

Back to the Feature: Classical 3D Features are (Almost) All You Need for 3D Anomaly Detection

Eliahu Horwitz, Yedid Hoshen
 School of Computer Science and Engineering
 The Hebrew University of Jerusalem, Israel

http://www.vision.huji.ac.il/3d_ads/
 {elياهو.horwitz, yedid.hoshen}@mail.huji.ac.il

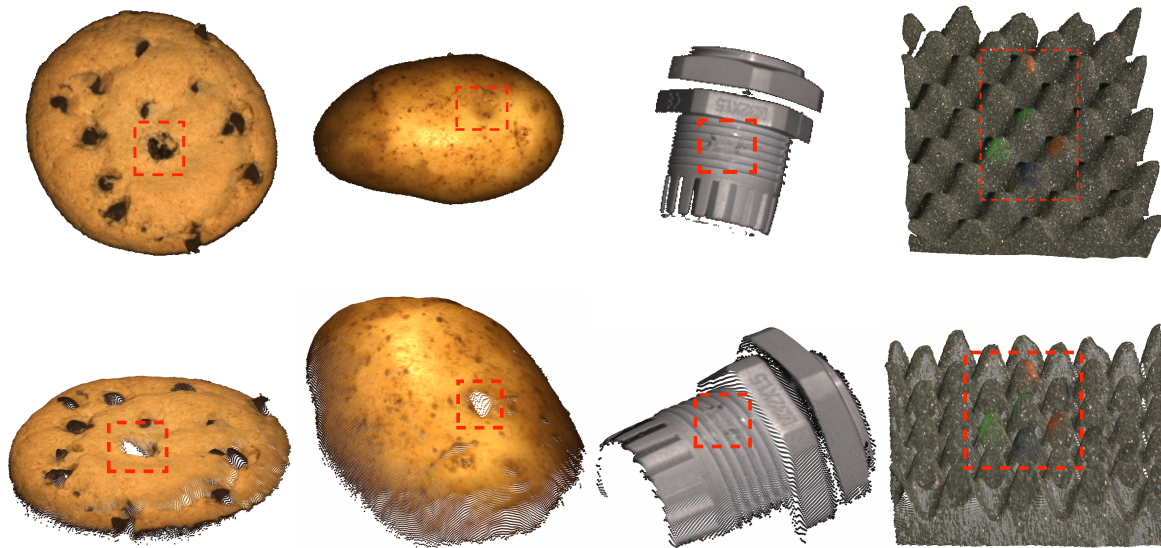


Figure 1. **Color and 3D - better together:** Objects can be anomalous in both shape and texture. Some shape anomalies are easily detected as sharp deformations in the 3D shape (marked in red, two leftmost objects - cookie and potato). In such cases, color is ineffective; the anomalies cannot be detected in top-row views¹. Conversely, 3D information often cannot identify texture anomalies e.g., deformation in cable gland, color of foam (two rightmost objects). These anomalies are easily detected in the 2D color images

Abstract

Despite significant advances in image anomaly detection and segmentation, few methods use 3D information. We utilize a recently introduced 3D anomaly detection dataset to evaluate whether or not using 3D information is a lost opportunity. First, we present a surprising finding: standard color-only methods outperform all current methods that are explicitly designed to exploit 3D information. This is counter-intuitive as even a simple inspection of the dataset shows that color-only methods are insufficient for images containing geometric anomalies. This motivates the question: how can anomaly detection methods effectively use 3D information? We investigate a range of shape representations including hand-crafted and deep-learning-based; we demonstrate that rotation invariance plays the leading

role in the performance. We uncover a simple 3D-only method that beats all recent approaches while not using deep learning, external pre-training datasets, or color information. As the 3D-only method cannot detect color and texture anomalies, we combine it with color-based features, significantly outperforming previous state-of-the-art. Our method, dubbed BTF (Back to the Feature) achieves pixel-wise ROCAUC: 99.3% and PRO: 96.4% on MVTEC 3D-AD.

1. Introduction

Although 3D understanding is fundamental to computer vision, it has typically not been considered by image anomaly detection and segmentation approaches, probably because of the lack of suitable datasets. To encour-

age research into 3D anomaly detection and segmentation (AD&S), the MVTEC 3D-AD [6] dataset was recently introduced alongside several baseline methods for 3D AD&S. However, despite the existence of a 3D AD&S dataset, the role of 3D information, as opposed to color-only, is still unclear. We conduct a careful study seeking answers to several questions:

1. Do current 3D AD&S methods truly outperform state-of-the-art 2D methods on 3D data?
2. Is 3D information *potentially* useful for AD&S?
3. What are the key properties of successful 3D AD&S representations?
4. Are there complementary benefits from using 3D shape and color modalities?

As very few previous image AD&S methods have used 3D information, we conducted a preliminary investigation of baseline methods on the MVTEC-3D dataset. Perhaps surprisingly, color-only methods (e.g. PatchCore [33]) outperform all current 3D AD&S methods by a wide margin. Next, we ask whether 3D information is potentially useful for AD&S. Encouragingly, we find that several types of anomalies go undetected when using color-only information (see Fig. 1 leftmost two, top row). In the bottom row, we present another view of the same objects, rendered using the 3D point cloud, where the anomalies are easily detected¹.

