

TransFusion – A Transparency-Based Diffusion Model for Anomaly Detection

Matic Fučka, Vitjan Zavrtnik, and Danijel Skočaj

University of Ljubljana, Faculty of Computer and Information Science, Slovenia
{matic.fucka, vitjan.zavrtnik, danijel.skocaj}@fri.uni-lj.si

Abstract. Surface anomaly detection is a vital component in manufacturing inspection. Current discriminative methods follow a two-stage architecture composed of a reconstructive network followed by a discriminative network that relies on the reconstruction output. Currently used reconstructive networks often produce poor reconstructions that either still contain anomalies or lack details in anomaly-free regions. Discriminative methods are robust to some reconstructive network failures, suggesting that the discriminative network learns a strong normal appearance signal that the reconstructive networks miss. We reformulate the two-stage architecture into a single-stage iterative process that allows the exchange of information between the reconstruction and localization. We propose a novel transparency-based diffusion process where the transparency of anomalous regions is progressively increased, restoring their normal appearance accurately while maintaining the appearance of anomaly-free regions using localization cues of previous steps. We implement the proposed process as TRANSparency DiffUSION (TransFusion), a novel discriminative anomaly detection method that achieves state-of-the-art performance on both the VisA and the MVTEC AD datasets, with an image-level AUROC of 98.5% and 99.2%, respectively. Code: https://github.com/MaticFuc/ECCV_TransFusion

Keywords: Anomaly detection · Diffusion model · Industrial inspection

1 Introduction

The primary objective of surface anomaly detection is the identification and localization of anomalies in images. In the standard problem setup, only anomaly-free (normal) images are used to learn a normal appearance model and any deviations from the learned model are classified as anomalies. Surface anomaly detection is commonly used in various industrial domains [6, 7, 45] where the limited availability of abnormal images, along with their considerable diversity, makes training supervised models impractical.

Many of the recent surface anomaly detection methods follow the discriminative [21, 40, 41, 43, 44] paradigm. Discriminative methods are trained to localize simulated anomalies. Discriminative methods typically follow a two-stage architecture: a normal-appearance model followed by a discriminative network. The

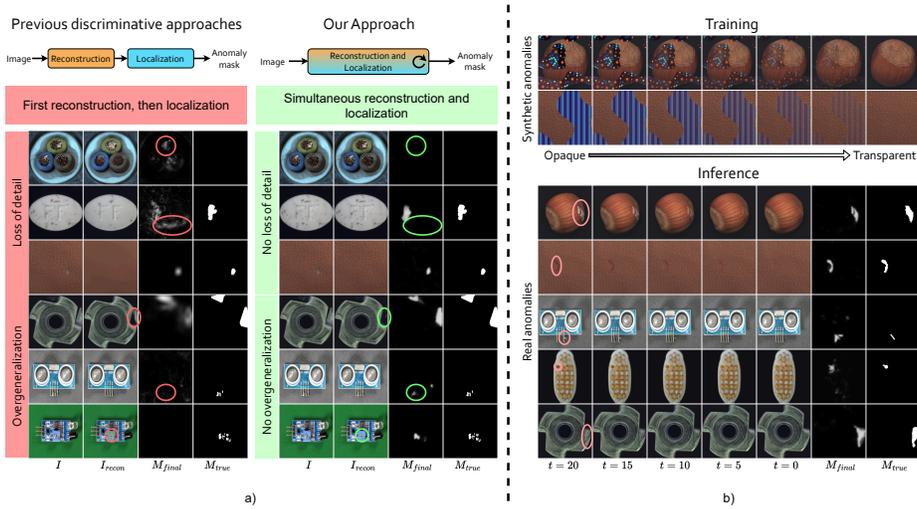


Fig. 1: a) Different than previous discriminative approaches, the proposed approach simultaneously reconstructs and localizes the anomalies through an iterative process, which results in a more potent normality model capable of detecting harder near-distribution anomalies. b) The reformulated diffusion model iteratively erases the anomalous regions during the reverse process. Training on synthetic anomalies (top) generalizes well to real anomalies (marked with red circles) seen at inference (bottom), leading to accurate output masks M_{final} that closely match the ground truth M_{true} .

normal-appearance model learns the anomaly-free object appearance and enables the detection of visual deviations. The discriminative network accurately localizes the anomalies and provides the per-pixel anomaly segmentation mask using the rich signal of the output of the normal-appearance model. Typically, the normal-appearance model is implemented as a reconstructive network. While the discriminative paradigm hailed the best performances in the past, it started to lag behind with the introduction of more challenging datasets [45].

The failures in reconstructive networks that hurt discriminative methods’ downstream anomaly detection capability can be characterised by two core issues. First, reconstructive methods may *overgeneralize*, which causes them to reconstruct even anomalous regions, leading to false negative detections. Second, due to the limited image generation capabilities of the commonly used reconstructive architectures, fine-grained details in normal regions tend to be erased, leading to *loss-of-detail* in normal regions, causing false positive detections. Some samples of these failures can be seen in Figure 1 a). In an attempt to address these two problems, the autoencoder-based reconstructive network of DRÆM [40] has been replaced with a diffusion model in previous works [43]. While the quality of reconstructions was somewhat improved, the loss-of-detail and overgeneralization problems remain in many cases, suggesting that simply

replacing the reconstructive subnetwork with a more powerful image generation model is insufficient.

In some cases, discriminative methods can successfully localize an anomaly despite the reconstruction network’s failure. This suggests that the discriminative network has the ability to learn the normal appearance signals that the reconstruction network misses. Similarly, discriminative methods can fail to localize an anomaly despite the reconstruction network’s success. Interaction between the reconstruction and localization normal-appearance signals might enable the extraction of additional information that is complementary to the information provided by each network and improves downstream anomaly detection performance. This interaction is not done with the current two-stage architecture most discriminative methods follow.

To address the problems of discriminative methods, we propose a novel *transparency-based diffusion process* reformulated explicitly for surface anomaly detection. Through the proposed diffusion process, the transparency of anomalies is iteratively increased so that they are gradually replaced with the corresponding normal appearance (Figure 1 b), effectively erasing the anomalies. Throughout the proposed process, the anomalies are simultaneously localized and restored to their anomaly-free appearance. This enables a precise anomaly-free reconstruction of the anomalous regions – addressing *overgeneralization*. Additionally, localization information is used to keep the anomaly-free regions intact – addressing the *loss-of-detail* problem (Figure 1 a). To implement the transparency-based diffusion process, we propose *TransFusion (TRANSparency DifFUSION)*, a surface anomaly detection method that integrates the powerful appearance modelling capabilities of diffusion models in the discriminative anomaly detection paradigm. Compared to the previously used reconstructive networks inside discriminative methods that attempted to implicitly detect and restore the anomaly-free appearance of anomalous regions in a single step [40, 41, 44], TransFusion can maintain more accurate restorations of anomalous regions without the overgeneralization problem and without loss-of-detail in the anomaly-free regions. Due to the iterative nature of the reformulated diffusion process, TransFusion is able to focus on various visual characteristics of anomalies at various time-steps, even potentially addressing the regions previous iterations may have missed. Additionally, the localization information of previous steps can be used as a cue in the reconstruction process, highlighting the potentially anomalous regions. This enables high-fidelity anomaly-free reconstructions and improves the downstream anomaly detection performance significantly, compared to previous discriminative approaches.

