



HAL
open science

Distillation-based fabric anomaly detection

Simon Thomine, Hichem Snoussi

► **To cite this version:**

Simon Thomine, Hichem Snoussi. Distillation-based fabric anomaly detection. Textile Research Journal, 2023, 10.1177/00405175231206820 . hal-04372602

HAL Id: hal-04372602

<https://hal.science/hal-04372602v1>

Submitted on 4 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distillation-based fabric anomaly detection

Thomine Simon^{1,2*} and Snoussi Hichem¹

¹*LIST3N, University of technology of Troyes, Troyes, 10000, France.

²Aquilae, Troyes, 10000, France.

*Corresponding author(s). E-mail(s): simon.thomine@utt.fr;

Contributing authors: hichem.snoussi@utt.fr;

Abstract

Unsupervised texture anomaly detection has been a concerning topic in a vast amount of industrial processes. Patterned textures inspection, particularly in the context of fabric defect detection, is indeed a widely encountered use case. This task involves handling a diverse spectrum of colors and textile types, encompassing a wide range of fabrics. Given the extensive variability in colors, textures, and defect types, fabric defect detection poses a complex and challenging problem in the field of patterned textures inspection. In this article, we propose a knowledge distillation-based approach tailored specifically for addressing the challenge of unsupervised anomaly detection in textures resembling fabrics. Our method aims to redefine the recently introduced reverse distillation approach, which advocates for an encoder-decoder design to mitigate classifier bias and to prevent the student from reconstructing anomalies. In this study, we present a new reverse distillation technique for the specific task of fabric defect detection. Our approach involves a meticulous design selection that strategically highlights high-level features. To demonstrate the capabilities of our approach both in terms of performance and inference speed, we conducted a series of experiments on multiple texture datasets, including MVTEC AD, AITEX, and TILDA, alongside conducting experiments on a dataset acquired from a textile manufacturing facility. The main contributions of this paper are the following: a robust texture anomaly detector utilizing a reverse knowledge-distillation technique suitable for both anomaly detection and domain generalization and a novel dataset encompassing a diverse range of fabrics and defects.

Keywords: unsupervised, anomaly, texture, pattern, domain-generalized, fabric, knowledge distillation

1 Introduction

Fabric defect detection has been a really worrisome problem for the textile industry as the quality of the garment is a priority for the customer. Due to the absence of superior alternatives, inspection is currently carried out mainly by humans, which represents a significant labor cost and does not guarantee flawless detection due to visual fatigue and distraction. Currently, the accuracy of humans in detecting defects is around 60% and decreases for visually difficult fabrics such as striped fabrics. Several constraints are relevant to consider when dealing with this problem: industrial processes are well oiled and cannot suffer from stoppage or slowing of any type, the variety of defects/-fabric type is colossal and evolving over time with new collections, defects can be tiny and therefore require an industrial camera with high resolution. As a result, we should aim for a method that does not require defective samples in the training, while offering generalization capability and a fast inference time. Recently, a predominant trend in unsupervised anomaly detection methods has emerged. It involves the utilization of features extracted from pre-trained networks to learn descriptive representations of the analyzed objects or textures [1–8, 8–12] with several approaches such as normalizing flow [13], memory bank and knowledge distillation [14]. Our approach is based on a knowledge distillation approach, primarily due to its advantages in terms of inference speed and its well-established capability in detecting defects on various types of textures [12]. Within the framework of knowledge distillation, knowledge transfer occurs between a teacher-student pair during the training. In the context of unsupervised anomaly detection, where the student model is trained solely on normal samples, the student model generates disparate representations compared to the teacher model when confronted with an anomalous sample. Our approach draws inspiration from the reverse distillation architecture [10]. This architecture follows a student-teacher paradigm, where the teacher model serves as the encoder and the student model as the decoder. By employing this architecture, we aim to leverage the knowledge and guidance provided by the teacher model to enhance the reconstruction capabilities of the student model in the context of anomaly detection. By design, this architecture emphasizes the deeper selected layers and therefore neglects the shallower reconstructed layers. In the domain of texture defect detection, the utmost significance lies in the high-level attributes, as the contextual information has less significance [12]. Since we aim at detecting defects on fabrics-like textures, we want to keep shallower information and to that extent we introduced residual connections between the teacher and the student along with custom attention layers. Another crucial component of the reverse distillation approach is the embedding bottleneck, which is responsible for gathering features from specific layers of the teacher model and creating a feature map that serves as input for the student model. In our bottleneck embedding architecture, we strive to design a texture-specific approach that selects the most representative features from the teacher model without altering the descriptive capabilities. To that extent, we carefully designed a bottleneck based mainly on 1×1 convolutions along with a custom attention block and an SSPCAB layer [15]. In industrial deployment, interrupting the production chain for training a model on a specific type of texture is not feasible. Hence, we propose a slight modification of our model to introduce a

domain-generalized model capable of detecting defects in any plain fabric. The primary idea behind this approach is to enable defect detection while a more precise and specific model is being trained. By adopting this strategy, we can ensure that the production chain remains operational, albeit with a concession in terms of precision. Our approach offers outstanding results for defect detection on plain textures, along with a fast inference time and domain generalization capabilities compared to the other state-of-the-art methods.

The primary contributions of this paper are outlined as follows:

- A robust texture anomaly detector utilizing a reverse knowledge-distillation technique. This approach enhances the model ability to detect anomalies in fabric-like textures with improved accuracy and reliability. A slight modification of this approach is also proposed for domain generalization purposes.
- Introduction of a novel dataset encompassing a diverse range of fabrics and defects. This dataset provides a valuable resource for training and evaluating fabric anomaly detection models, facilitating research and advancements in the textile domain.

Following the introduction, the subsequent section of this paper is dedicated to reviewing existing literature concerning both feature extraction and deep learning approaches for unsupervised anomaly detection. Subsequently, we present our novel reverse distillation approach with a precise description of the model. The following two sections focus on introducing a new dataset and conducting experiments to evaluate the effectiveness of our method. The final section provides a conclusion for the paper.

