

CFLOW-AD: Real-Time Unsupervised Anomaly Detection with Localization via Conditional Normalizing Flows

Denis Gudovskiy¹

¹Panasonic AI Lab, USA

denis.gudovskiy@us.panasonic.com

Shun Ishizaka²

²Panasonic Technology Division, Japan

{ishizaka.shun, kozuka.kazuki}@jp.panasonic.com

Kazuki Kozuka²

Abstract

Unsupervised anomaly detection with localization has many practical applications when labeling is infeasible and, moreover, when anomaly examples are completely missing in the train data. While recently proposed models for such data setup achieve high accuracy metrics, their complexity is a limiting factor for real-time processing. In this paper, we propose a real-time model and analytically derive its relationship to prior methods. Our CFLOW-AD model is based on a conditional normalizing flow framework adopted for anomaly detection with localization. In particular, CFLOW-AD consists of a discriminatively pre-trained encoder followed by a multi-scale generative decoders where the latter explicitly estimate likelihood of the encoded features. Our approach results in a computationally and memory-efficient model: CFLOW-AD is faster and smaller by a factor of $10\times$ than prior state-of-the-art with the same input setting. Our experiments on the MVTec dataset show that CFLOW-AD outperforms previous methods by 0.36% AUROC in detection task, by 1.12% AUROC and 2.5% AUPRO in localization task, respectively. We open-source our code with fully reproducible experiments¹.

1. Introduction

Anomaly detection with localization (AD) is a growing area of research in computer vision with many practical applications *e.g.* industrial inspection [4], road traffic monitoring [20], medical diagnostics [44] *etc.* However, the common supervised AD [32] is not viable in practical applications due to several reasons. First, it requires labeled data which is costly to obtain. Second, anomalies are usually rare long-tail examples and have low probability to be acquired by sensors. Lastly, consistent labeling of anomalies is subjective and requires extensive domain expertise as illustrated in Figure 1 with industrial cable defects.

¹Our code is available at github.com/gudovskiy/cflow-ad

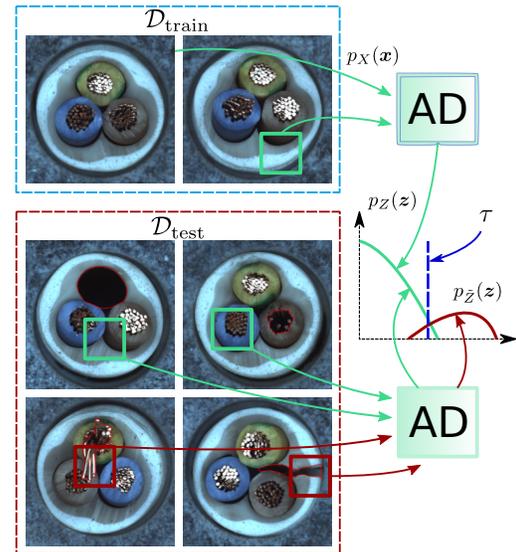


Figure 1. An example of the proposed out-of-distribution (OOD) detector for anomaly localization trained on anomaly-free $\mathcal{D}_{\text{train}}$ (top row). Sliced cable images are from the MVTec dataset [4], where the bottom row illustrates $\mathcal{D}_{\text{test}}$ ground truth masks for anomalies (red) and the middle row shows examples of anomaly-free patches (green). The OOD detector learns the distribution of anomaly-free patches \mathbf{x} with $p_X(\mathbf{x})$ density and transforms it into a Gaussian distribution with $p_Z(\mathbf{z})$ density. Threshold τ separates in-distribution patches from the OOD patches with $p_{\bar{Z}}(\mathbf{z})$ density.

With these limitations of the supervised AD, a more appealing approach is to collect only *unlabeled anomaly-free images* for train dataset $\mathcal{D}_{\text{train}}$ as in Figure 1 (top row). Then, any deviation from anomaly-free images is classified as an anomaly. Such data setup with low rate of anomalies is generally considered to be *unsupervised* [4]. Hence, the AD task can be reformulated as a task of *out-of-distribution detection* (OOD) with the AD objective.

While OOD for low-dimensional industrial sensors (*e.g.* power-line or acoustic) can be accomplished using a common k -nearest-neighbor or more advanced clustering methods [10], it is less trivial for high-resolution images. Re-

cently, convolutional neural networks (CNNs) have gained popularity in extracting semantic information from images into downsampled feature maps [4]. Though feature extraction using CNNs has relatively low complexity, the post-processing of feature maps is far from real-time processing in the state-of-the-art unsupervised AD methods [8].

To address this complexity drawback, we propose a CFLOW-AD model that is based on conditional normalizing flows. CFLOW-AD is agnostic to feature map spatial dimensions similar to CNNs, which leads to a higher accuracy metrics as well as a lower computational and memory requirements. We present the main idea behind our approach in a toy OOD detector example in Figure 1. A distribution of the anomaly-free image patches \mathbf{x} with probability density function $p_X(\mathbf{x})$ is learned by the AD model. Our translation-equivariant model is trained to transform the original distribution with $p_X(\mathbf{x})$ density into a Gaussian distribution with $p_Z(\mathbf{z})$ density. Finally, this model separates in-distribution patches \mathbf{z} with $p_Z(\mathbf{z})$ from the out-of-distribution patches with $p_{\bar{Z}}(\mathbf{z})$ using a threshold τ computed as the Euclidean distance from the distribution mean.

2. Related work

We review models² that employ the data setup from Figure 1 and provide experimental results for popular MVTEC dataset [4] with factory defects or Shanghai Tech Campus (STC) dataset [22] with surveillance camera videos. We highlight the research related to a more challenging task of *pixel-level anomaly localization (segmentation)* rather than a more simple *image-level anomaly detection*.

Napoletano *et al.* [25] propose to use CNN feature extractors followed by a principal component analysis and k -mean clustering for AD. Their feature extractor is a ResNet-18 [13] pretrained on a large-scale ImageNet dataset [16]. Similarly, SPADE [7] employs a Wide-ResNet-50 [43] with multi-scale pyramid pooling that is followed by a k -nearest-neighbor clustering. Unfortunately, clustering is slow at test-time with high-dimensional data. Thus, parallel convolutional methods are preferred in real-time systems.

Numerous methods are based on a natural idea of *generative modeling*. Unlike models with the discriminatively-pretrained feature extractors [25, 7], generative models learn distribution of anomaly-free data and, therefore, are able to estimate a proxy metrics for anomaly scores even for the unseen images with anomalies. Recent models employ generative adversarial networks (GANs) [35, 36] and variational autoencoders (VAEs) [3, 38].