The main contributions of our work are as follows:

- We propose a novel transparency-based diffusion process reformulated explicitly for the problem of surface anomaly detection. The proposed diffusion process iteratively increases the transparency of anomalies and simultaneously provides their explicit localization.
- We propose TransFusion - A strong discriminative anomaly detection model that implements the transparency-based diffusion process. TransFusion di-

rectly addresses the overgeneralization and loss-of-detail problems of recent discriminative anomaly detection methods, leading to a strong anomaly detection performance even in difficult near-in-distribution scenarios.

We perform extensive experiments on two challenging datasets and show that TransFusion achieves state-of-the-art results in anomaly detection on two standard challenging datasets – VisA [45] and MVTEC AD [6], with an AUROC of 98.5% and 99.2%, respectively. TransFusion sets a new state-of-the-art in anomaly detection in terms of the mean across both datasets, achieving a 98.9% AUROC.

2 Related Work

Surface anomaly detection has been a subject of intense research in recent years, and various approaches have been proposed to address this task. Methods can be divided into three main paradigms: reconstructive, embedding-based, and discriminative.

Reconstructive methods train an autoencoder-like network [9, 28, 42] or a generative model [1, 22, 30, 37] and assume that anomalies will be poorly reconstructed compared to the normal regions making them distinguishable by reconstruction error. The poor reconstruction assumption does not always hold, leading to poor performance.

Embedding-based methods use feature maps [16, 21] extracted with a pre-trained network to learn normality on these maps. Patchcore [25] creates a core-set memory bank out of the extracted normal features. Several normalizing-flow-based [13, 26, 33, 39] approaches have been proposed as well. Some methods utilize a student-teacher [8, 11, 27] network and assume that the student will not be able to produce meaningful features for the anomalies as it had not seen them during training. All these methods assume that the distribution of normal regions will be well represented in the training data and fail on rare normal regions unseen during training, producing false positives.

Discriminative methods use synthetically generated defects [18, 21, 38, 40, 41, 44] to train their model with the idea that the model can then generalize on real anomalies. In seminal works of this paradigm, such as DRÆM [40], a two-stage architecture was proposed. First, a reconstructive module is trained to restore the normal appearance, and then, a discriminative network is trained to segment synthetic anomalies. This idea has been followed by the vast majority of models [18, 38, 41, 43, 44] inside this paradigm. The normal appearance can also be modelled using pretrained features [21, 38, 44]. DiffAD [43] has tried to improve DRAEM [40] by exchanging the reconstructive subnetwork with a more powerful appearance modelling model, a diffusion model, but the overgeneralization and loss-of-detail problems remained. This suggests that the standard two-stage approach is not optimal for harder near-distribution anomalies.

Diffusion models recently emerged as state-of-the-art in image generation [14]. They have been extended to various domains, such as audio [15, 17] and text generation [3, 19]. Methods have also been proposed that tackle problems such as

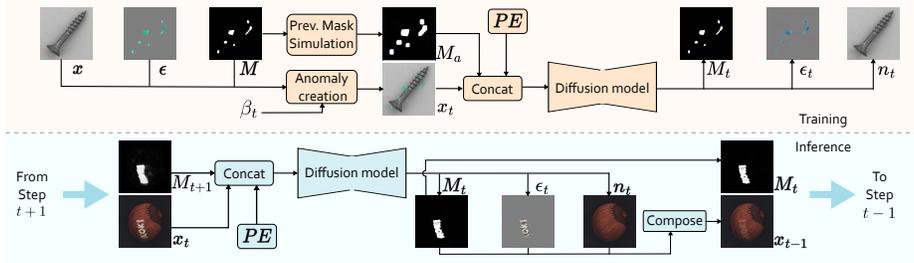


Fig. 2: TransFusion’s **training** and **inference** pipelines. **Training** examples are created from normal images x by generating the anomaly mask M and the anomaly appearance ϵ and imposing them on x according to the transparency schedule β_t . The resulting image x_t contains synthetic anomalies. TransFusion is guided by an augmented mask M_a . TransFusion outputs the estimated anomaly mask M_t , the anomaly appearance ϵ_t , and the normal appearance n_t . At **inference**, TransFusion infers M_t , ϵ_t , and n_t from the input image and constructs the next step image according to Eq. 4. The predicted mask M_t and the constructed x_{t-1} are used as the input in the next step.

semantic segmentation [2, 36] and object detection [10]. It has also been shown that the Gaussian noise-based diffusion process is not necessary for all problems [4, 10].

Diffusion-based anomaly detection Wyatt *et. al.* [37] proposed AnoDDPM which is based on a standard diffusion architecture [14]. AnoDDPM was applied to a medical image dataset and achieved state-of-the-art results. Lu *et. al.* [22] proposed using a DDPM to simultaneously predict the noise and to generate features that mimic the features extracted from a pretrained convolutional neural network. DiffAD [43] exchanged the autoencoder from DRÆM [40] with a latent diffusion model to limited success. All recent diffusion approaches face problems with loss-of-detail in the normal regions. As a result, they exhibit a high rate of false positives. This suggests that naively applying the standard diffusion process is insufficient for surface anomaly detection.

3 TransFusion

Discriminative anomaly detection approaches attempt to reconstruct the normal visual appearance of anomalies and localize them based on the output of the reconstruction module. An appropriate diffusion model is defined to reformulate this two-stage approach as an iterative one-stage process in order to achieve better detection robustness and reconstruction capability. Previous work [4] has established that a variety of iterative processes can be used to achieve the desired diffusion effect. In the proposed transparency-based diffusion process reformulation, images are thought of as a composition of anomalous and normal components, partitioned by the anomaly mask M . The anomalous regions are expressed as a linear interpolation between the anomalous and the normal appearance at each step to frame the anomaly localisation and restoration as an

iterative process. This equates to the transparency of the anomalous regions increasing throughout the diffusion process (Figure 1 b). In this section, we describe TransFusion in detail.

