

Real-IAD Variety: Pushing Industrial Anomaly Detection Dataset to a Modern Era

Wenbing Zhu^{a,d,1}, Chengjie Wang^{b,c,1}, Bin-Bin Gao^{b,1}, Jiangning Zhang^b,
Guannan Jiang^e, Jie Hu^f, Zhenye Gan^b, Lidong Wang^a, Ziqing Zhou^a, Jianghui
Zhang^g, Linjie Cheng^a, Yurui Pan^a, Bo Peng^g, Mingmin Chi^{a,*}, Lizhuang Ma^c

^a*College of Computer Science and Artificial Intelligence, Fudan University*

^b*Youtu Lab, Tencent*

^c*School of Computer Science, Shanghai Jiao Tong University*

^d*Rongcheer Co., Ltd*

^e*Department of Computer Science, City University of Hong Kong*

^f*Electrical and Computer Engineering, National University of Singapore*

^g*School of Information Science, Shanghai Ocean University*

Abstract

Industrial Anomaly Detection (IAD) is a cornerstone for ensuring operational safety, maintaining product quality, and optimizing manufacturing efficiency. However, the advancement of IAD algorithms is severely hindered by the limitations of existing public benchmarks. Current datasets often suffer from restricted category diversity and insufficient scale, leading to performance saturation and poor model transferability in complex, real-world scenarios. To bridge this gap, we introduce Real-IAD Variety, the largest and most diverse IAD benchmark. It comprises 198,950 high-resolution images across 160 distinct object categories. The dataset ensures unprecedented diversity by covering 28 industries, 24 material types, 22 color variations, and 27 defect types. Our extensive experimental analysis highlights the substantial challenges posed by this benchmark: state-of-the-art multi-class unsupervised anomaly detection methods suffer significant performance degradation (ranging from 10% to 20%) when scaled from 30 to 160 categories. Conversely, we demonstrate that zero-shot and few-shot IAD models exhibit remarkable robustness to category scale-up,

*Corresponding author

Email address: mmchi@fudan.edu.cn (Mingmin Chi)

¹Equal contribution.

maintaining consistent performance and significantly enhancing generalization across diverse industrial contexts. This unprecedented scale positions Real-IAD Variety as an essential resource for training and evaluating next-generation foundation IAD models.

Keywords: Multi-Class Unsupervised Learning, Zero-Shot and Few-Shot Anomaly Detection, Foundation Models

1. Introduction

The field of anomaly detection has undergone a paradigm shift towards fully unsupervised [1, 2, 3], multi-class unsupervised (unified) [4, 5, 6], and zero-/few-shot [7, 8, 9] learning frameworks, **ushering in a transformative era for model architecture**. Within the specialized domain of Industrial Anomaly Detection (IAD), substantial progress has been achieved in addressing the inherent complexity and variability of industrial processes.




Figure 1: **Visual illustration of diverse anomaly patterns in the Real-IAD Variety dataset.** We present representative samples across 27 defect types, where anomaly regions are highlighted with red boundaries. Each sample is captured from five distinct viewpoints, demonstrating that certain defects are only detectable from specific angles while remaining invisible in others. Text annotations follow the format of “object (defect type)”, covering both structural damages (*e.g.*, scratch) and logical inconsistencies (*e.g.*, missing part).




Figure 2: **Industry distribution of the proposed Real-IAD Variety dataset.** The dataset encompasses 8 major industrial groups (denoted as *c*): electrical, transport, cultural products, metal, general, electronics, rubber plastic, and other manufacturing sectors. These major categories are further subdivided into 28 industrial subcategories (denoted as *s*), and complete details are provided in Appendix.

Recent IAD research has predominantly concentrated on developing unified models capable of detecting anomalies across multiple industrial objects using a single model [4, 10, 11, 12, 5]. Concurrently, zero-shot and few-shot anomaly detection approaches [13, 7, 8, 9] have emerged to identify anomalies in previously unseen domains without requiring **task-specific fine-tuning**. These advances are fundamentally reshaping IAD enabling more robust and adaptable frameworks. **However, the progress is increasingly bottlenecked by the limitations of existing benchmarks.**

Despite the contributions of **established** benchmarks such as MVTec [14],

VisA [15], PAD [16], and Real-IAD [17], these benchmarks exhibit critical limitations in category diversity and scenario coverage. **For instance**, Real-IAD [17], offers the largest category count, **is limited to 30 categories**, while the widely adopted MVTec and VisA datasets comprise only 15 and 12 categories, respectively. Therefore, it constrains the comprehensive evaluation of contemporary unified and zero-/few-shot IAD models. Furthermore, while synthetic data generation via pre-trained diffusion models has been explored [18, 19, 20], these approaches still face challenges in generating authentic and diverse anomalous patterns.