2 Related work

To handle the problem of texture defect detection, we first have to extract the main characteristics that describe the texture. As mentioned in [16], we can distinguish seven categories of approaches to handle this problem namely statistical approaches, structural approaches, transform-based approaches, model-based approaches, graph-based approaches, learning-based approaches and entropy-based approaches. Each of these approaches offers different advantages and drawbacks and some of them offer competitive results while being really fast with recent computing resources such as GLCM [17] and local binary pattern [18] [19].

More recently, deep learning based feature extraction brought outstanding results for extracting the most important features characterizing a texture. An example of supervised training approaches consists of training a deep texture classifier and then extracting the last layer to obtain informative features. [20]

Autoencoder-based approaches are the most common unsupervised approaches to extract features [21] but still suffer from high generalization capabilities which can prevent accurate defect detection.

To tackle the problem of defect detection in the industrial case, gathering every defect of a given object or texture is a laborious task and can lead to poor performances if every type of defect was not considered[22]. Certain methods still utilize supervised training methods like mobile-unet [23] which integrates classification and

segmentation techniques, or the utilization of SSD-based approaches [24]. Unsupervised anomaly detection deals with the problem of detecting defects on an object or a texture without any prior on the type of possible defects. Consequently, many methods emerged proposing different types of algorithms such as autoencoders [25] and variational autoencoder variants [26] [27]. Another common way of detecting anomalies is Generative Adversarial Networks (GAN), introduced by [28], adapted to unsupervised anomaly detection such as Ano-GAN [29], G2D [30] and OCR-GAN [31]. More recently, approaches using a classifier pretrained on natural images as feature extractor made their way in industrial anomaly detection, offering outstanding performance. These methods, such as normalizing flow, knowledge distillation and memory banks, use selected layers output features to obtain descriptive feature maps. Then, a training is conducted based on the prior that this feature maps describe the most important characteristics of the given object. The normalizing flow approach consists of a flow training based on relevant features of good samples from a pretrained network such as AlexNet [1], Resnet [2], efficient-net [3] and even vision transformers [32] [33] trained on imageNet. Different strategies were used to enhance performances, such as a 2D flow [4] or a cross-scale flow [5]. The use of a memory bank has the advantage of offering a description over the whole class to compare and detect if there is an anomaly [6] [7].

Distilling the knowledge from a teacher to a student [14] was recently adapted for unsupervised anomaly detection [8] [9]. This concept involves training a student network on good samples using the output features of a teacher network that has already been pre-trained for classification purposes. In the test phase, the student network will be capable of replicating the output features of the teacher network when provided with good samples, but when confronted with anomalous samples, the student will not be able to extract descriptive features and therefore a relevant anomaly score can be computed.

Several methods used this principle with different strategies such as a multi-layer feature selection [9], a reverse distillation approach [10], a coupled-hypersphere-based feature adaptation [11] and a mixed-teacher approach [12].

3 Residual Reverse distillation using specifically selected embeddings

This part is dedicated to the use of reverse distillation to detect defects in textures. To this extent, we divided this chapter in four subsections: the reverse distillation approach along with the residual connections, the description of the bottleneck embedding that transforms the teacher’s output into input for the student model, the loss and score calculation strategy and the adaptation for domain generalization.

3.1 Residual reverse distillation

Approach: The idea of reverse distillation was introduced in [10] with the purpose of proposing an encoder-decoder version of the knowledge distillation paradigm applied to unsupervised anomaly detection. An important problem of unsupervised anomaly detection is the generalization capability, which is problematic since our purpose is to

be specific to a texture for detecting any kind of defect. To tackle this specific problem, we propose a residual reverse distillation with a reconstructed embedding emphasizing high-level features. The main idea is to extract the most relevant features describing the texture using a limited encoding and residual connections, allowing a selective and descriptive embedding for input to the decoder student.

Problem definition: Given a training dataset of images without anomaly $D = [I_1, I_2, \dots, I_n]$, our goal is to extract the information of the L bottom layers. For an image $I_k \in R^{w \times h \times c}$ where w is the width, h the height, c the number of channels and l the l th bottom layer, the teacher outputs features $F_t^l(I_k) \in R^{w_l \times h_l \times c_l}$. The decoder output features are defined as $F_s^{L-l}(I_k) \in R^{w_l \times h_l \times c_l}$. During the training phase, the student model is trained to reproduce the teacher features on good samples. In testing, when inferring an anomaly, the teacher and the student will output significantly different features and thus will give us an anomaly score. In the anomaly detection setting, normal samples follow the same distribution in both F_t and F_s . Out-of-distribution samples are considered anomalies.

Our reverse distillation backbone draws inspiration from the one outlined in reverse distillation [10], differing in terms of layer selection, selective embedding, and residual connections. Given the texture specific architecture described in [12], we have decided to put an emphasis on the first layers of the network, since they describe high-level features that are primordial in patterned textures. To this purpose, we employed adaptive loss functions to assign varying significance to individual layers, not to mention that with identical loss factors for every layer, we determine that the loss is more impacted by deeper layers with random weight initialization in the student. Also, we include the initial ResNet layer before the first residual block to guarantee a high-level representation of the texture.

The goal of the teacher encoder T in the reverse distillation paradigm is to extract thorough representations so that the decoder student S can reconstruct the different selected teacher’s features. As in previous work [10] [11], we used ResNet pretrained on ImageNet as our backbone for T. As in any knowledge distillation training, the teacher T is frozen to avoid converging to a trivial solution for T and S. In [10], they used the same exact number of layer for the student and the teacher to match the intermediate representation. We demonstrate that a shallower student with less residual blocks between each dimension is capable of offering a good representation, which also prevents the student to generalize in case of defective sample and therefore increase our defect detection capability in textures.

As part of its inherent design, the traditional reverse distillation approach highlights the deeper selected layers, owing to their dimensional proximity with the embedding representation. As mentioned earlier, for textures anomaly detection, shallower layers are more descriptive and less subject to misinterpretation. In this respect, we had some residual connections between the teacher and the student to try to compensate of the loss of information caused by the embedding. However, adding a direct connection between the equivalent layers of the student and the teacher would most likely result in identity or zero-weights depending on the selected method of connection. To counteract this issue, we choose to connect the teacher layer l to the student

layer $L - l - 1$ with a convolution layer and a pooling layer as well as a self-attention layer, shown in figure 1.

