# BATCHED POINT CLOUD ANOMALY DETECTION VIA STRUCTURAL POSITION-AWARE COMPARISON

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Point Cloud Anomaly Detection (PCAD) is critical for industrial inspection. We propose Structural position-aware point cloud Anomaly Detection (SpotAD), the novel batched PCAD framework that performs anomaly detection directly on point cloud. Our SpotAD enables position-aware comparison through deterministic PCA-based alignment. To ensure both efficiency and structural consistency, we introduce object-aware prototype selection to restrict comparisons to structurally similar objects, position-aware prototype filtering for geometrically constrained comparison, and local prototype score smoothing to propagate prototype-level anomaly scores to local neighborhoods, enforcing spatial continuity. This strategy preserving geometric correspondence. Extensive evaluations on four benchmark datasets demonstrate state-of-the-art performance under the batched PCAD setting, achieving 87.0% object-level AUROC and 86.3% point-level AUROC, while outperforming existing zero-shot and full-shot PCAD methods.

## 1 INTRODUCTION

Anomaly Detection (AD) plays a critical role in industrial inspection by enabling early identification of defective products and process failures, which directly impacts production yield, safety, and maintenance cost. With the increasing adoption of 3D sensors in manufacturing, industrial Point Cloud Anomaly Detection (PCAD) has become particularly important, as point clouds provide rich geometric information and are less sensitive to illumination or texture variations compared to RGB images. However, in real-world industrial environments, collecting and annotating abnormal 3D samples from production lines is often impractical, making unsupervised PCAD a practical and widely adopted setting.

Traditional unsupervised PCAD (Liu et al., 2023; Zhao et al.; Li et al., 2024a; Zhou et al., 2024b; Zhu et al., 2024; Wang et al., 2025; Zheng et al., 2025; Liang et al., 2025a; Ye et al., 2025; Liang et al., 2025c; Yu et al., 2025) aims to learn with normal point clouds and identifies anomalies as deviations from the learned normality of geometric information during inference. Despite their effectiveness, unsupervised PCAD typically requires category-specific training, which increases computational cost and memory overhead as the number of categories increases. This requirement also limits applicability in cold-start deployments where normal point cloud data is unavailable. To overcome the need for category-specific training, recent studies have explored zero-shot PCAD framework (Cheng et al., 2025a; Zhou et al., 2024a) that aims to generalize anomaly detection to unseen objects. While zero-shot PCAD reduces the need for category-specific training, existing approaches still rely on additional supervision through auxiliary annotated training data.

Motivated by these limitations of unsupervised and zero-shot PCAD, a batched framework is needed that performs anomaly detection without relying on any training data. While batched anomaly detection has been actively studied in the image domain, directly transferring such 2D methods to PCAD is non-trivial, as point clouds are first rendered into 2D images to be compatible with image-based methods. This rendering process inevitably causes geometric information loss and creates blind spots, where occluded or self-hidden structures are not captured, resulting in unreliable anomaly localization for inherently 3D defects. To extend this paradigm to the 3D domain, we propose Structural position-aware point cloud Anomaly Detection (SpotAD), a novel batched PCAD framework that operates directly on point clouds without any training. Our key insight is that anomalies can be identified by comparing spatially corresponding regions across point cloud samples,



Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

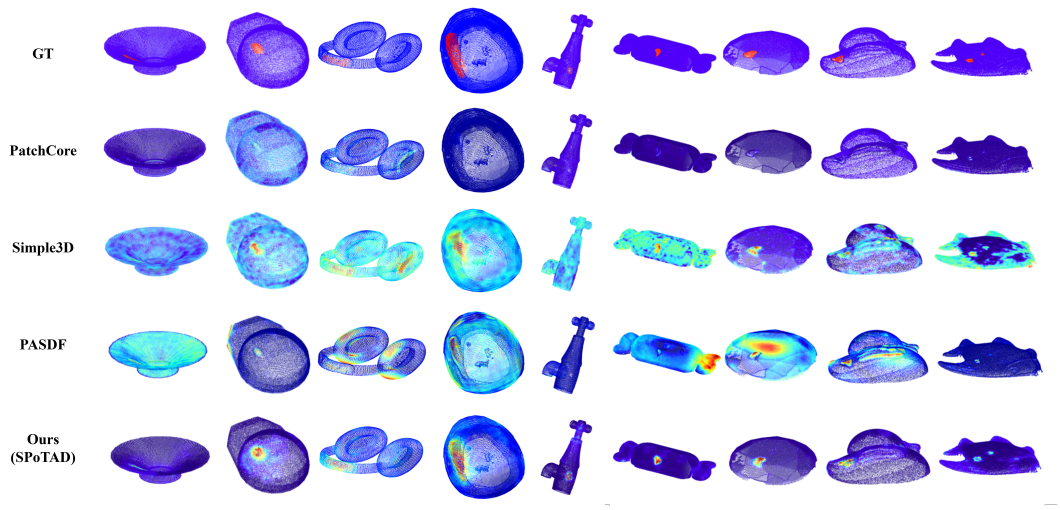


Figure 1: Qualitative results comparing prior methods with the proposed SpotAD, demonstrating its improved anomaly localization and overall performance across multiple benchmark datasets.

rather than performing exhaustive position-agnostic comparison. As shown in Fig. 1, SpotAD highlights localized defective regions while preserving the shared normal geometric structure across samples. This formulation is particularly well suited to industrial PCAD, where objects exhibit strong geometric consistency within the same category, while defects are sparse and confined to limited regions. This geometric regularity ensures that even batches containing anomalous samples largely preserve shared normality of geometric information, allowing meaningful reference structure to be obtained directly from the test set.