A fully-generative models [35, 36, 3, 38] are directly applied to images in order to estimate pixel-level probability density and compute per-pixel reconstruction errors as

anomaly scores proxies. These fully-generative models are unable to estimate the *exact data likelihoods* [6, 24] and do not perform better than the traditional methods [25, 7] according to MVTEC survey in [4]. Recent works [34, 15] show that these models tend to capture only low-level correlations instead of relevant semantic information. To overcome the latter drawback, a hybrid DFR model [37] uses a pretrained feature extractor with multi-scale pyramid pooling followed by a convolutional autoencoder (CAE). However, DFR model is unable to estimate the exact likelihoods.

Another line of research proposes to employ a student-teacher type of framework [5, 33, 41]. Teacher is a pretrained feature extractor and student is trained to estimate a scoring function for AD. Unfortunately, such frameworks underperform compared to state-of-the-art models.

Patch SVDD [42] and CutPaste [19] introduce a self-supervised pretraining scheme for AD. Moreover, Patch SVDD proposes a novel method to combine multi-scale scoring masks to a final anomaly map. Unlike the nearest-neighbor search in [42], CutPaste estimates anomaly scores using an efficient Gaussian density estimator. While the self-supervised pretraining can be helpful in uncommon data domains, Schirrmeister *et al.* [34] argue that large natural-image datasets such as ImageNet can be a more representative for pretraining compared to a small application-specific datasets *e.g.* industrial MVTEC [4].

The state-of-the-art PaDiM [8] proposes surprisingly simple yet effective approach for anomaly localization. Similarly to [37, 7, 42], this approach relies on ImageNet-pretrained feature extractor with multi-scale pyramid pooling. However, instead of slow test-time clustering in [7] or nearest-neighbor search in [42], PaDiM uses a well-known Mahalanobis distance metric [23] as an anomaly score. The metric parameters are estimated for each feature vector from the pooled feature maps. PaDiM has been inspired by Rippel *et al.* [29] who firstly advocated to use this measure for anomaly detection without localization.

DifferNet [30] uses a promising class of generative models called *normalizing flows* (NFWs) [9] for image-level AD. The main advantage of NFLOW models is ability to estimate the exact likelihoods for OOD compared to other generative models [35, 36, 3, 38, 37]. In this paper, we extend DifferNet approach to pixel-level anomaly localization task using our CFLOW-AD model. In contrast to RealNVP [9] architecture with global average pooling in [30], we propose to use conditional normalizing flows [2] to make CFLOW-AD suitable for low-complexity processing of multi-scale feature maps for localization task. We develop our CFLOW-AD with the following contributions:

- Our theoretical analysis shows why multivariate Gaussian assumption is a justified prior in previous models and why a more general NFLOW framework objective converges to similar results with the less compute.

²For comprehensive review of the existing AD methods we refer readers to Ruff *et al.* [31] and Pang *et al.* [26] surveys.

- We propose to use conditional normalizing flows for unsupervised anomaly detection with localization using computational and memory-efficient architecture.
- We show that our model outperforms previous state-of-the-art in both detection and localization due to the unique properties of the proposed CFLOW-AD model.

3. Theoretical background

3.1. Feature extraction with Gaussian prior

Consider a CNN $h(\lambda)$ trained for classification task. Its parameters λ are usually found by minimizing Kullback-Leibler (D_{KL}) divergence between joint train data distribution $Q_{x,y}$ and the learned model distribution $P_{x,y}(\lambda)$, where (x, y) is an input-label pair for supervised learning.

Typically, the parameters λ are initialized by the values sampled from the Gaussian distribution [12] and optimization process is regularized as

$$\arg \min_{\lambda} D_{KL} [Q_{x,y} \| P_{x,y}(\lambda)] + \alpha R(\lambda), \quad (1)$$

where $R(\lambda)$ is a regularization term and α is a hyperparameter that defines regularization strength.

The most popular CNNs [13, 43] are trained with L_2 weight decay [17] regularization ($R(\lambda) = \|\lambda\|_2^2$). That imposes multivariate Gaussian (MVG) prior not only to parameters λ , but also to the feature vectors z extracted from the feature maps of $h(\lambda)$ [11] intermediate layers.

3.2. A case for Mahalanobis distance

With the same MVG prior assumption, Lee *et al.* [18] recently proposed to model distribution of feature vectors z by MVG density function and to use Mahalanobis distance [23] as a confidence score in CNN classifiers. Inspired by [18], Rippel *et al.* [29] adopt Mahalanobis distance for anomaly detection task since this measure determines a distance of a particular feature vector z to its MVG distribution. Consider a MVG distribution $\mathcal{N}(\mu, \Sigma)$ with a density function $p_Z(z)$ for random variable $z \in \mathbb{R}^D$ defined as

$$p_Z(z) = (2\pi)^{-D/2} \det \Sigma^{-1/2} e^{-\frac{1}{2}(z-\mu)^T \Sigma^{-1}(z-\mu)}, \quad (2)$$

where $\mu \in \mathbb{R}^D$ is a mean vector and $\Sigma \in \mathbb{R}^{D \times D}$ is a covariance matrix of a true anomaly-free density $p_Z(z)$.

Then, the Mahalanobis distance $M(z)$ is calculated as

$$M(z) = \sqrt{(z - \mu)^T \Sigma^{-1}(z - \mu)}, \quad (3)$$

Since the true anomaly-free data distribution is unknown, mean vector and covariance matrix from (3) are replaced by the estimates $\hat{\mu}$ and $\hat{\Sigma}$ calculated from the empirical train dataset $\mathcal{D}_{\text{train}}$. At the same time, density function $p_{\tilde{Z}}(z)$ of anomaly data has different $\tilde{\mu}$ and $\tilde{\Sigma}$ statistics, which allows to separate out-of-distribution and in-distribution feature vectors using $M(z)$ from (3).

This framework with MVG distribution assumption shows its effectiveness in image-level anomaly detection task [29] and is adopted by the state-of-the-art PaDiM [8] model in pixel-level anomaly localization task.