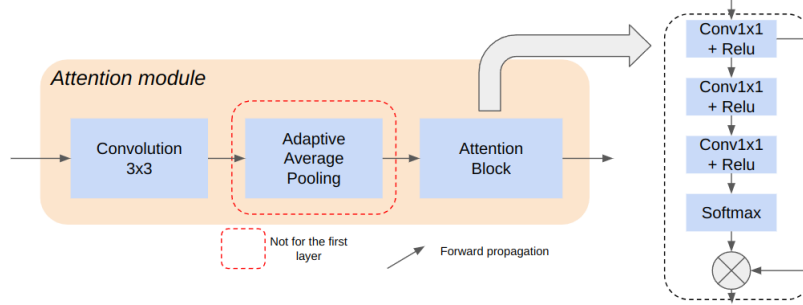


Fig. 1: Attention module used for residual connection

With the purpose of enriching the feature representation and modulating the feature importance, we introduce a custom attention block, described in figure 1, which is similar to a spatial attention but uniquely relies on 1×1 convolutions. The inclusion of this layer enables the introduction of non-linearity, leading to improved model performance with minimal additional computational expense.

The complete architecture is shown in figure 2

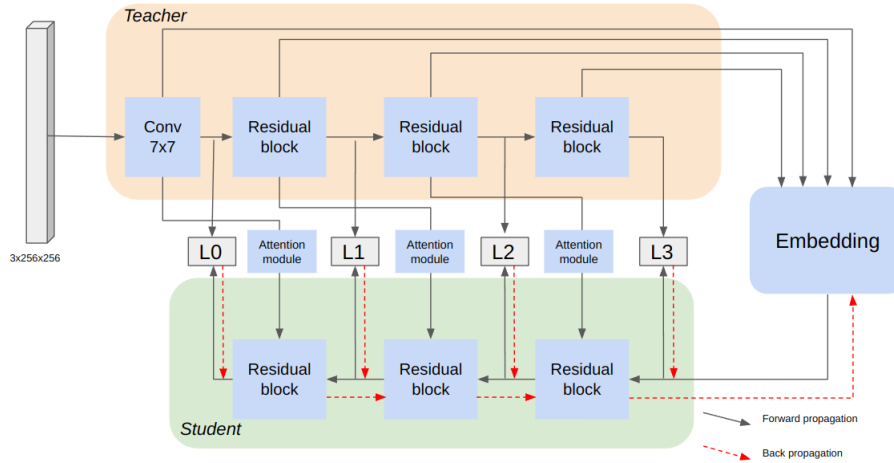


Fig. 2: Reverse residual distillation architecture.

3.2 Bottleneck Embedding

One of the main contributions of the reverse distillation paradigm is the embedding representation. In other knowledge-distillation anomaly detector models, the main drawback was the student capacity which could eventually generalize to anomalous defects. MixedTeacher [12] proposed a smaller student to counteract this problem and manage to increase anomaly detection on textures by a huge factor compared to the original STPM model [9]. Our approach goes even further with the selective embedding encoding approach that prevents the student from any unnecessary parameters. This approach is specifically outlined for texture defect detection; however, for object defect detection, the size of the embedded minimal representation could potentially vary. This issue revolves around striking a balance between the model’s capacity and its capability to accurately reconstruct correct elements, all the while ensuring it remains incapable of reconstructing defective elements.

First, to aggregate the features from the selected teacher’s layers, we used convolution of size 1 kernel as described in figure 3 to get all the features at the size of the last layer divided by 2 to prevent the prevalence of the last layer over the others. The obtained feature map is our bottleneck input and the bottleneck embedding will then learn to reconstruct this representation before giving to the student decoder.

To increase model performance, we have decided to include a SSPCAB layer [15] as the first layer of our bottleneck network since this specific type of layer shows increased performance for a wide variety of unsupervised anomaly detector. To make sure the local information was preserved, we decided to keep a relatively shallow encoding embedding. Each layer is constituted of a convolutional layer with a kernel size of 1 and a stride of 1 followed by a ReLU and a BatchNorm layer [34] except for the last which have a stride of 2 to fit the decoder input. Employing convolutional operations with a filter size of 1 confers the benefit of computational efficiency, alongside the capacity to introduce non-linearities and trainable parameters.

3.3 Loss and Anomaly Score Calculation

Mathematically, given the problem definition parameter, the pixel-loss function along the channel axis is defined as:

$$M^l(I_k)_{ij} = \frac{1}{2} \| \text{norm}(F_t^l(I_k)_{ij}) - \text{norm}(F_s^{L-l}(I_k)_{ij}) \| \quad (1)$$

with $M^l \in \mathbb{R}^{H_l \times W_l}$, the layer l loss function as :

$$\text{loss}^l(I_k) = \frac{1}{w_l h_l} \sum_{i=1}^{w_l} \sum_{j=1}^{h_l} M^l(I_k)_{ij} \quad (2)$$

and the global loss is written as:

$$\text{loss}(I_k) = \sum^l \alpha_l \text{loss}^l(I_k) \quad (3)$$

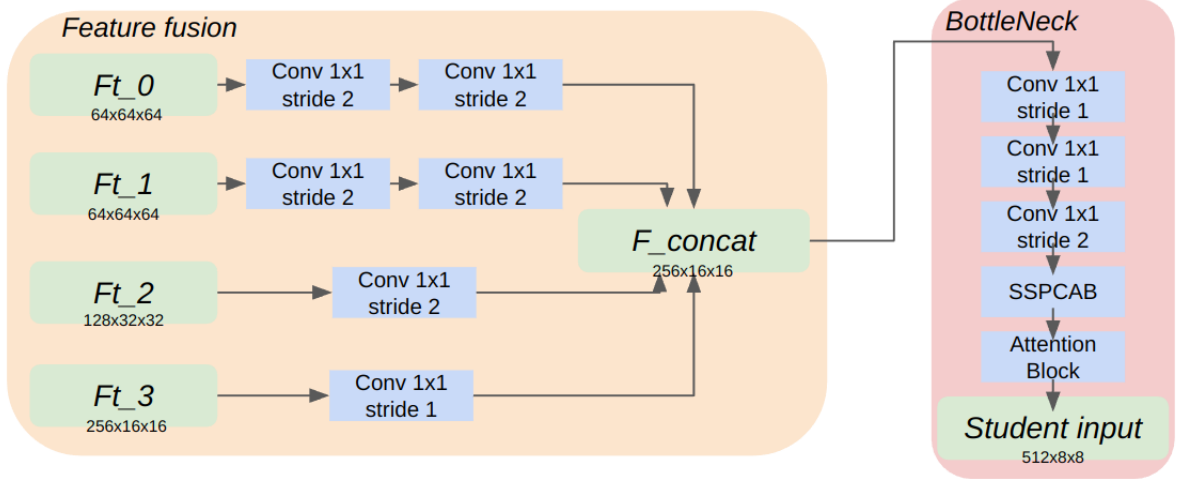


Fig. 3: Feature fusion and bottleneck architecture

with α_l the loss factor for the layer l .

At the testing stage, we want an anomaly score and an anomaly localization. To begin with, let's address the anomaly localization problem. The paradigm of knowledge distillation for anomaly detection and localization stipulates that, in case of anomalous sample, the teacher is able to represent the defect while the student lacks generalization capability as it has learnt to generate features only for defect-free samples. This means that the student and the teacher features will be different and this will reflect on the anomaly map calculated in equation 1. To account for the anomaly maps of all the features, we regroup all these features maps with upsampling in a single localization map denoted L_{map} :

$$L_{map} = \sum_{i=1}^l UpSample(M^l) \quad (4)$$

A Gaussian filter is then used on L_{map} to smooth the defect localization.

For the image anomaly score or defect detection, to avoid the prevalence of huge defects over small defects, we decided not to use the mean of the anomaly map but the max value instead based on the paradigm that a non-defective anomaly map will have no significant out-of-distribution value especially after the application of the Gaussian filter.

3.4 Reverse distillation for domain generalization

Domain generalization in deep learning refers to the task of training a model on data from multiple source domains, and then generalizing this model to perform well on a new target domain. The goal of domain generalization is to improve the generalization performance of deep learning models, especially in situations where the test data comes from a different distribution than the training data. In the specific domain of fabric defect detection, our objective is to develop an anomaly detection system capable of

identifying defects in all types of textiles. However, it is important to note that our design and testing procedures are specifically tailored for homogeneous textures and may not be directly applicable to fabric designs that exhibit more intricate patterns or variations. As demonstrated in FABLE [35], the knowledge distillation architecture has outstanding domain-generalized performance for homogeneous textures.

The domain generalization approach for fabric anomaly detection can be viewed as a strategy that provides additional training samples to the model. By incorporating a diverse range of fabric samples during training, the model becomes more adept at recognizing fabric-specific characteristics, such as thread patterns and lattice structures. This broader exposure helps the model generalize better across different fabric types, enhancing its ability to detect anomalies across a variety of textile materials. Indeed, while the idea of learning inherent textile patterns is appealing, achieving effective fabric anomaly detection is a complex task. There are various challenges associated with the intricacies of fabric textures, including variations in weave patterns, colors, textures, and defect types. Additionally, factors such as lighting conditions, scale, and resolution can impact the detection accuracy. Developing a robust and accurate fabric anomaly detection model requires careful consideration of these complexities and the implementation of sophisticated techniques to account for them.

As previously discussed, simpler patterns like thread and lattices are often captured by the shallower layers of the model architecture. This observation motivated the inclusion of a residual connection, which places emphasis on these particular layers. By doing so, the model can effectively capture and learn these fundamental fabric patterns, thereby enhancing its ability to detect anomalies associated with thread and lattice structures.

On the contrary, in the case of domain generalization, we believe that all deeper layers are a key to understanding the intricacies of such a descriptive problem. Consequently, we have made the decision to eliminate the residual connection mentioned in Section 3.1.1, as we consider it to be less relevant in this context.

4 Industrial textile dataset

In this part, we introduce the industrial fabrics dataset containing 2150 images with 6 types/colors of textile. We constitute two nomenclature for both domain-generalized and classic unsupervised anomaly detection training. We decide to use the same folder architecture as MVTEC AD dataset [36].

To constitute the dataset, we placed 2 cameras from different brand in front of a visiting machine with strong and nearly uniform light source.

4.1 Industry constraints

Industrial deployments are subject to certain constraints pertaining to the positioning of cameras and the intensity of light sources to prevent any disruption to the machine operator’s activities. Our aim was to carefully choose an ideal pairing of light intensity and shutter duration to guarantee the capture of images that are sharp and clearly

defined. However, due to variations in fabric color and the velocity of the visiting machine, our chosen light source did not consistently provide adequate brightness. As a result, certain images within the dataset exhibit a slight degree of blurriness, contributing to the creation of a more demanding and authentic dataset. This reflects the fact that achieving flawless conditions, even in an industrial setting, is frequently unachievable.

4.2 Variety of defects

As previously stated, defects within the textile industry exhibit a wide range of variations. Furthermore, our dataset will continue to expand as we accumulate additional instances of defects through the ongoing use of our cameras. In addition, it is worth noting that certain anomalies present in our database may consist of dust particles that have settled on the textile, which may not typically be classified as defects in most textile industries. However, due to our inability to definitively determine whether these anomalies are genuine defects or simply dust, we have chosen to include them in the dataset. Figure 4 showcases examples of such anomalies.



Fig. 4: Images extracted from our proposed dataset with defect that could be dust

4.3 Dataset description

The dataset is composed of patches of size 256x256 extracted from 2 industrial cameras of resolution 3088x2076 and 2464x2056 respectively. The defects included in the dataset are not artificially generated but rather originate from real-world instances within the textile industry. However, we added some non-defect from the industry like writing on the fabric or the sewing attaching two pieces of fabric. These added "defects" aim to diversify the dataset and provide a broader range of defect types for training and analysis purposes. The dataset contains 6 classes with between 199 and 588 images for training and 30 and 114 for testing. The precise distribution is presented in table 1. It is important to acknowledge that the dataset poses a potential long-tail distribution challenge, especially for domain generalization training, but remains suitable for "classic" unsupervised anomaly detection where distinct models are trained for each class. An overview of the dataset can be seen in figure 5.