Exploiting this property, SpotAD identifies anomalies as localized deviations through direct test-time comparison. To enable position-aware comparison, SpotAD first aligns point clouds via a deterministic PCA-based alignment strategy, achieving rotational invariance without learnable parameters. Each point cloud is then represented by object-level prototypes obtained through farthest point sampling and neighborhood aggregation. Object-aware Prototype Selection (OPS) selects structurally similar objects for comparison, and Position-aware Prototype Filtering (PPF) restricts comparisons to spatially corresponding prototype locations, improving geometric correspondence. Local Prototype Score Smoothing (LPSS) further propagates prototype-level anomaly scores to local neighborhoods to enforce spatial continuity. Extensive evaluations on four benchmark datasets demonstrate state-of-the-art (SoTA) performance under the batched PCAD setting, achieving 87.0% object-level AUROC and 86.3% point-level AUROC, while outperforming existing zero-shot and full-shot PCAD methods. Our key contributions are summarized as follows:

- We propose the first novel batched framework for point cloud anomaly detection using only 3D data.
- We propose a position-aware comparison framework that models defects as region-level structural deviations and uses non-learnable PCA-based alignment, significantly reducing memory and training overhead.
- Our SpotAD is simple and efficient that requires no training or category-specific learning, and enables scalable deployment under cold-start and industrial environments.
- Extensive experiments across four benchmarks demonstrate SoTA performance across all settings.

# 2 Related Works

## 2.1 3D Anomaly Detection

Compared to 2D anomaly detection, 3D point clouds are robust to blind spot and illumination changes. With benchmarks such as Real3D-AD (Liu et al., 2023) and Anomaly-ShapeNet (Li et al., 2024a), prior work has largely followed two lines, namely memory bank-based and reconstruction-based methods. Memory bank-based methods store normal feature representations of geometry features and compute anomaly scores through feature-space distances at inference. Among them, M3DM (Wang et al., 2023), CPMF (Cao et al., 2024) fuse handcrafted 3D and 2D features, and Simple3D (Cheng et al., 2025b) utilizes handcrafted 3D features, while Reg3D-AD (Liu et al., 2023) builds memory bank of normal samples using pretrained 3D feature extractor. Among reconstruction-based methods, IMRNet (Li et al., 2024a), R3D-AD (Zhou et al., 2024b), and PASDF (Zheng et al., 2025) learn to reconstruct normal point cloud geometry, identifying anomalies as regions with high reconstruction error.

Despite these advances, due to large training overhead, zero-shot 3D anomaly detection has emerged using pretrained vision-language models (VLMs). To utilize 3D point cloud data, PointAD (Zhou et al., 2024a) and MVP-PCLIP (Cheng et al., 2025a) project point clouds into multiple 2D renderings to exploit pretrained VLMs. However, the projection process inevitably leads to the loss of intricate internal structures and fine-grained geometric details. Consequently, the development of zero-shot 3D anomaly detection capable of directly operating on raw point clouds, without relying on two-dimensional surrogates, remains a critical research objective. To address this limitation, we extend the batched zero-shot paradigm to directly operate on 3D point clouds, eliminating both the training overhead of unsupervised methods and the geometric information loss inherent in 2D projection-based approaches.

## 2.2 Batched Zero-shot Anomaly Detection

Batched zero-shot anomaly detection constitutes a distinct inference paradigm in which predictions are made not on a single sample but over an entire batch, typically the full test set, in a training-free manner. This setting exploits the statistical property that, in most real-world inspection scenarios, normal samples constitute the overwhelming majority of any given batch. ACR (Li et al., 2023) integrates batch normalization into anomaly scoring network such that the normal-majority statistics geometrically separate normal and anomalous embeddings without any test-time optimization. MuSc (Li et al., 2024b) exploits the same property by scoring each patch through mutual comparison across the batch. MuSc-V2 (Li et al., 2026) extends this framework to multimodal data with multi-scale patch aggregation and cross-modal score fusion. Unlike previous work, we extend this paradigm to 3D point clouds with the insight that normal geometry dominates not only at the object-level but also at the point-level within each sample. Furthermore, unlike MuSc-V2, our approach operates directly on raw point clouds without relying on any 2D images.

# 3 Method

## 3.1 Overview

We propose SpotAD, a batched framework for PCAD. Given a test set comprising both normal and abnormal samples, SpotAD directly computes object-level and point-level anomaly scores without requiring training data, auxiliary annotations, or any additional learning procedures. Our approach is inspired by the spot-the-difference analogy. In such tasks, humans first align two images to establish spatial correspondences, and subsequently detect both shared structural patterns and region-level deviations. As illustrated in Fig. 2, SpotAD explicitly integrates this comparison-based strategy into the PCAD. Unlike prior training-based methods, our framework leverages shared normality of geometric information across test samples and identifies anomalies as position-wise structural deviations.

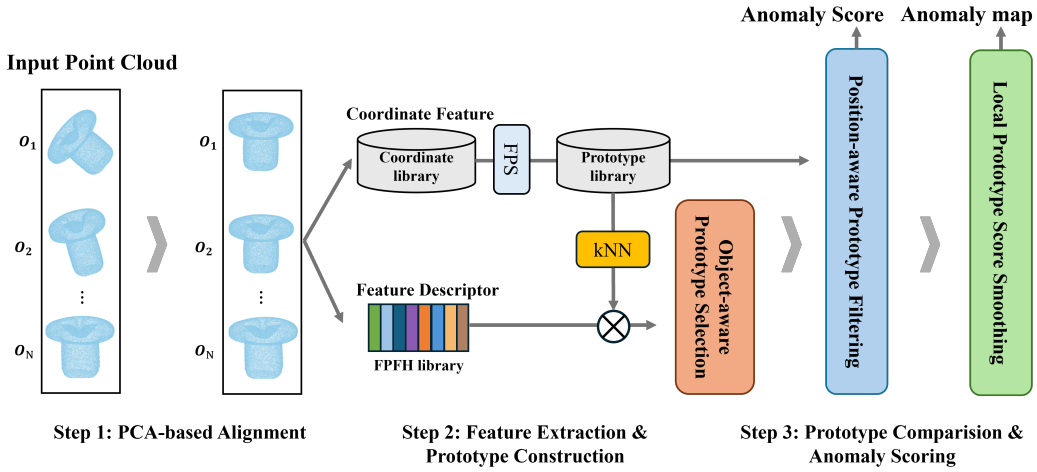


Figure 2: Framework of SpotAD. The framework consists of three main stages: (1) PCA-based alignment, (2) feature extraction and prototype construction, and (3) prototype comparison and anomaly scoring. Test point clouds are aligned and encoded into geometric and coordinate features, followed by position-aware prototype comparison and anomaly score computation.

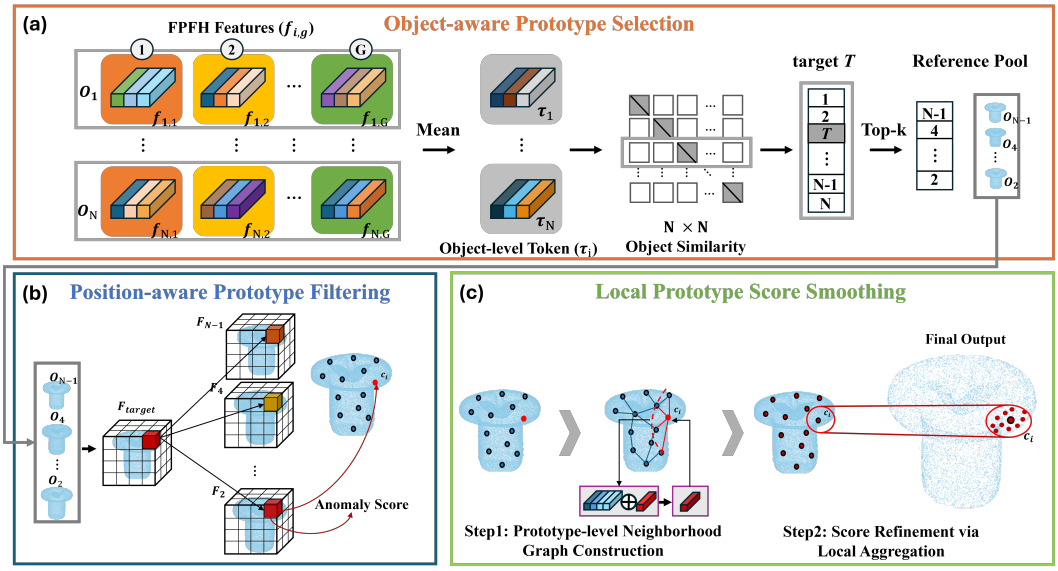


Figure 3: Detailed illustration of Object-aware Prototype Selection, Position-aware Prototype Filtering and Local Prototype Score Smoothing. (a) Object-aware Prototype Selection retrieves semantically relevant object prototypes. (b) Position-aware Prototype Filtering enable comparison between spatially corresponding regions. (c) Local Prototype Score Smoothing then propagates prototype-level scores to point-level predictions, producing the final anomaly map.

### 3.2 PCA Alignment-based Feature Extraction

#### 3.2.1 Canonical Pose Normalization

Let $P_i \in \mathbb{R}^{n\times 3}$ represents the input point cloud of the $i$-th sample, where each row corresponds to the coordinates of a point. Point clouds are typically captured under arbitrary orientations, which hinders reliable position-aware comparison across samples. To address this challenge, a deterministic

PCA-based pose normalization alignment strategy is adopted to establish a canonical coordinate frame. Given $P_i \in \mathbb{R}^{n\times 3}$, the principal axes are computed through eigen-decomposition of the covariance matrix and the point cloud is aligned as follows:

$$\tilde{P}_i = (P_i - \bar{p}_i)V, \tag{1}$$

where $\bar{p}_i$ represents the centroid and $V$ represents the principal directions. The principal axes are disambiguated to ensure consistent orientation across samples. This canonical alignment enables direct position-aware comparison across samples.

