

Anomaly Factory 3D: A Modular Framework for Diverse Pseudo-Anomaly Synthesis in Unsupervised 3D Anomaly Detection

Ali Balapour^{1,2,3}  and Faraz Hach^{2,3,4} *

¹ Department of Computer Science, The University of British Columbia, Vancouver, V6T 1Z4 BC, Canada

² Department of Urologic Sciences, The University of British Columbia, Vancouver, V5Z 1M9 BC, Canada

³ Vancouver Prostate Centre, M.H. Mohseni Institute of Urologic Sciences, Vancouver, V6H 3ZB BC, Canada

⁴ The Reproductive and Sexual Medicine Centre, M.H. Mohseni Institute of Urologic Sciences, Vancouver, V6H 3ZB BC, Canada

Abstract. Detecting and localizing defects in 3D point clouds is challenging because abnormal samples are scarce and diverse, while training is often limited to normal data. We propose Anomaly Factory 3D (AF3AD), a modular framework that synthesizes diverse pseudo-anomalies from normal point clouds to expand the training data for unsupervised 3D anomaly detection methods that rely on pseudo-anomalies. AF3AD uses a center-conditioned parametric deformation model defined in local PCA frames, with kernel-controlled spatial falloff, anisotropy, directional gating, and normal/tangential displacement fields, enabling a broad set of geometric defect presets. We demonstrate its ease-of-use and effectiveness by integrating AF3AD with an offset-prediction detector and a reconstruction-based anomaly detection method, showing that AF3AD transfers across detection paradigms. Experiments on AnomalyShapeNet and Real3D-AD show consistent improvements in object- and point-level detection and localization, supported by ablations on preset groups and robustness under noise. AF3AD is designed as a standalone synthesis tool to facilitate adoption across different 3D anomaly detection paradigms. Code is available at github.com/vpc-ccg/AF3AD.

Keywords: 3D Anomaly Detection · Unsupervised Learning · Data Synthesis

1 Introduction

Anomaly detection is a critical task across multiple domains, including natural language, audio, and image data [12]. However, the 3D shape modality remains less explored, despite being essential for detailed analysis in scientific and engineering applications. This gap exists largely because 3D data are difficult to

* Corresponding author: faraz.hach@ubc.ca.

capture and computationally intensive to process. Furthermore, 3D anomaly detection (3DAD) is fundamentally more challenging than its 1D (time-series) or 2D (images) counterparts due to structural complexities such as geometric sparsity, viewpoint dependence, and sensor noise [3, 9].

Despite these technical challenges, performing anomaly detection on 3D data offers distinct advantages over traditional 2D images. These include a richer understanding of an object’s spatial geometry, reduced reliance on a single viewpoint, and lower sensitivity to the illumination variations that frequently affect 2D data [9]. One primary use case for 3DAD is found in industrial settings [1], particularly for identifying defects in manufactured objects.

The core task of 3DAD involves identifying regions on the surface of a 3D object, particularly within point cloud representations, that deviate from normal patterns. These deviations can range from simple surface bulges to complex structural deformations. To identify these areas, Artificial Intelligence (AI) based methods have proven particularly effective for both detection and localization. Recent advancements in these methods have enabled faster and more accurate detection, with some approaches now achieving real-time defect detection on 3D data [6, 17]. AI-based methods for 3DAD can generally be grouped into three main paradigms: (1) reconstruction-based approaches, which learn a model of normal geometry and flag deviations via reconstruction residuals; (2) memory-bank (feature-based) methods, which store representations of normal samples and detect anomalies as feature mismatches at inference time; and (3) multi-view approaches, which render 3D data into a set of 2D views to leverage mature 2D anomaly detection pipelines.

Due to the scarcity of anomalous data, training 3DAD methods directly on defective samples is often impractical. As a result, many works adopt an unsupervised setting, where only normal samples are available during training and the model is not exposed to anomalous observations. This paradigm is commonly known as Unsupervised 3D Anomaly Detection (U3DAD) [3]. Within this paradigm, an increasingly adopted strategy is to synthesize pseudo-anomalies from normal data to create surrogate supervision. Recent methods such as R3D-AD [28] and PO3AD [25] follow this direction and demonstrate that learning to distinguish normal patterns from generated defects can improve detection at inference time. However, existing synthesis pipelines are often limited in the diversity and structural complexity of the generated defects, which may restrict the range of abnormal cues the model learns. This motivates a more structured pseudo-anomaly generation framework that can produce richer and more realistic defect patterns to better support robust 3D anomaly detection.

We propose **Anomaly Factory for 3D Anomaly Detection (AF3AD)**, a method designed to synthesize a diverse set of pseudo-anomalies using geometric kernels. These kernels are controlled by probabilistic distributions to introduce greater variation into the synthesis process. Using parameterized kernels, we design a repertoire of pseudo-anomalies ranging from simple bulges and dents to directional drags and elongated ridges. These are applied to normal, non-anomalous 3D point cloud samples to prepare them for unsupervised anomaly

detection training. Our framework is a standalone, reusable module that supports a wide range of parameter settings, such as radius and displacement rate, enabling greater diversity in the generated anomalies. To validate the effectiveness of our synthesis framework, we instantiate and evaluate it with an offset-prediction detector based on PO3AD [25] and a reconstruction-based detector based on R3D-AD [28]. We show that using more diverse and complex anomalies increases the accuracy and precision on established benchmarks. We evaluate our method on two prominent datasets, AnomalyShapeNet [15] and Real3D-AD [18], demonstrating that our integration achieves competitive performance. Furthermore, we discuss the advantages that synthetic diversity provides for the task of unsupervised anomaly detection.

In summary, our main contributions are as follows: (1) We propose AF3AD, a modular synthesis framework that generates diverse geometric pseudo-anomalies with explicit parametric control, designed to enhance 3D anomaly detection methods that rely on pseudo-anomaly synthesis. (2) When integrated with an offset-prediction detector, we achieve competitive performance, reaching 91.5% O-AUROC on AnomalyShapeNet (+5.5 points) and 85.2% on Real3D-AD (+7.1 points). (3) We further provide extensive ablations demonstrating that increased synthesis diversity consistently improves detection performance, remains robust under noise, and is supported by an empirical feature-space alignment analysis.

2 Related Works

2.1 2D Anomaly Detection

Unsupervised 2D visual anomaly detection, where only normal samples are available during training, has advanced considerably in recent years and can be grouped into four main paradigms [3]. Memory-bank methods store feature representations of normal data and detect anomalies by measuring their deviation from learned patterns [22]. Reconstruction-based approaches identify anomalies through errors in reconstructing normal samples or features [24]. Knowledge-distillation methods use teacher–student discrepancies as anomaly signals [8], while flow-based models detect anomalies via low likelihood under an explicitly learned distribution of normal features [10]. Despite their success in 2D image-based settings, directly applying these paradigms to 3D data remains challenging because point clouds are unstructured, unordered, and sparse, unlike the regular grids of images.

2.2 3D Anomaly Detection

3D anomaly detection aims to detect and localize geometric irregularities (e.g., structural deformations or surface defects) in unstructured point clouds [1]. Earlier methods relied on handcrafted descriptors like FPFH [11] to capture local geometry invariant to rotation. With the advent of deep learning, teacher–student architectures (e.g., TS-AD [21]) emerged, distilling features from pre-trained

Table 1: Comparison of pseudo-anomaly synthesis methods for 3D anomaly detection.

Method	Deformation Diversity	Anisotropic Support	Directional Control	Modular Design
CLIP3D-AD	Low (1 type)	✗	Normal only	✗
Patch-Gen (R3D-AD)	Low (3 types)	✗	Random	✗
GLFM	Low (1 type)	✗	Normal only	✗
Norm-AS (PO3AD)	Low (2 types)	✗	Normal only	✗
AF3AD (Ours)	High (11 types)	✓	Normal+Tangential	✓

encoders like PointNet [20]. Recent approaches leverage large-scale pre-trained transformers; for instance, PointMAE [19] and PointBERT [26] are frequently used as backbones to extract robust semantic features from raw geometry. Building on these representations, recent works have diversified: Reg3D-AD [18] utilizes registration to align test samples with a normal template, while Reg2Inv [27] explicitly learns registration-invariant representations to handle pose variations. Other strategies include bridging the modality gap by fusing 2D multi-view images with 3D descriptors (CPMF [2]), or explicitly modeling anomaly magnitude via per-point offset regression using sparse backbones (PO3AD [25]).

2.3 Pseudo-Anomaly Synthesis

To mitigate the scarcity of anomalous training data, a body of work synthesizes pseudo-anomalies directly on 3D geometry. While earlier multimodal methods, such as CLIP3D-AD [30], applied 2D noise, geometric methods focus on structural manipulation. Patch-Gen in R3D-AD [28] introduces patch-based rotation and translation, as well as bulges, dents, and random damage, while GLFM [5] stretches points along normal vectors to simulate protrusions. Addressing unrealistic overlaps, Norm-AS in PO3AD [25] constrains displacements such as bulges and dents along surface normals. Despite these advances, current synthesis methods often lack the structural diversity required for robust generalization, motivating the need for richer generation strategies (see the summary in Table 1). MiniShift [6] also uses geometric synthesis, but for building a fixed benchmark rather than generating training supervision. Therefore, it is complementary to AF3AD.

3 Methods

3.1 Overview

Our framework is modular and operates independently of specific detection architectures. AF3AD is designed to address a key limitation of prior pseudo-anomaly synthesis methods: tight coupling with specific detection paradigms. For instance, Norm-AS [25] is designed for offset regression, whereas Patch-Gen [28]




Fig. 1: Overview of AF3AD pipeline for geometric pseudo-anomaly synthesis. From a non-anomalous point cloud and its surface normals, AF3AD selects a local segment and computes a PCA-based local frame. Within this segment, each point is transformed into local coordinates, evaluated through an anisotropic distance function, and displaced according to preset-specified gating, direction, and deformation parameters. This modular design enables the controlled synthesis of diverse geometric pseudo-anomalies, producing an anomalous point cloud along with preset-specific anomaly patterns and ground-truth (GT) masks.

is tailored to diffusion-based reconstruction. Such coupling reduces reusability and limits systematic study of synthesis strategies. In contrast, AF3AD acts as a standalone synthesis module that transforms normal point clouds into diverse pseudo-anomalous samples with corresponding anomaly masks. These synthetic samples can be used to supervise detectors trained with pseudo-anomaly injection, including reconstruction-, registration-, and offset-based methods. We instantiate AF3AD with both offset-prediction (Section 3.3) and reconstruction-based (Section 3.4) detectors, although the synthesis module itself (Section 3.2) is designed to remain standalone.

3.2 Pseudo-Anomaly Synthesis via Parametric Deformations

Pseudo-anomalies are synthesized by applying localized, parametric deformations to a non-anomalous 3D point cloud. Let $P = \{\mathbf{p}_i\}_{i=1}^M \subset \mathbb{R}^3$ denote a point cloud with M points and $N = \{\hat{\mathbf{n}}_i\}_{i=1}^M$ the corresponding surface normals. The point cloud is partitioned into $K = 64$ patches by randomly sampling K centers and assigning each point to its nearest center (single pass, no k -means); a random patch center $c \in \mathbb{R}^3$ then defines the deformation region. This produces a deformed point cloud $P' = \{\mathbf{p}'_i\}_{i=1}^M$ by applying a displacement field to each point in that segment.

General Deformation Formulation In the generated pseudo-anomalous sample P' , each point is displaced as

$$\mathbf{p}'_i = \mathbf{p}_i + \Delta(\mathbf{p}_i), \quad \Delta(\mathbf{p}_i) = s B w_i \mathbf{d}_i, \quad (1)$$

where $s \in \{+1, -1\}$ controls the deformation sign (e.g., inward vs. outward), $B > 0$ controls the maximum displacement magnitude, $w_i \in [0, 1]$ is a spatial falloff weight, and $\mathbf{d}_i \in \mathbb{R}^3$ denotes the unit deformation direction vector. This parameterized formulation provides control and flexibility to build a diverse set of pseudo-anomalies. Points outside the selected segment remain unchanged, i.e., $\mathbf{p}'_i = \mathbf{p}_i$ for points not assigned to the deformed segment.

Local Coordinate Frame and Anisotropic Support To define anisotropic deformations, we estimate a local orthonormal frame at the anomaly center $\mathbf{c}$ using local Principal Component Analysis (PCA) [13]:

$$U(\mathbf{c}) = [\mathbf{u} \ \mathbf{v} \ \mathbf{n}] \in \mathbb{R}^{3 \times 3}, \quad U^\top U = I, \quad (2)$$

where $\mathbf{u}$ and $\mathbf{v}$ span the principal tangent directions of the local neighborhood, and $\mathbf{n}$ is the local mean normal. Points are then expressed in this local frame as

$$\mathbf{x}_i = U^\top (\mathbf{p}_i - \mathbf{c}). \quad (3)$$

Let R denote a nominal anomaly radius, and let (r_u, r_v, r_n) be dimensionless anisotropy multipliers. We define an ellipsoidal normalized distance

$$t_i = \left(\frac{x_{i,u}^2}{(r_u R)^2} + \frac{x_{i,v}^2}{(r_v R)^2} + \frac{x_{i,n}^2}{(r_n R)^2} \right)^{1/2}, \quad (4)$$

which induces anisotropic spatial support around the anomaly center.

Spatial Falloff Kernels Each point is assigned a spatial weight through a kernel function

$$w_i = K(t_i), \quad w_i \in [0, 1]. \quad (5)$$

We employ smooth kernels such as cosine, Gaussian, or polynomial functions, optionally parameterized by shape parameters (e.g., q or σ), to ensure spatially coherent deformations and avoid isolated point displacements.

To support asymmetric deformations, we optionally apply a soft gating function $g_i \in [0, 1]$ and set $w_i \leftarrow w_i g_i$, where g_i is either a half-space gate in a chosen direction or a normal-alignment gate restricting deformation to similarly oriented surfaces. We choose $\mathbf{d}_i \in \{\mathbf{n}, \mathbf{u}, \mathbf{v}\}$ (from principal directions). When gating is disabled, we set $g_i = 1$ for all points.

Stochastic Parameter Sampling The sign parameter $s \in \{+1, -1\}$ determines the direction of the deformation (e.g., $+1$ for an outward bulge and -1 for an inward dent). This parameter is either fixed for a specific anomaly preset or sampled stochastically during synthesis. Both the deformation magnitude B and the nominal anomaly radius R are sampled from bounded Beta distributions to introduce controlled stochasticity while preserving geometric plausibility. Specifically, for a parameter $z \in \{B, R\}$, we sample

$$z = z_{\min} + (z_{\max} - z_{\min}) \cdot \text{Beta}(\alpha_z, \beta_z), \quad (6)$$

where $\text{Beta}(\alpha_z, \beta_z)$ denotes the Beta distribution with shape parameters (α_z, β_z) and support on $[0, 1]$. The bounds $(z_{\min}, z_{\max})$ and distribution parameters (α_z, β_z) are specified per anomaly preset, allowing precise control over the typical scale and severity of synthesized defects while preventing extreme or implausible deformations.

Combining all components, each point is transformed as

$$\mathbf{p}'_i = \mathbf{p}_i + s B K(t_i) g_i \mathbf{d}_i. \quad (7)$$

Preset Taxonomy This formulation unifies isotropic and anisotropic bulges and dents, elongated ridges and trenches, skewed impact craters, tangential shear and drag anomalies, and compound micro-defect fields under a single-center-conditioned deformation model. A summary of the defined presets is provided in Table 2, and the overall procedure is outlined in Figure 1. We have grouped these anomalies based on their features (e.g., direction, anisotropy, etc.) for further analysis. Additionally, AF3AD is not limited to these specific presets; additional deformation types can be easily incorporated into the framework.

3.3 Instantiation: Offset-Prediction Detector

To validate the effectiveness of our modular synthesis framework, we integrate it with an offset-prediction detection architecture. We build upon PO3AD [25] with architectural refinements to better leverage our diverse pseudo-anomalies, replacing the Norm-AS module proposed in that work.

After synthesizing pseudo-anomalies on training samples using AF3AD, we train a modified version of the PO3AD model. In this approach, the model predicts an offset for every point in the point cloud, representing its displacement from its expected normal position (i.e., O_x, O_y, O_z). An offset of $O_x = 0, O_y = 0, O_z = 0$ indicates that a point is in its original, non-anomalous position, while any non-zero value suggests displacement in 3D space by the corresponding amount.

Anomaly Score Calculation. For each point in the point cloud, we compute an individual anomaly score using the L_1 norm of its predicted offset (i.e., $|O_x| + |O_y| + |O_z|$). This provides a straightforward measure of point-wise deviation,

Table 2: Pseudo-anomaly preset taxonomy.

Anomaly Type	Aniso.	Sym.
Group A — <i>Normal-direction, isotropic single-lobe</i>		
Basic Bulge	Iso	1-side
Basic Dent	Iso	1-side
Group B — <i>Normal-direction, anisotropic structured marks</i>		
Elongated Ridge	Aniso	1-side
Trench / Groove	Aniso	1-side
Elliptic Flat Spot	Mild	1-side
Skewed Impact Crater	Aniso	Dir.
Group C — <i>Normal-direction, multi-frequency/periodic</i>		
Double-Sided Ripple	Iso	2-side
Micro-Dimple Field	Iso	1-side
Group D — <i>Tangential-direction surface shifts</i>		
Shear / Slip (u)	Iso	2-side
Shear / Slip (v)	Iso	2-side
Directional Drag / Stretch	Aniso	1-side
<i>Note: Aniso. = anisotropy (Iso: isotropic, Aniso: anisotropic, Mild: mildly anisotropic); Sym. = symmetry (1-side: one-sided, 2-side: two-sided, Dir.: direction-gated).</i>		

where a higher score indicates a greater likelihood of an anomaly at that specific location. To determine the overall anomaly score for the entire point cloud (i.e., for object-level detection), we simply calculate the sum of these individual point-level scores. Detailed training procedures are provided in the supplementary material (Section S1).

3.4 Instantiation: Reconstruction-based Anomaly Detection

To demonstrate AF3AD’s modularity and ease of integration, we instantiate it within R3D-AD [28] as a second reference framework for anomaly detection. R3D-AD reconstructs the geometry of 3D point clouds using an autoencoder architecture with a PointNet-based encoder and a diffusion-based decoder. During training, it leverages pseudo-anomalous samples generated from normal data via a module called Patch-Gen. Patch-Gen synthesizes three anomaly types: bulges, dents, and damages (random point perturbations), by selecting a random view of the point cloud and extracting local segments using a k -nearest-neighbors (k NN) search. In contrast, AF3AD is geometry-driven and view-independent, operating directly on the 3D structure of the point cloud. In this instantiation, we replace Patch-Gen with AF3AD to generate pseudo-anomalous training samples while keeping the remaining R3D-AD architecture and training pipeline unchanged.

4 Experiments

4.1 Experimental Settings

Datasets We evaluated our method on two standard benchmark datasets: AnomalyShapeNet [15] and Real3D-AD [18]. AnomalyShapeNet, derived from ShapeNet [4], contains 40 categories and 1,600 samples with synthetically introduced anomalies. The training set includes only four normal samples per category, while the test set contains both normal and anomalous samples. Real3D-AD contains 12 categories of high-resolution point clouds scanned from real objects. Its training set includes four normal samples per category, and its test set consists of 100 normal and anomalous samples per category. A key difference between the datasets is scan coverage: Real3D-AD uses full 360-degree scans for training but single-view scans for testing, whereas AnomalyShapeNet uses full scans for both. This train-test discrepancy in Real3D-AD introduces a domain shift that makes evaluation more challenging.

Evaluation Metrics To evaluate 3D anomaly detection performance, it is common practice to use the Area Under the Receiver Operating Characteristic Curve (AUROC). We utilize object-level AUROC (O-AUROC) to assess anomaly detection for the entire object, and point-level AUROC (P-AUROC) to evaluate anomaly localization on the point cloud, independent of a specific decision threshold. Furthermore, to provide a more comprehensive analysis, we report the Area Under the Precision-Recall Curve (AUPR) at both levels (O-AUPR and P-AUPR) whenever possible. Finally, when comparing with existing methods in the literature, we provide the average rank as an additional metric to facilitate detailed multi-method comparisons.

4.2 Implementation Details

The anomaly detector uses a MinkUNet34C [7] backbone and a three-head offset MLP with hidden size 128 and PReLU activations. We matched training exposure across datasets to 150,000 effective epochs and trained using AdamW at an initial learning rate of 5×10^{-4} . For Real3D-AD, we used cosine annealing with 50 warmup epochs and a two-stage downsampling strategy (voxel grid sampling followed by farthest point sampling (FPS)), which improved efficiency and occasionally improved performance (Table S1, Section S2). Input point clouds were centered, normalized, randomly rotated, and processed using 64 patches, with normal and tangent directions extracted from the provided `.obj` and `.ply` files. Profiling on AnomalyShapeNet shows that AF3AD is lightweight, requiring 0.96ms per call, comparable to Norm-AS (1.02ms) and faster than Patch-Gen (3.24ms), with negligible impact on training throughput.

4.3 Baseline Methods

We compare our method with multiple prominent works in the literature: BTF [11] (using raw and FPFH features), M3DM [23], PatchCore [22] (using FPFH fea-

Table 3: Summary of mean AUROC and average rank on AnomalyShapeNet and Real3D-AD. AF3AD is reported as mean $\pm$ std (%) over three different seeds. Higher AUROC is better ($\uparrow$); lower rank is better ($\downarrow$). The results of the offset-prediction instantiation are reported. Best is **bold**, second-best is underlined.