Fig. 5: Images of our proposed datasets, (upper) defect-free images, (lower) defective images. Some images are a bit blurry due to motion blur. Other kind of defective samples are presented in the experiment part.

Category	total	train	testGood	testDefective
type1cam1	386	272	28	86
type2cam2	257	199	19	39
type3cam1	689	588	54	47
type4cam2	229	199	19	11
type5cam2	298	199	19	80
type6cam2	291	199	19	73
total	2150	1656	158	336

Table 1: Industrial textile dataset distribution

4.4 Existing datasets

In order to provide contextualization and to position our dataset within the broader landscape of existing datasets, it is essential to present other relevant datasets that specifically address the problem of texture anomaly detection.

One of the widely used industrial unsupervised anomaly detection datasets is the MVTEC AD dataset [36], which comprises 10 objects and 5 textures. While 4 of these textures align well with the problem of textile anomaly detection, one texture is significantly different. Other datasets, such as TILDA [37], Aitex [38] and DAGM [39], offer grayscale fabrics-like datasets for anomaly detection purposes. In addition, datasets like BTech [40], NEU-CLS [41], KolektorSDD [42], Magnetic Tile Dataset [43] and RSDDs [44] provide supplementary datasets for surface defect detection, which are not specific to textiles.

Our dataset strives to offer color images of fabrics captured under authentic conditions, showcasing a wide range of real defects. With an emphasis on diversity, our dataset encompasses an extensive array of defect types, ensuring its applicability in addressing various anomaly detection challenges within the textile industry.

Category	CFA	PatchCore[6]	FastFlow[4]	Reverse distillation[10]	MixedTeacher[12]	Ours
carpet	97.3	98.7	99.4	98.9	99.8	100
tile	99.4	98.7	100	99.3	100	99.7
wood	99.7	99.2	99.2	99.2	99.6	99.8
leather	100	100	99.9	100	100	100
Mean	99.1	99.1	99.6	99.3	99.8	99.9

Table 2: Anomaly detection results with AUROC on MVTEC fabric-like textures

Category	PatchCore[6]	FastFlow[4]	Reverse distillation[10]	MixedTeacher[12]	Ours
carpet	98.9	99.1	98.9	99.0	98.8
tile	95.6	96.6	95.6	95.9	93.7
wood	95	94.1	95.3	94.9	92.4
leather	99.3	99.6	99.4	99.6	98.7
Mean	97.2	97.3	97.3	97.4	95.9

Table 3: Anomaly localization results with AUROC on MVTEC fabric-like textures

5 Experiments

This part showcases performances of our knowledge distillation approach both in terms of AUROC score and inference time since these are the most concerning application requirements in anomaly detection.

We have also added a coverage measure to determine the actual benefit of this approach in real case scenario.

Training and inference were done on an RTX 3080ti. To ensure efficient inference for industrial applications, we employ ResNet18 and ResNet34 as the backbone architectures for our anomaly detection system. The teacher model, which serves as the reference model, is a ResNet34 trained on the ImageNet dataset. In contrast, the student model consists of a ResNet18 backbone. It is worth noting that the student model is structured in a backward manner, as explained in the part 3. We used ADAM optimizer[45], a learning rate of 0.005 with scheduler to reduce the learning rate when hitting a plateau during the training. The training was done for 100 epochs with a batch size of 4. We resized all the images to a size of 256x256 keeping respectively 70% and 30% for training and validation. We kept the checkpoint with the lowest validation loss and introduce early stopping.

5.1 State of the art AUROC comparison

To evaluate our architecture, we use several datasets including MVTEC AD textures [36], TILDA [37] dataset, AITEX dataset [38] and our dataset presented in section 4. For MVTEC AD, we selected the four fabric-like textures and compared our results to the other state-of-the-art methods. The results are presented in table 2 for anomaly detection and table 3 for anomaly localization.

For the TILDA dataset, we restructured the data to conform to the MVTEC nomenclature. To ensure a fair comparison with state-of-the-art methods, the anomaly library [46] was employed. The results are shown in table 4. Examples of input




Fig. 6: Input image and anomaly map in the 4 fabric-like MVTEC AD textures

Tilda	fabric1	fabric2	fabric3	fabric4	fabric5	fabric6	fabric7	fabric8	mean
reverseDistillation[10]	94.8	88.2	91.4	59.6	67.4	78.7	78.6	84.5	80.4
ours	96.9	95.8	92.5	75.0	87.2	88.6	70.7	61.9	83.6

Table 4: Anomaly detection results with AUROC on Tilda textures

images and results can be seen in figure 7

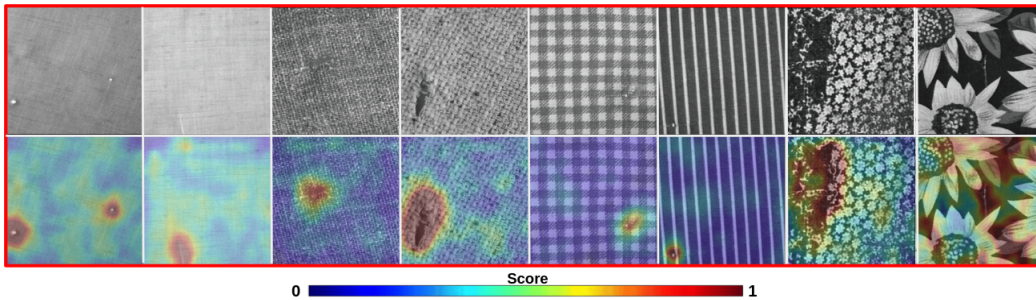


Fig. 7: Input image and anomaly map in all 8 TILDA textile. As expected, the model is too restrictive to handle textures with patterns like flowers

For the aitex dataset, we separate each 4096x256 images into 16 patches of size 256x256 and used the MVTEC AD nomenclature. The results of our experiment are presented in table 5. Examples of input images and their corresponding results are visually depicted in figure 8

Subsequently, we conduct assessments on the dataset for which our model has been meticulously tailored. This dataset is derived from an authentic industrial process, aligning with our model’s primary objective of optimizing performance specifically for this designated task. Results are depicted in table 6 and images and anomaly maps are