Having shown that 3D information is often needed for AD&S, our goal is to identify effective 3D representations for AD&S. We investigate a broad range of hand-crafted and deep representations and find that rotation invariance is key for 3D AD&S. *Our surprising result: a classical, hand-crafted 3D point cloud descriptor outperforms all other current methods, including learning-based representations.*

Notwithstanding the previous results, it is clear that color information is helpful. E.g., we present some examples from MVTEC 3D-AD where the anomaly is much clearer in the color than in the shape (Fig. 1, rightmost two examples). This motivates our final approach, *BTF* (Back to the Feature) which combines 3D and color to achieve the best-recorded result on the MVTEC 3D-AD dataset by a very wide margin (99.3% *Pixel-wise ROCAUC*, 96.4% *PRO*, and 87.3% *Image ROCAUC*).

Our main contributions in this paper are:

- Conducting a thorough analysis of the important and unexplored field of anomaly detection and segmentation for images with color and 3D information.

- Identifying that current 2D representations significantly outperform 3D representations on 3D data.
- Discovering that rotation invariant representations are key for 3D AD&S.
- Proposing *BTF*, a method that combines handcrafted 3D representations (FPFH) with a deep, color-based method (PatchCore), outperforming the state-of-the-art by a wide margin.

2. Related Work

Anomaly detection and segmentation. Anomaly detection methods have been researched for several decades, most approaches are based either on density estimation or out-of-domain generalization ideas. Classical approaches include: k-Nearest-Neighbors (kNN) [15], KDE [24], GMM [18], PCA [23], one class SVM (OCSVM) [39], and isolation forests [26]. With the advent of deep learning, these methods were extended with deep representations including: DAGMM [44] extending PCA, and DeepSVDD [34] extending OCSVM. A novel line of work extends self-supervised approaches to anomaly detection, including Golan and El-Yaniv [19], and Hendrycks et al. [22] that extend RotNet [17] and CSI [41] who extend contrastive methods [10, 20, 21]. We follow another line of works that assumes the availability of pre-trained representations and combines them with a kNN scoring function. Such works include Perera and Patel [28], and PANDA [31]. These works have been extended to anomaly segmentation including SPADE [11], PADIM [13] and PatchCore [33]. Very recent works have used more advanced density estimation models on the extracted representation, an example is FastFlow [43]. Other approaches for anomaly segmentation include Student-Teacher autoencoder approaches [5] as well as self-supervised methods that synthesize anomalies such as CutPaste [25] and NSA [38].

Anomaly detection and segmentation with 3D information. In contrast to the large amount of research on 2D anomaly detection approaches, 3D anomaly detection has not been extensively researched. In medical imaging research, work was performed to adapt anomaly detection methods to voxel data. Simarro et al. [40] extend f-Anogan [36, 37] to 3D. Bengs et al. [3] presented a 3D autoencoder approach for medical voxel data. Voxel data is significantly different from point cloud 3D data. Bergmann et al. [6] recognized that a dataset for anomaly segmentation in 3D point cloud data is missing and introduced MVTEC 3D-AD [6]. We expect this to be a critical contribution to the development of 3D anomaly detection and segmentation. Concurrently to our work, Bergmann and Sattlegger [7] introduced a 3D point cloud based approach dubbed *3D - ST₁₂₈* for anomaly detection, we include this work in our investigation.

¹Note that the black back-plane of the images was removed for visualization purposes. In some cases, this removal is only possible given the 3D information (e.g. We cannot tell apart the hole from the chocolate by looking only at the color image, 3D information is needed).

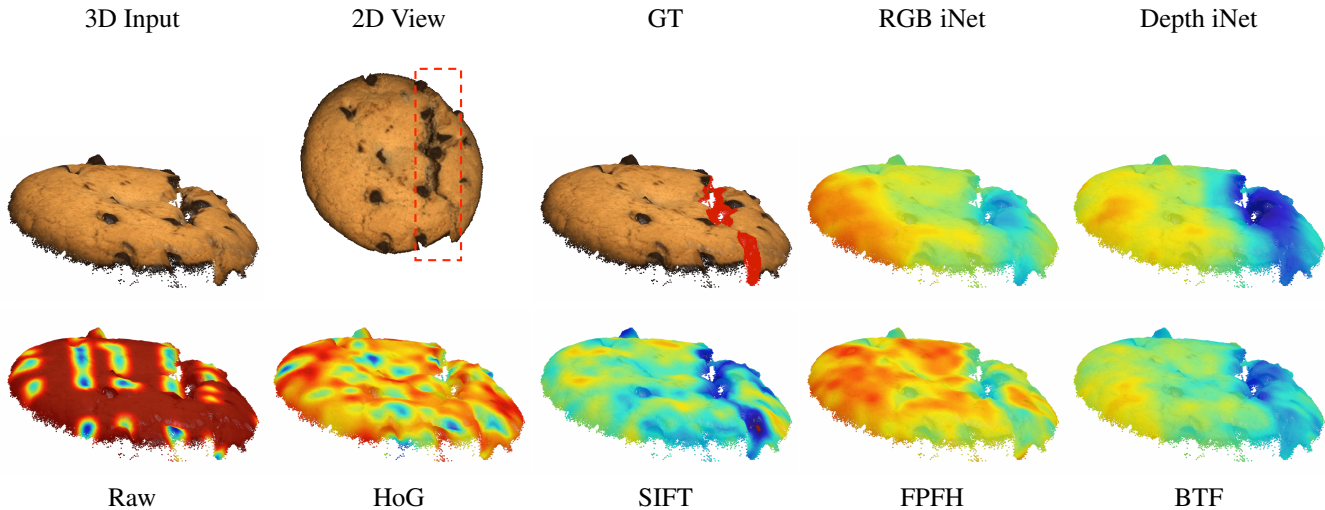


Figure 2. *Anomaly Heatmaps (Pixel-wise)*: The distance heatmaps are shown for each method. The *PRO* and *P-ROC* metrics use these distances for the final score. Blue colors indicate large distances (i.e., more anomalous) while red colors are smaller distances (i.e., less anomalous). Anomalies are indicated in red. “iNet” indicates ImageNet pre-trained features

3. Problem Definition

3.1. Setting

We assume a set of input training samples $x_1, x_2 \dots x_N$ that are all normal. At test time, we are given a test sample y . The goal of anomaly detection is to learn a sample-level scoring function σ_d , such that $\sigma_d(y) > 0$ for anomalous samples and $\sigma_d(y) \leq 0$ for normal ones. The goal of anomaly segmentation is to learn a pixel-level scoring function σ_s , which satisfies $\sigma_s(y, i) > 0$ if pixel i of sample y is anomalous, and $\sigma_s(y, i) \leq 0$ if it is normal.

Many current state-of-the-art methods (e.g. SPADE [11], PatchCore [33]) follow the following stages: i) Extracting representation of local regions ii) Estimating the probability density of normal local regions. For example, PatchCore and SPADE perform the density estimation by the nearest-neighbor distance to the normal training dataset.

Representation. We first compute a representation of each local region that may consist of one or more pixels. The representation of region j of image x is denoted $\phi(x, j)$. In this paper we focus on the representation stage, particularly, our goal is to find learned or handcrafted representations for 3D AD&S.

Anomaly scoring. Given the representations for every local region j of every training image x , we can train a model $\sigma_s(y, j)$ which computes the likelihood of a new representation $\phi(y, j')$. Although some approaches train parametric models for the density of the representations, non-parametric approaches are much simpler and require no training. Specifically, we use the k-Nearest-Neighbor distance of representation $\phi(y, j')$ to the set of all training representations $S = \{\phi(x, j) \forall x \forall j\}$. Despite their sim-

ilarity, such approaches are very accurate, require no training, and can be significantly sped up.

3.2. 3D Representations

Although RGB images are the default modality, they lack explicit 3D information. Other representations contain direct 3D information e.g., depth maps, *organized* point clouds, *unorganized* point clouds, and voxels. Both organized and unorganized point clouds represent the XYZ location of points in 3D space. However, Organized point clouds retain spatial relation and can thus be treated as images, allowing the use of RGB-based methods (e.g. CNN). Unorganized point clouds, in contrast, do not retain spatial relation and thus require specific methods and models. Finally, voxels are derived from point clouds and can be thought of as a 3D extension of pixels. For brevity we use the term “pixel” throughout the paper, however, depending on the context, it may refer to any of the above representations.

3.3. Benchmark

Our investigation uses the recently published MVTEC 3D Anomaly Detection dataset [6]. It contains over 4000 high-resolution 3D scans of industrially manufactured products across 10 categories. Each sample is represented by an *organized* point cloud and a corresponding RGB image with a one-to-one mapping between the pixels in the point cloud and those in the RGB image. Five of the classes in the dataset exhibit natural variations (*bagel*, *carrot*, *cookie*, *peach*, and *potato*). The classes *cable gland* and *dowel* are of rigid bodies, while the classes *foam*, *rope*, and *tire* are “man-made” but deformable. Bergmann et al.

Table 1. *MVTec Baselines vs. 2D RGB*: Average metrics across all classes, best performing MVTEC 3D-AD methods are shown

Voxel GAN		Voxel + RGB GAN		Point Cloud 3D - ST ₁₂₈		RGB PatchCore	
<i>PRO</i>	<i>I-ROC</i>	<i>PRO</i>	<i>I-ROC</i>	<i>PRO</i>	<i>I-ROC</i>	<i>PRO</i>	<i>I-ROC</i>
0.583	0.537	0.639	0.517	0.833	-	0.876	0.785

[6] introduced three baselines for the dataset: GAN-based, Autoencoder-based (AE), and Variation Model (VM) - a simple baseline based on per-pixel mean and standard deviation. These models operate either on the depth images or in voxel space with additional variants that operate on 3D+RGB information.

3.4. Evaluation Metrics

We use several evaluation metrics. Image-level anomaly detection is measured using image-level ROCAUC [9] (denoted *I-ROC*). Two pixel-level metrics are used for anomaly segmentation: i) pixel-wise ROCAUC, an extension of the standard ROCAUC for the pixel level which simply treats each pixel in the dataset as a sample and computes the ROCAUC over all pixels in the dataset (denoted *P-ROC*). ii) The *PRO* [4] metric, defined as the average relative overlap of the binary prediction *P* with each ground truth connected component C_k where K denotes the number of ground truth components. The final metric is computed by integrating this curve up to some false positive rate and normalizing

$$PRO = \frac{1}{K} \sum_{k=1}^K \frac{|P \cap C_k|}{|C_k|},$$

Following common practice, unless otherwise stated, we use the integration limit of 0.3.

4. An Empirical Investigation of 3D AD&S

4.1. Do current 3D methods beat 2D methods?

We begin our investigation by evaluating if current 3D AD&S methods are actually better than the SoTA 2D methods when applied on 3D data. To represent 3D methods, we test two approaches: i) Voxel GAN [6], a generative method proposed as a baseline for 3D AD&S. While it has several variants, we use the best performing ones, which are “Voxel” and “Voxel + RGB”. ii) 3D-ST [8], a concurrent method that uses a point cloud student-teacher model to learn 3D representations. We use PatchCore [33] to represent color-based image AD&S methods. Importantly, PatchCore uses features that were pre-trained on the ImageNet [14] dataset, which has been shown to be highly effective for image AD&S. In contrast, 3D-ST used Model-




Figure 3. *3D-Aware Preprocessing*: A nuisance artifact in the fabric

Net10 [42] for pre-training their teacher model. We present the results in Tab. 1. Surprisingly, PatchCore, which does not use 3D information, outperforms all previous methods.

Conclusion. Currently, state-of-the-art methods for image AD&S that use only color information, outperform 3D AD&S methods that use 3D or 3D + color information.

4.2. Is 3D information potentially useful for AD&S?

Provided the results of Sec. 4.1, we are faced with a second question: “Is 3D information *potentially* useful for AD&S?”. Below we present two cases in which 3D information is indeed useful for AD&S.

Ambiguous geometry. Frequently, we are unable to determine the underlying geometry of an object by only looking at the color information of the object. In such cases, 3D information may reveal the true geometry. We present several examples of such cases in the left half of Fig. 1-top row, the anomaly in each object cannot be detected from color information only. In the bottom row, using the 3D information, we present another view of the same objects where the anomalies are easily detected. In the case of the cookie¹, looking at the color-only image, the hole blends in with the rest of the chocolate chips, making it hard to visually identify the image as anomalous. Using the 3D information, we visualize the cookie from a different angle, making it easy to spot the anomaly. Looking at the image of the potato, it is hard to infer the geometry of the dent from shadow and texture. However, viewing the potato from different angles (by using the 3D information), the different texture reveals the dent.

Background variation. Curated datasets usually contain synthetic conditions such as centered objects and clean backgrounds, but the reality is seldom so simple. Many methods mistakenly classify cluttered image backgrounds as anomalous. Although background segmentation is not trivial, it is far easier when the 3D information is provided. We found cases in the MVTEC-3D datasets where backgrounds nuisance artifacts triggered false-positive alerts.

Table 2. *Summary of Our Findings*: Average metrics across all classes, “iNet“ indicates ImageNet pre-trained, PC indicates point cloud

Modality Method	RGB iNet	Depth iNet	Depth NSA	Depth Raw	Depth HoG	Depth SIFT	PC FPFH	PC + RGB PointNeXt	PC SpinNet	RGB + PC BTF
PRO	0.876	0.586	0.572	0.191	0.614	0.866	0.924	0.380	0.654	0.964
I-ROC	0.785	0.637	0.696	0.528	0.560	0.714	0.753	0.587	0.524	0.865
P-ROC	0.966	0.821	0.817	0.548	0.845	0.954	0.980	0.687	0.873	0.993

We demonstrate such a case in Fig. 3, the background fabric contains “wave“ like patterns which are hard to detect given the dark background color.

Conclusion. 3D information is often required to identify anomalies, even when color is available.

4.3. What are the key properties of successful 3D AD&S representations?

Having shown that 3D information is under-utilized by current methods, and having established the necessity of 3D information for image AD&S, we now seek to answer a third question: “What are the key properties of successful 3D AD&S representations?”. We distinguish among several categories.

Learning-based representations designed for images.

We adapt the two most popular learning-based image AD&S paradigms to 3D data: i) ImageNet pre-trained features ii) Self-supervised methods.

Depth-only ImageNet features. Motivated by the impressive results of ImageNet pre-trained features on color images (Sec. 4.1), we apply PatchCore on depth images.

NSA. A different class of learning-based methods approaches AD&S from a generative perspective. CutPaste and NSA [25, 38] are recent works that try to mimic anomalies by pasting image patches at different image locations. Specifically, NSA uses Poisson blending [29] to make these augmentations appear more natural.

Results. ImageNet pre-trained features significantly outperform NSA on depth images (Tab. 2). Both approaches underperform PatchCore applied to color images.

Handcrafted Image Representations. Depth patterns are often much simpler than color patterns. We hypothesize that a simple, handcrafted descriptor should suffice. The following depth representations do not require external data or training.

Raw Depth Values. Here, we test perhaps the simplest possible representation, the raw depth values of a patch.

Histogram of Oriented Gradients (HoG). HoG [12] considers image gradients and uses histograms to capture the distribution of gradient orientations in a patch. This is potentially more powerful than raw values as the descriptor encodes the spatial structure of the data while being invari-

ant to small translations. On the other hand, HoG is not invariant to global rotations, a much-desired property for 3D representations. Additionally, the small context of HoG makes it invariant to local geometric changes. This is counterproductive to our goal of detecting anomalies - usually manifested as local geometric changes.

Dense Scale-Invariant Feature Transform (D-SIFT). In contrast to HoG, SIFT [27] is rotation, scale, and shift-invariant as it is rotated to align the most dominant direction to the base orientation. This reduces the rotation ambiguity allowing matches between rotated images.

Results. HoG significantly boosts pixel-level accuracy, achieving better results than raw and learning-based features. These strong results are obtained despite HoG not being specifically designed for 3D information. Finally, the D-SIFT descriptor is able to surpass all previous depth-based results (including learning-based ones) on all three metrics.

3D rotation-Invariant Representations. Rotation-invariant features were very effective on depth maps. We now ask if rotation-invariant 3D features can do better?