#### 3.2.2 PROTOTYPE CONSTRUCTION

Local geometric features are extracted using handcrafted descriptors and aggregated to form region-level prototypes. Specifically, $G$ representative centers are selected through farthest point sampling and aggregate neighboring point features by mean pooling. This process produces a prototype set $\mathcal{F}_i = \{(f_{i,g}, c_{i,g})\}_{g=1}^{G}$ for each object $\tilde{P}_i$. Here, $f_{i,g}$ and $c_{i,g}$ represents the prototype-level geometric descriptor and spatial center of the $g$-th prototype of object $\tilde{P}_i$, respectively. All prototypes obtained from the test samples are aggregated into a unified reference prototype library. For each object, an object-level token is computed by averaging the corresponding prototype features, and the resulting token is subsequently utilized for object-aware prototype selection. Mathematically, the reference prototype library is defined as follows:

$$\mathcal{L} = \{(f_{i,g}, c_{i,g}) \mid i = 1, \ldots, N,\ g = 1, \ldots, G\}. \tag{2}$$

where $N$ represents the number of test samples used to construct the reference prototype library.

### 3.3 OBJECT-AWARE PROTOTYPE SELECTION

In SpotAD, all test samples are included in the reference pool. Prototypes from geometrically dissimilar objects can be incorporated into the comparison, leading to unreliable anomaly scores. To alleviate this challenges, we propose Object-aware Prototype Selection (OPS), which constrains prototype comparison at the object-level by restricting the reference pool to similar objects. As shown in Fig. 3 (a), for each object $\tilde{P}_i$, an object-level token is obtained by averaging prototype features as follows:

$$\tau_i = \frac{1}{G}\sum_{g=1}^{G} f_{i,g}, \tag{3}$$

where $f_{i,g}$ represents the $g$-th prototype feature of object $i$. Given a query object $\tilde{P}_i$, cosine similarity is computed between the token $\tau_i$ and all other object-level tokens. The top-$K_{\text{obj}}$ most similar objects are retained, and the reference prototype library is restricted to prototypes derived from these selected objects.

In a batched setting, the reference pool cannot be explicitly partitioned into normal and abnormal samples. Consequently, directly comparing a normal query with a geometrically dissimilar abnormal reference may lead to unreliable anomaly scores, as global shape discrepancies can be misinterpreted as defects. To mitigate this challenges, OPS restricting comparison to structurally similar objects, which improves correspondence consistency and suppresses spurious mismatches while reducing the search space.

This constraint also preserves the validity of abnormal-abnormal comparisons, as defects typically occupy a limited surface area and vary in scale, morphology, and spatial location. Therefore, the comparisons are dominated by shared normality of geometric information rather than being biased by localized defect patterns. This strategy enables reliable anomaly estimation without requiring an explicit reference pool of normal samples. By filtering the reference pool at the object-level, OPS prevents comparisons with geometrically irrelevant objects and reduces the candidate prototype set before position-aware comparison.

### 3.4 POSITION-AWARE PROTOTYPE FILTERING

Normality of geometric information within an object category are inherently position-dependent. Comparing prototypes across unrelated spatial regions introduces geometric irrelevant and degrades

anomaly reliability. Therefore, we introduce Position-aware Prototype Filtering (PPF) to enforce spatial correspondence prior to feature distance computation.

For each object, let $\{c_{i,g}\}_{g=1}^{G}$ represents the prototype centers. As shown in Fig. 3 (b), the relative spatial configuration of prototypes is encoded by normalizing each center with respect to the object centroid $\bar{c}$:

$$\hat{c}_{i,g} = \frac{c_{i,g} - \bar{c}_i}{\|c_{i,g} - \bar{c}_i\|_2}. \tag{4}$$

This removes scale dependency while preserving directional positional information. Given a query prototype $g$, cosine similarity is computed between the normalized center $\hat{c}_{i,g}$ and normalized candidate centers from the filtered by OPS reference pool. Only the top-$K_{\text{pos}}$ positionally consistent candidates are retained for feature comparison. The anomaly score of prototype $g$ is then computed as the minimum feature distance within this restricted candidate set as follows:

$$s_{i,g} = \min_{(j,h)\in\mathcal{M}_{i,g}} \|f_{i,g} - f_{j,h}\|_2, \tag{5}$$

where $\mathcal{M}_{i,g}$ represents the selected position-consistent candidates. By constraining feature comparisons to spatially corresponding regions, PPF mitigates inter-part mismatches and enables geometrically consistent anomaly scoring.

### 3.5 Local Prototype Score Smoothing

While PPF yields discriminative prototype-level anomaly scores, these scores are computed independently for each prototype. In real-world environments, defects typically appear as region-level structural deviations rather than isolated point-wise irregularities. As a result, anomaly scores should exhibit local consistency instead of abrupt fluctuations between adjacent prototypes. To alleviate noisy outliers caused by descriptor variation, sparse sampling, or imperfect alignment, we introduce Local Prototype Score Smoothing (LPSS). As shown in Fig. 3 (c), LPSS enhances intra-object spatial consistency by refining each prototype score based on the local neighborhood, without any learnable parameters. Given a set of prototype centers $\{c_{i,g}\}_{g=1}^{G}$ for an object $\tilde{P}_i$, an intra-object proximity graph is constructed based on the Euclidean distance in 3D space. For each prototype $g$, we define $\mathcal{N}_{\text{local}}(g)$ as the set of its $K_{\text{local}}$ nearest neighboring prototype centers in 3D Euclidean space. This graph captures the local geometric adjacency of the object surface, ensuring that smoothing is restricted to spatially consistent regions.

Let $s_{i,g}$ represents the anomaly score of prototype $g$ obtained after PPF. LPSS improves the anomaly score by combining the original score with the mean anomaly score of neighboring prototypes:

$$\tilde{s}_{i,g} = \alpha \cdot s_{i,g} + (1-\alpha)\frac{1}{K_{\text{local}}} \sum_{g'\in\mathcal{N}_{\text{local}}(g)} s_{i,g'} \tag{6}$$

where $\alpha \in [0,1]$ controls the trade-off between the original prototype score and the neighborhood consensus score. The final object-level anomaly score is obtained by propagating the refined prototype scores $\{\tilde{s}_{i,g}\}_{g=1}^{G}$ to points via nearest-center assignment and averaging the top-$K$ resulting point-level scores. Since real defects tend to generate high anomaly scores across multiple spatially adjacent prototypes, LPSS improves region-level consistency without blurring structural boundaries.

## 4 Experiments

### 4.1 Experimental setups

We evaluate SpotAD on four PCAD benchmarks: Anomaly-ShapeNet (Li et al., 2024a), Real3D-AD (Liu et al., 2023), MulSen-AD (Li et al., 2025), and MiniShift (Cheng et al., 2025b). For evaluation, we adopt the Area Under the Receiver Operating Characteristic Curve (AUROC). Specifically, Point-level AUROC (P-AUROC) measures point-level anomaly localization performance, while Object-level AUROC (O-AUROC) evaluates object-level anomaly detection performance. We compare SpotAD, a batched method with baselines from existing paradigms: unsupervised PCAD (PatchCore-FPFH (Roth et al., 2022), M3DM (Wang et al., 2023), ISMP (Liang et al., 2025b), and Simple3D (Cheng et al., 2025b), PO3AD (Ye et al., 2025), PASDF (Zheng et al., 2025), R3D-AD (Zhou et al., 2024b)) , zero-shot PCAD (PointAD (Zhou et al., 2024a), MVP-PCLIP (Cheng et al.,

Table 1: **Comparison of Object-level and Point-level AUROC (%) across different datasets.** The best and second-best results are highlighted in **bold** and underlined, respectively.

**(a) Object-level AUROC (↑)**

| Dataset | Full-shot | | | | | | | Zero-shot | | Batched |
|---|---|---|---|---|---|---|---|---|---|---|
| | PatchCore -FPFH | M3DM | ISMP | PO3AD | PASDF | R3D -AD | Simple3D | PointAD | MVP -PCLIP | **Ours** |
| A-ShapeNet | 77.3 | 60.3 | 75.7 | 83.9 | 90.8 | 73.1 | 86.0 | 54.4 | 49.0 | **91.3** |
| Real3D-AD | 59.4 | 53.5 | 76.7 | 76.5 | **80.6** | 52.1 | 80.4 | 74.1 | 61.8 | 77.8 |
| Mulsen | 85.6 | 73.7 | 59.5 | 74.9 | 47.0 | 66.3 | 88.2 | 59.2 | 50.6 | **89.1** |
| MiniShift | 64.7 | 57.7 | 50.4 | 48.3 | 50.6 | 73.9 | 68.6 | 50.5 | 50.3 | **89.9** |
| **Average** | 71.8 | 61.3 | 65.6 | 70.9 | 67.3 | 66.4 | 80.8 | 59.6 | 52.9 | **87.0** |

**(b) Point-level AUROC (↑)**

| Dataset | Full-shot | | | | | | | Zero-shot | | Batched |
|---|---|---|---|---|---|---|---|---|---|---|
| | Patchcore -FPFH | M3DM -PMAE | ISMP | PO3AD | PASDF | R3D -AD | Simple3D | PointAD | MVP -PCLIP | **Ours** |
| A-ShapeNet | 82.4 | 66.2 | 69.1 | 89.8 | 89.7 | 59.2 | 92.9 | 73.9 | 39.8 | **97.5** |
| Real3D-AD | 58.2 | 62.8 | 83.6 | – | 75.6 | 63.4 | **92.3** | 61.6 | 86.5 | 86.8 |
| Mulsen | 75.1 | 64.6 | 50.3 | 64.1 | 53.8 | 53.8 | 80.3 | 61.6 | 38.5 | **87.0** |
| MiniShift | 64.3 | 58.2 | 48.3 | 52.0 | 49.1 | 51.8 | 66.2 | 60.2 | 27.2 | **73.7** |
| **Average** | 70.0 | 63.0 | 62.8 | 68.6 | 67.1 | 57.1 | 82.9 | 64.3 | 48.0 | **86.3** |

2025a)). All experiments are conducted on a single NVIDIA RTX 4090 GPU. To extract geometric features, we employ Fast Point Feature Histograms (FPFH) (Rusu et al., 2009). For computational efficiency, voxel grid downsampling was applied to dense real-world datasets, with voxel sizes set to 0.09 for Real3D-AD.

### 4.2 Main Results

Table 1 presents quantitative comparisons with state-of-the-art PCAD methods across four benchmarks. SpotAD achieves the best overall performance, 87.0% O-AUROC and 86.3% P-AUROC, consistently outperforming all competing approaches. Compared with the training-based baseline Simple3D, SpotAD improves O-AUROC by +6.2% (80.8 → 87.0) and P-AUROC by +3.4% (82.9 → 86.3). Notably, these gains are achieved in a batched setting. Compared with zero-shot PCAD methods such as PointAD and MVP-PCLIP, SpotAD demonstrates substantial improvements. In particular, SpotAD surpasses MVP-PCLIP by +34.1% in O-AUROC and +38.3% in P-AUROC. These results highlight the advantage of directly modeling geometric features rather than relying on 2D projections or cross-modal alignment. Overall, the results validate that spatially corresponding, position-aware comparison is more effective and scalable than full-shot and zero-shot paradigms. Fig. 1 illustrates qualitative comparisons on Anomaly-ShapeNet, and Real3D-AD. Compared with competing methods, SpotAD produces sharper and more spatially coherent anomaly maps. While memory bank-based and reconstruction-based methods tend to generate scattered false positives and over-smooth abnormal regions, SpotAD consistently focuses on localized geometric inconsistencies at spatially corresponding positions.

### 4.3 Effect of Object-aware Prototype Selection and Position-aware Prototype Filtering

Fig. 4 qualitatively illustrates the effect of the two-stage prototype selection process consisting of Object-aware Prototype Selection (OPS) and Position-aware Prototype Filtering (PPF). As shown in Fig. 4 (a), OPS first retrieves structurally similar reference objects from the reference pool based on cosine similarity between object-level tokens. Despite the absence of supervision, normal targets consistently retrieve normal references, while abnormal targets retrieve structurally similar abnormal references. This indicates that localized defects exert limited influence on global object similarity, allowing OPS to suppress mismatches caused by large structural discrepancies and reduce the

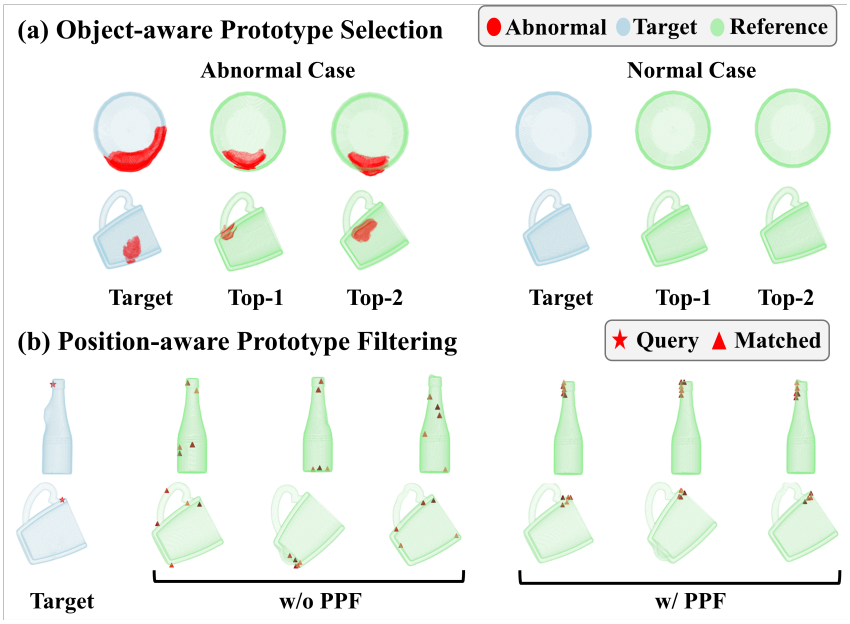


Figure 4: Qualitative visualization of OPS and PPF. (a) OPS retrieves structurally similar reference objects based on object-level similarity for both abnormal and normal targets. (b) PPF restricts prototype matching to positionally corresponding regions. Without PPF, query prototypes are matched to positionally inconsistent library prototypes, whereas PPF suppresses such cross-location mismatches and enables structurally consistent prototype comparison.

Table 2: Robustness analysis of the batched evaluation protocol. (a) Comparison between transductive and inductive evaluation protocols. (b) Sensitivity to the normal-to-abnormal composition of the transductive batch, averaged over all datasets.

**(a) Transductive vs. Inductive Evaluation**

| | Transductive | | Inductive | |
|---|---|---|---|---|
| **Dataset** | **O-AUROC** | **P-AUROC** | **O-AUROC** | **P-AUROC** |
| Anomaly-ShapeNet | 91.3 | 97.5 | 89.7 | 95.0 |
| Real3D-AD | 77.8 | 86.8 | 69.3 | 90.4 |
| MulSen-AD | 89.1 | 87.0 | 91.0 | 86.0 |
| MiniShift | 89.9 | 73.7 | 86.4 | 73.6 |
| **Average** | **87.0** | **86.3** | **84.1** | **86.3** |

**(b) Batch Composition (Avg.)**

| Normal:Abnormal | O-AUROC | P-AUROC |
|---|---|---|
| 9:1 | 86.0 | 87.1 |
| 8:2 | **88.9** | **89.7** |
| 7:3 | 86.9 | 88.2 |
| 6:4 | 86.4 | 87.2 |
| 4:6 | 84.9 | 86.1 |
| 3:7 | 82.8 | 85.4 |
| 2:8 | 81.7 | 85.1 |
| 1:9 | 77.6 | 85.3 |

candidate set before prototype-level comparison. Further analyses on robustness and reproducibility, including multi-class evaluation and robustness to increasing deformation severity, are provided in Appendix **??**.

