wang, Lequan Yu, Caizi Li, Chi-Wing Fu, and Pheng-Ann Heng. Learning from extrinsic and intrinsic supervisions for domain generalization. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX*, pages 159–176. Springer, 2020.
- [78] Peng Wu, Jing Liu, and Fang Shen. A deep one-class neural network for anomalous event detection in complex scenes. *IEEE transactions on neural networks and learning systems*, 31(7):2609–2622, 2019.
- [79] Xudong Yan, Huaidong Zhang, Xuemiao Xu, Xiaowei Hu, and Pheng-Ann Heng. Learning semantic context from normal samples for unsupervised anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 3110–3118, 2021.
- [80] Xufeng Yao, Yang Bai, Xinyun Zhang, Yuechen Zhang, Qi Sun, Ran Chen, Ruiyu Li, and Bei Yu. Pcl: Proxy-based contrastive learning for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7097–7107, 2022.
- [81] Jihun Yi and Sungroh Yoon. Patch svdd: Patch-level svdd for anomaly detection and segmentation. In *Proceedings of the Asian Conference on Computer Vision*, 2020.
- [82] Chris Yoon, Ghassan Hamarneh, and Rafeef Garbi. Generalizable feature learning in the presence of data bias and domain class imbalance with application to skin lesion classification. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part IV* 22, pages 365–373. Springer, 2019.
- [83] Muhammad Zaigham Zaheer, Jin-ha Lee, Marcella Astrid, and Seung-Ik Lee. Old is gold: Redefining the adversarially learned one-class classifier training paradigm. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14183–14193, 2020.
- [84] Vitjan Zavrtanik, Matej Kristan, and Danijel Škočaj. Reconstruction by inpainting for visual anomaly detection. *Pattern Recognition*, 112:107706, 2021.
- [85] Ling Zhang, Xiaosong Wang, Dong Yang, Thomas Sanford, Stephanie Harmon, Baris Turkbey, Bradford J Wood, Holger Roth, Andriy Myronenko, Daguang Xu, et al. Generalizing deep learning for medical image segmentation to unseen domains via deep stacked transformation. *IEEE transactions on medical imaging*, 39(7):2531–2540, 2020.
- [86] Yabin Zhang, Minghan Li, Ruihuang Li, Kui Jia, and Lei Zhang. Exact feature distribution matching for arbitrary style transfer and domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8035–8045, 2022.
- [87] Shanshan Zhao, Mingming Gong, Tongliang Liu, Huan Fu, and Dacheng Tao. Domain generalization via entropy regularization. *Advances in Neural Information Processing Systems*, 33:16096–16107, 2020.
- [88] Yuyang Zhao, Zhun Zhong, Fengxiang Yang, Zhiming Luo, Yaojin Lin, Shaozi Li, and Nicu Sebe. Learning to generalize unseen domains via memory-based multi-source meta-learning for person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6277–6286, 2021.
- [89] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle.