

Towards High-Resolution 3D Anomaly Detection: A Scalable Dataset and Real-Time Framework for Subtle Industrial Defects

Yuqi Cheng^{*1}, Yihan Sun^{*1}, Hui Zhang², Weiming Shen^{† 1}, Yunkang Cao^{† 2}

¹ State Key Laboratory of Intelligent Manufacturing Equipment and Technology,
Huazhong University of Science and Technology

² School of Artificial Intelligence and Robotics, Hunan University

{yuqicheng, yihansun}@hust.edu.cn, zhanghuihy@126.com, {wshen, caoyunkang}@ieee.org

Abstract

In industrial point cloud analysis, detecting subtle anomalies demands high-resolution spatial data, yet prevailing benchmarks emphasize low-resolution inputs. To address this disparity, we propose a scalable pipeline for generating realistic and subtle 3D anomalies. Employing this pipeline, we developed **MiniShift**, the inaugural high-resolution 3D anomaly detection dataset, encompassing 2,577 point clouds, each with 500,000 points and anomalies occupying less than 1% of the total. We further introduce **Simple3D**, an efficient framework integrating Multi-scale Neighborhood Descriptors (MSND) and Local Feature Spatial Aggregation (LFSA) to capture intricate geometric details with minimal computational overhead, achieving real-time inference exceeding 20 fps. Extensive evaluations on **MiniShift** and established benchmarks demonstrate that **Simple3D** surpasses state-of-the-art methods in both accuracy and speed, highlighting the pivotal role of high-resolution data and effective feature aggregation in advancing practical 3D anomaly detection.

HomePage —

<https://hustcyq.github.io/MiniShift-Simple3D>

Code — <https://github.com/hustCYQ/MiniShift-Simple3D>

Datasets —

<https://huggingface.co/datasets/ChengYuQi99/MiniShift>

Introduction

Despite the exceptional precision of modern manufacturing processes, subtle imperfections often persist, eluding conventional inspection techniques and introducing significant safety risks in downstream applications (Cheng et al. 2024a; Yu, Guo, and Li 2025; Lin et al. 2024). Unsupervised point cloud anomaly detection emerges as a compelling strategy for identifying these hard-to-detect flaws (Cao et al. 2024; Wang et al. 2023; Ye et al. 2025). However, the efficacy of such methods is curtailed by their reliance on the widely utilized low-resolution evaluation framework. This dependency creates a pronounced gap between cutting-edge research and the rigorous demands of industrial implementation.

^{*}These authors contributed equally.

[†]corresponding author.

Conventionally, prevalent approaches for 3D anomaly detection begin by selecting a sparse subset of points (typically around 1k) from the original point cloud. These methods encode local neighborhood features to form point groups and assign an anomaly score to each group. Reconstruction-based techniques, such as IMRNet (Li and Xu 2024) and R3D-AD (Zhou et al. 2024), transform each group into representative tokens, which are then reconstructed via a point transformer architecture. The anomaly score is derived by computing the feature-wise discrepancy between the original and reconstructed tokens. In contrast, prototype-based methods like ISMP (Liang et al. 2025) and GLFM (Cheng et al. 2025b) utilize a pre-trained Point Transformer to extract discriminative features from the point groups, and quantify anomalies based on their deviation from learned prototype representations. To generate dense anomaly maps, these sparse group-level scores are interpolated across the full point cloud. However, such interpolation inevitably compromises the spatial granularity necessary for detecting fine-grained defects. For example, an imperfection measuring just 1mm × 1mm on a 20mm × 20mm surface becomes almost imperceptible when the point cloud is reduced to 1k points. As illustrated in Figure 1 (a), a subtle protrusion is clearly discernible in the full-resolution 500k point cloud and remains detectable at 8k points, yet it vanishes almost entirely after downsampling to 1k points. This observation highlights the critical need for high-resolution anomaly detection frameworks that preserve spatial fidelity to effectively identify subtle industrial defects.

To detect such subtle anomalies, industrial inspection systems often capture hundreds of thousands to over a million points per object (Cheng et al. 2024b). Although datasets like MVTec 3D-AD (Bergmann et al. 2022), Real3D-AD (Liu et al. 2024), Anomaly-ShapeNet (Li and Xu 2024), and MulSenAD (Li et al. 2024) offer high-resolution point clouds, their anomalies tend to be relatively prominent, contrasting sharply with the exceedingly subtle irregularities encountered in real-world contexts. To bridge this divide between public datasets and practical applications, we propose an innovative pipeline for synthesizing realistic, subtle anomalies within high-resolution point clouds. Additionally, to systematically evaluate detection performance across varying scales, we introduce a multi-level difficulty protocol, classifying anomalies into three tiers—easy, medium,

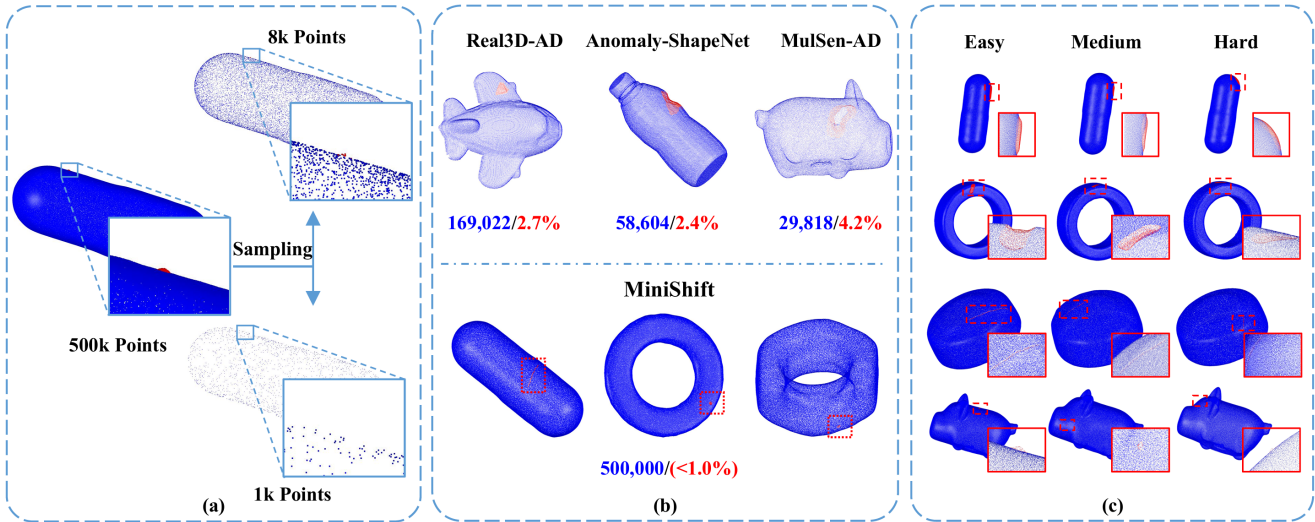


Figure 1: (a) Detecting subtle defects requires high-resolution point clouds to ensure precise localization and accurate identification. (b) The proposed **MiniShift** dataset features high-density point clouds with approximately 500k points per sample, where anomalies occupy less than 1% of the total surface area—posing a significantly more challenging scenario than existing benchmarks. Blue numbers indicate the average number of points per sample, while red numbers denote the anomaly occupancy rate. (c) Visualization of representative samples from MiniShift. Each row corresponds to a different defect category: Areal, Striate, Scratch, and Sphere.

and hard—based on their geometric perturbations and visual detectability. Applying this pipeline to a subset of normal samples from MulSenAD (Li et al. 2024), we establish a novel high-resolution benchmark, **MiniShift**. As illustrated in Figure 1 (b), MiniShift features point clouds with 500k points, far surpassing the resolution of existing datasets (typically below 170k points). Crucially, its anomalies constitute less than 1% of the surface area, necessitating the use of high-resolution spatial features and posing a markedly greater challenge than prior benchmarks.

We conducted an extensive evaluation of state-of-the-art (SOTA) methods on MiniShift using point clouds with 8k points, revealing that these approaches fall short in both performance and efficiency. This inadequacy stems primarily from the limitations of CNN- (He et al. 2016) and ViT- (Zhao et al. 2021) based backbones in applying high-resolution point cloud detection: (i) prohibitive computational complexity that scales with input group quantity, and (ii) inaccurate representation of local geometric information. In response, we revisit the utility of handcrafted point cloud descriptors and present a streamlined yet robust baseline, **Simple3D**, tailored for high-resolution anomaly detection. Simple3D integrates two novel components: the Multi-scale Neighborhood Descriptor (MSND) and Local Feature Spatial Aggregation (LFSA). By computing multi-scale descriptors for each point and hierarchically aggregating local spatial features through progressive downsampling, Simple3D achieves a detailed characterization of local geometry. Leveraging lightweight descriptors, it delivers substantially improved accuracy and efficiency over existing methods. Comprehensive experiments affirm that Simple3D attains SOTA performance not only on MiniShift but also across

traditional benchmarks, while sustaining real-time inference speeds exceeding 20 FPS. In summary, our contributions are as follows:

- We present MiniShift, a high-resolution 3D anomaly detection dataset comprising 2,577 point clouds, each with 500k points, accompanied by a triple-level difficulty framework to thoroughly assess detection methods.
- We introduce Simple3D, a high-resolution anomaly detection approach with two core components—MSND and LFSA—that markedly enhances the identification of subtle geometric anomalies while preserving computational efficiency.
- Through rigorous experimentation, we demonstrate that Simple3D consistently surpasses existing SOTA methods across MiniShift, and other existing datasets including Real3D-AD, Anomaly-ShapeNet, and MulSenAD, achieving superior accuracy and real-time performance.
- Our systematic analysis underscores the critical role of high-resolution input in precise 3D anomaly detection, revealing that judiciously aggregated local geometric features serve as effective representations for this purpose.

MiniShift Dataset

Pipeline

The development of the MiniShift dataset is illustrated in Figure 2. We commence by extracting 3D data from 12 distinct categories of typical industrial components within the MulSen-AD dataset. These data undergo dense sampling to produce high-resolution point clouds. Subsequently, we propose the Anchor-Guided Geometric Anomaly Synthesis (AG-GAS) method, which facilitates the creation of

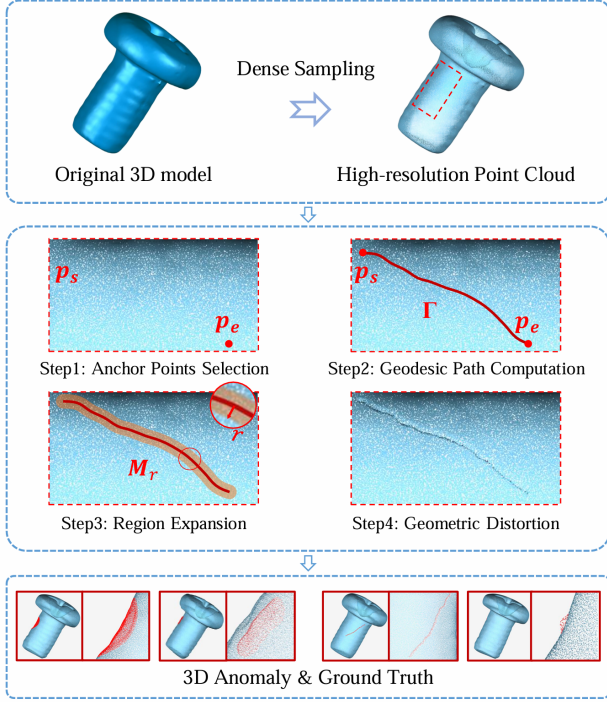


Figure 2: **The construction pipeline of MiniShift.** The 3D models are utilized to generate various defects through the proposed AG-GAS framework.

four anomaly types—namely Areal, Striate, Scratch, and Sphere—alongside their associated defect masks. This synthesis is achieved through precise manipulation of anchor point positions and geometric deformation parameters. Additionally, we implement a difficulty protocol, categorizing the dataset into three subsets—easy, medium, and hard—to replicate real-world industrial scenarios characterized by varying detection complexities.

Anchor-Guided Geometric Anomaly Synthesis

We introduce an automated framework engineered to synthesize realistic and diverse 3D defects, as depicted in Figure 2. The process begins with the random selection of two anchor points, followed by the determination of connecting path points. These path points are then expanded into a local neighborhood, delineating the deformation region wherein a stretching operation generates the defect. Unlike anomaly synthesis methods in R3D (Zhou et al. 2024) and GLFM (Cheng et al. 2025b), which predominantly simulate basic protrusions or depressions within localized circular zones, our framework affords the flexibility to craft defects of diverse shapes and scales through minimal parameter adjustments. This adaptability enhances the simulation of the randomness and intricacy characteristic of industrial defects. The framework is organized into four sequential stages:

Anchor Points Selection: Two anchor points, denoted p_s and $p_e \in P$, are selected from the input point cloud $P \in \mathbb{R}^{n \times 3}$, comprising n points, to regulate the position and span of the synthetic anomaly.

Geodesic Path Computation: As illustrated in Figure 2, we calculate the geodesic path points between the anchor points to support subsequent anomaly generation. Initially, we construct an undirected graph $G = (\{p_i\}, \{e_{ij}\})$ by defining edges $e_{ij} = \langle p_i, p_j \rangle$, where $p_i \in P$ and p_j represents one of the k nearest neighbors of p_i within P . Dijkstra’s algorithm (Frana and Misa 2010) is then applied to ascertain the shortest path π^* between p_s and p_e , yielding the geodesic path point set Γ :

$$\pi^* = \arg \min_{\pi \in \Pi(p_s, p_e)} \sum_{e_{ij} \in \pi} w(e_{ij}) \quad (1)$$

$$\Gamma = \{\{p_i, p_j\} \mid e_{ij} \in \pi^*\} \quad (2)$$

where $\Pi(p_s, p_e)$ denotes all possible paths from p_s to p_e , π is a path within this set, and $w(e_{ij})$ is the Euclidean distance between p_i and p_j .

Region Expansion: Utilizing the geodesic path Γ as the central axis, we define the mask region M_r through expansion with a control radius r , encompassing points in P that contribute to anomaly generation:

$$M_r = \{p_j \in P \mid \|p_j - p_i\| < r, \forall p_i \in \Gamma\} \quad (3)$$

Geometric Distortion: The distortion direction is established by computing the average normal vector \bar{n}_{avg} of points within M_r :

$$\bar{n}_{avg} = \frac{\bar{n}}{\|\bar{n}\|_2}, \quad \bar{n} = \frac{1}{|M_r|} \sum_{p_j \in M_r} n_j \quad (4)$$

where n_j is the normal vector of point p_j , and $|M_r|$ indicates the number of points in M_r . The distortion, controlled by a parameter $dir \in \{1, -1\}$ (1 for protrusion, -1 for depression), shifts points in M_r by a distance d along \bar{n}_{avg} . To ensure smoothness, the stretching distance is scaled according to proximity to the central axis, yielding the anomaly point:

$$p'_j = p_j + dir \cdot \bar{n}_{avg} \cdot \left(1 - \frac{d_j}{d_{\max}}\right) \cdot d, \quad d_{\max} = \max_{p_i \in M_r} d_i \quad (5)$$

where $d_j = \min_{p_m \in \Gamma} \|p_j - p_m\|$.

Through precise parameterization, our framework synthesizes diverse defect types for the MiniShift dataset, selecting four representative categories—Areal, Striate, Scratch, and Sphere—reflecting prevalent industrial defect patterns. Areal and Striate defects, marked by subtle distortions, commonly result from uneven mechanical stress or thermal gradients in machining and casting. Conversely, Scratch and Sphere anomalies, with minimal surface coverage, typically arise from impacts or friction during milling, handling, or assembly. We generate 30 unique samples per category and defect type.

Difficulty Protocol

To tackle the dual objectives of industrial inspection realism and algorithmic performance evaluation, we devise a difficulty protocol that classifies synthesized anomalies

into three tiers—easy, medium, and hard—based on geometric saliency and perceptual visibility. Large-scale deformations and expansive defects are readily detectable by human inspectors and automated systems, whereas subtle anomalies (e.g., minor protrusions or micro-scratches) exhibit higher oversight rates. The easy subset features anomalies with prominent distortions or broad spatial extent, facilitating identification, while the hard subset includes minute or nearly imperceptible defects, mirroring critical industrial edge cases prone to missed detection. Figure 1 (c) showcases anomaly samples across difficulty levels. This protocol augments the dataset’s hierarchical representation of real-world defect complexity and establishes a benchmark for assessing algorithmic robustness and generalization.

The parameters governing anomaly difficulty are outlined in Table 1, adjusted via the ranges of α , β , and γ , defined as:

$$\alpha = l/D, \quad \beta = r/D, \quad \gamma = d/D, \quad (6)$$

where $D = \|\max(\mathbf{P}) - \min(\mathbf{P})\|_2$ represents the point cloud’s bounding box diagonal, l is the geodesic distance between \mathbf{p}_s and \mathbf{p}_e .

Protocol		Easy	Medium	Hard
Areal	α	0.05 ~ 0.1	0.05 ~ 0.1	0.05 ~ 0.1
	β	0.04 ~ 0.07	0.04 ~ 0.07	0.04 ~ 0.07
	γ	$5 \times 10^{-3} \sim 7 \times 10^{-3}$	$3 \times 10^{-3} \sim 5 \times 10^{-3}$	$1 \times 10^{-3} \sim 3 \times 10^{-3}$
Striate	α	0.02 ~ 0.1	0.02 ~ 0.1	0.02 ~ 0.1
	β	0.01 ~ 0.03	0.01 ~ 0.03	0.01 ~ 0.03
	γ	$5 \times 10^{-3} \sim 7 \times 10^{-3}$	$3 \times 10^{-3} \sim 5 \times 10^{-3}$	$1 \times 10^{-3} \sim 3 \times 10^{-3}$
Scratch	α	0.3 ~ 0.4	0.2 ~ 0.3	0.1 ~ 0.2
	β	$3 \times 10^{-3} \sim 4 \times 10^{-3}$	$2 \times 10^{-3} \sim 3 \times 10^{-3}$	$1 \times 10^{-3} \sim 2 \times 10^{-3}$
	γ	$1 \times 10^{-3} \sim 5 \times 10^{-3}$	$1 \times 10^{-3} \sim 5 \times 10^{-3}$	$1 \times 10^{-3} \sim 5 \times 10^{-3}$
Sphere	α	$< 1 \times 10^{-3}$	$< 1 \times 10^{-3}$	$< 1 \times 10^{-3}$
	β	$7 \times 10^{-3} \sim 9 \times 10^{-3}$	$5 \times 10^{-3} \sim 7 \times 10^{-3}$	$3 \times 10^{-3} \sim 5 \times 10^{-3}$
	γ	$1 \times 10^{-3} \sim 5 \times 10^{-3}$	$1 \times 10^{-3} \sim 5 \times 10^{-3}$	$1 \times 10^{-3} \sim 5 \times 10^{-3}$

Table 1: **Difficulty Protocol**. Color-filled cells highlight the key indicators that determine the difficulty of the category.

Simple3D Method

Overview

To address the challenges posed by high-resolution detection, we present Simple3D, a simple yet effective framework. The overall architecture of this framework is illustrated in Figure 3. Unlike existing approaches that commonly utilize CNNs or ViTs for point cloud feature extraction, resulting in prohibitive computational costs for high-resolution detection, our method focuses on extracting local point features using handcrafted descriptors. These features are subsequently aggregated across spatial hierarchies to enhance the detection of anomalies. Simple3D employs a prototype-based detection mechanism, where anomaly scores are assigned by quantifying the deviations of features from established normal prototypes.

Multi-Scale Neighborhood Description (MSND)

Let the input point cloud be denoted by $\mathbf{P} \in \mathbb{R}^{n \times 3}$, where n represents the total number of points. For each point $\mathbf{p}_i \in \mathbf{P}$, we identify its k nearest neighbors to construct a point set

\mathbf{R}_i . To harness the complementary information inherent in local point distributions across various neighborhood scales, we employ multiple neighbor counts k_1, k_2, \dots, k_m , thereby generating a sequence of point sets $\mathbf{R}_{i1}, \mathbf{R}_{i2}, \dots, \mathbf{R}_{im}$. Subsequently, a local feature operator \mathbf{f} is applied to extract feature descriptors from each point set. These descriptors are concatenated to form the multi-scale local feature representation \mathbf{F}_i of \mathbf{p}_i :

$$\mathbf{F}_i = \text{Concat}(\mathbf{R}_{i1}, \mathbf{R}_{i2}, \dots, \mathbf{R}_{im}) \quad (7)$$

Local Feature Spatial Aggregation (LFSA)

To further refine the feature representation, we randomly sample t points from \mathbf{P} and aggregate the MSND within their respective neighborhoods. This process yields enhanced point features that possess both an expanded receptive field and enriched local geometric information. Specifically, let the sampled points be denoted as \mathbf{p}_{s_j} for $0 \leq j \leq t$ and $0 \leq s_j \leq n$, with the neighbor count specified as k_L . The enhanced point feature $\mathbf{F}_{s_j}^A$ for \mathbf{p}_{s_j} is computed as:

$$\mathbf{F}_{s_j}^A = \frac{1}{k_L} \sum_{\mathbf{F} \in \mathbf{R}_{s_j k_L}} \mathbf{F} \quad (8)$$

Anomaly Detection

The enhanced point features derived from normal point clouds are utilized to establish a normal prototype set $\mathcal{S} = \{\mathbf{F}_{s_1}^A, \mathbf{F}_{s_2}^A, \dots, \mathbf{F}_{s_t}^A\}$. For the enhanced point features $\mathbf{F}_{\text{test}}^A$ of test point clouds, deviations from this normal prototype set are interpreted as anomaly scores. The point-wise anomaly scores \mathbf{A} and the object-wise anomaly score ξ are determined by:

$$\mathbf{A} = \|\mathbf{F}_{\text{test}}^A - \mathbf{F}^*\| \quad (9)$$

$$\xi = \max(\mathbf{A}) \quad (10)$$

where $\mathbf{F}^* = \min_{\mathbf{s} \in \mathcal{S}} \|\mathbf{s} - \mathbf{F}_{\text{test}}^A\|$.

Experiment

Experimental Setups

Datasets: Our experimental evaluation encompasses the proposed **MiniShift** dataset alongside three established public benchmarks: **Real3D-AD** (Liu et al. 2024), **Anomaly-ShapeNet** (Li and Xu 2024), and **MulSen-AD** (Li et al. 2024).

Implementation Details: For point cloud feature extraction, our approach utilizes the handcrafted descriptor FPFH (Rusu, Blodow, and Beetz 2009) due to its efficient CUDA-accelerated implementation available in Open3D (Zhou, Park, and Koltun 2018). The multi-scale ranges are set to 40, 80, and 120 by default, and the number of aggregation points is configured to 128. All experiments are conducted on NVIDIA A100 GPUs.

Evaluation Metrics: To assess the efficacy of our method, we employ the widely recognized Area Under the Receiver

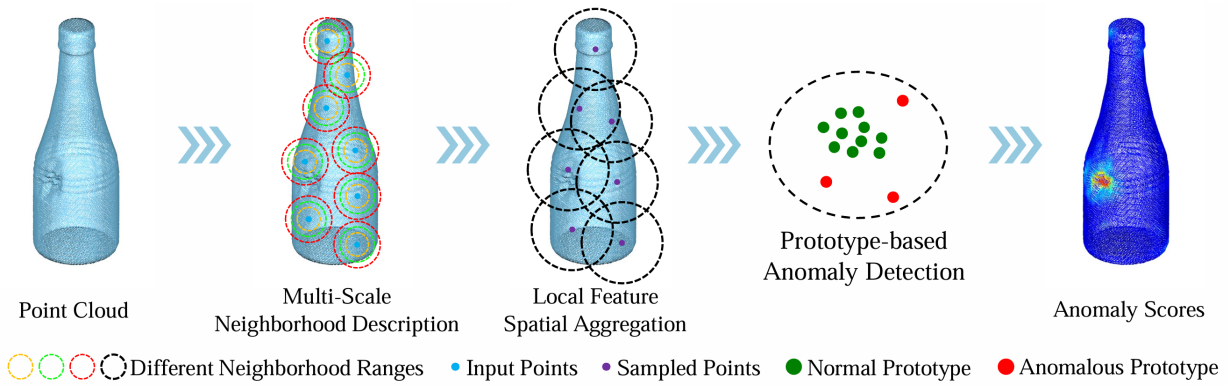


Figure 3: **Framework of the proposed method, Simple3D.** High-resolution point clouds undergo initial processing to derive multi-scale neighborhood descriptors. These descriptors are then spatially aggregated to yield enhanced point cloud features. Simple3D utilizes a prototype-based approach to facilitate both object-wise and point-wise anomaly detection.

Method →	PatchCore-FP	PatchCore-PM	R3D-AD	GLFM	Simple3D
Level ↓	CVPR'22	CVPR'22	ECCV'24	TASE'25	Ours
Easy	<u>68.3</u> /56.5	57.2/57.3	56.3/48.3	57.7/ <u>67.7</u>	75.6/77.3
Medium	<u>65.6</u> /54.1	55.8/52.6	50.8/50.7	55.0/ <u>57.3</u>	68.6/65.5
Hard	<u>61.4</u> /51.1	54.3/52.0	50.1/50.6	52.8/ <u>52.7</u>	61.6/56.3
ALL	<u>65.1</u> /53.7	55.9/54.1	53.6/49.8	55.8/ <u>58.7</u>	68.6/66.2

Table 2: **Quantitative Results on MiniShift.** The results are presented in O-ROC%/P-ROC%. The best performance is in **bold**, and the second best is underlined.

Operating Characteristic curve (AUROC) metric. Specifically, we calculate both object-wise AUROC (O-ROC) and point-wise AUROC (P-ROC) to gauge performance across the four datasets.

Comparison Methods: Our proposed method is rigorously benchmarked against a diverse array of recent SOTA approaches. These include prototype-based methods such as PatchCore (Roth et al. 2022; Horwitz and Hoshen 2023), BTF (Horwitz and Hoshen 2023), M3DM (Wang et al. 2023), CPMF (Cao, Xu, and Shen 2024), Reg3D-AD (Liu et al. 2024), Group3AD (Zhu et al. 2024), and GLFM (Cheng et al. 2025b); reconstruction-based methods like R3D-AD (Zhou et al. 2024), IMRNet (Li and Xu 2024), and MC3D-AD (Cheng et al. 2025a); and the regression-based method PO3AD (Ye et al. 2025). Notably, M3DM and PatchCore are evaluated with various feature types, denoted as “-FP” for FPFH (Rusu, Blodow, and Beetz 2009), “-PM” for PointMAE (Pang et al. 2023), “-PB” for Point-BERT (Yu et al. 2022).

Benchmarking Results on MiniShift

The quantitative assessment of all competing methods on MiniShift is delineated in Table 2. Our approach consistently surpasses all baseline methods across diverse difficulty tiers, underscoring its exceptional generalization and robustness. As the dataset complexity escalates, existing methods manifest a marked deterioration in performance, particularly in the most arduous scenarios, where their efficacy in both O-ROC and P-ROC metrics approaches futility. Conversely,

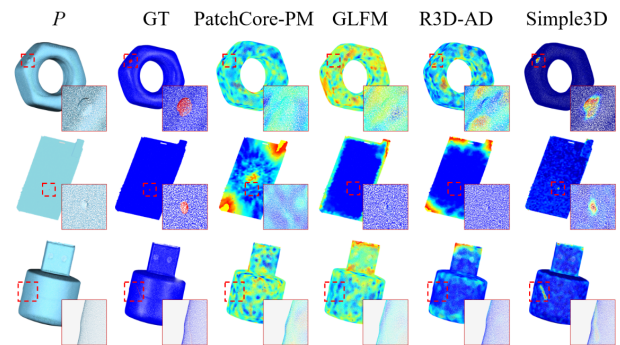


Figure 4: **Visualization of our point-wise anomaly detection results on MiniShift.** From top to bottom: input point clouds, ground truths, anomaly maps.

while our method also incurs a modest decline in performance, it sustains a decisive and substantial advantage over all rival techniques. This superior performance is chiefly ascribed to two pivotal components of our methodology, which enable the extraction of highly discriminative and expressive features. As a result, our approach adeptly discerns even the most subtle and intricate anomalies that elude detection by other methods.

The qualitative analysis, as depicted in Figure 4, corroborates these quantitative outcomes. Our method, Simple3D, produces anomaly maps of remarkable sharpness and precision across multiple categories, even amidst intricate and challenging conditions. In stark contrast, the majority of baseline methods struggle to delineate defects from the background effectively, thereby underscoring the critical necessity of high-resolution representations and meticulous local feature descriptions for proficient anomaly detection.

Benchmarking Results on Existing Datasets

To further ascertain the efficacy of Simple3D, we extend our evaluation to three prominent public benchmarks: Real3D-AD, Anomaly-ShapeNet, and MulSen-AD. The outcomes are encapsulated in Table 3. Simple3D uniformly sets new

Method → Real3D-AD	CPMF PR'24	Reg3D-AD NeurIPS'23	Group3AD ACM MM'24	IMRNet CVPR'24	ISMP AAAI'25	GLFM TASE'25	PO3AD CVPR'25	MC3D-AD IJCAI'25	Simple3D Ours
O-ROC/P-ROC	62.5/75.9	70.4/70.5	75.1/73.5	72.5/-	76.7/ <u>83.6</u>	75.0/76.7	76.5/-	<u>78.2</u> /76.8	80.4/92.3
Method → Anomaly-ShapeNet	M3DM-PM CVPR'23	CPMF PR'24	Reg3D-AD NeurIPS'23	IMRNet CVPR'24	ISMP AAAI'25	GLFM TASE'25	PO3AD CVPR'25	MC3D-AD IJCAI'25	Simple3D Ours
O-ROC/P-ROC	51.1/54.9	55.9/-	57.2/-	66.1/65.0	75.7/69.1	61.9/74.5	83.9/ <u>89.8</u>	<u>84.2</u> /74.8	86.0/92.9
Method → MulSen-AD	M3DM-PM CVPR'23	M3DM-PB CVPR'23	PatchCore-FP CVPR'22	PatchCore-FP-R CVPR'22	PatchCore-PM CVPR'22	IMRNet CVPR'24	Reg3D-AD NeurIPS'23	GLFM TASE'25	Simple3D Ours
O-ROC/P-ROC	62.8/58.7	70.5/61.1	<u>86.0</u> /64.0	83.3/62.0	84.0/60.5	60.1/46.7	74.9/64.1	78.5/66.5	88.2/80.3

Table 3: **Quantitative Results on Real3D-AD, Anomaly-ShapeNet and MulSen-AD.** The results are presented in O-ROC%/P-ROC%. The best performance is in **bold**, and the second best is underlined.




Figure 5: **Visualization of our point-wise anomaly detection results on Real3D-AD, Anomaly-ShapeNet, and MulSen-AD.** From top to bottom: input point clouds, ground truths, anomaly maps.

SOTA benchmarks across all datasets for both O-ROC and P-ROC metrics. Moreover, our method exhibits robust per-category performance and precise localization of fine-grained anomalies, thereby affirming its generalizability and discriminative prowess across a spectrum of 3D anomaly types.

Real3D-AD: On the Real3D-AD dataset, Simple3D attains scores of 80.4% and 92.3% for O-ROC and P-ROC, respectively, thereby eclipsing prior SOTA methods such as MC3D-AD and ISMP by margins of 2.2% and 8.7% in the respective metrics. The comprehensive per-category analysis, available at arXiv, reveals that Simple3D secures the highest object-wise or point-wise performance in 10 out of 12 categories, thereby attesting to its efficacy. Particularly noteworthy is the near-perfect detection achieved in the Diamond (100%/99.0%) and Car (98.1%/99.2%) categories.

Anomaly-ShapeNet: For the Anomaly-ShapeNet dataset, Simple3D records O-ROC and P-ROC scores of 86.0% and 92.9%, respectively. Significantly, our method not only outperforms MC3D-AD by 1.8% in O-ROC but also demonstrates a substantial enhancement of 18.1% in P-ROC. The per-category breakdown, available at arXiv, indicates that Simple3D achieves superior performance in 35 out of 40 categories and attains complete detection in 5 categories, further solidifying its preeminence.

MulSen-AD: On the MulSen-AD dataset, Simple3D establishes a new SOTA benchmark with scores of 88.3% and

80.3% for O-ROC and P-ROC, respectively, thereby surpassing the previous best by 2.2% and 16.2%. The per-category evaluation, available at arXiv, reveals that our method attains the highest point-wise performance in 12 out of 15 categories.

Qualitative Results and Analysis: The qualitative results, as showcased in Figure 5, exemplify the adeptness of Simple3D in precisely localizing subtle anomalies across the Real3D-AD, Anomaly-ShapeNet, and MulSen-AD datasets. These visualizations accentuate the method’s superior detection capabilities, particularly when juxtaposed with the shortcomings exhibited by competing approaches.

Ablation Study

To elucidate the impact of MSND and LFSA, we undertake a thorough ablation study across all four datasets. The results, as tabulated in Table 4, demonstrate that both components invariably augment detection performance, yielding substantial gains across all evaluated metrics.