Fig. 4 (b) further shows the effect of PPF. Without PPF, query prototypes are matched globally in feature space and can therefore correspond to geometrically unrelated regions, such as the body or bottom of a bottle for a query located near the neck. PPF restricts candidate prototypes according to positional consistency in the PCA-normalized coordinate system, thereby suppressing such cross-location mismatches. The cup examples further show that corresponding regions are retained despite moderate rotational variations and imperfect alignment. Overall, OPS and PPF progressively reduce object-level and position-level mismatches and enable structurally consistent prototype comparison for reliable anomaly scoring.

Table 3: Ablation study of different modules on Object-level / Point-level AUROC (↑).

| PPF | LPSS | PCA | OPS | Anomaly ShapeNet | Real3D-AD | MulSen-AD | MiniShift | **Mean** |
|---|---|---|---|---|---|---|---|---|
| ✓ | ✓ | | ✓ | 68.6 / 95.7 | 71.6 / 82.7 | 63.6 / 82.6 | 33.1 / 67.8 | 59.2 / 82.2 |
| ✓ | | ✓ | ✓ | 90.3 / 96.0 | 70.0 / 78.8 | 89.5 / 86.4 | 82.4 / 65.4 | 83.1 / 81.7 |
| | ✓ | ✓ | ✓ | 90.3 / 96.0 | 70.0 / 78.8 | 85.8 / 86.1 | 65.4 / 68.5 | 77.9 / 82.4 |
| ✓ | ✓ | ✓ | | 90.8 / 97.5 | 77.8 / 86.0 | 88.7 / 87.0 | 86.8 / 74.2 | 86.0 / 86.2 |
| ✓ | ✓ | ✓ | ✓ | 91.3 / 97.5 | 77.8 / 86.8 | 89.1 / 87.0 | 89.9 / 73.7 | 87.0 / 86.3 |

### 4.4 Robustness to Evaluation Protocol and Batch Composition

Since SpotAD follows a batched framework, reference library is constructed from samples available at inference time. This raises two practical questions regarding whether the reported performance depends on the transductive use of the test set and how sensitive the method is to the normal-to-anomaly composition of the inference batch. To examine the first question, we construct an inductive variant in which the reference pool is formed from the training split rather than the test set, while test samples are used only as queries. As shown in Table 2 (a), the inductive setting achieves a mean O-AUROC/P-AUROC of 84.1/86.3, compared with 87.0/86.3 under the transductive protocol. In particular, the mean point-level performance remains unchanged, indicating that the localization capability is not solely attributable to transductive access to the test-set structure.

We further vary the normal-to-anomaly ratio within the transductive batch from 9:1 to 1:9. As shown in Table 2 (b), object-level performance gradually decreases as anomalous samples become dominant, whereas point-level performance remains relatively stable. Even under the extreme 1:9 setting, SpotAD retains a P-AUROC of 85.3. These results indicate that batch composition mainly affects global object-level discrimination, while position-aware local comparison remains robust to substantial changes in the normal-to-anomaly ratio. Overall, these results demonstrate that the effectiveness of SpotAD is not solely explained by the transductive evaluation protocol or a strongly normal-dominated test batch.

### 4.5 Ablation Study

As shown in Table 3, we evaluate the contribution of each module. The full configuration, which progressively restricts the candidate set, achieves the best overall performance and validates the synergy between all components. Specifically, OPS effectively filters irrelevant category-level noise. PPF serves as a core mechanism for structural consistency. LPSS ensures precise anomaly localization by leveraging spatial neighborhood consensus, and PCA provides a necessary foundation for PPF.

## 5 Conclusion

This paper introduced SpotAD, a batched framework for PCAD that identifies defects via structural, position-aware comparisons. By leveraging PCA-based alignment, our method establishes reliable correspondences between point clouds without learnable parameters, enabling position-aware comparison under geometric consistency constraints. While the approach is validated primarily on rigid industrial objects, extending its applicability to highly deformable geometries and non-industrial domains such as medical 3D imaging remains an important direction for future work.

**Limitation** SpotAD extends naturally to the multi-class setting, but its behavior under severe class imbalance remains unexplored. Since each sample is scored by comparison with others, classes represented by only one or two samples may yield unstable comparisons and anomaly scores that are poorly calibrated relative to well-represented classes. We leave a systematic study of this regime to future work.

### AI use statement

In this work, we used generative AI tools to assist with language editing, sentence-level rewriting, improving the clarity of the manuscript, and identifying relevant literature. All technical content,

including the proposed method, mathematical formulations, experimental design, results, and AI-assisted code, was carefully reviewed and verified. We take full responsibility for the final content of this work.

## REFERENCES


Yunkang Cao, Xiaohao Xu, and Weiming Shen. Complementary pseudo multimodal feature for point cloud anomaly detection. *Pattern Recognition*, 156:110761, 2024.

Yuqi Cheng, Yunkang Cao, Guoyang Xie, Zhichao Lu, and Weiming Shen. Toward zero-shot point cloud anomaly detection: A multiview projection framework. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2025a.