3.3. Relationship with the flow framework

Dinh *et al.* [9] introduce a class of generative probabilistic models called normalizing flows. These models apply change of variable formula to fit an arbitrary density $p_Z(z)$ by a tractable base distribution with $p_U(u)$ density and a bijective invertible mapping $g^{-1} : Z \rightarrow U$. Then, the log-likelihood of any $z \in Z$ can be estimated by

$$\log \hat{p}_Z(z, \theta) = \log p_U(u) + \log |\det \mathbf{J}|, \quad (4)$$

where a sample $u \sim p_U$ is usually from standard MVG distribution ($u \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$) and a matrix $\mathbf{J} = \nabla_z g^{-1}(z, \theta)$ is the Jacobian of a bijective invertible flow model ($z = g(u, \theta)$ and $u = g^{-1}(z, \theta)$) parameterized by vector θ .

The flow model $g(\theta)$ is a set of basic layered transformations with tractable Jacobian determinants. For example, $\log |\det \mathbf{J}|$ in RealNVP [9] coupling layers is a simple sum of layer's diagonal elements. These models are optimized using stochastic gradient descent by maximizing log-likelihood in (4). Equivalently, optimization can be done by minimizing the reverse $D_{KL} [\hat{p}_Z(z, \theta) \| p_Z(z)]$ [27], where $\hat{p}_Z(z, \theta)$ is the model prediction and $p_Z(z)$ is a target density. The loss function for this objective is defined as

$$\mathcal{L}(\theta) = \mathbb{E}_{\hat{p}_Z(z, \theta)} [\log \hat{p}_Z(z, \theta) - \log p_Z(z)]. \quad (5)$$

If $p_Z(z)$ is distributed according to Section 3.1 MVG assumption, we can express (5) as a function of Mahalanobis distance $M(z)$ using its definition from (3) as

$$\mathcal{L}(\theta) = \mathbb{E}_{\hat{p}_Z(z, \theta)} \left[\frac{M^2(z) - E^2(u)}{2} + \log \frac{|\det \mathbf{J}|}{\det \Sigma^{-1/2}} \right], \quad (6)$$

where $E^2(u) = \|u\|_2^2$ is a squared Euclidean distance of a sample $u \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ (detailed proof in Appendix A).

Then, the loss in (6) converges to zero when the likelihood contribution term $|\det \mathbf{J}|$ of the model $g(\theta)$ (normalized by $\det \Sigma^{-1/2}$) compensates the difference between a squared Mahalanobis distance for z from the target density and a squared Euclidean distance for $u \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

This normalizing flow framework can estimate the exact likelihoods of any arbitrary distribution with p_Z density, while Mahalanobis distance is limited to MVG distribution only. For example, CNNs trained with L_1 regularization would have Laplace prior [11] or have no particular prior in the absence of regularization. Moreover, we introduce conditional normalizing flows in the next section and show that they are more compact in size and have fully-convolutional parallel architecture compared to [7, 8] models.

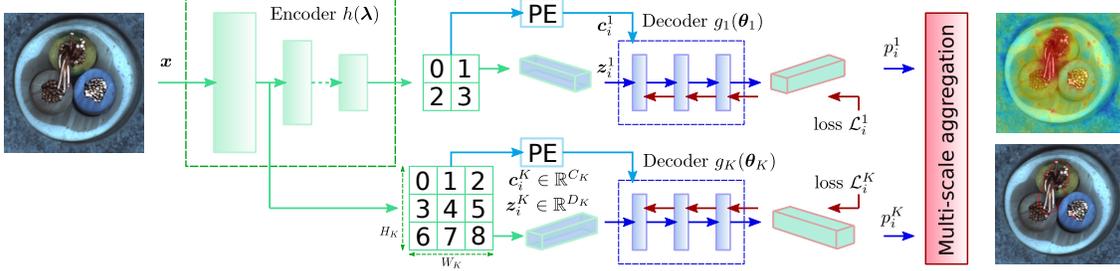


Figure 2. Overview of our CFLOW-AD with a fully-convolutional translation-equivariant architecture. Encoder $h(\lambda)$ is a CNN feature extractor with multi-scale pyramid pooling. Pyramid pooling captures both global and local semantic information with the growing from top to bottom receptive fields. Pooled feature vectors z_i^k are processed by a set of decoders $g_k(\theta_k)$ independently for each k th scale. Our decoder is a conditional normalizing flow network with a feature input z_i^k and a conditional input c_i^k with spatial information from a positional encoder (PE). The estimated multi-scale likelihoods p_i^k are upsampled to the input size and added up to produce anomaly map.

4. The proposed CFLOW-AD model

4.1. CFLOW encoder for feature extraction

We implement a feature extraction scheme with multi-scale feature pyramid pooling similar to recent models [7, 8]. We define the discriminatively-trained CNN feature extractor as an encoder $h(\lambda)$ in Figure 2. The CNN encoder maps image patches x into a feature vectors z that contain relevant semantic information about their content. CNNs accomplish this task efficiently due to their translation-equivariant architecture with the shared kernel parameters. In our experiments, we use ImageNet-pretrained encoder following Schirrmeyer *et al.* [34] who show that large natural-image datasets can serve as a representative distribution for pretraining. If a large application-domain unlabeled data is available, the self-supervised pretraining from [42, 19] can be a viable option.

One important aspect of a CNN encoder is its *effective receptive field* [21]. Since the effective receptive field is not strictly bounded, the size of encoded patches x cannot be exactly defined. At the same time, anomalies have various sizes and shapes, and, ideally, they have to be processed with the variable receptive fields. To address the ambiguity between CNN receptive fields and anomaly variability, we adopt common multi-scale feature pyramid pooling approach. Figure 2 shows that the feature vectors $z_i^k \in \mathbb{R}^{D_k}$, $i \in \{H_k \times W_k\}$ are extracted by K pooling layers. Pyramid pooling captures both local and global patch information with small and large receptive fields in the first and last CNN layers, respectively. For convenience, we number pooling layers in the last to first layer order.

4.2. CFLOW decoders for likelihood estimation

We use the general normalizing flow framework from Section 3.3 to estimate log-likelihoods of feature vectors z . Hence, our generative decoder model $g(\theta)$ aims to fit true density $p_Z(z)$ by an estimated parameterized density $\hat{p}_Z(z, \theta)$ from (1). However, the feature vectors are as-

sumed to be independent of their spatial location in the general framework. To increase efficacy of distribution modeling, we propose to incorporate spatial prior into $g(\theta)$ model using conditional flow framework. In addition, we model $\hat{p}_Z^k(z, \theta)$ densities using K independent decoder models $g_k(\theta_k)$ due to multi-scale feature pyramid pooling setup.