Fast Point Feature Histograms (FPFH) [35]. The method first computes the k-Nearest-Neighboring points to the region center point. It then computes a histogram-based representation as a function of the surface normals and vector distance to the nearest neighbors. We choose this as the representative due to its time-tested excellent performance.

Point-cloud specific learning-based representations.

PointNeXt [30]. A U-Net [32] architecture in which the encoder hierarchically abstracts the point cloud features while the decoder gradually interpolates the abstracted features.

SpinNet [1]. A rotation invariant, learning-based representation learning method. A transformation and voxelization phase make the model rotation invariant.

Results. Compared to most methods mentioned above, PointNeXt falls short. SpinNet performs better than PointNeXt (another indicator of the importance of the rotation invariance), yet fails to surpass the proposed rotation invariant, handcrafted methods. See Tab. 3, Fig. 5, and the supplementary materials (SM) for the results.

Conclusion. FPFH outperforms all methods that use color, depth, or both (Tab. 2). The results show that strong, handcrafted, rotation-invariant 3D representations are extremely effective for AD&S when 3D information is available. Fur-

Table 3. **Detailed PRO Results:** Top half are current state-of-the-art, bottom half are methods investigated by us. Many of our methods outperform all current methods by a wide margin. “iNet” indicates ImageNet pre-trained

Method	Bagel	Cable Gland	Carrot	Cookie	Dowel	Foam	Peach	Potato	Rope	Tire	Mean		
Previous Methods	Voxel GAN	0.440	0.453	0.825	0.755	0.782	0.378	0.392	0.639	0.775	0.389	0.583	
	+ RGB	0.664	0.620	0.766	0.740	0.783	0.332	0.582	0.790	0.633	0.483	0.639	
	Voxel AE	0.260	0.341	0.581	0.351	0.502	0.234	0.351	0.658	0.015	0.185	0.348	
	+ RGB	0.467	0.750	0.808	0.550	0.765	0.473	0.721	0.918	0.019	0.170	0.564	
	Voxel VM	0.453	0.343	0.521	0.697	0.680	0.284	0.349	0.634	0.616	0.346	0.492	
	+ RGB	0.510	0.331	0.413	0.715	0.680	0.279	0.300	0.507	0.611	0.366	0.471	
	Depth GAN	0.111	0.072	0.212	0.174	0.160	0.128	0.003	0.042	0.446	0.075	0.143	
	+ RGB	0.421	0.422	0.778	0.696	0.494	0.252	0.285	0.362	0.402	0.631	0.474	
	Depth AE	0.147	0.069	0.293	0.217	0.207	0.181	0.164	0.066	0.545	0.142	0.203	
	+ RGB	0.432	0.158	0.808	0.491	0.841	0.406	0.262	0.216	0.716	0.478	0.481	
	Depth VM	0.280	0.374	0.243	0.526	0.485	0.314	0.199	0.388	0.543	0.385	0.374	
	+ RGB	0.388	0.321	0.194	0.570	0.408	0.282	0.244	0.349	0.268	0.331	0.335	
	3D – ST_{128}	0.950	0.483	0.986	0.921	0.905	0.632	0.945	0.988	0.976	0.542	0.833	
	Our Findings	RGB iNet	0.898	0.948	0.927	0.872	0.927	0.555	0.902	0.931	0.903	0.899	0.876
		Depth iNet	0.701	0.544	0.791	0.835	0.531	0.100	0.800	0.549	0.827	0.185	0.586
		NSA	0.724	0.228	0.716	0.856	0.320	0.432	0.712	0.655	0.818	0.258	0.572
Raw		0.040	0.047	0.433	0.080	0.283	0.099	0.035	0.168	0.631	0.093	0.191	
HoG		0.518	0.609	0.857	0.342	0.667	0.340	0.476	0.893	0.700	0.739	0.614	
SIFT		0.894	0.722	0.963	0.871	0.926	0.613	0.870	0.973	0.958	0.873	0.866	
FPFH		0.972	0.849	0.981	0.939	0.963	0.693	0.975	0.981	0.980	0.949	0.928	
PointNext		0.425	0.294	0.365	0.772	0.227	0.151	0.408	0.101	0.771	0.295	0.380	
SpinNet		0.635	0.316	0.922	0.780	0.870	0.380	0.585	0.699	0.955	0.400	0.654	
BTF		0.976	0.967	0.979	0.974	0.971	0.884	0.976	0.981	0.959	0.971	0.964	

thermore, as anomalies are usually local and “fine-grained”, using only a small subset of points (as required by many deep-learning-based methods) reduces performance.

4.4. Are there complimentary benefits from using both 3D and color modalities?

While the best depth-only representation outperformed existing color-only representations, we hypothesize that combining them might achieve the best of both worlds. In some cases, geometry alone does not suffice for detecting anomalies. Two examples are fine textures and color-based anomalies. The “cable gland” in Fig. 1-right is slightly scraped. While this anomalous texture is clearly observed in the color image, it is virtually impossible to detect with the current resolution of the 3D information. This is even more apparent in the foam example, wherein the anomaly is manifested as a change in color. As our exclusive focus on 3D fails to account for certain anomalies, it is necessary to combine 3D and color information.

BTF - A Combined color + 3D Approach. We take a combined color + 3D approach. To this end, color repre-

sentations are extracted using the ImageNet-based method discussed in Sec. 4.1 and 3D representations are extracted using FPFH as discussed in Sec. 4.3. We concatenate these two representations, forming a color + 3D representation which we dub *BTF* (Back to the Feature).