Method	AnomalyShapeNet				Real3D-AD			
	O-AUROC $\uparrow$	P-AUROC $\uparrow$	O-Rank $\downarrow$	P-Rank $\downarrow$	O-AUROC $\uparrow$	P-AUROC $\uparrow$	O-Rank $\downarrow$	P-Rank $\downarrow$
BTF(RAW)	49.3	55.0	9.70	9.76	63.5	57.1	9.21	9.42
BTF(FPFH)	52.8	62.8	9.00	7.55	60.3	73.3	9.25	6.33
M3DM	55.2	61.6	8.78	7.88	55.2	63.1	11.62	8.75
PatchCore(FPFH)	56.8	58.0	8.34	8.78	59.3	68.2	10.00	6.92
PatchCore(PMAE)	56.2	57.7	8.39	9.00	59.4	62.0	10.58	9.00
CPMF	55.9	57.3	8.31	9.35	58.6	75.9	10.58	5.29
Reg3D-AD	57.2	66.8	8.43	6.25	70.4	70.5	7.50	6.83
Group3AD	–	–	–	–	75.1	73.5	5.83	5.17
IMRNet	65.9	64.9	5.88	6.86	72.5	–	6.58	–
ISMP	–	69.1	–	6.05	75.7	83.6	5.42	3.21
R3D-AD	74.9	–	4.06	–	73.4	–	6.17	–
PO3AD	83.9	<u>89.9</u>	2.94	<u>2.35</u>	76.6	–	5.33	–
Reg2Inv	<u>86.0</u>	88.2	<u>2.56</u>	2.52	<u>78.1</u>	87.9	<u>4.42</u>	<u>2.58</u>
PO3AD+AF3AD (Ours)	91.5$\pm$2.1	92.5$\pm$1.6	1.62	1.65	85.2$\pm$1.2	<u>86.1$\pm$1.9</u>	2.50	2.50

tures or the PointMAE [19] backbone), ISMP [16], CPMF [2], Reg3D-AD [18], IMRNet [15], R3D-AD [28], Group3AD [29], PO3AD [25], and Reg2Inv [27]. Note that we only include Group3AD on Real3D-AD. The results for these methods were obtained either from the original papers or from publicly available implementations.

4.4 Main Results

Results on AnomalyShapeNet Table 3 summarizes the results of our proposed framework on the AnomalyShapeNet dataset in comparison to baseline methods. We report the performance of AF3AD as the mean and standard deviation over three random seeds. The average rank for AF3AD is computed using the per-category mean over three seeds. AF3AD achieves the best performance overall at both object and point levels, with 91.5% O-AUROC and 92.5% P-AUROC. It outperforms the best competing methods by 5.5 points in O-AUROC (vs. Reg2Inv at 86.0%) and 2.6 points in P-AUROC (vs. PO3AD at 89.9%). Compared to PO3AD, AF3AD improves O-AUROC by 7.6 points and P-AUROC by 2.6 points. Detailed class-wise results are provided in Supplementary Tables S4 and S5.

Results on Real3D-AD Table 3 summarizes the results of our proposed framework on the Real3D-AD dataset compared to baseline methods. To address the train–test discrepancy in scan coverage (full scans during training versus partial scans during testing), we additionally investigated different augmentation settings during preprocessing to mimic partial scans in training data, with details provided in Section S4, Table S3. Our method outperforms the second-best approach in O-AUROC by a margin of 7.1 points and achieves competitive results at the point level (second in raw P-AUROC, best in average point-level rank).

Table 4: Mean performance across 15 AnomalyShapeNet categories (index-0 split) (%).

Method	O-AUROC $\uparrow$	P-AUROC $\uparrow$	O-AUPR $\uparrow$	P-AUPR $\uparrow$
R3D-AD	58.7	59.9	63.2	7.5
R3D-AD + Patch-Gen	62.7	60.7	67.7	8.4
R3D-AD + AF3AD	67.2	60.7	71.3	9.2

Notably, our framework achieves an 8.6 O-AUROC points gain over PO3AD, which serves as the foundational method for our work. Detailed class-wise results are provided in Supplementary Tables S6 and S7.

4.5 Results of R3D-AD Instantiation on AnomalyShapeNet

We used the public GitHub implementation of R3D-AD by the authors, but found that in the released training script/configuration, Patch-Gen was not enabled by default during the preparation of training data; as a result, the released training pipeline uses only augmented normal samples (e.g., random rotations) rather than pseudo-anomalies. To enable a fair comparison, we evaluate three settings for this instantiation: (i) base R3D-AD (no Patch-Gen; exactly as provided), (ii) R3D-AD + Patch-Gen, and (iii) R3D-AD + AF3AD (ours). We run all three settings on 15 classes of AnomalyShapeNet (index-0 split) and report the results in Table 4. Overall, AF3AD improves performance by 4.5 O-AUROC points over the stronger Patch-Gen baseline and by 8.5 points over base R3D-AD. Nevertheless, the instantiation based on offset prediction remains stronger and is selected as our main result. This instantiation mainly shows AF3AD’s modularity across paradigms: object-level detection improves from 62.7 to 67.2 O-AUROC, while P-AUROC stays the same at 60.7. This is likely due to R3D-AD’s downsampling and diffusion-decoder residuals limit point-level localization, regardless of the synthesis method.

4.6 Ablation Study

Effect of Offset Prediction Architecture We study the interaction between the offset-head design and the synthesis module under a fixed training budget to verify that AF3AD’s gains come from synthesis diversity rather than the modified prediction head. We compare a baseline head against a multi-head variant for both Norm-AS and AF3AD, and sweep the hidden dimension $h \in \{32, 64, 128\}$. The baseline predicts the full 3D offset (O_x, O_y, O_z) using a single MLP, whereas the multi-head design uses separate prediction heads for each coordinate. Table 5 shows two consistent trends. First, the multi-head design improves performance relative to the baseline, with the largest gains appearing under AF3AD training. Second, AF3AD can underperform Norm-AS at low head capacity (e.g., baseline at $h \in \{32, 64\}$), which is expected because AF3AD draws from a larger

Table 5: Architecture $\times$ synthesis ablation on 15 categories of AnomalyShapeNet (index 0, mean over categories, in %). Best is in bold and second-best is underlined (per column), computed over the 15-category runs.