3.1 Transparency-based diffusion model

In the transparency-based diffusion process reformulation, each image I is expressed as a composition of the normal appearance N , the anomaly appearance A , the anomaly mask M , and the blending factor between the anomalous and the normal appearance β , i.e., the transparency level of the anomaly:

$$I = \overline{M} \odot N + \beta(M \odot A) + (1 - \beta)(M \odot N), \quad (1)$$

where M is a binary mask where the anomalous pixels are set to 1 and \overline{M} is the inverse of M . The anomalous region is an interpolation between the anomaly appearance A and the normal appearance N in the region specified by the anomaly mask M . The transparency of the anomalous region is defined by β . The restoration of the normal appearance from an anomalous image I can be modelled as an iterative process of gradually increasing the anomaly transparency until only the normal appearance remains. This is not a trivial task, since the accurate localization M , normal appearance N , and anomaly appearance A must be inferred from the input image I .

During training, images containing synthetic anomalies and their corresponding anomaly masks are used. For each step in the *forward process*, the value of β is gradually increased, thus decreasing the transparency of anomalies, and increasing their prominence. Let x_t denote the anomalous image I at time step t . The transparency schedule is denoted as $\beta_0 < \beta_1 < \dots < \beta_{T-1} < \beta_T$, where $\beta_0 = 0$ and $\beta_T = 1$. Eq. (1) is rewritten to correspond to timestep t by substituting the variables A with ϵ_t , M with M_t , and N with n_t :

$$x_t = \overline{M}_t \odot n_t + \beta_t(M_t \odot \epsilon_t) + (1 - \beta_t)(M_t \odot n_t). \quad (2)$$

The image with more transparent anomalies x_{t-1} at iteration $t - 1$ is then computed:

$$x_{t-1} = \overline{M}_{t-1} \odot n_{t-1} + \beta_{t-1}(M_{t-1} \odot \epsilon_{t-1}) + (1 - \beta_{t-1})(M_{t-1} \odot n_{t-1}). \quad (3)$$

β_t decreases between steps t and $t-1$, while the correct values of M_t , n_t and ϵ_t are predefined and remain constant throughout the forward process. We can thus write $M_t = M_{t-1} = \dots = M$, $\epsilon_t = \epsilon_{t-1} = \dots = A$ and $n_t = n_{t-1} = \dots = N$. After substituting M_{t-1} for M_t , ϵ_{t-1} for ϵ_t and n_{t-1} for n_t in Eq. (3), subtracting it from Eq. (2) and then rearranging it, the transition between steps x_t and x_{t-1} is computed:

$$x_{t-1} = x_t - (\beta_t - \beta_{t-1})(M_t \odot \epsilon_t) + (\beta_t - \beta_{t-1})(M_t \odot n_t). \quad (4)$$

At each time step in the *reverse process*, the value of x_t moves towards the anomaly-free x_0 by an amount influenced by $\beta_t - \beta_{t-1}$. The anomaly's transparency is therefore gradually increased, reconstructing the normal appearance

until the final anomaly-free restoration x_0 is reached. This requires an accurate estimation of the anomaly mask M_t , the normal appearance n_t and the anomaly appearance ϵ_t at each time step.

3.2 Architecture

The architecture of TransFusion, depicted in Figure 2, is based on ResUNet [12], which is commonly used in diffusion models. TransFusion has three prediction heads, which output the anomaly appearance ϵ_t , anomaly mask M_t , and the normal appearance n_t , enabling the generation of the image in the next reverse step according to Eq. (4). The anomaly and normal appearance heads consist of a single convolutional layer, while the anomaly mask head consists of a BatchNorm, SiLU and a convolutional layer.

The input to the diffusion model at each timestep consists of four elements: the current reconstruction estimate x_t , the mask estimate M_t , the 2D sinusoidal positional encoding PE [34], and the timestep t . All the elements are channel-wise concatenated except for the timestep embedding, which is added to the features. PE helps the model to learn the global composition of some objects. During training, images containing synthetic anomalies are generated from an anomaly-free image x , the anomaly mask M , and the anomaly appearance ϵ . The input image x_t is generated according to Eq. 2, where $n_t = x$, $\epsilon_t = \epsilon$ and $M_t = M$, and the β schedule for the sampled timestep t . Losses for the prediction head outputs n_t , M_t and ϵ_t are calculated using x , M and ϵ as ground truth values, respectively.

Separate loss functions are used for each prediction head. The *normal appearance prediction head* uses the structural similarity (SSIM) loss [35] and the \mathcal{L}_1 loss:

$$\mathcal{L}_n = SSIM(n_t, x) + \mathcal{L}_1(n_t, x). \quad (5)$$

The *anomaly mask head* uses the focal loss [20] and the Smooth \mathcal{L}_1 loss, commonly used in discriminative anomaly detection [38, 40]:

$$\mathcal{L}_m = \alpha \mathcal{L}_{foc}(M_t, M) + \mathcal{L}_{1Smooth}(M_t, M). \quad (6)$$

The weighting parameter α is set to 5 in all experiments. The *anomaly appearance prediction head* employs the standard \mathcal{L}_2 reconstruction loss:

$$\mathcal{L}_a = \mathcal{L}_2(\epsilon_t, \epsilon). \quad (7)$$

To ensure the consistency between diffusion steps, where x_{t-1} is computed from the estimated M_t , ϵ_t , n_t and the previous step x_t using Eq. (4), an additional *consistency loss* function \mathcal{L}_c is employed. \mathcal{L}_c compares the predicted x_{t-1} with the ground truth \tilde{x}_{t-1} computed using the ground truth M , ϵ , and x :

$$\mathcal{L}_c = \mathcal{L}_2(x_{t-1}, \tilde{x}_{t-1}). \quad (8)$$

The complete TransFusion loss is then given as:

$$\mathcal{L} = \mathcal{L}_n + \mathcal{L}_m + \mathcal{L}_a + \mathcal{L}_c. \quad (9)$$

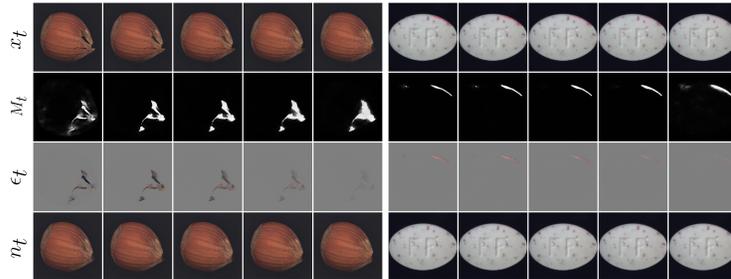


Fig. 3: TransFusion inference. For every fourth timestep, the input image x_t and the predictions for the mask M_t , anomaly appearance ϵ_t and normal appearance n_t are shown. As seen in the top row, TransFusion first reconstructs larger anomalies and inpaints the details near the end of the reconstruction process.

3.3 Synthetic anomaly generation

We directly adopt the synthetic anomaly generation from MemSeg [38], which is an extension to the synthetic anomaly generation proposed by DRÆM [40]. Synthetic anomalies are generated by pasting out-of-distribution regions on the anomaly-free inputs, outputting the image containing synthetic anomalies I and the anomaly mask M . M is generated using *Perlin noise* [24]. Synthetic anomalous examples are shown in the top part of Figure 1 b). Depending on the timestep used, anomalies are generated at different transparency levels.

It would be unrealistic to expect that the current mask estimate M_t would be perfect during inference. To mimic this observation during training, the previous mask estimate imperfection is simulated. The *simulated previous mask estimate* is obtained by thresholding the Perlin noise map used for generating M , resulting in a reduction or an expansion of the size of M_t . M_t is also dropped during training in 25% of training samples.

3.4 Inference

At *inference*, Figure 2, the starting mask estimate is initialized to all zero values. Then, the reverse process of T time steps is performed. T is set to 20 in all experiments unless stated otherwise. At each time step t , the current approximation of the reconstructed image x_t is channel-wise concatenated with the binarized previous mask estimate M_{t+1} and positional encoding PE . This composite input and the current timestep t are fed into the diffusion model. The model’s output consists of the current mask estimate M_t , an anomaly appearance estimation ϵ_t , and a normal appearance estimation n_t (Figure 2, bottom middle). Based on these outputs, the next step x_{t-1} is predicted using Eq. (4) (Figure 2, bottom right). Anomaly mask M_t is binarized by thresholding and used in the next step. An example of the inference process is visualized in Figure 3. The reverse process iteratively reduces the transparency of the anomalous regions, progressively

restoring the anomaly-free appearance of the image. At time step 0, the result is a fully reconstructed anomaly-free image x_0 .

The *final anomaly mask* M_{final} is derived from M_{disc} , the pixel-wise mean of anomaly masks M_t , with t going from 1 to T , produced throughout the reverse process and from M_{recon} , the reconstruction error between the initial image x and the diffusion model output x_0 .

To obtain the final mask M_{final} , a weighted combination of M_{disc} and M_{recon} is performed:

$$M_{final} = (\lambda M_{disc} + (1 - \lambda) M_{recon}) * f_n, \quad (10)$$

where the influence of M_{disc} and M_{recon} is weighted by λ ($\lambda=0.95$ in all experiments), f_n is a mean filter of size $n \times n$ (in our case 7×7) and $*$ is the convolution operator. The mean filter smoothing is performed to aggregate the local anomaly map responses for a robust image-level score estimation. The image-level anomaly score AS is obtained by the maximum value of M_{final} :

$$AS = \max(M_{final}). \quad (11)$$

Including both M_{disc} and M_{recon} gives the final mask M_{final} a balanced anomaly representation, allowing it to benefit from both discriminative and reconstructive cues.

4 Experiments

4.1 Datasets

Experiments are performed on two standard anomaly detection datasets: the VisA dataset [45] and the MVTec AD dataset [6]. The VisA dataset is comprised of 10,821 images distributed across 12 object categories, while the MVTec AD dataset contains 5,354 images encompassing 5 texture categories and 10 object categories. Notably, both datasets provide pixel-level annotations for the test images, enabling accurate evaluation and analysis.

4.2 Evaluation metrics

Standard anomaly detection evaluation metrics are used. The image-level anomaly detection performance is evaluated by the Area Under the Receiver Operator Curve (AUROC), while for the pixel-level anomaly localization the Area Under the Per Region Overlap (AUPRO) is utilized.

4.3 Implementation details

During both training and inference, 20 steps ($T = 20$) are used in the diffusion process with a linear transparency (β) schedule ranging from 0 to 1. The model was trained for 1500 epochs using the AdamW optimizer with a batch size of 8. The learning rate was set to 10^{-5} and was multiplied by 0.1 after 800 epochs.

Synthetic anomalies were added to half of the training batch. Rotation augmentation was used following DRÆM [40]. A standard preprocessing approach is employed to ensure experimental consistency. Each image is resized to dimensions of 256×256 and subsequently center-cropped to 224×224 following recent literature [21, 25, 38]. The image is then linearly scaled between -1 and 1 following recent diffusion model literature [14, 31]. Following the standard protocol for unsupervised anomaly detection, a separate model was trained for each category, and the same hyperparameters were set across both datasets and all categories.

4.4 Experimental results

Anomaly detection results on VisA are shown in Table 1. TransFusion achieves the best results on 5 out of the 12 categories and outperforms the previous best state-of-the-art method by 0.4 percentage points in terms of the mean AUROC performance. On the MVTec AD dataset, TransFusion achieves state-of-the-art results with a mean anomaly detection AUROC of 99.2%. Results are shown in Table 2.

Due to the significant differences in anomaly types between the VisA and MVTec AD datasets, very few recent methods exhibit the generalization capability necessary to achieve top results for both datasets. Table 3 shows results on both VisA and MVTec AD. Additionally, the average scores across both datasets are shown. TransFusion outperforms all recent methods in terms of the average anomaly detection AUROC by 0.3 percentage points and, more notably, outperforms the next best discriminative method by a significant margin of 4.0 percentage points, reducing the error by 78.5%.

TransFusion also achieves the second highest score in *anomaly localization* when averaged across both datasets, achieving an AUPRO of 91.6%. In terms of anomaly detection, TransFusion outperforms competing methods significantly on the VisA dataset and achieves state-of-the-art performance on MVTec AD. TransFusion also outperforms other diffusion-based methods, AnoDDPM [37], DiffAD [43] and AnomDiff [22], by a significant margin, which suggests that simply relying on a standard diffusion process for reconstruction may not be sufficient for anomaly detection. TransFusion also significantly outperforms previous state-of-the-art discriminative methods, such as DRÆM [40], DSR [41], DiffAD [43] and SimpleNet [21], on the VisA dataset in terms of anomaly detection. This suggests that simultaneous localization and reconstruction provide a more potent normality model in comparison to the previous two-stage paradigm.