To bridge this gap and propel IAD research toward unified and zero-/few-shot paradigms, **there is an urgent requirement for benchmarks** that encompass substantially broader industrial categories and defect typologies. In this paper, we introduce Real-IAD Variety, a large-scale benchmark that represents a significant advancement in the field. As illustrated in Figs. 1, 2 and 4, this dataset substantially expands the scope of IAD data collection, featuring 160 categories spanning 28 industries, 24 material types, 22 color variations and **27 defect types, culminating in** a total of **198,950** images with high-fidelity pixel-level annotations. The category count of Real-IAD Variety is approximately 5.3 times that of **its predecessor**, Real-IAD [17], making it the first truly multi-industry benchmark that overcomes the constraints of prior datasets for unified and zero-/few-shot evaluation.

Beyond scale, Real-IAD Variety offers unprecedented diversity in materials (*e.g.*, composites, metals, ceramics) and industries (*e.g.*, electrical manufacturing, transport, electronics), addressing the practical needs of real-world deployment. We leverage this benchmark to establish rigorous evaluations across three critical settings: multi-class unsupervised anomaly detection (MUAD), multi-view anomaly detection, and zero-/few-shot anomaly detection.

Our extensive experiments yield two pivotal insights, **scalability bottlenecks** and **paradigm robustness**. We observe that an increase in category diversity leads to a significant performance degradation (10%-20%) in state-of-the-art MUAD methods, revealing their limited scalability when tasked with han-

dling multiple category distributions. In contrast, zero-/few-shot approaches, particularly those leveraging large-scale vision-language models, demonstrate remarkable robustness. Their performance remains consistent regardless of category count, eventually surpassing MUAD methods once a certain complexity threshold is reached.

In summary, our contributions are threefold:

- We present Real-IAD Variety, a large-scale IAD dataset featuring unprecedented category diversity, extensive industry coverage, and high-quality pixel-level annotations. It expands the category count to 160 across 28 industries, 24 material types, 22 color variations and **27 defect types** with **198,950** images, representing a substantial advancement over existing IAD benchmarks.
- We establish comprehensive benchmarks on Real-IAD Variety across three critical IAD settings: multi-class unsupervised anomaly detection, multi-view anomaly detection, and zero-/few-shot anomaly detection, facilitating systematic evaluation of current methodologies.
- We reveal that increasing category count significantly impairs the performance of multi-class and multi-view unsupervised anomaly detection methods, while zero-/few-shot approaches exhibit minimal sensitivity to category variations, suggesting a fundamental shift in scalability characteristics across different learning paradigms.

2. Related Work

2.1. IAD Datasets

In Industrial Anomaly Detection (IAD), dataset diversity and scale are critical determinants of model effectiveness. Current 2D IAD benchmarks, including MVTec AD [14], VisA [15] and Real-IAD [17], have demonstrated considerable success in establishing foundational evaluation frameworks. However, their categorical scope remains limited. For example, Real-IAD [17], the largest existing

dataset, encompasses 30 categories, while MVTec AD [14] and VisA [15] contain 15 and 12 categories, respectively. Additional datasets such as MTD [21], MPDD [22], BTAD [23], and KolektorSDD [24] contribute valuable resources but suffer from restricted scale and low categorical diversity.

Recent efforts have extended these benchmarks toward multi-modal or large-scale settings. For instance, MANTA [25] introduce descriptive textual annotations for multi-modal IAD focusing on tiny objects across 38 categories. In the logistics domain, Kaputt [26] presents a large-scale collection of 48,000 objects; however, its utility is constrained by the lack of pixel-level mask annotations, which are essential for precise defect localization. Furthermore, some recent benchmarks like COCO-AD [27] and ADNet [28] are constructed through the re-purposing or integration of existing datasets. COCO-AD [27] manually re-defines categories from the COCO dataset as anomalies, which fundamentally differs from the authentic "defects" encountered in real-world industrial production. ADNet [28] primarily focuses on the consolidation of existing open-source benchmarks. Despite these contributions, existing datasets often exhibit limited category variability, insufficient scale, or a gap between synthetic anomalies and real-world industrial scenarios. In contrast, our proposed Real-IAD Variety addresses these limitations by substantially enhancing category diversity, industrial scope, and defect typology, thereby advancing the field toward more comprehensive and realistic evaluation scenarios.

2.2. Unsupervised Anomaly Detection

Traditional unsupervised anomaly detection (UAD) algorithms typically focuses on single-class, single-view 2D images, where models are trained independently for each object category using only normal samples. These methods are generally categorized into three paradigms: (1) Discriminative methods (*e.g.*, CutPaste [29], DRAEM [30], and SimpleNet [11]), which learn decision boundaries by distinguishing normal data from synthetically generated anomalies; (2) Embedding-based methods (*e.g.*, PatchCore [31], CFA [10], and CSFlow [32]), which model the distribution of normal feature representations in a frozen or

learned embedding space. While these approaches have achieved notable success in single-class scenarios, their scalability to multi-class and cross-domain settings remains limited, motivating the development of unified detection frameworks; (3) Reconstruction-based methods (*e.g.*, DAE [33], **OCR-GAN** [34], and RD [35]), which operate under the assumption that anomalous regions exhibit higher reconstruction errors compared to normal patterns.

2.3. Multi-Class Unsupervised AD

The **Multi-class Unsupervised Anomaly Detection (MUAD)** setting, pioneered by UniAD [4], requires a single unified model to handle training and inference across multiple categories simultaneously. MUAD benchmarks encompass both traditional UAD methods, such as DRAEM [30], SimpleNet [11], CFA [10], CFLOW-AD [36], and RD [35], explicitly designed for unified training, including UniAD [4], OneNIP [5], MambaAD [12], DesTSeg [37], LGC [38] and Dinomaly [6]. However, existing MUAD evaluations have been conducted on datasets with limited category counts (typically ≤ 30 classes), leaving the scalability of these methods to large-scale, highly diverse industrial scenarios largely unexplored. Our Real-IAD Variety dataset, with its 160 categories, **provides the first rigorous testing ground for MUAD scalability.**

2.4. Multi-View Anomaly Detection

Beyond single-image analysis, Multi-View Anomaly Detection (MVAD) leverages complementary information from multiple viewpoints to enhance inspection accuracy and robustness. This approach originates from the stringent precision requirements of industrial quality control, where anomalies may be occluded or imperceptible from a single perspective. Representative methods such as MVAD [39] exploit multi-view consistency constraints and cross-view feature fusion to improve detection performance, which is critical for minimizing false negatives in practical deployment scenarios. Nevertheless, existing MVAD datasets predominantly focus on limited object types and controlled imaging conditions,

constraining the evaluation of cross-view generalization in diverse industrial contexts. Real-IAD Variety addresses this gap by providing multi-view annotations across 160 categories, facilitating rigorous MVAD benchmarking.

2.5. Zero-/Few-Shot AD Models

Vision-Language Models (VLMs) have emerged as a promising direction for anomaly detection, enabling semantic reasoning and cross-domain generalization. Recent advances have explored zero-shot (ZSAD) and few-shot (FSAD) anomaly detection using large-scale pre-trained models such as CLIP [40]. WinCLIP [41] introduces dual-class textual prompts and patch-wise window operation for ZSAD and FSAD. AnomalyCLIP [7] learns object-agnostic prompt embeddings to capture generalizable normal and abnormal representations through an auxiliary anomaly detection dataset. AdaCLIP [42] introduces learnable hybrid prompts and regional anomaly feature extraction to improve detection precision, while VCP-CLIP [43] designs Pre-VCP and Post-VCP modules to optimize textual embeddings with visual contextual prompts, promoting cross-modal information interaction. Recently, AdaptCLIP [8] treats CLIP as a foundational service, proposing alternating and comparative learning strategies based on three lightweight adapters to support zero- and few-shot generalization across domains.

Unlike VLM-based ZSAD and FSAD, MetaUAS [9] is the first to reframe anomaly detection from a change detection perspective, achieving a pure-vision one-shot anomaly detection paradigm that eliminates the reliance on specific anomaly detection datasets for pre-training. Despite these methodological advances, existing ZSAD and FSAD approaches have been validated primarily on small-scale datasets with limited category diversity (typically ≤ 30 classes), leaving their scalability and robustness under large-scale, multi-industry conditions unverified. Our Real-IAD Variety dataset enables the first large-scale evaluation of ZSAD and FSAD methods across 160 categories, thereby facilitating systematic investigation of their generalization capabilities in realistic industrial settings.