Synthesis	Offset head	h	O-AUROC $\uparrow$	P-AUROC $\uparrow$	O-AUPR $\uparrow$	P-AUPR $\uparrow$
Norm-AS	Baseline	32	90.9	86.8	89.4	46.6
Norm-AS	Baseline	64	91.2	87.5	91.7	49.8
Norm-AS	Baseline	128	90.5	88.8	91.4	49.3
Norm-AS	Multi-head	32	91.7	87.0	89.8	46.1
Norm-AS	Multi-head	64	91.1	<u>90.3</u>	92.3	50.3
Norm-AS	Multi-head	128	91.6	88.9	90.7	47.1
AF3AD	Baseline	32	88.4	86.2	91.0	48.1
AF3AD	Baseline	64	87.5	87.5	91.1	51.3
AF3AD	Baseline	128	90.3	87.8	91.8	50.7
AF3AD	Multi-head	32	87.1	86.4	90.0	52.4
AF3AD	Multi-head	64	<u>92.0</u>	90.2	<u>94.0</u>	<u>61.0</u>
AF3AD	Multi-head	128	97.0	93.6	97.1	67.7

preset repertoire and therefore reduces per-preset exposure under a fixed budget. However, once the head is sufficiently expressive, AF3AD becomes competitive in O-AUROC and improves localization, ultimately achieving the best O-AUROC, P-AUROC, and AUPRs at multi-head $h=128$. Crucially, under the matched multi-head $h=128$ setting, Norm-AS reaches 91.6/88.9 O/P-AUROC while AF3AD reaches 97.0/93.6, showing that the improvement is not explained by the architectural change alone but requires the diverse synthesis.

Ablation on Pseudo-Anomaly Presets To isolate the effect of pseudo-anomaly diversity, we evaluate all combinations of preset groups $\{A, B, C, D\}$ on 15 AnomalyShapeNet categories (index-0 split). For each configuration, we enforce equal per-preset exposure by balancing the sampling frequency of the selected presets during training, ensuring a fair comparison across different group combinations.

Overall, performance improves consistently as more groups are combined, with the full set (ABCD) yielding the overall strongest results, while Group B remains the most effective single-group baseline (see Table 6). These findings indicate that maintaining a diverse repertoire of pseudo-anomalies is critical for robust detection and localization.

Robustness to Noisy Data We evaluate robustness by adding Gaussian noise ($\sigma \in \{0.001, 0.002, 0.003\}$) to 15 AnomalyShapeNet categories (index-0 split), following [25] but with more classes. Each setting is repeated five times with different random seeds, and we report the mean over runs. As shown in Table 7, performance remains largely unchanged under mild perturbations ($\sigma = 0.001$), demonstrating robustness to jitter. While stronger noise gradually reduces accuracy at $\sigma = 0.002$ and $\sigma = 0.003$, the method retains meaningful detection performance, with the largest impact observed on fine-grained localization (P-AUPR).

Table 6: Preset-group ablation with equal per-preset exposure on AnomalyShapeNet (index-0 split). Mean across categories shared across all configurations (%).

Grp	O-AUROC	P-AUROC	O-AUPR	P-AUPR	Grp	O-AUROC	P-AUROC	O-AUPR	P-AUPR
A	62.5	62.4	68.0	15.3	BC	80.0	77.6	84.0	34.0
B	74.0	72.0	79.1	27.1	BD	83.3	80.0	87.7	37.9
C	64.6	63.2	71.8	17.7	CD	79.1	74.7	83.8	32.7
D	68.3	67.9	72.8	22.2	ABC	86.6	80.8	89.9	41.2
AB	81.5	78.5	83.9	37.7	ABD	88.3	<u>84.8</u>	<u>91.6</u>	<u>47.5</u>
AC	73.8	71.5	79.3	28.1	ACD	83.1	78.9	87.0	36.9
AD	75.7	73.4	80.4	28.8	BCD	89.0	83.7	91.4	47.3
					All	<u>88.9</u>	86.1	91.6	53.6

Table 7: Mean performance (% , across 15 classes) as a function of Gaussian noise level.

σ	O-AUROC	P-AUROC	O-AUPR	P-AUPR
0	93.3	97.9	84.8	59.0
0.001	93.3	97.4	84.8	58.0
0.002	82.3	88.0	77.7	34.4
0.003	76.7	80.0	79.7	29.6

Empirical Feature-space Alignment We further evaluate how well the synthesized anomalies align with real anomalies in feature space using k NN-based coverage and Maximum Mean Discrepancy (MMD) on Real3D-AD. As shown in Table 8, our synthesis achieves substantially higher coverage of test anomalies while yielding lower MMD than Norm-AS, indicating closer distributional alignment to real anomalous samples. Although this offline analysis requires access to test anomalies and is therefore not applicable in real deployment settings, it provides additional insight into the quality of the synthesized anomalies. Note that the coverage and MMD values are only proxies for the distribution of real anomalies, not proof of physical realism. They suggest that increasing geometric diversity makes synthesized anomalies closer to real anomaly regions, but fully modeling real defects remains an open problem (Further details in Section S3).

4.7 Qualitative Results

Figure 2 displays anomaly score maps that illustrate the effectiveness of our localization method on various test samples from the Real3D-AD and AnomalyShapeNet datasets. Our approach effectively identifies and highlights anomalous regions while assigning appropriate scores to normal points, a critical requirement in the 3DAD task. However, we observe a few failure cases, which are detailed in Supplementary Figure S2.

Table 8: Real3D-AD anomaly-feature alignment (mean $\pm$ std over 12 categories). Original: computed in feature space. UMAP/t-SNE: computed in 2D projections to support Figure S1.