Our conditional normalizing flow (CFLOW) decoder architecture is presented in Figure 2. We generate a conditional vector c_i^k using a 2D form of conventional positional encoding (PE) [39]. Each $c_i^k \in \mathbb{R}^{C_k}$ contains sin and cos harmonics that are unique to its spatial location $(h_k, w_k)_i$. We extend unconditional flow framework to CFLOW by concatenating the intermediate vectors inside decoder coupling layers with the conditional vectors c_i as in [2].

Then, the k th CFLOW decoder contains a sequence of conventional coupling layers with the additional conditional input. Each coupling layer comprises of fully-connected layer with $(D_k + C_k) \times (D_k + C_k)$ kernel, softplus activation and output vector permutations. Usually, the conditional extension does not increase model size since $C_k \ll D_k$. For example, we use the fixed $C_k = 128$ in all our experiments. Our CFLOW decoder has translation-equivariant architecture, because it slides along feature vectors extracted from the intermediate feature maps with kernel parameter sharing. As a result, both the encoder $h(\lambda)$ and decoders $g_k(\theta_k)$ have convolutional translation-equivariant architectures.

We train CFLOW-AD using a maximum likelihood objective, which is equivalent to minimizing loss defined by

$$\mathcal{L}(\theta) = D_{KL} [p_Z(z) \| \hat{p}_Z(z, \mathbf{c}, \theta)] \approx \frac{1}{N} \sum_{i=1}^N \left[\frac{\|u_i\|_2^2}{2} - \log |\det J_i| \right] + \text{const}, \quad (7)$$

where the random variable $u_i = g^{-1}(z_i, c_i, \theta)$, the Jacobian $J_i = \nabla_z g^{-1}(z_i, c_i, \theta)$ for CFLOW decoder and an expectation operation in D_{KL} is replaced by an empirical train dataset $\mathcal{D}_{\text{train}}$ of size N . For brevity, we drop the k th scale notation. The derivation is given in Appendix B.

Table 1. Complexity estimates for SPADE [7], PaDiM [8] and our CFLOW-AD. We compare train and test complexity as well as memory requirements. All models use the same encoder $h(\lambda)$ setup, but diverge in the post-processing. SPADE allocates memory for a train gallery \mathbb{G} used in k -nearest-neighbors. PaDiM keeps large matrices $(\Sigma_i^k)^{-1}$, $i \in \{H_k \times W_k\}$ for Mahalanobis distance. Our model employs trained decoders $g_k(\theta_k)$ for post-processing.

Model	Train	Test	Memory
SPADE [7]	\mathbb{G}	$\sum_{\mathbf{v} \in \mathbb{G}} \ \mathbf{v} - \mathbf{z}_i\ _2^2$	$\lambda + \mathbb{G}$
PaDiM [8]	Σ_i^{-1}	$M(\mathbf{z}_i)$	$\lambda + \Sigma_i^{-1}$
Ours	$\mathcal{L}(\theta)$	$\log \hat{p}_Z(\mathbf{z}_i, \mathbf{c}_i, \theta)$	$\lambda + \theta$

After training the decoders $g_k(\theta_k)$ for all K scales using (7), we estimate test dataset $\mathcal{D}_{\text{test}}$ log-likelihoods as

$$\log \hat{p}_Z(\mathbf{z}_i, \mathbf{c}_i, \hat{\theta}) = -\frac{\|\mathbf{u}_i\|_2^2 + D \log(2\pi)}{2} + \log |\det \mathbf{J}_i|. \quad (8)$$

Next, we convert log-likelihoods to probabilities $p_i^k = e^{\log \hat{p}_Z(\mathbf{z}_i^k, \mathbf{c}_i^k, \hat{\theta}_k)}$ for each k th scale using (8) and normalize them to be in $[0 : 1]$ range. Then, we upsample p_i^k to the input image resolution ($H \times W$) using bilinear interpolation $\mathbf{P}_k = b(p_i^k) \in \mathbb{R}^{H \times W}$. Finally, we calculate anomaly score maps $\mathbf{S}^{H \times W}$ by aggregating all upsampled probabilities as $\mathbf{S} = \max_{\sum_{k=1}^K} \mathbf{P}_k - \sum_{k=1}^K \mathbf{P}_k$.

4.3. Complexity analysis

Table 1 analytically compares complexity of CFLOW-AD and recent state-of-the-art models with the same pyramid pooling setup *i.e.* $\mathbf{z}^{k \in \{1 \dots K\}} = h(\mathbf{x}, \lambda)$.

SPADE [7] performs k -nearest-neighbor clustering between each test point \mathbf{z}_i^k and a gallery \mathbb{G} of train data. Therefore, the method requires large memory allocation for gallery \mathbb{G} and a clustering procedure that is typically slow compared to convolutional methods.

PaDiM [8] estimates train-time statistics *i.e.* inverses of covariance matrices $(\Sigma_i^k)^{-1} \in \mathbb{R}^{D_k \times D_k}$, $i \in \{H_k \times W_k\}$ to calculate $M(\mathbf{z}_i^k)$ at test-time. Hence, it has low computational complexity, but it stores in memory $H_k \times W_k$ matrices of $D_k \times D_k$ size for every k th pooling layer.

Our method optimizes generative decoders $g_k(\theta_k)$ using (7) during the train phase. At the test phase, CFLOW-AD simply infers data log-likelihoods $\log \hat{p}_Z(\mathbf{z}_i, \mathbf{c}_i, \hat{\theta})$ using (8) in a fully-convolutional fashion. Decoder parameters $\theta_{k \in \{1 \dots K\}}$ are relatively small as reported in Table 6.

5. Experiments

5.1. Experimental setup

We conduct unsupervised anomaly detection (image-level) and localization (pixel-level segmentation) experiments using the MVTEC [4] dataset with factory defects and the STC [22] dataset with surveillance camera videos. The

code is in PyTorch [28] with the FrEIA library [1] used for generative normalizing flow modeling.

Industrial MVTEC dataset comprises 15 classes with total of 3,629 images for training and 1,725 images for testing. The train dataset contains only anomaly-free images without any defects. The test dataset contains both images containing various types of defects and defect-free images. Five classes contain different types of textures (carpet, grid, leather, tile, wood), while the remaining 10 classes represent various types of objects. We resize MVTEC images without cropping according to the specified image resolution (*e.g.* $H \times W = 256 \times 256, 512 \times 512$ *etc.*) and apply $\pm 5^\circ$ augmentation rotations during training phase only.

STC dataset contains 274,515 training and 42,883 testing frames extracted from surveillance camera videos and divided into 13 distinct university campus scenes. Because STC is significantly larger than MVTEC, we experiment only with 256×256 resolution and apply the same pre-processing and augmentation pipeline as for MVTEC.