4.5 Qualitative comparisons

A qualitative comparison with the state-of-the-art methods DRÆM [40] and EfficientAD [5] can be seen in Figure 4. Note that TransFusion outputs very precise anomaly masks and does not produce significant false positives in the background as opposed to other state-of-the-art methods (Columns 3, 4, 14). Due to being a discriminative network, TransFusion outputs masks (Columns 1-14) that are much sharper than those of EfficientAD, which outputs a feature reconstruction

Method	AnoDDPM [37]	AnomDiff [22]	PatchCore [25]	RD4AD [11]	AST [27]	EfficientAD [5]	DiffAD [43]	DRÆM [40]	DSR [41]	SimpleNet [21]	TransFusion -
Candle	64.9	64.9	98.1	92.2	99.4	98.4	90.4	94.4	98.8	95.6	98.3
Capsules	76.5	80.0	85.7	90.1	85.4	93.5	87.6	76.3	99.1	76.7	99.6
Cashew	94.4	90.9	98.5	99.6	95.1	97.2	81.4	90.7	97.6	91.7	93.7
Chewing gum	91.3	98.1	99.0	99.7	100	99.9	94.0	94.2	93.8	99.1	99.6
Fryum	81.5	89.2	97.2	96.6	99.1	96.5	87.1	97.4	82.9	95.3	98.3
Macaroni1	58.8	77.8	95.7	98.4	93.9	99.4	87.6	95.0	87.3	90.8	99.4
Macaroni2	74.5	61.0	78.1	97.6	72.1	96.7	90.7	96.2	83.4	65.2	96.5
PCB1	42.1	86.7	98.3	97.6	99.2	98.5	75.0	54.8	90.5	60.1	98.9
PCB2	90.7	76.5	97.2	91.1	98.4	99.5	94.6	77.8	96.6	93.3	99.7
PCB3	92.3	80.4	96.2	95.5	97.4	98.9	94.7	94.5	94.8	94.9	99.2
PCB4	98.3	93.8	99.0	96.5	99.6	98.9	97.7	93.4	93.5	98.2	99.6
Pipe fryum	72.5	89.4	99.4	97.0	99.4	99.6	92.7	99.4	97.5	93.3	99.6
Average	78.2	83.7	94.3	96.0	94.9	98.1	89.5	88.7	91.6	87.9	98.5

Table 1: Comparison of TransFusion in anomaly detection (AUROC) with SOTA on VisA. **First**, **second** and **third** place are marked. The names of all previous discriminative approaches are typeset in **bold**.

Method	AnoDDPM [37]	AnomDiff [22]	PatchCore [25]	RD4AD [11]	AST [27]	EfficientAD [5]	DiffAD [43]	DRÆM [40]	DSR [41]	SimpleNet [21]	TransFusion -
Carpet	93.5	99.9	98.7	95.3	99.1	99.3	98.3	97.0	100	97.5	99.2
Grid	93.8	99.7	98.2	100	98.7	99.9	100	99.9	100	99.1	100
Leather	99.5	100	100	97.1	100	100	100	100	100	100	100
Tile	99.4	98.0	98.7	99.3	99.1	99.9	100	99.6	100	100	99.8
Wood	99.0	98.1	99.2	99.2	99.2	100	100	99.1	96.3	100	99.4
Bottle	98.4	99.3	100	100	100	99.9	100	99.2	100	100	100
Cable	52.7	91.2	99.5	95.0	98.5	95.2	94.6	91.8	93.8	99.9	97.9
Capsule	89.0	84.1	98.1	96.3	99.7	97.9	97.5	98.5	98.1	97.7	98.5
Hazelnut	84.5	97.9	99.9	100	100	99.4	100	100	95.6	100	100
Metal nut	92.8	99.2	100	100	98.5	99.6	99.5	98.7	98.5	100	100
Pill	80.9	64.7	96.6	96.6	99.1	98.6	97.7	98.9	97.5	99.0	98.3
Screw	20.3	89.9	98.1	97.0	99.7	96.9	97.2	93.9	96.2	98.2	97.2
Toothbrush	86.4	96.9	100	90.8	96.6	100	100	100	99.7	99.7	100
Transistor	65.0	92.3	100	96.7	99.3	99.9	96.1	93.1	97.8	100	98.3
Zipper	98.2	85.5	99.4	98.5	99.1	99.7	100	100	100	99.5	100
Average	83.5	93.1	99.1	98.5	99.2	99.1	98.7	98.0	98.2	99.6	99.2

Table 2: Comparison of TransFusion in anomaly detection (AUROC) with SOTA on MVTec AD.

Method	Venue	Disc.	VisA		MVTec AD		Average	
			Det.	Loc.	Det.	Loc.	Det.	Loc.
AnoDDPM	CVPRW'22		78.2	60.5	83.5	50.7	80.9	55.6
DRÆM	ICCV'21	✓	88.7	73.1	98.0	92.8	93.3	83.0
SimpleNet	CVPR'23	✓	87.9	68.9	99.6	89.6	93.8	79.3
DiffAD	ICCV'23	✓	89.5	71.2	98.7	84.8	94.1	78.0
DSR	ECCV'22	✓	91.6	68.1	98.2	90.8	94.9	79.5
Patchcore	CVPR'22		94.3	79.7	99.1	92.7	97.0	86.2
AST	WACV'23		94.9	81.5	99.2	81.2	97.1	81.4
RD4AD	CVPR'22		96.0	70.9	98.5	93.9	97.3	82.4
EfficientAD	WACV'24		98.1	94.0	99.1	93.5	98.6	93.7
TransFusion	ECCV'24	✓	98.5	88.8	99.2	94.3	98.9	91.6

Table 3: Results in anomaly detection (AUROC) and anomaly localization (AUPRO) on both VisA and MVTec AD.

error. DRÆM is unable to accurately detect small near-in-distribution anomalies (Columns 3, 5, 6, 14) mostly present in the VisA [45] Dataset. We hypothesize that this is due to the problems encountered in two-stage discriminative approaches, such as DRÆM.

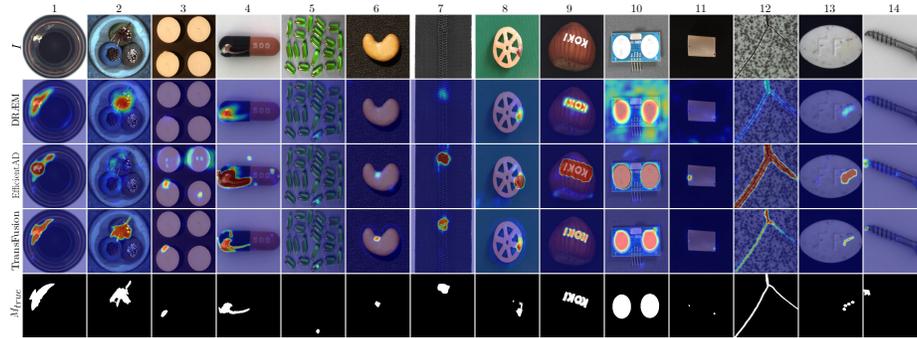


Fig. 4: Qualitative comparison of the masks produced by TransFusion and three other state-of-the-art methods. The anomalous images are shown in the first row. The middle four rows show the anomaly mask generated by DRÆM [40], EfficientAD [5] and TransFusion, respectively. The last row shows the ground truth anomaly mask.