Yuqi Cheng, Yihan Sun, Hui Zhang, Weiming Shen, and Yunkang Cao. Towards high-resolution 3d anomaly detection: A scalable dataset and real-time framework for subtle industrial defects. *arXiv preprint arXiv:2507.07435*, 2025b.

Aodong Li, Chen Qiu, Marius Kloft, Padhraic Smyth, Maja Rudolph, and Stephan Mandt. Zero-shot anomaly detection via batch normalization. *Advances in Neural Information Processing Systems*, 36:40963–40993, 2023.

Wenqiao Li, Xiaohao Xu, Yao Gu, Bozhong Zheng, Shenghua Gao, and Yingna Wu. Towards scalable 3d anomaly detection and localization: A benchmark via 3d anomaly synthesis and a self-supervised learning network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 22207–22216, 2024a.

Wenqiao Li, Bozhong Zheng, Xiaohao Xu, Jinye Gan, Fading Lu, Xiang Li, Na Ni, Zheng Tian, Xiaonan Huang, Shenghua Gao, et al. Multi-sensor object anomaly detection: Unifying appearance, geometry, and internal properties. In *Proceedings of the computer vision and pattern recognition conference*, pp. 9984–9993, 2025.

Xurui Li, Ziming Huang, Feng Xue, and Yu Zhou. Musc: Zero-shot industrial anomaly classification and segmentation with mutual scoring of the unlabeled images. In *The Twelfth International Conference on Learning Representations*, 2024b.

Xurui Li, Feng Xue, and Yu Zhou. Musc-v2: Zero-shot multimodal industrial anomaly classification and segmentation with mutual scoring of unlabeled samples. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2026.

Hanzhe Liang, Aoran Wang, Jie Zhou, Xin Jin, Can Gao, and Jinbao Wang. Examining the source of defects from a mechanical perspective for 3d anomaly detection. *arXiv preprint arXiv:2505.05901*, 2025a.

Hanzhe Liang, Guoyang Xie, Chengbin Hou, Bingshu Wang, Can Gao, and Jinbao Wang. Look inside for more: Internal spatial modality perception for 3d anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 5146–5154, 2025b.

Hanzhe Liang, Jie Zhou, Xuanxin Chen, Tao Dai, Jinbao Wang, and Can Gao. Fence theorem: Towards dual-objective semantic-structure isolation in preprocessing phase for 3d anomaly detection. *arXiv preprint arXiv:2503.01100*, 2025c.

Jiaqi Liu, Guoyang Xie, Ruitao Chen, Xinpeng Li, Jinbao Wang, Yong Liu, Chengjie Wang, and Feng Zheng. Real3d-ad: A dataset of point cloud anomaly detection. *Advances in Neural Information Processing Systems*, 36:30402–30415, 2023.

Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Schölkopf, Thomas Brox, and Peter Gehler. Towards total recall in industrial anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 14318–14328, 2022.

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*, pp. 3212–3217. IEEE, 2009.

Jiaxiang Wang, Haote Xu, Xiaolu Chen, Haodi Xu, Yue Huang, Xinghao Ding, and Xiaotong Tu. Exploiting point-language models with dual-prompts for 3d anomaly detection. *Expert Systems with Applications*, pp. 129758, 2025.

Yue Wang, Jinlong Peng, Jiangning Zhang, Ran Yi, Yabiao Wang, and Chengjie Wang. Multimodal industrial anomaly detection via hybrid fusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8032–8041, 2023.

Jianan Ye, Weiguang Zhao, Xi Yang, Guangliang Cheng, and Kaizhu Huang. Po3ad: Predicting point offsets toward better 3d point cloud anomaly detection. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 1353–1362, 2025.

Yuyang Yu, Zhengwei Chen, Xuemiao Xu, Lei Zhang, Haoxin Yang, Yongwei Nie, and Shengfeng He. Registration is a powerful rotation-invariance learner for 3d anomaly detection. *arXiv preprint arXiv:2510.16865*, 2025.

B Zhao, Q Xiong, X Zhang, J Guo, Q Liu, X Xing, and X PointCore Xu. Efficient unsupervised point cloud anomaly detector using local-global features. arxiv 2024. *arXiv preprint arXiv:2403.01804*.

Bozhong Zheng, Jinye Gan, Xiaohao Xu, Xintao Chen, Wenqiao Li, Xiaonan Huang, Na Ni, and Yingna Wu. Bridging 3d anomaly localization and repair via high-quality continuous geometric representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 27063–27072, 2025.

Qihang Zhou, Jiangtao Yan, Shibo He, Wenchao Meng, and Jiming Chen. Pointad: Comprehending 3d anomalies from points and pixels for zero-shot 3d anomaly detection. *Advances in Neural Information Processing Systems*, 37:84866–84896, 2024a.

Zheyuan Zhou, Le Wang, Naiyu Fang, Zili Wang, Lemiao Qiu, and Shuyou Zhang. R3d-ad: Reconstruction via diffusion for 3d anomaly detection. In *European conference on computer vision*, pp. 91–107. Springer, 2024b.

Hongze Zhu, Guoyang Xie, Chengbin Hou, Tao Dai, Can Gao, Jinbao Wang, and Linlin Shen. Towards high-resolution 3d anomaly detection via group-level feature contrastive learning. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pp. 4680–4689, 2024.