arXiv:2304.06433v1 [cs.CV] 13 Apr 2023

High-Fidelity Zero-Shot Texture Anomaly Localization Using Feature Correspondence Analysis

Andrei-Timotei Ardelean Tim Weyrich
Friedrich-Alexander-Universität Erlangen-Nürnberg
{timotei.ardelean, tim.weyrich}@fau.de

Abstract

We propose a novel method for Zero-Shot Anomaly Localization that leverages a bidirectional mapping derived from the 1-dimensional Wasserstein Distance. The proposed approach allows pinpointing the anomalous regions in a texture with increased precision by aggregating the contribution of a pixel to the errors of all nearby patches. We validate our solution on several datasets and obtain more than a 40% reduction in error over the previous state of the art on the MVTec AD dataset in a zero-shot setting.

1. Introduction

Anomaly Detection (AD) refers to discerning between elements that abide by a standard of normality and those which do not. Humans are generally able to perform this distinction without the need for an explicit guideline for the standard of normality simply by comparing them to items that agree to the standard [41]. Even further, we can often find anomalous regions from visual imagery without previous knowledge of how a certain object or material should look, by simply pinpointing what stands out in a single, isolated sample [28]. This motivates the search for an automatic system able to perform this task, *i.e.*, zero-shot anomaly localization (ZSAL).

Anomaly detection and localization has a wide range of applications. Automatically finding defects during manufacturing, identifying forgeries, detecting situations that require attention in medical imaging, and discovering inaccuracies in industrial machines are just a few of the domains where an anomaly detection system could bring considerable benefits.

The computer vision community has lately shown increased interest in solving the problem of anomaly detection and localization, encouraged by the success of deep learning methods on various tasks. The primary employed strategy is unsupervised learning, modeling normality from a collection of unblemished items. This alleviates the need for labeled anomalous data at training time, which can be difficult to acquire. However, in most existing systems the need for numerous normal samples still remains, and only recently has the more challenging task of few-shot and even zero-shot AL started to be addressed.

Figure 1: Anomaly localization. *Left:* input texture; *right:* anomaly map.

We develop a new system designed specifically for anomaly localization that works in a zero-shot setting, identifying the parts that break the homogeneity of a single textured sample (Figure 1). Our main contribution is a novel method for comparing the statistics between different patches in an image or feature map. To quantify the normality of a pixel location one could trivially compute the average of the nearby features and compare them to a global descriptor, however, as we show, the errors obtained by this method are too coarse for a pixel-level localization of anomalies. We analyze different methods for comparing the local statistics of a patch to a reference and show that one can use a bidirectional mapping that implicitly results from the Wasserstein distance to more precisely identify the offending pixels. This insight is the key element of our Feature Correspondance Analysis (FCA).

2. Related Work

The problem of anomaly detection can be posed for various types of data such as weather records [42], stock market and financial transactions [1], acoustic monitoring [16], video surveillance [24], medical imaging [20], manufacturing inspection [22, 26], etc. In this work, we address the detection of anomalies in images, more exactly detecting anomalous regions in otherwise homogeneous or stationary textures. This can be formulated as a multi-class segmentation and classification of anomalous pixels [11, 31], or in a simpler setting, as a binary separation between normal and anoma-

lous regions, as in [15, 17, 25, 34, 43, 46]. We focus on the latter, usually referred to as anomaly localization (AL) despite dealing with pixel-level segmentation (as opposed to localization understood in the context of object detection). AL can be considered a superclass of the anomaly detection task/classification over images, as an image label can be simply computed as the maximum of pixel-wise anomaly predictions [7]. Therefore, AL is more challenging and, for most purposes, more useful compared to image-level classification, making the result explainable and actionable [41].

In this section, we briefly address the most relevant methods and refer readers to a survey [27, 41] for a broader insight into anomaly localization literature. Most of the early machine learning methods for anomaly detection are reconstruction-based [41], using a (variational) autoencoder [3, 10, 44], or a generative adversarial network (GAN) [2, 5, 35] to learn to synthesize normal images. At inference, reconstruction errors reveal anomalies. These methods are intuitive and provide reasonable results on simple textures given enough normal samples; however, they do not incorporate any priors.

The evolution of AL methods, especially based on deep learning, has greatly accelerated with the introduction of the MVTec AD [8] dataset, which still is the *de facto* benchmarking standard. In the extended version of the MVTec AD paper [7] the authors compare several existing approaches including their own method, Uninformed Students [9], which tries to distill a larger descriptive teacher network into an ensemble of students on normal data only. At inference time, inconsistencies in embeddings from the students and the teacher indicate anomalies.

Uninformed Students is an example of an embedding-based approach. This class of methods generally relies on deep features extracted with the help of a larger network, pretrained on vast amounts of data, that serves as a prior. SPADE [15] demonstrates the effectiveness of neural features, surpassing previous methods without requiring a training stage. Features from a frozen Convolutional Neural Network (CNN), pretrained for classification, are used to find the nearest neighbors within the normal images database and pixel-level correspondences for localization of anomalies. PaDiM [17] uses neural features to characterize each pixel location. The features across the set of normal images are modeled as a multivariate Gaussian allowing for fast anomaly localization at inference using the Mahalanobis distance [29]. Similarly, PatchCore [33] uses extracted features, but compared to PaDiM it is less reliant on image alignment, creating a single memory bank accessible to all patches. Additionally, coreset reduction is used to limit the size of the memory bank and to retain a low computational cost. PatchCore also shows better performance in a few-shot (or low-shot [17]) setting compared to SPADE and PaDiM.

Few-shot anomaly detection was recently explicitly addressed with methods such as normalizing flows [34], hierarchical generative models [37], and feature registration [23]. These methods, however, rely on data augmentation which is problem specific and may require domain knowledge. Moreover, as observed in [4], they are not significantly better compared to, for example, PatchCore [33] which scales better with the number of samples.

Zero-shot anomaly localization represents the extreme case of anomaly detection where the anomalous regions are segmented without a set of unblemished textures to act as guidance. MAEDAY [36] (Nov 2022) claims to have introduced for the first time the task of zero-shot anomaly detection. The method pretrains a transformer-based network which is used to reconstruct a partially masked image at inference. By using this in-painting network, an anomaly score can be computed by identifying the differences between the unmasked image and the reconstructed output. Aota *et al*. [4] (Jan 2023) introduce a new method for zero-shot anomaly detection and localization that uses a large pretrained network as a feature extractor. For each pixel, the local features are averaged and compared to the K nearest neighbors in the input image. Our method is most similar to the latter as it compares local features with globally aggregated information, and it is designed to work on textures and not generic objects as [36]. While explicit efforts to solve ZSAL are a recent endeavor, the task bears similarities to texture stationarity analysis [30], image saliency [12], and weathering estimation [6].

3. Algorithm Design

This section describes the design decisions that went into building our method. We analyze how different components of a zero-shot patch-based anomaly localization system affect its performance, and we also introduce a novel procedure for estimating the anomaly degree at each spatial location.

We consider the following attributes of an AL method, identified as desirable: *high sensitivity at high specificity*, *ability to scale to higher resolutions*, and *fast running time*. Importantly, we focus on a zero-shot scenario and we are therefore mainly interested in textures, which are largely homogeneous, save for the anomalous regions themselves.