Results. Compared with the previous best method of combining 3D and RGB (“Voxel GAN + RGB”), our *BTF* improves the *PRO* (i.e. anomaly segmentation) metric by 32.5% and *I-ROC* (i.e. anomaly detection) by 33.6%. Compared to using only 3D information, our *BTF* improves on FPFH by 3.6% *PRO* and 12% *I-ROC*. Moreover, it achieves a score of 99.3% on *P-ROC*, a 1.3% improvement over FPFH (Tab. 3, Fig. 5). Other color and 3D combinations and extended results are found in the SM.

Conclusion. By combining color and 3D information, our *BTF* representation makes use of complementary attributes from both modalities, achieving the best results to date on the MVTEC 3D-AD dataset.

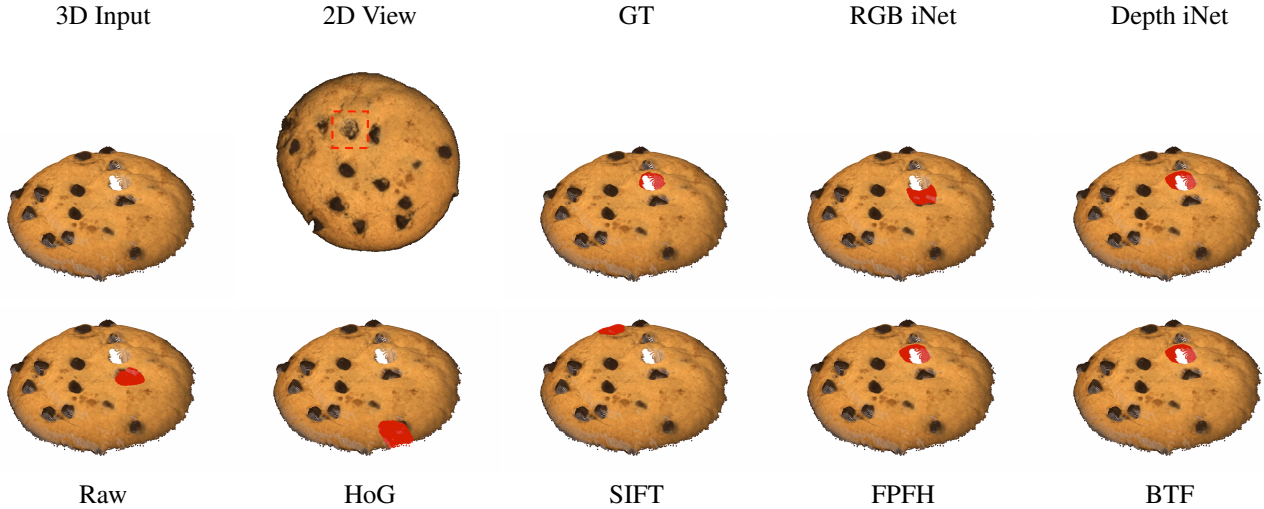


Figure 4. **Most Distant Patch (I-ROC)**: The patch with the largest kNN distance is shown in red for each representation. Anomaly indicated by a red square in the 2D view, “iNet” indicates ImageNet pre-trained

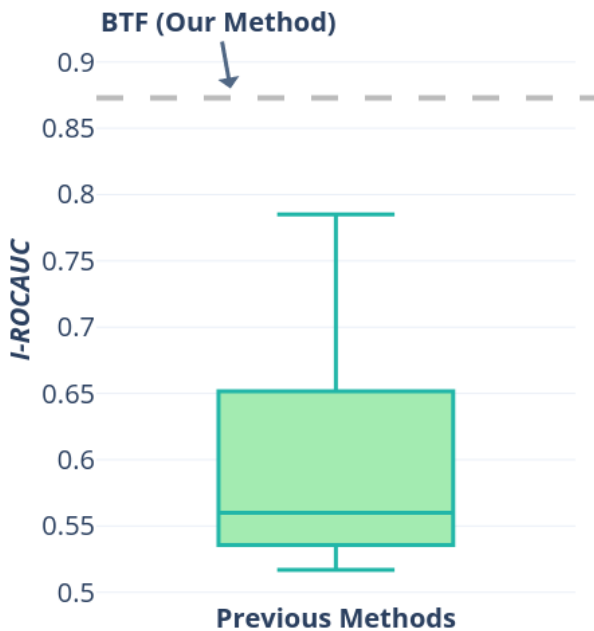


Figure 5. **Summary I-ROCAUC Results**: Our proposed BTF (top, dashed line) outperforms all 21 other methods (bottom, candle chart). For the numbers, see SM

4.5. Implementation Details

Unless otherwise stated, the original point clouds and color images are downsampled to 224×224 . For point clouds, we downsample the *organized* point cloud (i.e. im-

age downsampling) using nearest-neighbor interpolation, the color images are downsampled using bicubic interpolation. For *unorganized* point clouds, we reshape the *organized* point cloud from $n \times m \times 3$ into $n \cdot m \times 3$. We use the Z channel of the *organized* point cloud as our depth map. We extract $28 \cdot 28 = 784$ patches (features) from each sample, the feature dimension varies based on the representations used. When the representation is extracted at a different resolution, we use average pooling to match $28 \cdot 28 = 784$. For nonsquare classes (i.e. *rope* and *tire*), we pad the color and 3D images with zeros. For PointNeXt, we use the PointNeXt-XL architecture pre-trained on S3DIS [2] with a segmentation objective. We report the results on area1 as it performs best. For further details see SM.