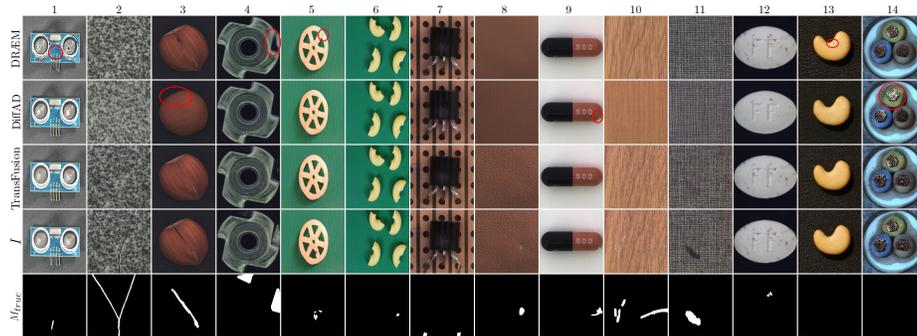


Fig. 5: Qualitative reconstruction results. TransFusion better restores anomalies to their normal appearance and better preserves the details in the normal regions than competing methods DRÆM [40] and DiffAD [43]. A few of the larger differences are highlighted in red.

TranFusion exhibits a strong reconstructive ability. A qualitative comparison can be seen in Figure 5. Compared to DRÆM [40], TransFusion outputs higher-quality reconstructions and even produces realistic results in difficult reconstruction cases, such as strong deformations, while maintaining fine-grained details in normal regions. TransFusion better addresses the loss-of-detail problem than the previously proposed method DiffAD [43]. The reconstructions suggest that simultaneous reconstruction and localization produce more powerful normal-appearance models.

4.6 Ablation study

The results of the evaluation of individual components of TransFusion and its training process are shown in Table 4.

Group	Condition	VisA		MVTec AD	
		Det.	Loc.	Det.	Loc.
Input strategies	w/o PE	-1.4	-0.1	-1.9	-3.2
	w/o Simulated Mask	-1.2	-8.6	-2.4	-6.2
Loss Function	w/o \mathcal{L}_n	-31.8	-45.3	-26.0	-43.2
	w/o \mathcal{L}_m	-1.5	-3.2	-2.2	-2.7
	w/o \mathcal{L}_a	-1.1	-1.5	-1.0	-1.0
	w/o \mathcal{L}_c	-0.9	-0.3	-0.7	-0.8
Final mask calc.	Only M_1	-1.5	-0.3	-1.4	-2.6
	Only M_{disc}	-0.1	+0.1	-0.2	+0.1
	Only M_{recon}	-0.9	-13.5	-0.7	-5.9
Diffusion step num.	5 steps	-1.0	-4.3	-0.7	-1.2
	10 steps	-0.5	-1.1	-0.7	-1.0
	50 steps	-0.3	+0.7	-0.5	-0.8
Transparency sched.	Quadratic	-2.0	+1.2	-2.0	+0.4
	Root	-1.7	-2.7	-0.7	-1.6
TransFusion	Linear, 20 steps	98.5	88.8	99.2	94.3

Table 4: Ablation study results. Detection results are reported in AUROC and localization results are reported in AUPRO. In each row, the difference to the actual model is shown. The highest discrepancy for each experiment group is marked in blue.

Input strategies. In addition to the image x_t , the Positional Encoding (PE) and the simulated previous mask are input during training. The impact of PE and the simulated previous mask is evaluated by excluding each individually from the architecture. *Excluding PE* leads to a 1.4 percentage points (p. p.) drop on VisA and a 1.9 p. p. drop on MVTec AD. *Excluding the simulated previous mask* leads to a 1.2 p. p. drop on VisA and a 2.4 p. p. drop on MVTec AD, showing the benefit of the localization information gained from the previous step. There is also a significant drop (8.6 p. p. on VisA and 6.2 p. p. on MVTec AD) in localization when excluding the simulated mask, highlighting its importance for precise localization.

Importance of loss functions. The importance of each loss function was evaluated by excluding one loss function at a time and training the model. Removing L_a , L_c or L_m reduces the overall anomaly detection performance by approximately 1 p. p. on VisA and MVTec AD, demonstrating their usefulness. Notably, removing \mathcal{L}_n leads to a major drop in performance (31.8 p. p. AUROC on VisA, 26 p. p. AUROC on MVTec AD), showing the necessity of learning a strong normal appearance model of the object. Without \mathcal{L}_n , TransFusion may focus on learning the synthetic anomaly appearance, leading to poor generalization.

Final mask calculation. The anomaly mask calculation methods using either only the last mask estimate M_1 , the discriminative mask M_{disc} , or the reconstruction mask M_{recon} are evaluated. Using only M_{recon} leads to a 0.9 p. p. drop on VisA and a 0.7 p. p. MVTec AD in terms of AUROC. M_{disc} can accurately localize the anomalies even without M_{recon} , leading to only a 0.1 and 0.2 p. p. drop on VisA and MVTec AD, respectively. The impact of mask averaging throughout the diffusion process is significant since using only the last estimated mask (Last Mask Est.) causes a 1.5 and 1.4 p. p. drop in anomaly detection performance on the VisA and MVTec AD, respectively.

Number of diffusion steps. The impact of the number of diffusion steps on the anomaly detection performance is evaluated. Although a lower number of steps leads to a poorer normal appearance restoration, TransFusion remains robust

Method	DR.EM [40]	Patchcore [25]	DiffAD [43]	<i>TransFusion</i>
Inference [s]	0.05	0.22	1.00	0.34

Table 5: Results for average inference time of a single sample with NVIDIA A100 GPU. Inference times are reported in seconds.

across various time-step settings, achieving similar results across both VisA and MVTec AD, even achieving state-of-the-art results on VisA at only 5 timesteps. A higher number of diffusion steps also increases the result in localization on VisA.

Transparency schedule. The impact of replacing the linear β schedule with alternative schedules is evaluated. The Root and the Quadratic schedule are examined, where the β values change from 0 to 1 using a quadratic or a square-root function, respectively. Using a Quadratic schedule causes a 2 p. p. drop in performance on both VisA and MVTec AD. The Root schedule leads to a 1.7 and a 0.7 p. p. drop on the VisA and the MVTec AD, respectively. Interestingly, using a quadratic schedule improves anomaly localization by a 1.2 p. p. on VisA and a 0.4 p. p. on MVTec AD.