As a generic function for the zero-shot anomaly localization problem, we propose the following self-similarity formulation for computing the anomaly map A, given an image I:

$$A(x, y; F, C, R) = \sum_{F_r \in R(F(I))} S(x, y, F(I), F_r) \ . \tag{1}$$

Such an AL system is defined by three different components: feature extraction (F), patch statistics comparison (S), and reference selection (R). Simply put, the anomaly score A at location (x, y) is computed as the sum of the costs when comparing features within one or more patches P containing

PRO ↑ / AUROC ↑	Colors	RandProj	Steerable	VGG
Moments	46.51 / 75.62	40.66 / 73.33	64.21 / 80.78	61.96 / 83.82
Histogram	50.43 / 77.80	53.74 / 80.48	70.64 / 84.43	73.17 / 88.44
SWW	58.62 / 83.62	62.21 / 85.89	**73.08 / 87.77**	77.40 / 91.44
FCA (ours)	**63.30 / 85.76**	**66.28 / 87.62**	71.75 / 86.99	**81.08 / 92.58**

Table 1: Preliminary experiment, comparing our patch statistics method to different baselines. Compared in terms of two metrics: PRO(0.3) and AUROC. The best results are highlighted in bold.

(x, y) with a set of references. We note that the proposed definition is a superset of the discrete form of the stationarity measure introduced in [30]. While not explicitly designed for anomaly localization, by isolating the influence of each spatial location to the stationarity measure in [30] one can use it as an anomaly localization score. The main difference is that Moritz *et al*. assume the reference set R consists of all patches in $F(I)$, which, as we show, is suboptimal.

3.1. Feature Extraction

We evaluate the effect of different feature extractors $F(I) \rightarrow \mathbb{R}^{H \times W \times C}$ and confirm the findings of previous work that neural networks pretrained on ImageNet [18] provide useful features for AL.

The performance of four different options for the feature extractor function F are presented in Table 1. The metrics used for evaluation are detailed in Section 4. We consider using the colors directly ($F(I) = I$), random projections by convolving the image I with a set of random kernels, Steerable Filters [21], and neural features from a simple pretrained VGG19 network [39]. In this preliminary experiment, all feature extractors operate on a single resolution and have a relatively small receptive field, *i.e*., the images are scaled to 256×256, from which feature maps of the same resolution are extracted, with $C \approx 128$ channels (except for colors, where $C = 3$). The simple RGB colors obviously have a receptive field of 1; the random projections are inspired by [19], where they are used in the context of the Sliced Wasserstein Distance, and consist of normalized random 5×5 kernels; we use only one level of steerable filters (full spatial resolution); finally, we use the output of the first and second (3×3) convolutional layer of a pretrained VGG network, having an effective receptive field of 5×5. We use the same patch size of 25×25 for all stationarity measures, which is, by inspection, large enough to capture the difference in appearance between normal and anomalous regions.

As shown in Table 1, the embeddings obtained from the VGG network consistently outperform the other types of features.

3.2. Patch Statistics Comparison

The function $S(x, y, F(I), F_r)$ evaluates the degree of anomaly at position (x, y) given its local context in the feature maps $F(I)$, by comparing with the reference F_r. The function should analyze how do the local statistics around (x, y) differ from the statistics in F_r. In this subsection we describe different options for S, together with their limitations, and introduce our Feature Correspondence Analysis (FCA) method for comparing patch statistics.

Moments. In general, only a small region around a certain location is needed to identify an anomaly. This leads to a trivial patch statistics comparison method, computed by averaging the features around (x, y), *i.e*.,

$$S(x, y) = \left\| \frac{1}{T^2} \sum_{(x', y') \in P_{xy}} F(I)(x', y') - \text{avg}(F_r) \right\|_2^2 , \qquad (2)$$

where $F(I)$ and F_r have been omitted from S for brevity, and P_{xy} denotes a patch of size $T \times T$ centered in (x, y). The definition can be easily extended to include spatial weighting (*e.g*., Gaussian) and moments of higher order, becoming equivalent to the method of moments from [30] when using RGB colors directly as features.

Histogram. Moritz *et al*. [30] propose another two options for computing the stationarity measure, which can be described in our conceptual framework as using a histogram-based patch statistics comparison over RGB colors, and steerable filters, respectively. The histogram-based algorithm can be described as:

$$S(x, y) = \text{hist}\Big(\bigcup_{(x', y') \in P_{xy}} F(I)(x', y') \Big) \ominus \text{hist}(F_r) , \qquad (3)$$

where $\ominus$ gives the earth mover's (Wasserstein) distance between the two histograms. As in the case of moments, when computing the histogram one can employ spatial weighting to increase the importance of the pixels closer to (x, y).

Sample-weighted Wasserstein (SWW). The previous methods have limited expressive powers, specifically because they consider the distribution inside a patch as a whole, unable to pinpoint "outlieredness" of individual samples.

That ability conveniently occurs in an efficient implementation of the 1-D Wasserstein distance when operating on individual samples drawn from distributions. If two sets of samples have the same size, the Wasserstein distance can be obtained by sorting the samples and then summing over the absolute differences between the elements corresponding to the same rank [19]. That comparison of samples of the same rank within a sorting can be seen as a bijective mapping

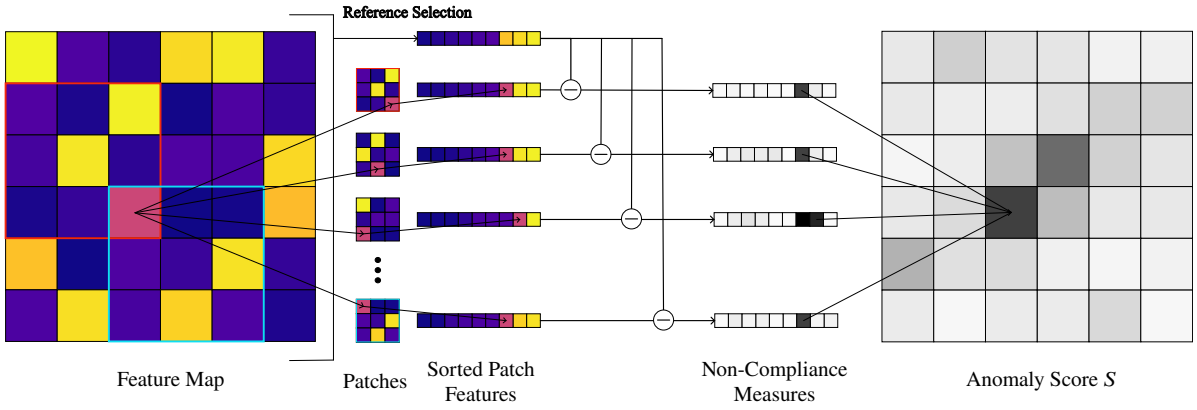

Figure 2: Depiction of our Feature Correspondance Analysis (FCA). The $\ominus$ denotes the absolute difference.

between two sample sets, and the difference between corresponding samples is an immediate measure for those samples' non-compliance with the respective other distribution. That resulting non-compliance calculation translates directly to a coarse anomaly measure of the corresponding coefficient channel for each pixel; summing them across channels results in an error score $M(x, y; P)$ for each pixel (x, y) in a patch P.