Method	Original (feature space)		UMAP (2D)		t-SNE (2D)	
	Cov($T \rightarrow \cdot$) $\uparrow$	MMD $\downarrow$	Hull IoU $\uparrow$	Cov($T \rightarrow \cdot$) $\uparrow$	Hull IoU $\uparrow$	Cov($T \rightarrow \cdot$) $\uparrow$
Norm-AS	0.087 $\pm$ 0.130	0.326 $\pm$ 0.146	0.157 $\pm$ 0.133	0.007 $\pm$ 0.013	0.191 $\pm$ 0.105	0.002 $\pm$ 0.006
Ours	0.380$\pm$0.187	0.105$\pm$0.079	0.483$\pm$0.242	0.242$\pm$0.107	0.521$\pm$0.220	0.196$\pm$0.118

5 Discussion

AF3AD functions as a modular synthesis framework capable of generating diverse pseudo-anomalies to train detection methods that rely on synthetic data. While we demonstrate integration with two different anomaly detection methods, the framework’s modular design enables adoption by other paradigms. Our kernel-based synthesis can easily replace simplistic generators, such as PO3AD’s Norm-AS or R3D-AD’s Patch-Gen. Across both benchmarks, AF3AD improves over competing methods at the object level and stays competitive at the point level, with all gains averaged over three seeds (Section 4.4). We observe that our "Group A" presets are similar to the anomalies generated by Norm-AS, the synthesis module used in PO3AD. As shown in Table 6, integrating additional types of anomalies beyond this baseline consistently leads to better results. These gains persist under realistic perturbations, such as random noise (Table 7), demonstrating the robustness of our proposed framework. In Table 8, we empirically show that utilizing diverse anomaly presets allows us to populate wider regions of the "anomaly space" by considering factors like scale, anisotropy, boundary sharpness, displacement ratio, and affected area. This prevents the model from overfitting to a narrow range of synthetic anomalies. Furthermore, our approach enables real-time, controllable anomaly generation without requiring labeled data. This diversity acts as an inductive bias, helping to train a more powerful anomaly detector—a benefit also observed in 2D anomaly detection studies [14]. AF3AD has two main limitations. First, because it synthesizes anomalies through smooth geometric displacement of points, its coverage is uneven across real defect types: it directly covers geometric surface deformations such as bulges, dents, ridges, trenches, and shear (Table 2); it only partially covers sharp or high-frequency defects such as cracks and fine scratches, which smooth kernels approximate but do not reproduce exactly; and it does not address non-geometric defects such as material flaws or purely appearance-based anomalies, since the synthesis modifies geometry rather than surface appearance. Second, the preset configurations are fixed and not tailored to a given dataset, so if the synthesized anomalies diverge too far from the real defect distribution, generalization may suffer. Incorporating minimal domain guidance and developing a principled way to predict which anomaly types are most effective for a specific dataset are therefore promising directions for future work.




Fig. 2: Samples of anomaly prediction using offset prediction method trained with AF3AD.

6 Conclusion

In this work, we proposed AF3AD, a modular framework for synthesizing diverse pseudo-anomalies to generate training signals for 3D anomaly detection. Our framework provides a systematic taxonomy of geometric deformations with explicit parametric control, operating independently of specific detection architectures. We demonstrated its effectiveness by integrating AF3AD with an offset-prediction detector, achieving superior performance on AnomalyShapeNet (91.5% O-AUROC) and Real3D-AD (85.2% O-AUROC), with gains of 5.5 and 7.1 points over the previous best methods, respectively. Extensive ablations confirm that diverse synthesis strategies consistently outperform limited baselines and remain robust under noise. Future work will explore integrating generative modeling for richer anomaly representations and implementing curriculum learning to progressively introduce complex anomalies.

References

1. Bergmann, P., Jin, X., Sattlegger, D., Steger, C.: The mvtec 3d-ad dataset for unsupervised 3d anomaly detection and localization. CoRR **abs/2112.09045** (2021)
2. Cao, Y., Xu, X., Shen, W.: Complementary pseudo multimodal feature for point cloud anomaly detection. Pattern Recognit. **156**, 110761 (2024).
3. Cao, Y., Xu, X., Zhang, J., Cheng, Y., Huang, X., Pang, G., Shen, W.: A survey on visual anomaly detection: Challenge, approach, and prospect. CoRR **abs/2401.16402** (2024).
4. Chang, A.X., Funkhouser, T.A., Guibas, L.J., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F.: Shapenet: An information-rich 3d model repository. CoRR **abs/1512.03012** (2015)
5. Cheng, Y., Cao, Y., Wang, D., Shen, W., Li, W.: Boosting global-local feature matching via anomaly synthesis for multi-class point cloud anomaly detection. IEEE Trans Autom. Sci. Eng. **22**, 12560–12571 (2025).
6. Cheng, Y., Sun, Y., Zhang, H., Shen, W., Cao, Y.: Towards high-resolution 3d anomaly detection: A scalable dataset and real-time framework for subtle industrial defects. In: Koenig, S., Jenkins, C., Taylor, M.E. (eds.) Fortieth AAAI Conference on Artificial Intelligence, Thirty-Eighth Conference on Innovative Applications of Artificial Intelligence, Sixteenth Symposium on Educational Advances in Artificial Intelligence, AAAI 2026, Singapore, January 20–27, 2026. pp. 3327–3334. AAAI Press (2026).
7. Choy, C.B., Gwak, J., Savarese, S.: 4d spatio-temporal convnets: Minkowski convolutional neural networks. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16–20, 2019. pp. 3075–3084. Computer Vision Foundation / IEEE (2019).
8. Deng, H., Li, X.: Anomaly detection via reverse distillation from one-class embedding. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18–24, 2022. pp. 9727–9736. IEEE (2022).
9. Du, J., Tao, C., Cao, X., Tsung, F.: 3d vision-based anomaly detection in manufacturing: A survey. Frontiers of Engineering Management **12**(2), 343–360 (Jun 2025).
10. Gudovskiy, D.A., Ishizaka, S., Kozuka, K.: CFLOW-AD: real-time unsupervised anomaly detection with localization via conditional normalizing flows. In: IEEE/CVF Winter Conference on Applications of Computer Vision, WACV 2022, Waikoloa, HI, USA, January 3–8, 2022. pp. 1819–1828. IEEE (2022).
11. Horwitz, E., Hoshen, Y.: Back to the feature: Classical 3d features are (almost) all you need for 3d anomaly detection. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023 - Workshops, Vancouver, BC, Canada, June 17–24, 2023. pp. 2968–2977. IEEE (2023).
12. Huang, H., Wang, P., Pei, J., Wang, J., Alexanian, S., Niyato, D.: Deep learning advancements in anomaly detection: A comprehensive survey. IEEE Internet Things J. **12**(21), 44318–44342 (2025).
13. Jolliffe, I.T.: Principal Component Analysis. Springer Series in Statistics, Springer (1986).
14. Kim, H., Lee, C.: Enhancing anomaly detection via generating diversified and hard-to-distinguish synthetic anomalies. In: Proceedings of the 33rd ACM International Conference on Information and Knowledge Management. p. 1089–1098. CIKM '24, Association for Computing Machinery, New York, NY, USA (2024).