Aitex	fabric1	fabric2	fabric3	fabric4	mean
reverseDistillation [10]	92.7	96.5	77.3	97.7	91.0
ours	97.8	99.2	89.1	100	96.5

Table 5: Anomaly detection results with AUROC on Aitex textures

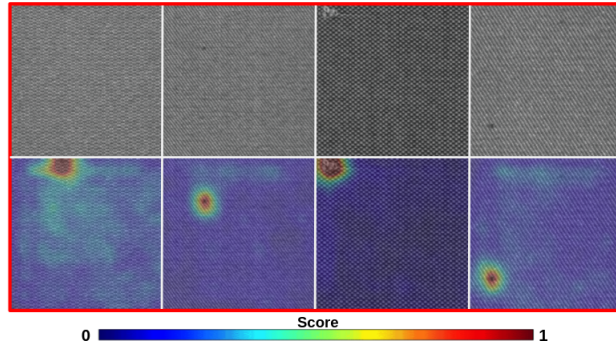


Fig. 8: Input image and anomaly map in all 4 Aitex textile

shown in figure 9. Our results on the dataset are exceptionally accurate as our model was specifically designed and optimized for it. However, we continue to record images using the camera and actively strive to gather a larger and more diverse dataset. By expanding the dataset, we aim to introduce greater challenges and complexities, which will further enhance the robustness and generalization capabilities of our model.




Fig. 9: Input image and anomaly map in industrial textile database. We showcase speed-blurry image on purpose to demonstrate our capability to deal with such quality differences

The domain generalization results are presented in table 7

Results interpretation : Results on the TILDA dataset demonstrates that our approach offers outstanding results on plain fabrics and simple patterns such as checkered or striped. The approach is therefore limited for more complex patterns that

ITD	type1cam1	type2cam2	type3cam1	type4cam2	type5cam2	type6cam2	mean
reverseDistillation [10]	100	100	98.3	100	100	100	99.7
ours	100	100	100	100	100	100	100

Table 6: Anomaly detection results with AUROC on Industrial textile dataset

Category	Epi-FRC+ [47]	EISNet+[48]	DGTSAD [49]	Fable [35]	Ours
carpet	91.6	98.2	94.3	98.5	99.3
tile	95.1	85.1	99.4	96.5	98.8
wood	94.1	97.9	96.2	99.9	99.7
leather	100	100	100	99.1	99.9
Mean	95.2	95.3	97.5	98.5	99.4

Table 7: Domain generalization comparison

Batch size	Inference time	FPS
16 : time per patch	0.7ms	1450
32 : time per patch	0.5ms	2000
whole image (96 patches)	45ms	22

Table 8: Inference speed

require more contextual information and a higher receptive field. That consideration establishes the limitation of such a restrictive architecture that emphasizes mostly high-level features along with 1×1 convolutions that do not increase the receptive field. Nonetheless, the results on MVTEC AD, AITEX and our dataset show that our model outperforms state-of-the-art approaches for “simple” textures.

5.2 Inference speed comparison

When designing the architecture, one of the primary concerns was the inference speed. Our approach of knowledge distillation offers the advantage of fast inference due to the utilization of “smaller” networks such as ResNet18 and ResNet34. Despite their reduced size, these networks can still deliver state-of-the-art performance, as demonstrated in the previous part. The efficacy of both the bottleneck embedding component and the residual connection component is significantly reliant upon the utilization of 1×1 convolutions, which exert minimal influence on the overall computational velocity. This combination allows for efficient and accelerated anomaly detection inference while maintaining high detection accuracy.

The added embedding and residual connections are even lighter and therefore our method offer fast inference time.

In our specific scenario, we process high resolution images of size 3000×2000 and therefore, we will have to divide this image in smaller patches of size 256×256 . The corresponding number of patches is 96 patches. To report industrial plausible inference time, we processed in batch of size 32 and reported the inference time for the whole image processing (3 batches) as shown in table 8.

Category	Coverage	Coverage +/- 0.02	Coverage +/- 0.05
carpet	100	97	91
tile	91	81	71
wood	99	96	91
leather	100	100	100
Mean	97.5	93.5	88.3

Table 9: Coverage measure with fixed precision and recall at 100 percent. The +/- indicates the score tolerance to consider that 2 samples are ambiguous. Base coverage is with score tolerance 0

5.3 Coverage Measure

Applications of this nature are primarily designed to assist humans in the task of detecting defects in the visiting machine. By leveraging automated anomaly detection methods, the goal is to provide support and assistance to human operators in their inspection and quality control tasks. These applications aim to improve the efficiency and accuracy of defect detection processes, enabling operators to identify and address anomalies more effectively. Consequently, we have made the decision to incorporate a metric that accurately reflects the practical impact of our method on assisting industrial processes. This metric will serve as an objective measure to assess the effectiveness and efficiency of our approach in real-world scenarios.

The coverage metric quantifies the level of confidence exhibited by the model while maintaining a fixed precision and recall. For example, setting a target of 80 percent coverage with a precision of 95 percent would imply that the model will classify 80 percent of defects with a precision of 95 percent, while external assistance would be required for the remaining 20 percent.

In an industrial scenario, the tolerance for false negatives is typically extremely low, indicating that the detection of all defects is essential. Therefore, a recall of 100 percent, ensuring that no defects are missed, is a critical requirement for fabric defect detection. Additionally, to prioritize precision and avoid false positives, we have decided to set the precision at 100 percent as well. This means that all identified defects are guaranteed to be genuine, providing a high level of confidence in the accuracy of the defect detection system. In conjunction with the outcomes related to inference speed and domain generalization performance, these results collectively support the potential applicability of our approach within a real-world industrial setting, potentially involving collaboration with a human operator.

The table 9 showcases coverage value for the tested MVTEC textures.