Subsequently, we aggregate the per-patch error map M into an anomaly measure for the center of the patch P. At this point, averaging $M(\cdot, \cdot; P)$ would yield the exact Wasserstein distance, and it would be equivalent to the previously defined histogram method (with bins $\to \infty$). Instead, we use a Gaussian-weighted average to increase the spatial sensitivity of the resulting anomaly score $S(x, y)$.

Upon cursory observation, this may resemble the weighted, sliding-window histogram calculations of Moritz *et al*. [30]; however, Moritz *et al*. compute weighted distributions before calculating their metric, whereas we preserve the original patch distribution but weight the influence of each sample's non-compliance score on the final anomaly score. The equation for this sample-weighted Wasserstein method is:

$$S(x, y) = \sum_{(x', y') \in P_{xy}} M(x', y'; P_{xy}) G_{\sigma_w}(x' - x, y' - y) \,, \quad (4)$$

where $G_{\sigma_w}(\Delta_x, \Delta_y)$ is a spatial weighting function, for which we use a Gaussian with variance σ_w^2.

FCA. The previous definition of SWW allows us to separate the context size and the amount of smoothing in the aggregation through the parameter of the Gaussian; however, the limitation is that the final anomaly score for any location (x, y) uses as context only the patch P_{xy}. We further leverage the bijective mapping from SWW by computing the anomaly score at location (x, y) as the sum of the matching errors for position (x, y) in the context of all surrounding patches, which gives:

$$S(x, y) = \sum_{(x', y') \in P_{xy}} M(x, y; P_{x'y'}) G_{\sigma_p}(x' - x, y' - y) \,. \quad (5)$$

Please note the change of parameters in M compared to the SWW equation (4). The main difference is that instead of considering one context patch P_{xy} when computing $S(x, y)$, we consider all patches that contain (x, y), and aggregate the contribution of the location (x, y) in all of these contexts. Anomalies are generally considered smooth and all available datasets present anomalies as binary blobs that mark anomalous regions, rather than continuous scores depicting the contribution of each pixel to the anomaly. To attend to this, we introduce Gaussian smoothing $\mathcal{G}_{\sigma_s}$ after matching errors, yielding the final formula:

$$S(x, y) = \sum_{(x', y') \in P_{xy}} \mathcal{G}_{\sigma_s}\big(M(\cdot, \cdot; P_{x'y'})\big)(x, y)\, G_{\sigma_p}(x' - x, y' - y) \,. \quad (6)$$

The workflow of the algorithm is illustrated in Figure 2. We name this novel method Feature Correspondence Analysis (FCA), as it computes the anomaly score based on the correspondence of features from patches to a reference.

In Figure 3, we showcase the effect of the proposed method on an artificial problem. We run FCA without smoothing (equation 5) to show how our formulation allows significantly better localization of the source of the error when comparing the patch features statistics to the reference. While running the histogram method with a small patch size would improve the first result, it would fail in the second example because it contains a contextual (also called conditional [40]) anomaly.

We also compare the Histogram method and our FCA on a real case example from MVTec AD [7], in Figure 4.

3.3. Reference Selection

We analyze several options for the set of references $R(F(I))$. An intuitive solution is to use all the patches in the image as references, however, this amounts to computing the pairwise distances between all patches in an image which can be very time-consuming, scaling poorly as the image resolution increases. Choosing a single patch at random is fast but is a poor approximation of the global statistics.

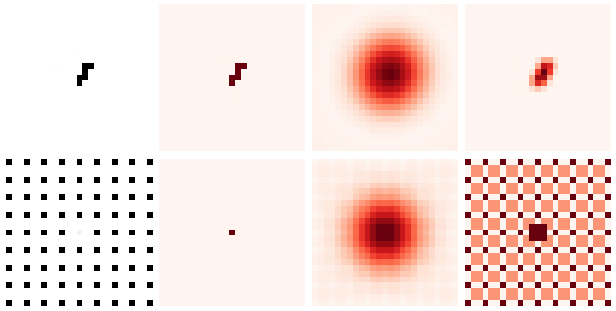

Figure 3: Anomaly localization maps when matching patch statistics with our FCA compared to histograms. Visualized on 2 synthetic examples.

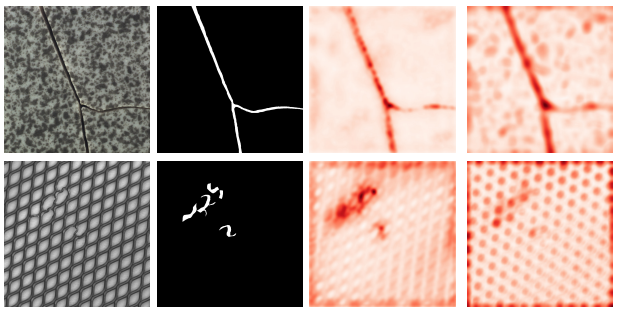

Figure 4: Anomaly localization maps when matching patch statistics with our FCA compared to histograms. Visualized on 2 textures from MVTec AD [8].

One can alternatively use one reference that aggregates the global information (*e.g.*, global average for moments, and the histogram over the whole feature map, for the histogram-based patch statistics comparison). In the case of SWW and FCA, we choose F_r as:

$$\underset{F_r}{\operatorname{argmin}} \sum_{(x,y)} A(x, y; \cdot, \cdot, R = \{F_r\}) \ . \tag{7}$$

The feature set that minimizes the Wasserstein distance across all patches has a closed-form solution, obtained by taking the median over the features at each sorted position individually, *i.e.*, compute the median for each order statistic for each feature channel. We analyze the performance of the global statistic aggregation method and the trade-off between the number of random patches used and performance in Table 2.

Using the median works well when the texture is homogeneous but struggles to capture the global statistics for multimodal textures (*e.g.*, structured textures with the period larger than the patch size). To avoid this issue, one can use the pairwise distances and discard the outliers by considering only the closest K distances. In this case, R selects the K

PRO / Time (s)	Hist	SWW	FCA
Random (1)	60.95 / 1.1	59.56 / 5.7	62.01 / 9.2
Random (3)	67.04 / 1.1	66.79 / 8.2	69.57 / 18.7
Random (10)	72.99 / 1.2	73.23 / 17.3	75.91 / 52.3
Random (100)	74.55 / 1.7	75.69 / 134	78.45 / 482
All	74.01 / 380	– / 84984	– / 314577
Mean/Median	73.17 / 1.1	77.40 / 5.7	81.08 / 9.2

Table 2: Analysis of the effect of the Reference Selection method. We report the PRO (0.3) metric as well as the running time (in seconds) per image. Variants that would be unreasonably slow to be used in practice were marked with "–", and only the time was reported.

nearest neighbors (KNN) over all patches in the feature maps, with respect to the cost $S(x, y; F(I), F_r)$. We only report results using KNN references in Section 4.3, when running on low-resolution feature maps, due to the high running time of this method. Notably, employing a WideResnet-50 [45] as feature extractor, using the first moment for patch statistics comparison, and taking the KNN for reference selection yields a system equivalent to Aota *et al.* [4].