Establishing a 3D-based preprocessing protocol. Preprocessing is sometimes required for removing nuisance artifacts. To handle such cases we developed a simple preprocessing method. We first remove the background plane by applying RANSAC [16] on the point cloud data. Once removed, we discard outliers and areas far from the plane by applying a connected-components-based algorithm (for implementation details see the SM). This preprocessing phase left the results of color-only methods mostly unaffected. More interestingly, it drastically improved results for the depth-based methods, while for 3D-based methods (i.e. FPFH) it slightly decreases results. We postulate this is caused by the difference in how depth and 3D methods handle missing sensor information². For point clouds, these missing values are all located at the origin (since their value is 0) and are easily ignored (since they are not in the spatial

²3D sensing methods are prone to sampling noise and missing information (e.g. occlusions). In MVtec 3D-AD it is common to have very noisy backgrounds, these areas are replaced by zeros by the dataset designers.




Figure 6. *Additional results of our method (BTF)*: All anomalies are correctly segmented (marked in red)

context of other points). In contrast, for depth images, these values are in the spatial context of other points and are thus taken into account. Removing these planes creates a similar situation to that of the point clouds and hence benefits depth-based methods. When using the preprocessed data, even the simplest feature (i.e. Raw) outperforms the original baselines [6], results are shown in Fig. 7.

4.6. Limitations.

Our proposed method BTF has several limitations:

Feature fusion. Both *cable gland* and *foam* perform poorly for all depth-based methods (Tab. 3 and SM). While the anomalies in these classes are easier to detect by using color than by using 3D (see Fig. 1), we expected the fusion of both modalities to improve performance. Unfortunately, for these classes, the fused features underperformed the color-only method. Future work should address this issue.

Image-level accuracy. While *BTF* establishes a new state-of-the-art on all metrics, the image level detection accuracy is far from perfect. It reaches an *I-ROC* of 86.5%, a large improvement compared to past methods, but still a relatively low score. Since we use PatchCore as the backbone for most of our experiments, the *I-ROC* score is determined by the image patch that is most distant from all training patches. We expect that better metrics can be devised for 3D data; investigating them is left for future work.

5. Conclusion

Our study was motivated by the outperformance of color-only approaches over all existing 3D methods on the

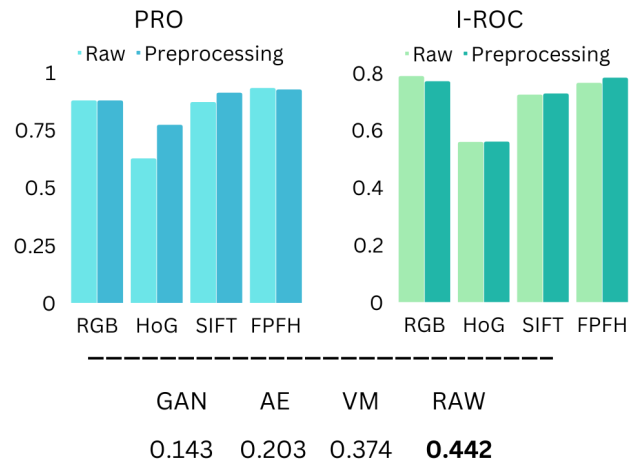


Figure 7. *Preprocessing*: Top: *PRO* vs. *I-ROC*. Bottom: Preprocessed MVTEC Baselines vs. Raw (*PRO*), average metrics across all classes are reported

MVTEC 3D-AD dataset. We conducted an extensive investigation of 3D representations and found that rotation-invariant representations achieves the best performance on 3D anomaly detection. We proposed *BTF*, a combination of 3D and color features that set a new state-of-the-art. As our method is simple, we expect it to serve as a strong baseline for future work.