6 Conclusion

In summary, we have introduced a novel unsupervised anomaly detection approach utilizing knowledge distillation principles, with a specific emphasis on reverse distillation. By using thoughtful layer selection and incorporating residual connections, we have achieved exceptional outcomes on the MVTEC anomaly detection dataset, surpassing the performance of previous state-of-the-art methods specifically for these targeted classes. In order to address the challenge of fabric anomaly detection, we conducted

extensive testing of our model using a diverse fabric defect database. Our findings validate the superior performance of our unsupervised fabric defect detection approach. However, the architectural design of our approach is specifically tailored for the problem of plain fabric defect detection and is limited when dealing with big patterns such as drawings. Due to the lack of existing databases tailored to our specific fabric anomaly detection problem, which involves high-resolution color images captured by industrial cameras in real-world scenarios, we have taken the initiative to create a novel dataset. It comprises 2150 images of diverse fabric types, encompassing several commonly encountered defect types prevalent in the fabric industry. Presently, this dataset is limited in terms of the pool of defects gathered since in a fabric factory some defects are rarer than others. It will, therefore, undergo regular updates to incorporate additional classes and to introduce new types of defects. This continuous expansion aims to enhance the dataset diversity and ensure its relevance to evolving fabric anomaly detection research and industry requirements. This dataset will be freely available for the scientific community. (<https://github.com/SimonThomine/IndustrialTextileDataset>)

References

- [1] He, K., Zhang, X., Ren, S., Sun, J.: Deep Residual Learning for Image Recognition. arXiv. <http://arxiv.org/abs/1512.03385> Accessed 2022-09-30
- [2] Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks **60**(6), 84–90 <https://doi.org/10.1145/3065386> . Accessed 2022-11-25
- [3] Tan, M., Le, Q.V.: EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. arXiv. <http://arxiv.org/abs/1905.11946> Accessed 2022-09-30
- [4] Yu, J., Zheng, Y., Wang, X., Li, W., Wu, Y., Zhao, R., Wu, L.: FastFlow: Unsupervised anomaly detection and localization via 2d normalizing flows [2111.07677](https://arxiv.org/abs/2111.07677). Accessed 2021-11-22
- [5] Rudolph, M., Wehrbein, T., Rosenhahn, B., Wandt, B.: Fully convolutional cross-scale-flows for image-based defect detection [2110.02855](https://arxiv.org/abs/2110.02855). Accessed 2021-11-29
- [6] Roth, K., Pemula, L., Zepeda, J., Schölkopf, B., Brox, T., Gehler, P.: Towards total recall in industrial anomaly detection [2106.08265](https://arxiv.org/abs/2106.08265). Accessed 2021-10-14
- [7] Bae, J., Lee, J.-H., Kim, S.: PNI : Industrial Anomaly Detection using Position and Neighborhood Information. arXiv. <http://arxiv.org/abs/2211.12634> Accessed 2023-04-24
- [8] Salehi, M., Sadjadi, N., Baselizadeh, S., Rohban, M.H., Rabiee, H.R.: Multiresolution knowledge distillation for anomaly detection, pp. 14902–14912. https://openaccess.thecvf.com/content/CVPR2021/html/Salehi_Multiresolution_Knowledge_Distillation_for_Anomaly_Detection_CVPR_2021_paper.html Accessed 2023-04-24

- [9] Wang, G., Han, S., Ding, E., Huang, D.: Student-teacher feature pyramid matching for anomaly detection [2103.04257](#). Accessed 2022-05-05
- [10] 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
- [11] Lee, S., Lee, S., Song, B.C.: CFA: Coupled-hypersphere-based Feature Adaptation for Target-Oriented Anomaly Localization. arXiv. <http://arxiv.org/abs/2206.04325> Accessed 2022-09-12
- [12] Thomine, S., Snoussi, H., Soua, M.: MixedTeacher: Knowledge distillation for fast inference textural anomaly detection, pp. 487–494. <https://www.scitepress.org/Link.aspx?doi=10.5220/0011633100003417> Accessed 2023-04-24
- [13] Rezende, D.J., Mohamed, S.: Variational Inference with Normalizing Flows. arXiv. <http://arxiv.org/abs/1505.05770> Accessed 2023-06-05
- [14] Hinton, G., Vinyals, O., Dean, J.: Distilling the Knowledge in a Neural Network. arXiv. <http://arxiv.org/abs/1503.02531> Accessed 2022-09-30
- [15] Ristea, N.-C., Madan, N., Ionescu, R.T., Nasrollahi, K., Khan, F.S., Moeslund, T.B., Shah, M.: Self-supervised predictive convolutional attentive block for anomaly detection [2111.09099](#). Accessed 2022-05-12
- [16] Humeau-Heurtier, A.: Texture feature extraction methods: A survey **7**, 8975–9000 <https://doi.org/10.1109/ACCESS.2018.2890743> . Conference Name: IEEE Access
- [17] Partio, M., Cramariuc, B., Gabbouj, M., Visa, A.: Rock texture retrieval using gray level co-occurrence matrix
- [18] Ojala, T., Pietikainen, M., Maenpaa, T.: Multiresolution gray-scale and rotation invariant texture classification with local binary patterns **24**(7), 971–987 <https://doi.org/10.1109/TPAMI.2002.1017623> . Accessed 2023-04-24
- [19] Mäenpää, T., Pietikäinen, M.: TEXTURE ANALYSIS WITH LOCAL BINARY PATTERNS, 3rd edn., pp. 197–216. WORLD SCIENTIFIC. <https://doi.org/10.1142/9789812775320.0011> . http://www.worldscientific.com/doi/abs/10.1142/9789812775320_0011 Accessed 2023-04-24
- [20] Simon, P., V, U.: Deep learning based feature extraction for texture classification **171**, 1680–1687 <https://doi.org/10.1016/j.procs.2020.04.180> . Accessed 2023-04-24
- [21] Kramer, M.A.: Nonlinear principal component analysis using autoassociative neural networks **37**(2)