Analysis on Detection Resolution To delve deeper into the influence of detection resolution on efficacy, we assess our method across a spectrum of resolutions ranging from 2^8 to 2^{13} , while maintaining a fixed aggregation point count of 16, utilizing the four benchmark datasets. The findings, depicted in the first row of Figure 6, reveal a consistent enhancement in all evaluation metrics concomitant with in-

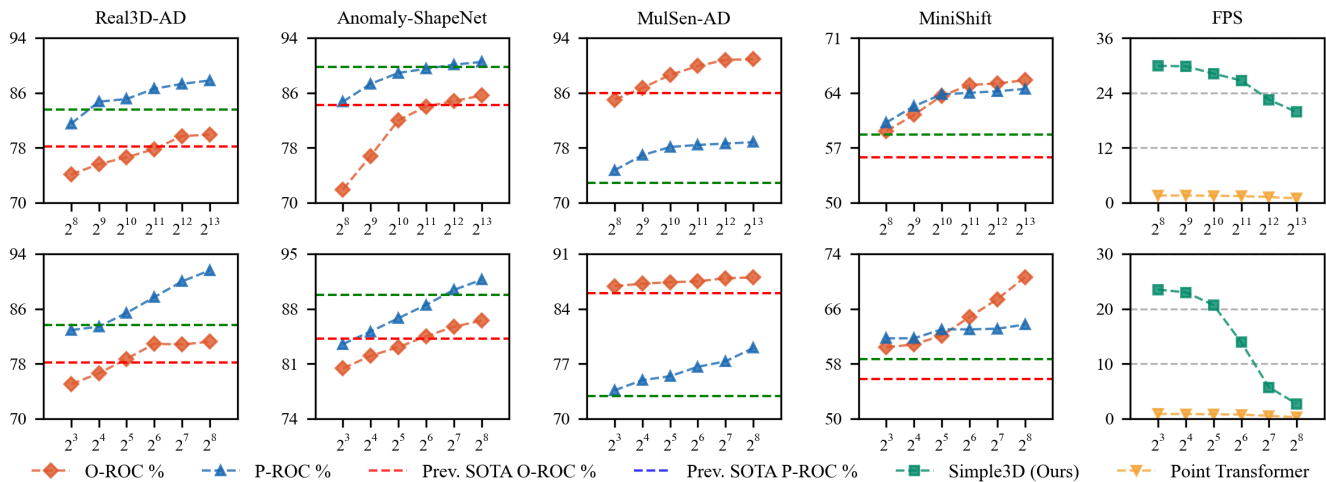


Figure 6: **Influence of detection resolution t (First Row) and aggregated neighborhood point number k_L (Second Row).** The left four figures show the performance on datasets Real3D-AD, Anomaly-ShapeNet, MulSen-AD, and MiniShift, while the last one shows the Frames Per Second (FPS), as the sampling resolution t varies from 256 (2^8) to 8192 (2^{13}) and aggregated neighborhood point number k_L varies from 8 (2^3) to 128 (2^8). The FPS is statistically analyzed on the Anomaly-ShapeNet.

MSND	✗	✓	✗	✓
LFSA	✗	✗	✓	✓
Real3D-AD	74.2/76.9	78.1/85.3	79.8/90.0	80.4/92.3
Anomaly-ShapeNet	78.8/77.0	83.8/85.6	85.7/90.4	86.0/92.9
MulSen-AD	85.1/66.2	86.8/70.7	87.1/77.3	88.2/80.3
MiniShift	67.2/56.7	67.8/57.9	67.4/63.1	68.6/66.2

Table 4: **Ablation results of MSND and LFSA between four datasets.**

creased resolution across all datasets. Notably, at resolutions exceeding 2^{12} , our method surpasses the performance of prior SOTA techniques. These observations underscore the pivotal importance of high-resolution representations in bolstering the accuracy and dependability of anomaly detection.

Analysis on Number of Aggregated Points Simple3D incorporates LFSA to expand the receptive field of local features for better feature distinguishability. To quantify the effect of the aggregated point count on detection performance, we systematically vary the neighborhood size from 2^3 to 2^8 points, while keeping the input resolution constant at 4k points. The outcomes, illustrated in the second row of Figure 6, demonstrate a uniform improvement in performance metrics across all datasets as the number of aggregated points increases. These findings accentuate the indispensable role of meticulous local feature description in 3D anomaly detection, with our approach occasionally outperforming features derived from learned 3D backbones.

Analysis on Runtime The runtime analysis, as presented in Figure 6, elucidates the influence of resolution and the number of aggregation points on detection speed. Notably, at a resolution of 4k points, Simple3D not only surpasses prior SOTA methods in performance but also sustains a detection speed exceeding 20 FPS, markedly outpacing all compet-

ing approaches. This real-time processing capability is of paramount significance for practical applications, such as industrial inspection.

Conclusion

We present **MiniShift**, a novel benchmark for high-resolution 3D anomaly detection, constructed via a controllable synthetic pipeline that scalably generates realistic and subtle surface defects. Our evaluation protocol classifies anomalies into three difficulty levels based on geometric saliency and perceptual visibility, ensuring a comprehensive assessment aligned with real-world industrial complexities. Current SOTA methods struggle to balance accuracy and efficiency in high-resolution contexts. To overcome this, we propose **Simple3D**, a lightweight yet potent method enabling real-time, fine-grained anomaly detection in dense point clouds. Rigorous experiments on **MiniShift** and datasets such as Real3D-AD, Anomaly-ShapeNet, and MulSenAD reveal that **Simple3D** achieves SOTA performance in both accuracy and speed, emphasizing the essential role of high-resolution representations in detecting minute defects. We assert that **MiniShift** and **Simple3D** lay robust groundwork for future advancements in high-resolution 3D anomaly detection.

Future Directions: Leveraging our anomaly generation pipeline, we aim to construct a large-scale, diverse dataset for 3D anomaly detection, facilitating the development of foundational models adept at identifying subtle anomalies across varied categories for seamless industrial integration.

Acknowledgments

This work was supported by Fundamental Research Funds for the Central Universities (HUST: 2021GCRC058) and was part by the HPC Platform of Huazhong University of

Science and Technology where the computation is completed.