3.4. Final Method and Implementation Details

In accordance with the observations made in this section, we design our final anomaly localization system to use neural features from a pretrained neural network, evaluate the local statistics with the newly introduced FCA, and use the median for reference selection as a balance between fidelity and speed. Following recent work on anomaly detection [15, 17, 33], we use a WideResnet-50 network [45] and extract the features computed by the second convolutional block, yielding feature maps with 512 channels. Because the output of this block has a resolution 8 times smaller than the input, and FCA can handle relatively large context sizes, we choose to run the method at full resolution and not resize it as a preprocessing step as done in previous work [4, 17, 33]. All patch statistics comparison variants, including our FCA, have been implemented in PyTorch [32], utilizing CUDA acceleration, and ran on an NVIDIA RTX A5000 GPU. We use the same hyperparameters for all experiments, setting $\sigma_p = 3.0, \sigma_s = 1.0$. The patch size T should be set depending on the size of the feature maps. We use $T = 9$ when running at full dataset resolution and $T = 3$ for consistency with Aota *et al.* [4] when running at 320×320.

4. Experiments

We compare our approach with state-of-the-art methods in zero-shot anomaly detection as well as a few other adapted baselines. Several datasets are considered in order to assess the robustness of the proposed approach.

4.1. Datasets

MVTec AD. Currently, the dataset most used in the context of anomaly classification and localization is the MVTec AD dataset [7, 8]. It contains 15 classes with about 5 different defects per class. We use the 5 texture classes, accumulating over 500 test images and their (manually annotated) ground truth binary segmentation masks. The resolution of these images ranges from 840×840 to 1024×1024 pixels. Past works [4, 7, 33, 36] propose various preprocessing and post-processing setups, consisting of resizing and cropping to various resolutions. For a fair evaluation, we compute the metrics at full resolution, following the original evaluation script from the dataset provider [7]. The only adaption performed is cropping to the center regions before evaluation to avoid computing metrics on the edges of the images where most methods do not provide reliable scores [4].

Woven Fabric Textures. Bergmann *et al*. [10] introduced a small dataset for the task of defect segmentation containing two woven fabric textures (denoted WFT from here on). For each of them, 50 test images with various types of anomalies are provided. The resolution of the images is 512×512 and ground truth segmentation masks are also included.

DTD-Synthetic. Aota *et al*. [4] constructed an artificial dataset to evaluate Anomaly Detection methods on more diverse data, including anisotropic textures. The dataset is based on the Describable Texture Dataset [14] on which various types of defects were artificially added. The textures are also randomly rotated and cropped, eventually yielding 1304 images of small resolution (180×180 to 384×384). Following [4], all images are resized to a fixed 320×320.

Aitex. The Aitex dataset [38] contains uniform fabric textures. The defects have been manually annotated in the original images of size 256×4096. Following standard practice, we split the images into square pieces. Additionally, we discard all frames that are not completely covered by the texture and images that do not contain any anomaly. For consistency with [4], we resize the images to 320×320 before evaluating both methods.

4.2. Metrics

The main metric for anomaly localization is the threshold-independent AUROC (area under the receiver operating characteristic curve). One limitation of this metric is that the size of the anomalous regions is not taken into consideration, making it less sensitive to spatially small anomalies. To account for this, [8] introduced the PRO metric which weights the size of each anomalous region. Moreover, they only compute the integral up to a False Positive Rate of 30% to avoid computing metrics on degenerate results. Since the purpose of the proposed method is to obtain a more accurate and detailed segmentation of the anomalies, the PRO metric is our most important indicator. Our main contribution deals with anomaly localization and does not focus on improving the image-level label (computed as the maximum across the anomaly scores map). Therefore, the anomaly classification metric, AUROC$_c$, is only reported on the MVTec AD dataset and omitted for the other experiments.

4.3. Results

We compare our final method against several existing methods for zero-shot and few-show anomaly localization.

The results on the MVTec AD dataset are reported in Table 3 below. As the zero-shot anomaly segmentation task is particularly new, there are very few methods designed explicitly for this task. Concretely we compare our system against MAEDAY, an image-reconstruction-based zero-shot AL [36], and [4] – based on a Wide-Resnet-50 feature extractor, using a simple average for patch statistics comparison, combined with a KNN search. We additionally adapt methods that were not explicitly designed for ZSAL but are related in scope, for a more complete comparison. Bellini *et al*. [6] propose a method for weathering arbitrary textures from a single image, and uses an age-estimation procedure as the first step in their pipeline. The age-estimation procedure targets the same goal, to highlight regions in an image that stray away from the pristine appearance. Saliency-RC [12] as a saliency detection method highlights parts of the image that stand out, which is related to anomaly localization. However, as mentioned by the authors, the method's performance on textures is limited.

	PRO (0.3)	AUROC	AUROC$_c$
Saliency [12]	23.88	60.40	46.00
Bellini *et al*. [6]	51.26	76.36	27.17
MAEDAY† [36]	–	75.2†	88.9†
Aota *et al*. [4]	93.82	97.47	**99.67**
PatchCore 1-shot	76.72	88.65	90.79
PatchCore all	92.42	97.61	99.44
PatchCore all†	93.64†	97.52†	98.96†
Ours$_{320}$	95.46	97.74	99.21
Ours$_{320}$ + KNN	95.58	97.77	99.17
Ours	**97.18**	**98.73**	99.58

Table 3: Quantitative comparison on the MVTec AD dataset. We note with † the results that are taken from different papers as reported by the authors and can differ slightly from the evaluation used for the other scores, as discussed in Section 4.1. The subscript $_{320}$ marks running our method at the lower resolution. The 1-shot results from PatchCore were averaged across 10 runs.

Please note that the proposed method improves the localization of anomalous textures significantly compared to the previous state-of-the-art zero-shot method of Aota *et al*. [4], thanks to the use of a more precise method for comparing patch statistics. Moreover, our system also outperforms one