- [22] Han, S., Hu, X., Huang, H., Jiang, M., Zhao, Y.: ADBench: Anomaly Detection Benchmark. arXiv. <http://arxiv.org/abs/2206.09426> Accessed 2022-12-01
- [23] Jing, J., Wang, Z., Rättsch, M., Zhang, H.: Mobile-unet: An efficient convolutional neural network for fabric defect detection **92**, 004051752092860 <https://doi.org/10.1177/0040517520928604>
- [24] Xie, H., Zhang, Y., Wu, Z.: An improved fabric defect detection method based on SSD **8**(1), 181–190 <https://doi.org/10.14504/ajr.8.S1.22> . Accessed 2023-06-05
- [25] Mei, S., Wang, Y., Wen, G.: Automatic fabric defect detection with a multi-scale convolutional denoising autoencoder network model **18**(4), 1064 <https://doi.org/10.3390/s18041064> . Accessed 2021-07-07
- [26] Nguyen, Q.P., Lim, K.W., Divakaran, D.M., Low, K.H., Chan, M.C.: GEE: A gradient-based explainable variational autoencoder for network anomaly detection **1903.06661**. Accessed 2021-06-16
- [27] Zavrtnik, V., Kristan, M., Skočaj, D.: DRAEM – A discriminatively trained reconstruction embedding for surface anomaly detection. arXiv. <http://arxiv.org/abs/2108.07610> Accessed 2022-11-21
- [28] Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks **1406.2661**. Accessed 2021-06-15
- [29] Schlegl, T., Seeböck, P., Waldstein, S.M., Langs, G., Schmidt-Erfurth, U.: f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks **54**, 30–44 <https://doi.org/10.1016/j.media.2019.01.010> . Accessed 2021-07-09
- [30] Pourreza, M., Mohammadi, B., Khaki, M., Bouindour, S., Snoussi, H., Sabokrou, M.: G2d: Generate to detect anomaly. In: 2021 IEEE Winter Conference on Applications of Computer Vision (WACV), pp. 2002–2011. IEEE. <https://doi.org/10.1109/WACV48630.2021.00205> . event-place: Waikoloa, HI, USA. <https://ieeexplore.ieee.org/document/9423181/> Accessed 2021-06-15
- [31] Liang, Y., Zhang, J., Zhao, S., Wu, R., Liu, Y., Pan, S.: Omni-frequency channel-selection representations for unsupervised anomaly detection **2203.00259**. Accessed 2022-05-05
- [32] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention Is All You Need. arXiv. <http://arxiv.org/abs/1706.03762> Accessed 2023-06-05
- [33] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby,

- N.: An image is worth 16x16 words: Transformers for image recognition at scale [2010.11929](https://arxiv.org/abs/2010.11929). Accessed 2021-11-25
- [34] Ioffe, S., Szegedy, C.: Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. arXiv. <http://arxiv.org/abs/1502.03167> Accessed 2023-04-24
- [35] Thomine, S., Snoussi, H., Soua, M.: FABLE : Fabric anomaly detection automation process, pp. 487–494
- [36] Bergmann, P., Fauser, M., Sattlegger, D., Steger, C.: MVTec AD — a comprehensive real-world dataset for unsupervised anomaly detection. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 9584–9592. IEEE. <https://doi.org/10.1109/CVPR.2019.00982> . <https://ieeexplore.ieee.org/document/8954181/> Accessed 2022-05-11
- [37] Deutsche Forschungsgemeinschaft, G.: TILDA Textile Texture-Database. <https://lmb.informatik.uni-freiburg.de/resources/datasets/tilda.en.html>
- [38] Silvestre-Blanes, J., Albero-Albero, T., Miralles, I., Pérez-Llorens, R., Moreno, J.: A public fabric database for defect detection methods and results **19**(4), 363–374 <https://doi.org/10.2478/aut-2019-0035>
- [39] Matthias Wieler, F.A.H. Tobias Hahn: Weakly Supervised Learning for Industrial Optical Inspection
- [40] Mishra, P., Verk, R., Fornasier, D., Piciarelli, C., Foresti, G.L.: VT-ADL: A vision transformer network for image anomaly detection and localization. In: 2021 IEEE 30th International Symposium on Industrial Electronics (ISIE), pp. 01–06. <https://doi.org/10.1109/ISIE45552.2021.9576231> . <http://arxiv.org/abs/2104.10036> Accessed 2022-09-16
- [41] Bao, Y., Song, K., Liu, J., Wang, Y., Yan, Y., Yu, H., Li, X.: Triplet-graph reasoning network for few-shot metal generic surface defect segmentation **70**, 1–11 <https://doi.org/10.1109/TIM.2021.3083561> . Conference Name: IEEE Transactions on Instrumentation and Measurement
- [42] Tabernik, D., Šela, S., Skvarč, J., Skočaj, D.: Segmentation-based deep-learning approach for surface-defect detection <https://doi.org/10.1007/s10845-019-01476-x>
- [43] Huang, Q.C.G.Y.W.X.\.Y.K. Y.: Surface Defect Saliency of Magnetic Tile. <https://doi.org/10.1109/coase.2018.8560423>
- [44] Lu, J., Liang, B., Lei, Q., Li, X., Liu, J., Liu, J., Xu, J., Wang, W.: SCueU-net: Efficient damage detection method for railway rail **8**, 125109–125120 <https://doi.org/10.1109/ACCESS.2020.3007603> . Conference Name: IEEE Access

- [45] Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. arXiv. <http://arxiv.org/abs/1412.6980> Accessed 2023-05-03
- [46] Akcay, S., Ameln, D., Vaidya, A., Lakshmanan, B., Ahuja, N., Genc, U.: Anomalib: A Deep Learning Library for Anomaly Detection. arXiv. <http://arxiv.org/abs/2202.08341> Accessed 2022-12-02
- [47] Li, D., Zhang, J., Yang, Y., Liu, C., Song, Y.-Z., Hospedales, T.: Episodic training for domain generalization. In: 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 1446–1455. IEEE. <https://doi.org/10.1109/ICCV.2019.00153> . <https://ieeexplore.ieee.org/document/9008109/> Accessed 2022-11-24
- [48] Wang, S., Yu, L., Li, C., Fu, C.-W., Heng, P.-A.: Learning from Extrinsic and Intrinsic Supervisions for Domain Generalization. arXiv. <http://arxiv.org/abs/2007.09316> Accessed 2022-11-24
- [49] Chen, S.-F., Liu, Y.-M., Lin, C.-C., Chen, T.P.-C., Wang, Y.-C.F.: Domain-Generalized Textured Surface Anomaly Detection. arXiv. <http://arxiv.org/abs/2203.12304> Accessed 2022-11-24