We compare CFLOW-AD with the models reviewed in Section 2 using MVTEC and STC datasets. We use widely-used threshold-agnostic evaluation metrics for localization: area under the receiver operating characteristic curve (AUROC) and area under the per-region-overlap curve (AUPRO) [4]. AUROC is skewed towards large-area anomalies, while AUPRO metric ensures that both large and small anomalies are equally important in localization. Image-level AD detection is reported by the AUROC only.

We run each CFLOW-AD experiment four times on the MVTEC and report mean (μ) of the evaluation metric and, if specified, its standard deviation ($\pm \sigma$). For the larger STC dataset, we conduct only a single experiment. As in other methods, we train a separate CFLOW-AD model for each MVTEC class and each STC scene. All our models use the same training hyperparameters: Adam optimizer with $2e-4$ learning rate, 100 train epochs, 32 mini-batch size for encoder and cosine learning rate annealing with 2 warm-up epochs. Since our decoders are agnostic to feature map dimensions and have low memory requirements, we train and test CFLOW-AD decoders with 8,192 (32×256) mini-batch size for feature vector processing. During the train phase 8,192 feature vectors are randomly sampled from 32 random feature maps. Similarly, 8,192 feature vectors are sequentially sampled during the test phase. The feature pyramid pooling setup for ResNet-18 and WideResnet-50 encoder is identical to PaDiM [8]. The effects of other architectural hyperparameters are studied in the ablation study.

5.2. Ablation study

Table 2 presents a comprehensive study of various design choices for CFLOW-AD on the MVTEC dataset using AUROC metric. In particular, we experiment with the input image resolution ($H \times W$), encoder architecture (ResNet-

Table 2. Ablation study of CFLOW-AD using localization AUROC metric on the MVTec [4] dataset, %. We experiment with input image resolution ($H \times W$), encoder architecture (ResNet-18 (R18), WideResnet-50 (WRN50) and MobileNetV3L (MNetV3)), type of normalizing flow (unconditional (UFLOW) and conditional (CFLOW)), number of coupling (# of CL) and pooling layers (# of PL).

Encoder	WRN50	WRN50	WRN50	WRN50	WRN50	R18	R18	MNetV3	MNetV3
# of CL	4	→ 8	8	8	8	8	8	8	8
# of PL	2	2	→ 3	3	3	3	3	3	3
H×W	256	256	256	→ 512	512	256	→ 512	256	→ 512
Type	CFLOW	CFLOW	CFLOW	CFLOW → UFLOW	CFLOW	CFLOW	CFLOW	CFLOW	CFLOW
Bottle	97.28 \pm 0.03	97.24 \pm 0.03	98.76 \pm 0.01	98.98 \pm 0.01	98.83 \pm 0.01	98.47 \pm 0.03	98.64 \pm 0.01	98.74	98.92
Cable	95.71 \pm 0.01	96.17 \pm 0.07	97.64 \pm 0.04	97.12 \pm 0.06	95.29 \pm 0.04	96.75 \pm 0.04	96.07 \pm 0.06	97.62	97.49
Capsule	98.17 \pm 0.02	98.19 \pm 0.05	98.98 \pm 0.00	98.64 \pm 0.02	98.40 \pm 0.12	98.62 \pm 0.02	98.28 \pm 0.05	98.89	98.75
Carpet	98.50 \pm 0.01	98.55 \pm 0.01	99.23 \pm 0.01	99.25 \pm 0.01	99.24 \pm 0.00	99.00 \pm 0.01	99.29 \pm 0.00	98.64	99.00
Grid	93.77 \pm 0.05	93.88 \pm 0.16	96.89 \pm 0.02	98.99 \pm 0.02	98.74 \pm 0.00	93.95 \pm 0.04	98.53 \pm 0.01	94.75	98.81
Hazelnut	98.08 \pm 0.01	98.13 \pm 0.02	98.82 \pm 0.01	98.89 \pm 0.01	98.88 \pm 0.01	98.81 \pm 0.01	98.41 \pm 0.01	98.88	99.00
Leather	98.92 \pm 0.02	99.00 \pm 0.06	99.61 \pm 0.01	99.66 \pm 0.00	99.65 \pm 0.00	99.45 \pm 0.01	99.51 \pm 0.02	99.50	99.64
Metal Nut	96.72 \pm 0.03	96.72 \pm 0.06	98.56 \pm 0.03	98.25 \pm 0.04	98.16 \pm 0.03	97.59 \pm 0.05	96.42 \pm 0.03	98.36	98.78
Pill	98.46 \pm 0.02	98.46 \pm 0.01	98.95 \pm 0.00	98.52 \pm 0.05	98.20 \pm 0.08	98.34 \pm 0.02	97.80 \pm 0.05	98.69	98.44
Screw	94.98 \pm 0.06	95.28 \pm 0.06	98.10 \pm 0.05	98.86 \pm 0.02	98.78 \pm 0.01	97.38 \pm 0.03	98.40 \pm 0.03	98.04	99.09
Tile	95.52 \pm 0.02	95.66 \pm 0.06	97.71 \pm 0.02	98.01 \pm 0.01	97.98 \pm 0.02	95.10 \pm 0.02	95.80 \pm 0.10	96.07	96.48
Toothbrush	98.02 \pm 0.03	97.98 \pm 0.00	98.56 \pm 0.02	98.93 \pm 0.00	98.89 \pm 0.00	98.44 \pm 0.02	99.00 \pm 0.01	98.09	98.80
Transistor	93.09 \pm 0.28	94.05 \pm 0.11	93.28 \pm 0.40	80.52 \pm 0.13	76.28 \pm 0.14	92.71 \pm 0.23	83.34 \pm 0.46	97.79	95.22
Wood	90.65 \pm 0.10	90.59 \pm 0.07	94.49 \pm 0.03	96.65 \pm 0.01	96.56 \pm 0.02	93.51 \pm 0.03	95.00 \pm 0.04	92.24	94.96
Zipper	96.80 \pm 0.02	97.01 \pm 0.05	98.41 \pm 0.09	99.08 \pm 0.02	99.06 \pm 0.01	97.71 \pm 0.06	98.98 \pm 0.01	97.50	99.07
Average	96.31	96.46	97.87	97.36	96.86	97.06	96.90	97.59	98.16

Table 3. The detailed comparison of PaDiM [8], SPADE [7], CutPaste [19] and our CFLOW-AD on the MVTec [4] dataset for every class using AUROC or, if available, a tuple (AUROC, AUPRO) metric, %. CFLOW-AD model is with the best hyperparameters from Section 5.2 ablation study. For fair comparison, we group together results with the same encoder architectures such as ResNet-18 and WideResNet-50.