Inference efficiency. Inference times of various methods can be seen in Table 5. Due to the complexity of diffusion models, TransFusion is slower than some competing methods however, it is faster than other diffusion-based methods. Additionally, reducing the number of inference steps does not drastically reduce performance (Table 4). Speeding up diffusion models is an active field [23, 29, 32] and may be helpful to increase the inference speed of TransFusion in the future.

5 Conclusion

A novel, transparency-based diffusion process is proposed, where the transparency of the anomalous regions is gradually increased, effectively removing them and restoring their normal appearance. TransFusion, a novel discriminative anomaly detection method that implements the transparency-based diffusion process, is proposed. With simultaneous localization and reconstruction, TransFusion is able to produce accurate anomaly-free reconstructions of anomalies while maintaining the appearance of normal regions, thus addressing both the overgeneralization and loss-of-detail problems of commonly used reconstructive models inside discriminative approaches. TransFusion achieves state-of-the-art results in anomaly detection on the standard VisA and MVTec AD datasets, achieving an AUROC of 98.5% and 99.2% for both datasets, respectively. The versatility of TransFusion and its robustness to near-in-distribution anomalies are further validated by the state-of-the-art performance across both datasets, where TransFusion achieves 98.9% mean AUROC, surpassing the previous state-of-the-art method by a significant margin of 0.3 percentage points. The results indicate that custom diffusion processes crafted specifically for surface anomaly detection are a promising direction for future research.

Acknowledgements This work was in part supported by the ARIS research project L2-3169 (MV4.0), research programme P2-0214 and the supercomputing network SLING (ARNES, EuroHPC Vega).

References

1. Akcay, S., Atapour-Abarghouei, A., Breckon, T.P.: GANomaly: Semi-supervised Anomaly Detection via Adversarial Training. In: Asian conference on computer vision. pp. 622–637. Springer (2018). https://doi.org/https://doi.org/10.1007/978-3-030-20893-6_39
2. Amit, T., Nachmani, E., Shaharbandy, T., Wolf, L.: SegDiff: Image Segmentation with Diffusion Probabilistic Models. arXiv preprint arXiv:2112.00390 (2021). <https://doi.org/https://doi.org/10.48550/arXiv.2112.00390>
3. Austin, J., Johnson, D.D., Ho, J., Tarlow, D., van den Berg, R.: Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems* **34**, 17981–17993 (2021)
4. Bansal, A., Borgnia, E., Chu, H.M., Li, J.S., Kazemi, H., Huang, F., Goldblum, M., Geiping, J., Goldstein, T.: Cold diffusion: Inverting arbitrary image transforms without noise. arXiv preprint arXiv:2208.09392 (2022). <https://doi.org/https://doi.org/10.48550/arXiv.2208.09392>
5. Batzner, K., Heckler, L., König, R.: EfficientAD: Accurate Visual Anomaly Detection at Millisecond-level Latencies. arXiv preprint arXiv:2303.14535 (2023). <https://doi.org/https://doi.org/10.48550/arXiv.2303.14535>
6. Bergmann, P., Batzner, K., Fauser, M., Sattlegger, D., Steger, C.: The MVTEC Anomaly Detection Dataset: A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection. *International Journal of Computer Vision* **129**(4), 1038–1059 (2021). <https://doi.org/https://doi.org/10.1007/s11263-020-01400-4>
7. Bergmann, P., Batzner, K., Fauser, M., Sattlegger, D., Steger, C.: Beyond Dents and Scratches: Logical Constraints in Unsupervised Anomaly Detection and Localization. *International Journal of Computer Vision* **130**(4), 947–969 (2022). <https://doi.org/https://doi.org/10.1007/s11263-022-01578-9>
8. Bergmann, P., Fauser, M., Sattlegger, D., Steger, C.: Uninformed Students: Student-Teacher Anomaly Detection with Discriminative Latent Embeddings. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4183–4192 (2020). <https://doi.org/https://doi.org/10.1109/CVPR42600.2020.00424>
9. Bergmann, P., Löwe, S., Fauser, M., Sattlegger, D., Steger, C.: Improving Unsupervised Defect Segmentation by Applying Structural Similarity to Autoencoders. In: Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2019) - Volume 5: VISAPP. pp. 372–380 (01 2019). <https://doi.org/https://doi.org/10.5220/0007364503720380>
10. Chen, S., Sun, P., Song, Y., Luo, P.: DiffusionDet: Diffusion Model for Object Detection. arXiv preprint arXiv:2211.09788 (2022). <https://doi.org/https://doi.org/10.48550/arXiv.2211.09788>
11. Deng, H., Li, X.: Anomaly Detection via Reverse Distillation From One-Class Embedding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9737–9746 (June 2022). <https://doi.org/https://doi.org/10.1109/CVPR52688.2022.00951>

12. Diakogiannis, F.I., Waldner, F., Caccetta, P., Wu, C.: ResUNet-a: A deep learning framework for semantic segmentation of remotely sensed data. *ISPRS Journal of Photogrammetry and Remote Sensing* **162**, 94–114 (2020). <https://doi.org/https://doi.org/10.1016/j.isprsjprs.2020.01.013>
13. Gudovskiy, D., Ishizaka, S., Kozuka, K.: CFLOW-AD: Real-Time Unsupervised Anomaly Detection with Localization via Conditional Normalizing Flows. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 98–107 (2022). <https://doi.org/https://doi.org/10.1109/WACV51458.2022.00188>
14. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* **33**, 6840–6851 (2020), <https://proceedings.neurips.cc/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf>
15. Huang, R., Lam, M.W., Wang, J., Su, D., Yu, D., Ren, Y., Zhao, Z.: FastDiff: A Fast Conditional Diffusion Model for High-Quality Speech Synthesis. In: *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22. International Joint Conferences on Artificial Intelligence Organization* (2022). <https://doi.org/https://doi.org/10.24963/ijcai.2022/577>
16. Jang, J., Hwang, E., Park, S.H.: N-Pad: Neighboring Pixel-Based Industrial Anomaly Detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. pp. 4364–4373 (June 2023). <https://doi.org/https://doi.org/10.1109/CVPRW59228.2023.00459>
17. Kong, Z., Ping, W., Huang, J., Zhao, K., Catanzaro, B.: DiffWave: A Versatile Diffusion Model for Audio Synthesis. In: *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021* (2021), <https://openreview.net/forum?id=a-xFK8Ymz5J>
18. Li, C.L., Sohn, K., Yoon, J., Pfister, T.: CutPaste: Self-Supervised Learning for Anomaly Detection and localization. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9664–9674 (2021). <https://doi.org/https://doi.org/10.1109/10.1109/CVPR46437.2021.00954>
19. Li, X.L., Thickstun, J., Gulrajani, I., Liang, P., Hashimoto, T.: Diffusion-LM Improves Controllable Text Generation. In: *Advances in Neural Information Processing Systems*. vol. 35, pp. 4328–4343 (2022)
20. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal Loss for Dense Object Detection. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence*. vol. 42, pp. 318–327 (2020). <https://doi.org/https://doi.org/10.1109/TPAMI.2018.2858826>
21. Liu, Z., Zhou, Y., Xu, Y., Wang, Z.: SimpleNet: A Simple Network for Image Anomaly Detection and Localization. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 20402–20411 (2023). <https://doi.org/https://doi.org/10.48550/arXiv.2303.15140>
22. Lu, F., Yao, X., Fu, C.W., Jia, J.: Removing anomalies as noises for industrial defect localization. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 16166–16175 (October 2023)
23. Meng, C., Rombach, R., Gao, R., Kingma, D., Ermon, S., Ho, J., Salimans, T.: On distillation of guided diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 14297–14306 (2023)
24. Perlin, K.: An image synthesizer. *ACM Siggraph Computer Graphics* **19**(3), 287–296 (1985). <https://doi.org/https://doi.org/10.1145/325165.325247>