15. Li, W., Xu, X., Gu, Y., Zheng, B., Gao, S., Wu, Y.: Towards scalable 3d anomaly detection and localization: A benchmark via 3d anomaly synthesis and A self-supervised learning network. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024. pp. 22207–22216. IEEE (2024).
16. Liang, H., Xie, G., Hou, C., Wang, B., Gao, C., Wang, J.: Look inside for more: Internal spatial modality perception for 3d anomaly detection. In: AAAI Conference on Artificial Intelligence (2024).
17. Lin, Y., Chang, Y., Tong, X., Yu, J., Liotta, A., Huang, G., Song, W., Zeng, D., Wu, Z., Wang, Y., Zhang, W.: A survey on rgb, 3d, and multimodal approaches for unsupervised industrial image anomaly detection. *Inf. Fusion* **121**, 103139 (2025).
18. Liu, J., Xie, G., Chen, R., Li, X., Wang, J., Liu, Y., Wang, C., Zheng, F.: Real3d-ad: A dataset of point cloud anomaly detection. In: Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S. (eds.) *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023* (2023)
19. Pang, Y., Wang, W., Tay, F.E.H., Liu, W., Tian, Y., Yuan, L.: Masked autoencoders for point cloud self-supervised learning. In: Avidan, S., Brostow, G.J., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part II. Lecture Notes in Computer Science*, vol. 13662, pp. 604–621. Springer (2022).
20. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017. pp. 77–85. IEEE Computer Society (2017).
21. Qin, J., Gu, C., Yu, J., Zhang, C.: Teacher-student network for 3d point cloud anomaly detection with few normal samples. *Expert Syst. Appl.* **228**, 120371 (2023).
22. Roth, K., Pemula, L., Zepeda, J., Schölkopf, B., Brox, T., Gehler, P.V.: Towards total recall in industrial anomaly detection. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022. pp. 14298–14308. IEEE (2022).
23. Wang, Y., Peng, J., Zhang, J., Yi, R., Wang, Y., Wang, C.: Multimodal industrial anomaly detection via hybrid fusion. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023. pp. 8032–8041. IEEE (2023).
24. Yang, J., Shi, Y., Qi, Z.: DFR: deep feature reconstruction for unsupervised anomaly segmentation. *CoRR* **abs/2012.07122** (2020)
25. Ye, J., Zhao, W., Yang, X., Cheng, G., Huang, K.: PO3AD: predicting point offsets toward better 3d point cloud anomaly detection. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025. pp. 1353–1362. Computer Vision Foundation / IEEE (2025).
26. Yu, X., Tang, L., Rao, Y., Huang, T., Zhou, J., Lu, J.: Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022. pp. 19291–19300. IEEE (2022).
27. Yu, Y., Chen, Z., Xu, X., Zhang, L., Yang, H., Nie, Y., He, S.: Registration is a powerful rotation-invariance learner for 3d anomaly detection. *CoRR* **abs/2510.16865** (2025).

28. Zhou, Z., Wang, L., Fang, N., Wang, Z., Qiu, L., Zhang, S.: R3D-AD: reconstruction via diffusion for 3d anomaly detection. In: Leonardis, A., Ricci, E., Roth, S., Russakovsky, O., Sattler, T., Varol, G. (eds.) *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part XXXVI*. Lecture Notes in Computer Science, vol. 15094, pp. 91–107. Springer (2024).
29. Zhu, H., Xie, G., Hou, C., Dai, T., Gao, C., Wang, J., Shen, L.: Towards high-resolution 3d anomaly detection via group-level feature contrastive learning. In: Cai, J., Kankanhalli, M.S., Prabhakaran, B., Boll, S., Subramanian, R., Zheng, L., Singh, V.K., César, P., Xie, L., Xu, D. (eds.) *Proceedings of the 32nd ACM International Conference on Multimedia, MM 2024, Melbourne, VIC, Australia, 28 October 2024 - 1 November 2024*. pp. 4680–4689. ACM (2024).
30. Zuo, Z., Dong, J., Wu, Y., Qu, Y., Wu, Z.: CLIP3D-AD: extending CLIP for 3d few-shot anomaly detection with multi-view images generation. *CoRR* **abs/2406.18941** (2024).