Task	Localization		Detection				Localization		
Encoder	ResNet-18		EffNetB4		WideResNet-50				
Class/Model	CutPaste	Ours	CutPaste	Ours	CutPaste	Ours	SPADE	PaDiM	Ours
Bottle	97.6	98.64	98.3	100.00	100.0	100.0	(98.4, 95.5)	(98.3, 94.8)	(98.98, 96.80)
Cable	90.0	96.75	80.6	97.62	96.2	97.59	(97.2, 90.9)	(96.7, 88.8)	(97.64, 93.53)
Capsule	97.4	98.62	96.2	93.15	95.4	97.68	(99.0, 93.7)	(98.5, 93.5)	(98.98, 93.40)
Carpet	98.3	99.29	93.1	98.20	100.0	98.73	(97.5, 94.7)	(99.1, 96.2)	(99.25, 97.70)
Grid	97.5	98.53	99.9	98.97	99.1	99.60	(93.7, 86.7)	(97.3, 94.6)	(98.99, 96.08)
Hazelnut	97.3	98.81	97.3	99.91	99.9	99.98	(99.1, 95.4)	(98.2, 92.6)	(98.89, 96.68)
Leather	99.5	99.51	100.0	100.00	100.0	100.0	(97.6, 97.2)	(98.9, 88.8)	(99.66, 99.35)
Metal Nut	93.1	97.59	99.3	98.45	98.6	99.26	(98.1, 94.4)	(97.2, 85.6)	(98.56, 91.65)
Pill	95.7	98.34	92.4	93.02	93.3	96.82	(96.5, 94.6)	(95.7, 92.7)	(98.95, 95.39)
Screw	96.7	98.40	86.3	85.94	86.6	91.89	(98.9, 96.0)	(98.5, 94.4)	(98.86, 95.30)
Tile	90.5	95.80	93.4	98.40	99.8	99.88	(87.4, 75.9)	(94.1, 86.0)	(98.01, 94.34)
Toothbrush	98.1	99.00	98.3	99.86	90.7	99.65	(97.9, 93.5)	(98.8, 93.1)	(98.93, 95.06)
Transistor	93.0	97.69	95.5	93.04	97.5	95.21	(94.1, 87.4)	(97.5, 84.5)	(97.99, 81.40)
Wood	95.5	95.00	98.6	98.59	99.8	99.12	(88.5, 97.4)	(94.9, 91.1)	(96.65, 95.79)
Zipper	99.3	98.98	99.4	96.15	99.9	98.48	(96.5, 92.6)	(98.5, 95.9)	(99.08, 96.60)
Average	96.0	98.06	95.2	96.75	97.1	98.26	(96.0, 91.7)	(97.5, 92.1)	(98.62, 94.60)

Table 4. Average AUROC and AUPRO on the MVTEC [4] dataset, %. Both the best detection and localization metrics are presented, if available. CFLOW-AD is with WideResNet-50 encoder.

Metric	AUROC		AUPRO
Model	Detection	Localization	
DifferNet [30]	94.9	-	-
DFR [37]	-	95.0	91.0
SVDD [42]	92.1	95.7	-
SPADE [7]	85.5	96.0	91.7
CutPaste [19]	97.1	96.0	-
PaDiM [8]	97.9	97.5	92.1
CFLOW-AD (ours)	98.26	98.62	94.60

Table 5. Average AUROC on the STC [22] dataset, %. Both the best available detection and localization metrics are showed. CFLOW-AD is with WideResNet-50 encoder.

Metric	AUROC	
Model	Detection	Localization
CAVGA [40]	-	85.0
SPADE [7]	71.9	89.9
PaDiM [8]	-	91.2
CFLOW-AD (ours)	72.63	94.48

18 [13], WideResnet-50 [43], MobileNetV3L [14]), type of normalizing flow (unconditional (UFLOW) or conditional (CFLOW)), number of flow coupling layers (# of CL) and pooling layers (# of PL).

Our study shows that the increase in number of decoder’s coupling layers from 4 to 8 gives on average 0.15% gain due to a more accurate distribution modeling. Even higher 1.4% AUROC improvement is achieved when processing 3-scale feature maps (layers 1, 2 and 3) compared 2-scale only (layers 2, 3). The additional feature map (layer 1) with larger scale ($H_1 \times W_1 = 32 \times 32$) provides more precise spatial semantic information. The conditional normalizing flow (CFLOW) is on average 0.5% better than the unconditional (UFLOW) due to effective encoding of spatial prior. Finally, larger WideResnet-50 outperforms smaller ResNet-18 by 0.81%. MobileNetV3L, however, could be a good design choice for both fast inference and high AUROC.

Importantly, we find that the optimal input resolution is not consistent among MVTEC classes. The classes with macro objects *e.g.* cable or pill tend to benefit from the smaller-scale processing (256×256), which, effectively, translates to larger CNN receptive fields. Majority of classes perform better with 512×512 inputs *i.e.* smaller receptive fields. Finally, we discover that the transistor class has even higher AUROC with the resized to 128×128 images. Hence, we report results with the highest performing input resolution settings in the Section 5.3 comparisons.

5.3. Quantitative comparison

Table 4 summarizes average MVTEC results for the best published models. CFLOW-AD with WideResNet-50 encoder outperforms state-of-the-art by 0.36% AUROC in detection, by 1.12% AUROC and 2.5% AUPRO in localization, respectively. Table 3 contains per-class comparison for the subset of models grouped by the task and type of encoder architecture. CFLOW-AD is on par or significantly exceeds the best models in per-class comparison with the same encoder setups.

Table 5 presents high-level comparison of the best recently published models on the STC dataset. CFLOW-AD outperforms state-of-the-art SPADE [7] by 0.73% AUROC in anomaly detection and PaDiM [8] by 3.28% AUROC in anomaly localization tasks, respectively.

Note that our CFLOW-AD models in Tables 3-4 use variable input resolution as discussed in the ablation study: 512×512 , 256×256 or 128×128 depending on the MVTEC class. We used fixed 256×256 input resolution in Table 5 for the large STC dataset to decrease training time. Other reference hyperparameters in Tables 4-5 are set as: WideResnet-50 encoder with 3-scale pooling layers, conditional normalizing flow decoders with 8 coupling layers.