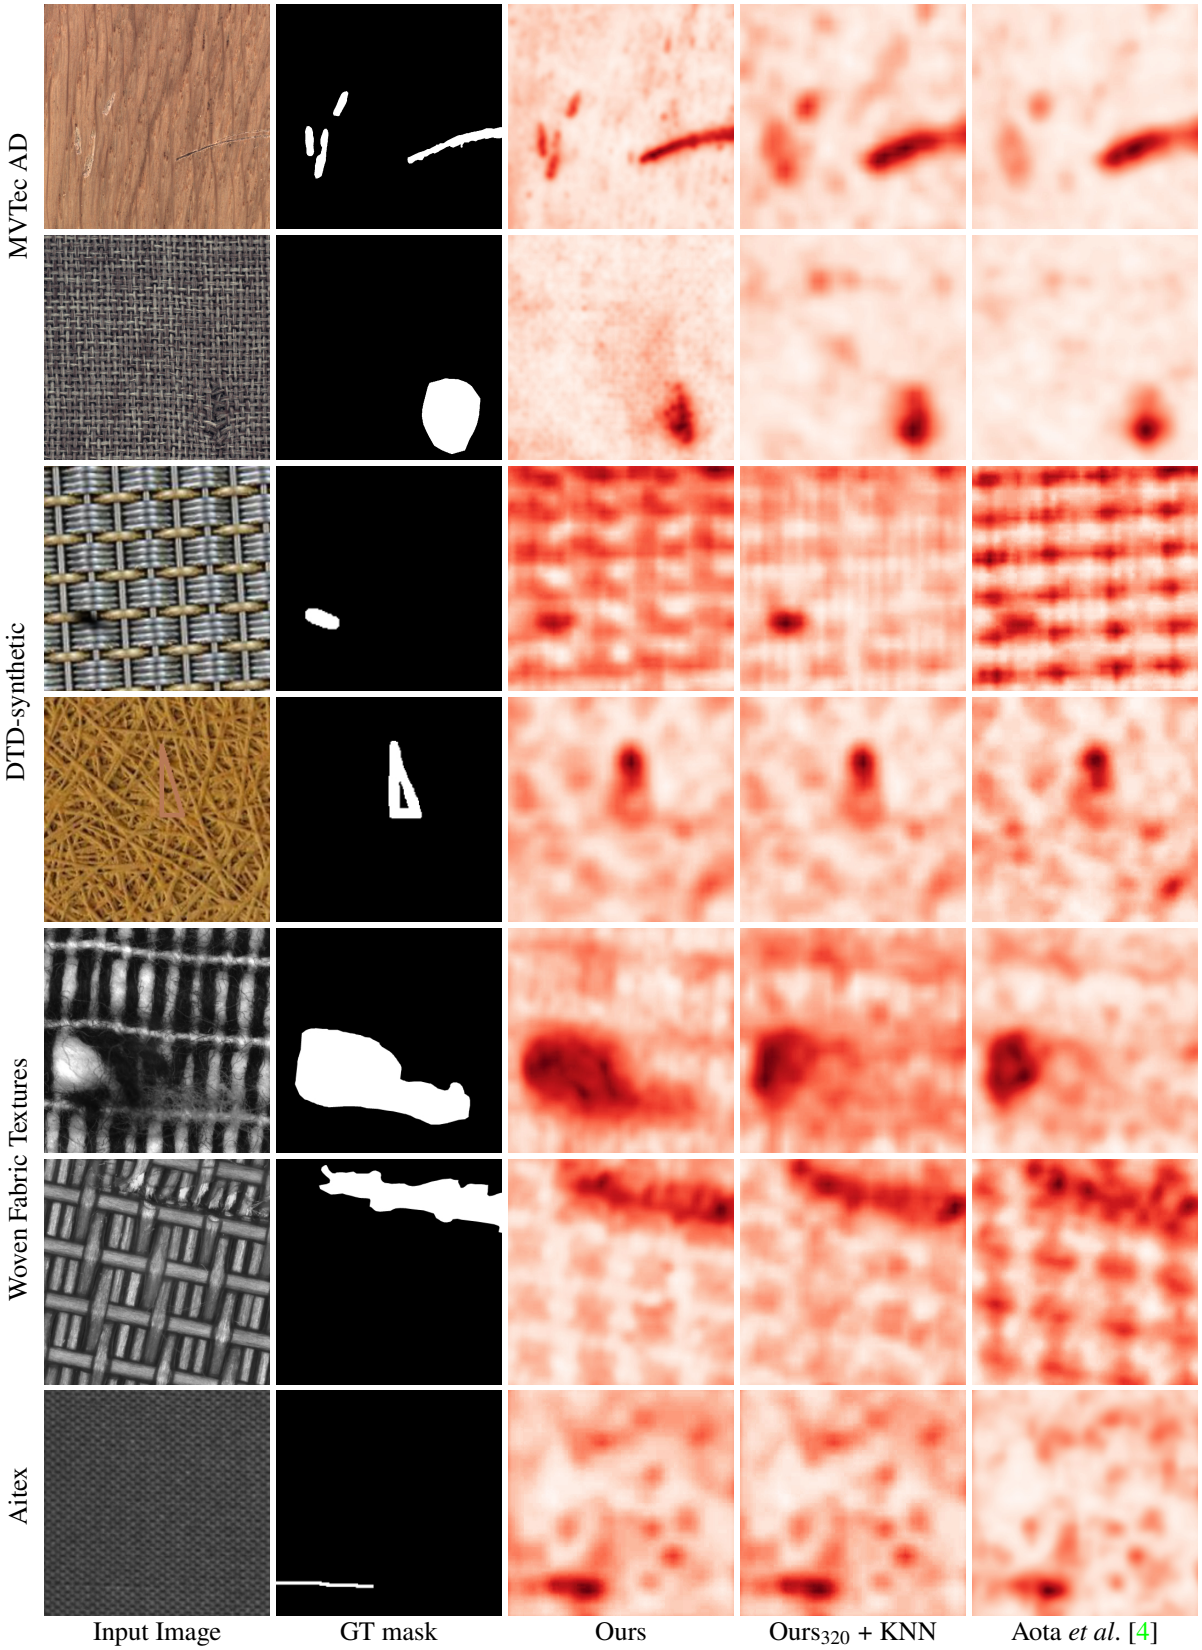

Figure 5: Qualitative comparison on challenging examples. The images are shown after cropping to center.

of the most prominent few-shot anomaly detection methods, PatchCore [33].

	PRO (0.3)	AUROC	Time [s]
DTD-Synthetic			
Aota *et al*. [4]	94.32	98.00	1.1
Ours	94.71	98.03	0.71
Ours + KNN	**95.93**	**98.51**	4.5
WFT			
Aota *et al*. [4]	84.59	96.11	1.1
Ours_{320}	73.00	92.82	0.71
Ours_{320} + KNN	86.24	96.19	4.5
Ours	**89.57**	**98.26**	0.88
Aitex			
Aota *et al*. [4]	87.11	96.70	1.1
Ours	91.07	97.51	0.71
Ours + KNN	**91.24**	**97.52**	4.5

Table 4: Quantitative comparison on DTD-Synthetic, Woven Fabric Textures (WFT), and Aitex.

We additionally run experiments on the DTD-Synthetic dataset from [4], Woven Fabric Textures from [10], and Aitex [38]. In these experiments we only compare with the identified leading method in ZSAL, *i.e*., [4]. The results are presented in Table 4 and show that our method consistently improves upon prior art in both PRO (0.3) and AUROC metrics. We find that using KNN for reference selection can, in some cases, significantly improve our results at the cost of a higher running time. On the other hand, using the median as reference allows our method to process images at higher resolution. Importantly, as the resolution of the images increases, the value added by our FCA also grows. This can be seen in the results on the WFT (512×512) and MVTec AD (1024×1024) datasets where running our method at full resolution outperforms the lower resolution + KNN variant.

In Figure 5 we present a qualitative comparison between our method and the zero-shot state-of-the-art method of Aota *et al*. [4]. We show the anomaly predictions on challenging samples from each dataset. In general, the anomaly maps produced by our method have high fidelity, with more precise localization (rows 1, 2, 4), fewer false positives (rows 3, 6), and more complete coverage of the anomalous regions (rows 5, 7) compared to [4]. For more visualizations please see the supplementary material.

5. Discussion

The results suggest the proposed method is effective in predicting anomaly scores with high fidelity. While FCA generally performs better compared to other methods, it also has a higher complexity. Computing local moments or histograms can be done efficiently thanks to the separability of the Gaussian kernel. This does not apply to SWW which requires sorting the values inside each sliding window.

Table 5 reports the computational complexity of various methods for anomaly localization. Due to the downsampling in feature extractors, the resolution of the maps is rather small, which in turn result in small values for $T < 10$ in our experiments. In practice, for most methods the running time is dominated by running the WideResnet-50 feature extractor. The notable exception is Ours + KNN which has a running time about 7 times larger compared to the single reference counterpart when run at the same resolution.