25. Roth, K., Pemula, L., Zepeda, J., Schölkopf, B., Brox, T., Gehler, P.: Towards Total Recall in Industrial Anomaly Detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14318–14328 (2022). <https://doi.org/https://doi.org/10.1109/CVPR52688.2022.01392>
26. Rudolph, M., Wehrbein, T., Rosenhahn, B., Wandt, B.: Fully Convolutional Cross-Scale-Flows for Image-based Defect Detection. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1088–1097 (2022). <https://doi.org/https://doi.org/10.1109/WACV51458.2022.00189>
27. Rudolph, M., Wehrbein, T., Rosenhahn, B., Wandt, B.: Asymmetric student-teacher networks for industrial anomaly detection. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2592–2602 (2023)
28. Sakurada, M., Yairi, T.: Anomaly Detection Using Autoencoders with Nonlinear Dimensionality Reduction. In: Proceedings of the MLSDA 2014 2nd workshop on machine learning for sensory data analysis. pp. 4–11 (2014). <https://doi.org/https://doi.org/10.1145/2689746.2689747>
29. Salimans, T., Ho, J.: Progressive distillation for fast sampling of diffusion models. In: International Conference on Learning Representations (2022), <https://openreview.net/forum?id=TIIXIpzhoI>
30. Schlegl, T., Seeböck, P., Waldstein, S.M., Langs, G., Schmidt-Erfurth, U.: f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks. *Medical image analysis* **54**, 30–44 (2019). <https://doi.org/https://doi.org/10.1016/j.media.2019.01.010>
31. Song, J., Meng, C., Ermon, S.: Denoising Diffusion Implicit Models. In: 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021 (2021), <https://openreview.net/forum?id=St1giarCHLP>
32. Song, Y., Dhariwal, P., Chen, M., Sutskever, I.: Consistency models. In: Proceedings of the 40th International Conference on Machine Learning. ICML'23, JMLR.org (2023)
33. Tailanian, M., Pardo, Á., Musé, P.: U-Flow: A U-shaped Normalizing Flow for Anomaly Detection with Unsupervised Threshold. arXiv preprint arXiv:2211.12353 (2022). <https://doi.org/https://doi.org/10.48550/arXiv.2211.12353>
34. Wang, Z., Liu, J.C.: Translating Math Formula Images to LaTeX Sequences Using Deep Neural Networks with Sequence-level Training. *International Journal on Document Analysis and Recognition (IJ DAR)* **24**(1–2), 63–75 (2021). <https://doi.org/https://doi.org/10.1007/s10032-020-00360-2>
35. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004). <https://doi.org/https://doi.org/10.1109/TIP.2003.819861>
36. Wu, J., FU, R., Fang, H., Zhang, Y., Yang, Y., Xiong, H., Liu, H., Xu, Y.: Med-SegDiff: Medical Image Segmentation with Diffusion Probabilistic Model. In: *Medical Imaging with Deep Learning* (2023), <https://openreview.net/forum?id=Jdwc2jG9>
37. Wyatt, J., Leach, A., Schmon, S.M., Willcocks, C.G.: AnoDDPM: Anomaly Detection With Denoising Diffusion Probabilistic Models Using Simplex Noise. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 650–656 (June 2022). <https://doi.org/https://doi.org/10.1109/CVPRW56347.2022.00080>

38. Yang, M., Wu, P., Feng, H.: MemSeg: A semi-supervised method for image surface defect detection using differences and commonalities. *Engineering Applications of Artificial Intelligence* **119**, 105835 (2023). <https://doi.org/https://doi.org/10.1016/j.engappai.2023.105835>
39. Yu, J., Zheng, Y., Wang, X., Li, W., Wu, Y., Zhao, R., Wu, L.: FastFlow: Unsupervised Anomaly Detection and Localization via 2D Normalizing Flows. arXiv preprint arXiv:2111.07677 (2021). <https://doi.org/https://doi.org/10.48550/arXiv.2111.07677>
40. Zavrtnik, V., Kristan, M., Skočaj, D.: DRAEM-A discriminatively trained reconstruction embedding for surface anomaly detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 8330–8339 (2021). <https://doi.org/https://doi.org/10.1109/ICCV48922.2021.00822>
41. Zavrtnik, V., Kristan, M., Skočaj, D.: DSR-A dual subspace re-projection network for surface anomaly detection. In: *European Conference on Computer Vision*. pp. 539–554. Springer (2022). https://doi.org/https://doi.org/10.1007/978-3-031-19821-2_31
42. Zavrtnik, V., Kristan, M., Skočaj, D.: Reconstruction by inpainting for visual anomaly detection. *Pattern Recognition* **112**, 107706 (2021). <https://doi.org/https://doi.org/10.1016/j.patcog.2020.107706>
43. Zhang, X., Li, N., Li, J., Dai, T., Jiang, Y., Xia, S.T.: Unsupervised surface anomaly detection with diffusion probabilistic model. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 6782–6791 (October 2023)
44. Zhang, X., Li, S., Li, X., Huang, P., Shan, J., Chen, T.: DeSTSeg: Segmentation Guided Denoising Student-Teacher for Anomaly Detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3914–3923 (2023)
45. Zou, Y., Jeong, J., Pemula, L., Zhang, D., Dabeer, O.: SPot-the-Difference Self-supervised Pre-training for Anomaly Detection and Segmentation. In: *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXX*. pp. 392–408. Springer (2022). https://doi.org/https://doi.org/10.1007/978-3-031-20056-4_23