5.4. Qualitative results

Figure 3 visually shows examples from the MVTEC and the corresponding CFLOW-AD predictions. The top row shows ground truth masks from $\mathcal{D}_{\text{test}}$ including examples with and without anomalies. Then, our model produces anomaly score maps (middle row) using the architecture from Figure 2. Finally, we show the predicted segmentation masks with the threshold selected to maximize F1-score.

Figure 4 presents an additional evidence that our CFLOW-AD model actually addresses the OOD task sketched in Figure 1 toy example. We plot distribution of output anomaly scores for anomaly-free (green) and anomalous feature vectors (red). Then, CFLOW-AD is able to distinguish in-distribution and out-of-distribution feature vectors and separate them using a scalar threshold τ .

5.5. Complexity evaluations

In addition to analytical estimates in Table 1, we present the actual complexity evaluations for the trained models using inference speed and model size metrics. Particularly, Table 1 compares CFLOW-AD with the models from Tables 4-5 that have been studied by Defard *et al.* [8].

The model size in Table 6 is measured as the size of all floating-point parameters in the corresponding model *i.e.* its encoder and decoder (post-processing) models. Because the encoder architectures are identical, only the post-processing models are different. Since CFLOW-AD decoders do not explicitly depend on the feature map dimensions (only on feature vector depths), our model is significantly smaller

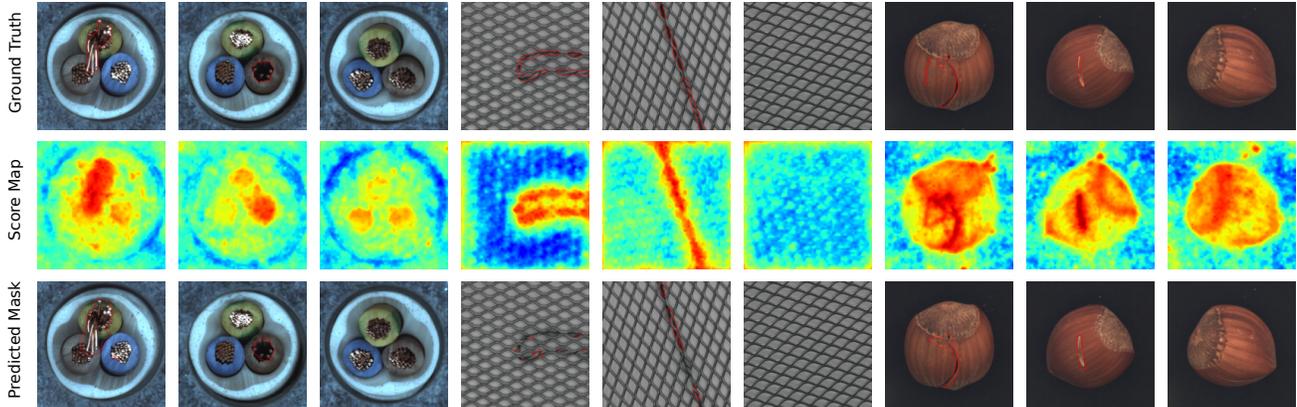


Figure 3. Examples of the input images with ground truth anomaly masks (top row) for various classes of the MVTEC. Our CFLOW-AD model from Table 4 estimates anomaly score maps (middle row) and generates segmentation masks (bottom row) for a threshold selected to maximize F1-score. The predicted segmentation mask should match the corresponding ground truth as close as possible.

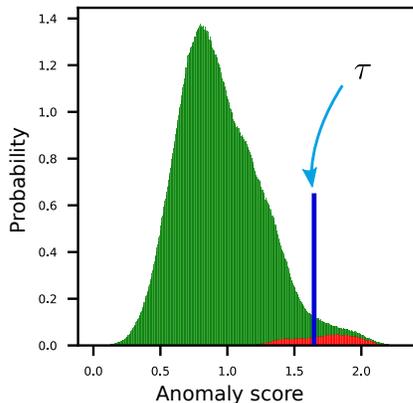


Figure 4. Distribution of anomaly scores for the cable class from the MVTEC learned by CFLOW-AD model from Table 4. Green density represent scores for the anomaly-free feature vectors, while region of red-color density shows scores for feature vectors with anomalies. The threshold τ is selected to optimize F1-score.

Table 6. Complexity comparison in terms of inference speed (fps) and model size (MB). Inference speed for CFLOW-AD models from Table 3 is measured for $(256 \times 256) / (512 \times 512)$ inputs.

Complexity metric and Model	Inference speed, fps	Model size, MB	
		STC	MVTEC
R18 encoder only	80 / 62	45	
PaDiM-R18 [8]	4.4	210	170
CFLOW-AD-R18	34 / 12	96	
WRN50 encoder only	62 / 30	268	
SPADE-WRN50 [7]	0.1	37,000	1,400
PaDiM-WRN50 [8]	1.1	5,200	3,800
CFLOW-AD-WRN50	27 / 9	947	
MNetV3 encoder only	82 / 61	12	
CFLOW-AD-MNetV3	35 / 12	25	

than SPADE and PaDiM. If we exclude the encoder parameters for fair comparison, CFLOW-AD is $1.7 \times$ to $50 \times$ smaller than SPADE and $2 \times$ to $7 \times$ smaller than PaDiM.

Inference speed in Table 6 is measured with INTEL I7 CPU for SPADE and PaDiM in Defard *et al.* [8] study with 256×256 inputs. We deduce that this suboptimal CPU choice was made due to large memory requirements for these models in Table 6. Thus, their GPU allocation for fast inference is infeasible. In contrast, our CFLOW-AD can be processed in real-time with $8 \times$ to $25 \times$ faster inference speed on 1080 8GB GPU with the same input resolution and feature extractor. In addition, MobileNetV3L encoder provides a good trade-off between accuracy, model size and inference speed for practical inspection systems.

6. Conclusions

We proposed to use conditional normalizing flow framework to estimate the exact data likelihoods which is infeasible in other generative models. Moreover, we analytically showed the relationship of this framework to previous distance-based models with multivariate Gaussian prior.

We introduced CFLOW-AD model that addresses the complexity limitations of existing unsupervised AD models by employing fully-convolutional translation-equivariant architecture. As a result, CFLOW-AD is faster and smaller by a factor of $10 \times$ than prior models with the same input resolution and feature extractor setup.

CFLOW-AD achieves new state-of-the-art for popular MVTEC with 98.26% AUROC in detection, 98.62% AUROC and 94.60% AUPRO in localization. Our new state-of-the-art for STC dataset is 72.63% and 94.48% AUROC in detection and localization, respectively. Our ablation study analyzed design choices for practical real-time processing including feature extractor choice, multi-scale pyramid pooling setup and the flow model hyperparameters.