Method	Complexity
Moments	$O(NTD)$
Histogram	$O(NTBD)$
Aota *et al*. [4]	$O(NTD + N^2D)$
Ours	$O(NT^2 \log(T)D)$
Ours + KNN	$O(N^2T^2 \log(T)D)$

Table 5: Complexity analysis. D: number of features; T^2: patch area; N: image pixel count; B: bins per histogram. For KNN, brute force complexity is used, as the number of neighbours K is in the order of N ($K = N \div 4$ in [4]).

Due to the large and varied datasets, our experiments arguably support these findings robustly; however, we observe that the manually defined ground-truth images of MVTec AD, WFT, and Aitex inevitably introduce a level of subjectivity to those ground-truth references. In some cases, as for instance shown in Figure 6, significantly different ground truth interpretations would have been possible, relativizing the accuracy score in such cases.

A limitation of our system is that, by design, it only works on textures. Generic objects can have very different feature statistics in different regions which would not be handled correctly by our method.

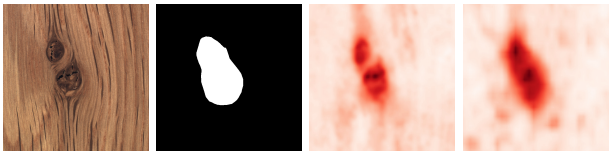

Input Image Mask Ours Aota *et al*. [4]

Figure 6: Manual ground-truth annotations remain subjective where multiple plausible interpretations exist.

6. Conclusion

In this work, we put forward a generic framework for performing zero-shot anomaly localization. We identify the importance of the different components and suggest a new approach that significantly improves upon prior art. The most important novelty is the proposed FCA for patch statistics comparison which enables high-fidelity anomaly localization that scales well with large textures. The performance of the method is validated on several datasets offering a comprehensive overview of the advantages of the method and the trade-off between running time and prediction quality.

Acknowledgement. This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 956585 (PRIME ITN).

References

[1] Mohiuddin Ahmed, Abdun Naser Mahmood, and Md Rafiqul Islam. A survey of anomaly detection techniques in financial domain. *Future Generation Computer Systems*, 55:278–288, 2016. 1

[2] Samet Akçay, Amir Atapour-Abarghouei, and T. Breckon. Skip-ganomaly: Skip connected and adversarially trained encoder-decoder anomaly detection. *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2019. 2

[3] Jinwon An and Sungzoon Cho. Variational autoencoder based anomaly detection using reconstruction probability. *Special lecture on IE*, 2(1):1–18, 2015. 2

[4] Toshimichi Aota, Lloyd Teh Tzer Tong, and Takayuki Okatani. Zero-shot versus many-shot: Unsupervised texture anomaly detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5564–5572, 2023. 2, 5, 6, 7, 8, 11, 12

[5] Christoph Baur, Benedikt Wiestler, Shadi Albarqouni, and Nassir Navab. Deep autoencoding models for unsupervised anomaly segmentation in brain mr images. In *BrainLes@MICCAI*, 2018. 2

[6] Rachele Bellini, Yanir Kleiman, and Daniel Cohen-Or. Time-varying weathering in texture space. *ACM Transactions on Graphics (TOG)*, 35(4):1–11, 2016. 2, 6

[7] Paul Bergmann, Kilian Batzner, Michael Fauser, David Sattlegger, and Carsten Steger. The mvtec anomaly detection dataset: a comprehensive real-world dataset for unsupervised anomaly detection. *International Journal of Computer Vision*, 129(4):1038–1059, 2021. 2, 4, 6

[8] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mvtec ad–a comprehensive real-world dataset for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9592–9600, 2019. 2, 5, 6, 11

[9] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4183–4192, 2020. 2

[10] Paul Bergmann, Sindy Löwe, Michael Fauser, David Sattlegger, and Carsten Steger. Improving unsupervised defect segmentation by applying structural similarity to autoencoders. In *VISIGRAPP*, 2018. 2, 6, 8, 11

[11] Rozi Bibi, Yousaf Saeed, Asim Zeb, Taher M Ghazal, Taj Rahman, Raed A Said, Sagheer Abbas, Munir Ahmad, and Muhammad Adnan Khan. Edge ai-based automated detection and classification of road anomalies in vanet using deep learning. *Computational intelligence and neuroscience*, 2021:1–16, 2021. 1

[12] Ming-Ming Cheng, Niloy J Mitra, Xiaolei Huang, Philip HS Torr, and Shi-Min Hu. Global contrast based salient region detection. *IEEE transactions on pattern analysis and machine intelligence*, 37(3):569–582, 2014. 2, 6

[13] N Chinchor. Muc-4 evaluation metrics in proc. of the fourth message understanding conference 22–29, 1992. 11

[14] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3606–3613, 2014. 6

[15] Niv Cohen and Yedid Hoshen. Sub-image anomaly detection with deep pyramid correspondences. *arXiv preprint arXiv:2005.02357*, 2020. 2, 5

[16] Clayton Cooper, Jianjing Zhang, Robert X Gao, Peng Wang, and Ihab Ragai. Anomaly detection in milling tools using acoustic signals and generative adversarial networks. *Procedia Manufacturing*, 48:372–378, 2020. 1

[17] Thomas Defard, Aleksandr Setkov, Angelique Loesch, and Romaric Audigier. Padim: a patch distribution modeling framework for anomaly detection and localization. In *Pattern Recognition. ICPR International Workshops and Challenges: Virtual Event, January 10–15, 2021, Proceedings, Part IV*, pages 475–489. Springer, 2021. 2, 5

[18] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 3

[19] Ariel Elnekave and Yair Weiss. Generating natural images with direct patch distributions matching. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVII*, pages 544–560. Springer, 2022. 3

[20] Tharindu Fernando, Harshala Gammulle, Simon Denman, Sridha Sridharan, and Clinton Fookes. Deep learning for medical anomaly detection–a survey. *ACM Computing Surveys (CSUR)*, 54(7):1–37, 2021. 1

[21] William T Freeman, Edward H Adelson, et al. The design and use of steerable filters. *IEEE Transactions on Pattern analysis and machine intelligence*, 13(9):891–906, 1991. 3

[22] Ruei-Jie Hsieh, Jerry Chou, and Chih-Hsiang Ho. Unsupervised online anomaly detection on multivariate sensing time series data for smart manufacturing. In *2019 IEEE 12th conference on service-oriented computing and applications (SOCA)*, pages 90–97. IEEE, 2019. 1

[23] Chaoqin Huang, Haoyan Guan, Aofan Jiang, Ya Zhang, Michael Spratling, and Yan-Feng Wang. Registration based few-shot anomaly detection. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIV*, pages 303–319. Springer, 2022. 2

[24] Phulpreet Kaur, M Gangadharappa, and Shalu Gautam. An overview of anomaly detection in video surveillance. In *2018 International Conference on Advances in Computing, Communication Control and Networking (ICACCCN)*, pages 607–614. IEEE, 2018. 1

[25] Chun-Liang Li, Kihyuk Sohn, Jinsung Yoon, and Tomas Pfister. Cutpaste: Self-supervised learning for anomaly detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9664–9674, 2021. 2

[26] Jiaqi Liu, Guoyang Xie, Jingbao Wang, Shangnian Li, Chengjie Wang, Feng Zheng, and Yaochu Jin. Deep industrial image anomaly detection: A survey. *arXiv e-prints*, pages arXiv–2301, 2023. 1