References

- [1] Sheng Ao, Qingyong Hu, Bo Yang, Andrew Markham, and Yulan Guo. Spinnet: Learning a general surface descriptor for 3d point cloud registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 5
- [2] Iro Armeni, Ozan Sener, Amir R. Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3d semantic parsing of large-scale indoor spaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 7
- [3] Marcel Bents, Finn Behrendt, Julia Krüger, Roland Opfer, and Alexander Schlaefer. Three-dimensional deep learning with spatial erasing for unsupervised anomaly segmentation in brain mri. *International journal of computer assisted radiology and surgery*, 16(9):1413–1423, 2021. 2
- [4] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Mvtec ad—a comprehensive real-world dataset for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9592–9600, 2019. 4
- [5] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. Uninformed students: Student-teacher anomaly detection with discriminative latent embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4183–4192, 2020. 2
- [6] Paul Bergmann, Xin Jin, David Sattlegger, and Carsten Steger. The mvtec 3d-ad dataset for unsupervised 3d anomaly detection and localization. *arXiv preprint arXiv:2112.09045*, 2021. 2, 3, 4, 8
- [7] Paul Bergmann and David Sattlegger. Anomaly detection in 3d point clouds using deep geometric descriptors. *arXiv preprint arXiv:2202.11660*, 2022. 2
- [8] Paul Bergmann and David Sattlegger. Anomaly detection in 3d point clouds using deep geometric descriptors. *arXiv preprint arXiv:2202.11660*, 2022. 4
- [9] Andrew P Bradley. The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern recognition*, 30(7):1145–1159, 1997. 4
- [10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020. 2
- [11] Niv Cohen and Yedid Hoshen. Sub-image anomaly detection with deep pyramid correspondences. *arXiv preprint arXiv:2005.02357*, 2020. 2, 3
- [12] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, volume 1, pages 886–893. Ieee, 2005. 5
- [13] Thomas Defard, Aleksandr Setkov, Angélique Loesch, and Romaric Audigier. Padim: a patch distribution modeling framework for anomaly detection and localization. In *International Conference on Pattern Recognition*, pages 475–489. Springer, 2021. 2
- [14] 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. 4
- [15] Eleazar Eskin, Andrew Arnold, Michael Prerau, Leonid Portnoy, and Sal Stolfo. A geometric framework for unsupervised anomaly detection. In *Applications of data mining in computer security*, pages 77–101. Springer, 2002. 2
- [16] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. 7
- [17] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Un-supervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018. 2
- [18] Michael Glodek, Martin Schels, and Friedhelm Schwenker. Ensemble gaussian mixture models for probability density estimation. *Computational Statistics*, 28(1):127–138, 2013. 2
- [19] Izhak Golan and Ran El-Yaniv. Deep anomaly detection using geometric transformations. In *NeurIPS*, 2018. 2
- [20] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. *arXiv preprint arXiv:1911.05722*, 2019. 2
- [21] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. 2
- [22] Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. Using self-supervised learning can improve model robustness and uncertainty. In *NeurIPS*, 2019. 2
- [23] Ian Jolliffe. *Principal component analysis*. Springer, 2011. 2
- [24] Longin Jan Latecki, Aleksandar Lazarevic, and Dragoljub Pokrajac. Outlier detection with kernel density functions. In *International Workshop on Machine Learning and Data Mining in Pattern Recognition*, pages 61–75. Springer, 2007. 2
- [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, 5
- [26] Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *2008 Eighth IEEE International Conference on Data Mining*, pages 413–422. IEEE, 2008. 2
- [27] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004. 5
- [28] Pramuditha Perera and Vishal M Patel. Learning deep features for one-class classification. *IEEE Transactions on Image Processing*, 28(11):5450–5463, 2019. 2
- [29] Patrick Pérez, Michel Gangnet, and Andrew Blake. Poisson image editing. *SIGGRAPH*, 2003. 5
- [30] Guocheng Qian, Yuchen Li, Houwen Peng, Jinjie Mai, Hasan Hammoud, Mohamed Elhoseiny, and Bernard Ghanem. Pointnext: Revisiting pointnet++ with improved training and scaling strategies. *arXiv:2206.04670*, 2022. 5

- [31] Tal Reiss, Niv Cohen, Liron Bergman, and Yedid Hoshen. Panda: Adapting pretrained features for anomaly detection and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2806–2814, 2021. [2](#)
- [32] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. [5](#)
- [33] Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Schölkopf, Thomas Brox, and Peter Gehler. Towards total recall in industrial anomaly detection. *arXiv preprint arXiv:2106.08265*, 2021. [2](#), [3](#), [4](#)
- [34] Lukas Ruff, Nico Gornitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Robert Vandermeulen, Alexander Binder, Emmanuel Müller, and Marius Kloft. Deep one-class classification. In *ICML*, 2018. [2](#)
- [35] Radu Bogdan Rusu, Nico Blodow, and Michael Beetz. Fast point feature histograms (fpfh) for 3d registration. In *2009 IEEE International Conference on Robotics and Automation*, pages 3212–3217, 2009. [5](#)
- [36] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Georg Langs, and Ursula Schmidt-Erfurth. f-anogan: Fast unsupervised anomaly detection with generative adversarial networks. *Medical image analysis*, 54:30–44, 2019. [2](#)
- [37] Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Ursula Schmidt-Erfurth, and Georg Langs. Unsupervised anomaly detection with generative adversarial networks to guide marker discovery. In *International Conference on Information Processing in Medical Imaging*, 2017. [2](#)
- [38] Hannah M Schlüter, Jeremy Tan, Benjamin Hou, and Bernhard Kainz. Self-supervised out-of-distribution detection and localization with natural synthetic anomalies (nsa). *arXiv preprint arXiv:2109.15222*, 2021. [2](#), [5](#)
- [39] Bernhard Scholkopf, Robert C Williamson, Alex J Smola, John Shawe-Taylor, and John C Platt. Support vector method for novelty detection. In *NIPS*, 2000. [2](#)
- [40] Jaime Simarro Viana, Ezequiel de la Rosa, Thijs Vande Vyvere, David Robben, Diana M Sima, et al. Unsupervised 3d brain anomaly detection. In *International MICCAI Brainlesion Workshop*, pages 133–142. Springer, 2020. [2](#)
- [41] Jihoon Tack, Sangwoo Mo, Jongheon Jeong, and Jinwoo Shin. Csi: Novelty detection via contrastive learning on distributionally shifted instances. *arXiv preprint arXiv:2007.08176*, 2020. [2](#)
- [42] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. [4](#)
- [43] Jiawei Yu, Ye Zheng, Xiang Wang, Wei Li, Yushuang Wu, Rui Zhao, and Liwei Wu. Fastflow: Unsupervised anomaly detection and localization via 2d normalizing flows. *arXiv preprint arXiv:2111.07677*, 2021. [2](#)
- [44] Bo Zong, Qi Song, Martin Renqiang Min, Wei Cheng, Cristian Lumezanu, Daeki Cho, and Haifeng Chen. Deep autoen-

coding gaussian mixture model for unsupervised anomaly detection. *ICLR*, 2018. [2](#)