References

- [1] Lynton Ardizzone, Jakob Kruse, Carsten Rother, and Ullrich Köthe. Analyzing inverse problems with invertible neural networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- [2] Lynton Ardizzone, Carsten Lüth, Jakob Kruse, Carsten Rother, and Ullrich Köthe. Guided image generation with conditional invertible neural networks. *arXiv:1907.02392*, 2019.
- [3] Christoph Baur, Benedikt Wiestler, Shadi Albarqouni, and Nassir Navab. Deep autoencoding models for unsupervised anomaly segmentation in brain MR images. *arXiv:1804.04488*, 2018.
- [4] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. MVTEC AD – a comprehensive real-world dataset for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [5] 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 (CVPR)*, 2020.
- [6] Paul Bergmann, Sindy Löwe, Michael Fauser, David Sattlegger, and Carsten Steger. Improving unsupervised defect segmentation by applying structural similarity to autoencoders. *Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, 2019.
- [7] Niv Cohen and Yedid Hoshen. Sub-image anomaly detection with deep pyramid correspondences. *arXiv:2005.02357v3*, 2021.
- [8] Thomas Defard, Aleksandr Setkov, Angélique Loesch, and Romaric Audigier. PaDiM: a patch distribution modeling framework for anomaly detection and localization. In *Proceedings of the International Conference on Pattern Recognition (ICPR) Workshops*, 2021.
- [9] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- [10] Markus Goldstein and Seiichi Uchida. A comparative evaluation of unsupervised anomaly detection algorithms for multivariate data. *PLOS ONE*, 2016.
- [11] Ian J. Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2015.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [14] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, Quoc V. Le, and Hartwig Adam. Searching for MobileNetV3. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [15] Polina Kirichenko, Pavel Izmailov, and Andrew Gordon Wilson. Why normalizing flows fail to detect out-of-distribution data. In *Advances in Neural Information Processing Systems*, 2020.
- [16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 2012.
- [17] Anders Krogh and John Hertz. A simple weight decay can improve generalization. In *Advances in Neural Information Processing Systems*, 1992.
- [18] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. In *Advances in Neural Information Processing Systems*, 2018.
- [19] 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 (CVPR)*, 2021.
- [20] Yingying Li, Jie Wu, Xue Bai, Xipeng Yang, Xiao Tan, Guanbin Li, Shilei Wen, Hongwu Zhang, and Errui Ding. Multi-granularity tracking with modularized components for unsupervised vehicles anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2020.
- [21] Wenjie Luo, Yujia Li, Raquel Urtasun, and Richard Zemel. Understanding the effective receptive field in deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, 2016.
- [22] Weixin Luo, Wen Liu, and Shenghua Gao. A revisit of sparse coding based anomaly detection in stacked RNN framework. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017.
- [23] Prasanta Chandra Mahalanobis. On the generalized distance in statistics. *Proceedings of the National Institute of Sciences (Calcutta)*, 1936.
- [24] Eric Nalisnick, Akihiro Matsukawa, Yee Whye Teh, Dilan Gorur, and Balaji Lakshminarayanan. Do deep generative models know what they don’t know? In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- [25] Paolo Napoletano, Flavio Piccoli, and Raimondo Schettini. Anomaly detection in nanofibrous materials by CNN-based self-similarity. *Sensors*, 2018.
- [26] Guansong Pang, Chunhua Shen, Longbing Cao, and Anton Van Den Hengel. Deep learning for anomaly detection: A review. *ACM Comput. Surv.*, 2021.
- [27] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 2021.
- [28] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban

- Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *Autodiff workshop at Advances in Neural Information Processing Systems*, 2017.
- [29] Oliver Rippel, Patrick Mertens, and Dorit Merhof. Modeling the distribution of normal data in pre-trained deep features for anomaly detection. In *Proceedings of the International Conference on Pattern Recognition (ICPR)*, 2021.
- [30] Marco Rudolph, Bastian Wandt, and Bodo Rosenhahn. Same same but DifferNet: Semi-supervised defect detection with normalizing flows. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2021.
- [31] Lukas Ruff, Jacob Reinhard Kauffmann, Robert Vandermeulen, Gregoire Montavon, Wojciech Samek, Marius Kloft, Thomas G. Dietterich, and Klaus-Robert Müller. A unifying review of deep and shallow anomaly detection. *Proc of the IEEE*, 2021.
- [32] Babak Saleh, Ali Farhadi, and Ahmed Elgammal. Object-centric anomaly detection by attribute-based reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [33] Mohammadreza Salehi, Niousha Sadjadi, Soroosh Baselizadeh, Mohammad Hossein Rohban, and Hamid R. Rabiee. Multiresolution knowledge distillation for anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [34] Robin Tibor Schirrmester, Yuxuan Zhou, Tonio Ball, and Dan Zhang. Understanding anomaly detection with deep invertible networks through hierarchies of distributions and features. In *Advances in Neural Information Processing Systems*, 2020.
- [35] 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*, 2019.
- [36] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *International Conference on Information Processing in Medical Imaging*, 2017.
- [37] Yong Shi, Jie Yang, and Zhiqian Qi. Unsupervised anomaly segmentation via deep feature reconstruction. *Neurocomputing*, 2021.
- [38] A. Vasilev, V. Golkov, M. Meissner, I. Lipp, E. Sgarlata, V. Tomassini, D.K. Jones, and D. Cremers. q-space novelty detection with variational autoencoders. *MICCAI 2019 International Workshop on Computational Diffusion MRI*, 2019.
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017.
- [40] Shashanka Venkataramanan, Kuan-Chuan Peng, Rajat Vikram Singh, and Abhijit Mahalanobis. Attention guided anomaly localization in images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [41] Guodong Wang, Shumin Han, Errui Ding, and Di Huang. Student-teacher feature pyramid matching for unsupervised anomaly detection. *arXiv:2103.04257*, 2021.
- [42] Jihun Yi and Sungroh Yoon. Patch SVDD: Patch-level SVDD for anomaly detection and segmentation. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, 2020.
- [43] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2016.
- [44] Kang Zhou, Yuting Xiao, Jianlong Yang, Jun Cheng, Wen Liu, Weixin Luo, Zaiwang Gu, Jiang Liu, and Shenghua Gao. Encoding structure-texture relation with P-Net for anomaly detection in retinal images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.