[27] Jiaqi Liu, Guoyang Xie, Jingbao Wang, Shangnian Li, Chengjie Wang, Feng Zheng, and Yaochu Jin. Deep industrial image anomaly detection: A survey, 2023. 2

[28] David Lowe. *Perceptual organization and visual recognition*, volume 5. Springer Science & Business Media, 2012. 1

[29] PC Mahalanobis. On the generalized distances in statistics: Mahalanobis distance. *Journal Soc. Bengal*, 26:541–588, 1936. 2

[30] Joep Moritz, Stuart James, Tom S.F. Haines, Tobias Ritschel, and Tim Weyrich. Texture stationarization: Turning photos into tileable textures. *Computer Graphics Forum (Proc. Eurographics)*, 36(2):177–188, 2017. 2, 3, 4

[31] Hugo S Oliveira, João F Teixeira, and Hélder P Oliveira. Lightweight deep learning pipeline for detection, segmentation and classification of breast cancer anomalies. In *Image Analysis and Processing–ICIAP 2019: 20th International Conference, Trento, Italy, September 9–13, 2019, Proceedings, Part II 20*, pages 707–715. Springer, 2019. 1

[32] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 5

[33] Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Schölkopf, Thomas Brox, and Peter Gehler. Towards total recall in industrial anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14318–14328, 2022. 2, 5, 6, 8

[34] Marco Rudolph, Bastian Wandt, and Bodo Rosenhahn. Same same but differnet: Semi-supervised defect detection with normalizing flows. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1907–1916, 2021. 2

[35] Thomas Schlegl, Philipp Seeböck, Sebastian M. Waldstein, Ursula Margarethe Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *Information Processing in Medical Imaging*, 2017. 2

[36] Eli Schwartz, Assaf Arbelle, Leonid Karlinsky, Sivan Harary, Florian Scheidegger, Sivan Doveh, and Raja Giryes. MAEDAY: MAE for few and zero shot AnomalY-detection. 2, 6

[37] Shelly Sheynin, Sagie Benaim, and Lior Wolf. A hierarchical transformation-discriminating generative model for few shot anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8495–8504, 2021. 2

[38] Javier Silvestre-Blanes, Teresa Albero-Albero, Ignacio Miralles, Rubén Pérez-Llorens, and Jorge Moreno. A public fabric database for defect detection methods and results. *Autex Research Journal*, 19(4):363–374, 2019. 6, 8, 11

[39] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3

[40] Xiuyao Song, Mingxi Wu, Christopher Jermaine, and Sanjay Ranka. Conditional anomaly detection. *IEEE Transactions on knowledge and Data Engineering*, 19(5):631–645, 2007. 4

[41] Xian Tao, Xinyi Gong, Xin Zhang, Shaohua Yan, and Chandranath Adak. Deep learning for unsupervised anomaly localization in industrial images: A survey. *IEEE Transactions on Instrumentation and Measurement*, 2022. 1, 2

[42] S Wibisono, M Anwar, Aji Supriyanto, and I Amin. Multivariate weather anomaly detection using dbscan clustering algorithm. *Journal of Physics: Conference Series*, 1869:012077, 04 2021. 1

[43] Jihun Yi and Sungroh Yoon. Patch svdd: Patch-level svdd for anomaly detection and segmentation. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, November 2020. 2

[44] Sanyapong Youkachen, Miti Ruchanurucks, Teera Phatrapomnant, and Hirohiko Kaneko. Defect segmentation of hot-rolled steel strip surface by using convolutional auto-encoder and conventional image processing. *2019 10th International Conference of Information and Communication Technology for Embedded Systems (IC-ICTES)*, pages 1–5, 2019. 2

[45] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In Edwin R. Hancock Richard C. Wilson and William A. P. Smith, editors, *Proceedings of the British Machine Vision Conference (BMVC)*, pages 87.1–87.12. BMVA Press, September 2016. 5

[46] Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Draem-a discriminatively trained reconstruction embedding for surface anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8330–8339, 2021. 2

Supplementary Material

A. Detailed Quantitative Comparison

In Table 6 we present the detailed breakdown of the metrics into texture classes for all datasets. We report the PRO(0.3) and AUROC metrics as introduced in the main paper. Additionally, the per-class F_1 score [13] corresponding to the (F_1-)optimal threshold is included.

B. Additional Qualitative Comparison

We visually compare the anomalous regions extracted by our method and Aota *et al*. [4] by thresholding the anomaly maps. We compute the optimal threshold with respect to the F_1 measure for each texture class individually and display the anomalous area in the original texture. This visualization is included in Figure 7.

	Ours			**Ours + KNN**			**Aota *et al*.**[4]		
MVTec AD [8]	PRO	AUROC	F1	PRO	AUROC	F1	PRO	AUROC	F1
carpet	95.44	98.30	**72.58**	**96.92**	98.81	71.53	96.13	**98.83**	69.97
grid	**98.07**	**99.46**	**61.62**	97.77	99.27	52.41	97.01	99.12	51.89
leather	98.90	99.45	**66.06**	**98.92**	**99.52**	60.26	98.13	99.47	58.74
tile	**96.33**	**98.22**	**82.16**	88.95	94.31	65.55	84.29	93.41	63.74
wood	**97.18**	**98.22**	**76.34**	95.22	96.98	62.42	93.56	96.54	58.68
DTD-Synthetic [4]	PRO	AUROC	F1	PRO	AUROC	F1	PRO	AUROC	F1
Blotchy_099	98.67	99.55	78.92	**98.73**	**99.57**	**79.50**	97.60	99.19	69.56
Marbled_078	97.97	99.33	76.25	**98.05**	**99.37**	**76.60**	96.51	98.80	66.11
Mesh_114	94.44	97.63	65.46	**95.91**	**98.16**	**66.95**	95.14	97.75	64.16
Stratified_154	98.78	99.15	66.67	**98.81**	99.19	**66.80**	98.53	**99.25**	64.48
Woven_068	97.24	98.86	70.51	**97.31**	**98.92**	**70.72**	95.34	98.29	65.61
Woven_125	98.51	99.52	77.03	**98.66**	**99.56**	**77.59**	97.00	98.99	67.50
Fibrous_183	96.95	98.96	72.82	**97.21**	**99.06**	**73.29**	94.42	98.20	65.42
Matted_069	**89.43**	99.33	76.15	88.89	**99.37**	**76.60**	89.34	99.17	68.64
Perforated_037	94.55	96.60	64.66	**96.62**	**97.76**	**68.78**	95.74	97.05	67.41
Woven_001	94.73	98.93	66.17	**97.79**	**99.59**	**68.96**	96.51	99.42	63.59
Woven_104	89.98	96.84	64.93	**90.40**	**96.96**	**66.05**	89.78	96.60	65.28
Woven_127	85.26	91.63	58.88	**92.82**	**94.62**	**69.58**	85.96	92.22	63.81
WFT [10]	PRO	AUROC	F1	PRO	AUROC	F1	PRO	AUROC	F1
texture_1	**92.38**	**97.98**	**80.34**	85.36	94.25	68.46	89.01	95.94	73.12
texture_2	86.77	**98.56**	**77.91**	**87.12**	98.13	77.05	80.16	96.28	71.02
Aitex [38]	PRO	AUROC	F1	PRO	AUROC	F1	PRO	AUROC	F1
t_00	**78.00**	**95.08**	49.41	75.19	94.08	**49.56**	62.91	90.21	45.33
t_01	76.85	91.28	69.62	**80.08**	**92.34**	**73.47**	70.71	91.31	71.04
t_02	96.34	**99.32**	56.30	**96.57**	**99.33**	**56.93**	92.92	99.02	53.45
t_03	88.77	97.94	**68.04**	**89.44**	**97.97**	67.75	87.43	97.59	65.83
t_04	**99.39**	**99.79**	**72.44**	**99.39**	**99.79**	72.23	97.89	99.75	72.04
t_05	**98.15**	**99.17**	**49.34**	98.04	99.11	47.66	97.95	99.05	43.41
t_06	**99.97**	**99.99**	71.60	99.97	**99.99**	70.73	**99.98**	**99.99**	**76.54**

Table 6: Metrics breakdown into texture classes.