References

- Bergmann, P.; Jin, X.; Sattlegger, D.; and Steger, C. 2022. The mvtec 3d-ad dataset for unsupervised 3d anomaly detection and localization. *Proceedings of the 17th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, 5: 202–213. Doi: 10.5220/0010865000003124.
- Cao, Y.; Xu, X.; and Shen, W. 2024. Complementary pseudo multimodal feature for point cloud anomaly detection. *Pattern Recognition*, 156: 110761.
- Cao, Y.; Zhang, J.; Frittoli, L.; Cheng, Y.; Shen, W.; and Boracchi, G. 2024. Adaclip: Adapting clip with hybrid learnable prompts for zero-shot anomaly detection. In *European Conference on Computer Vision*, 55–72. Springer.
- Cheng, J.; Gao, C.; Zhou, J.; Wen, J.; Dai, T.; and Wang, J. 2025a. MC3D-AD: A Unified Geometry-aware Reconstruction Model for Multi-category 3D Anomaly Detection. *IJCAI*.
- Cheng, Y.; Cao, Y.; Wang, D.; Shen, W.; and Li, W. 2025b. Boosting Global-Local Feature Matching via Anomaly Synthesis for Multi-Class Point Cloud Anomaly Detection. *IEEE Transactions on Automation Science and Engineering*, 1–1.
- Cheng, Y.; Cao, Y.; Xie, G.; Lu, Z.; and Shen, W. 2024a. Towards Zero-shot Point Cloud Anomaly Detection: A Multi-View Projection Framework. *arXiv preprint arXiv:2409.13162*.
- Cheng, Y.; Li, W.; Jiang, C.; Wang, D.; Xing, H.; and Xu, W. 2024b. MVGR: Mean-Variance Minimization Global Registration Method for Multi-view Point Cloud in Robot Inspection. *IEEE Transactions on Instrumentation and Measurement*, 1–1.
- Frana, P. L.; and Misa, T. J. 2010. An interview with edsgar w. dijkstra. *Communications of the ACM*, 53(8): 41–47.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Horwitz, E.; and Hoshen, Y. 2023. Back to the feature: classical 3d features are (almost) all you need for 3d anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2967–2976.
- Li, W.; and Xu, X. 2024. Towards Scalable 3D Anomaly Detection and Localization: A Benchmark via 3D Anomaly Synthesis and A Self-Supervised Learning Network. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Li, W.; Zheng, B.; Xu, X.; Gan, J.; Lu, F.; Li, X.; Ni, N.; Tian, Z.; Huang, X.; Gao, S.; et al. 2024. Multi-Sensor Object Anomaly Detection: Unifying Appearance, Geometry, and Internal Properties. *arXiv preprint arXiv:2412.14592*.
- Liang, H.; Xie, G.; Hou, C.; Wang, B.; Gao, C.; and Wang, J. 2025. Look Inside for More: Internal Spatial Modality Perception for 3D Anomaly Detection. *AAAI*.
- Lin, H.; Cai, D.; Xu, Z.; Wu, J.; Sun, L.; and Jia, H. 2024. Fabric4show: real-time vision system for fabric defect detection and post-processing. *Visual Intelligence*, 2(1): 13.
- Liu, J.; Xie, G.; Li, X.; Wang, J.; Liu, Y.; Wang, C.; Zheng, F.; et al. 2024. Real3D-AD: A Dataset of Point Cloud Anomaly Detection. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, volume 36.
- Pang, Y.; Tay, E. H. F.; Yuan, L.; and Chen, Z. 2023. Masked Autoencoders for 3D Point Cloud Self-supervised Learning. *World Scientific Annual Review of Artificial Intelligence*, 1: 2440001.
- Roth, K.; Pemula, L.; Zepeda, J.; Schölkopf, B.; Brox, T.; and Gehler, P. 2022. Towards total recall in industrial anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14318–14328. Doi: 10.1109/CVPR52688.2022.01392.
- Rusu, R. B.; Blodow, N.; and Beetz, M. 2009. Fast point feature histograms (FPFH) for 3D registration. In *2009 IEEE international conference on robotics and automation*, 3212–3217. IEEE.
- Wang, Y.; Peng, J.; Zhang, J.; Yi, R.; Wang, Y.; and Wang, C. 2023. Multimodal industrial anomaly detection via hybrid fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8032–8041.
- Ye, J.; Zhao, W.; Yang, X.; Cheng, G.; and Huang, K. 2025. Po3ad: Predicting point offsets toward better 3d point cloud anomaly detection. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 1353–1362.
- Yu, R.; Guo, B.; and Li, H. 2025. Anomaly Detection of Integrated Circuits Package Substrates Using the Large Vision Model SAIC: Dataset Construction, Methodology, and Application. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 22563–22574.
- Yu, X.; Tang, L.; Rao, Y.; Huang, T.; Zhou, J.; and Lu, J. 2022. Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 19313–19322.
- Zhao, H.; Jiang, L.; Jia, J.; Torr, P. H.; and Koltun, V. 2021. Point transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, 16259–16268.
- Zhou, Q.-Y.; Park, J.; and Koltun, V. 2018. Open3D: A modern library for 3D data processing. *arXiv preprint arXiv:1801.09847*.
- Zhou, Z.; Wang, L.; Fang, N.; Wang, Z.; Qiu, L.; and Zhang, S. 2024. R3D-AD: Reconstruction via Diffusion for 3D Anomaly Detection. In *European Conference on Computer Vision (ECCV)*.
- Zhu, H.; Xie, G.; Hou, C.; Dai, T.; Gao, C.; Wang, J.; and Shen, L. 2024. Towards High-resolution 3D Anomaly Detection via Group-Level Feature Contrastive Learning. *ACM Multimedia (ACM MM)*.