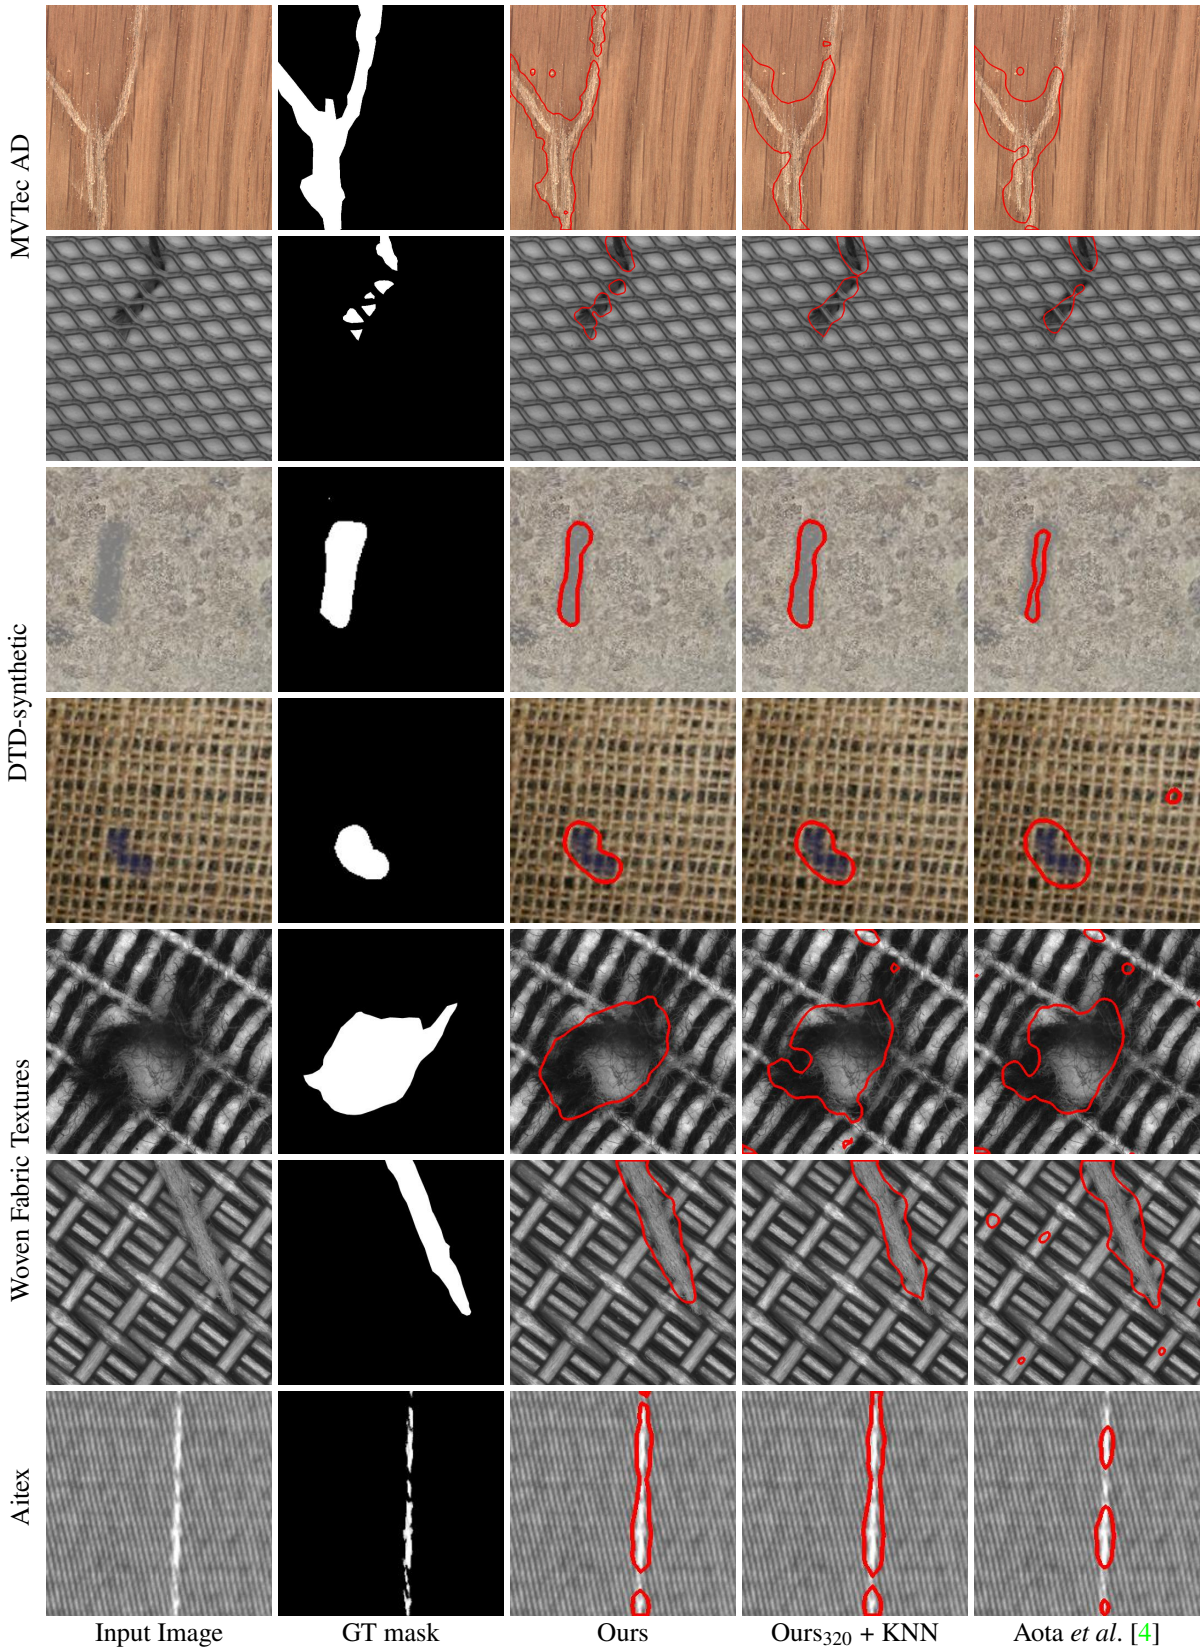

Figure 7: Qualitative comparison on challenging examples, displaying anomalous regions detected by each method via thresholding of the predicted anomaly maps. The respective thresholds are chosen to be F_1-optimal. The thresholded (binary) maps are represented through their enclosing contours; all images are shown after cropping to the center.