Figure 3: **Data collection and annotation pipeline for the proposed Real-IAD Variety.** The pipeline comprises a four-stage sequential process: (a) **Material Preparation:** This initial phase encompasses the assembly of a diverse array of materials, spanning 160 distinct categories sourced from 28 industrial domains and encompassing 24 material compositions. (b) **Acquisition Equipment Design:** The second phase involves the design of data capture apparatus, comprising one top-down camera for overhead views and four peripheral cameras to capture lateral perspectives. (c) **Data Collection and Annotation:** The third phase pertains to the data collection process, which includes meticulous pixel-level manual annotation, rigorous algorithmic cross-validation, and iterative refinement. **The process iterates until the variation in the model’s predicted Average Precision (AP) scores falls below a predetermined threshold**, following the methodology established in Real-IAD [17]. (d) **Defect Taxonomy:** The lower section illustrates distinct defect types alongside their characteristic visual representations. Zoom in for enhanced visibility of defect regions delineated in red.

3. Real-IAD Variety

3.1. Data Collection Pipeline

Inspired by the methodology established in Real-IAD [17], we construct a rigorous three-stage data collection pipeline to ensure dataset comprehensiveness and annotation quality, as illustrated in Fig. 3. The pipeline encompasses material preparation, acquisition equipment design, and data collection with iterative annotation refinement.

Stage 1: Diverse Material Preparation. To construct the Real-IAD Variety dataset, a team of 12 members dedicated 11,000 working hours to material selection and procurement, simulating a wide spectrum of real-world defects




Figure 4: **Statistical characteristics of Real-IAD Variety across multiple dimensions.** (a) **Anomalous region proportion:** Real-IAD Variety exhibits a broader and more balanced distribution of anomalous region proportions relative to total image area compared to Real-IAD [17], substantially increasing dataset complexity. (b) **Defect aspect ratio:** Real-IAD Variety provides diverse aspect ratios for minimum bounding rectangles of defects, comparable to Real-IAD, introducing additional diversity and detection challenges. Representative samples are shown for intuitive visualization. (c) **Material distribution:** Real-IAD Variety encompasses 24 material types for practical applications, imposing higher requirements on method robustness. (d) **Color distribution:** Real-IAD Variety captures a wide color spectrum, which is essential for color-based anomaly detection research.

Table 1: Comparison with popular 2D IAD datasets on different attributes. ✓: Satisfied. ✗: Unsatisfied.

Datasets	Classes	Number of Images			Image Resolution	Anomaly Masks	Multiple Views	Defect Types
		Normal	Anomaly	All				
MVTec AD [14]	15	4,096	1,258	5,354	700~1,024	✓	✗	20
VisA [15]	12	9,621	1,200	10,821	960~1,562	✓	✗	13
BTAD [23]	3	2,250	580	2,830	600~1,600	✓	✗	3
MPDD [22]	6	1,064	282	1,346	1,024~1,024	✓	✗	8
MAD-Real [16]	10	540	221	761	3,472~3,472	✓	✓	2
MAD-Sim [16]	20	4,838	4,951	9,789	800~800	✓	✓	3
Real-IAD [17]	30	99,721	51,329	151,050	2,000~5,000	✓	✓	8
MANTA [25]	38	652,455	34,235	686,690	162~3,365	✓	✓	30
Kaputt [26]	48,376	70,951	29,316	100,267	2,048~2,048	✗	✗	7
Real-IAD Variety	160	39,905	159,045	198,950	260~5,328	✓	✓	27

based on material characteristics. To enhance dataset representativeness, we assembled 160 object categories spanning 24 material types across 28 industrial domains. Representative samples are depicted in Fig. 3a. Given that industrial

Table 2: Detailed Specifications of the Multi-view Acquisition System.

Specification	Top-down Camera	Peripheral Cameras
Sensor Resolution	$5,328 \times 3,040$	$4,096 \times 3,000$
Quantity	1	4
Mounting Angle	90° (Vertical)	$40^\circ\text{--}45^\circ$ (Oblique)
Lateral Accuracy	0.01 mm/pixel	0.028 mm/pixel
Working Distance	185 mm	310 mm
Field of View (FoV)	53×30.4 mm	114×84 mm
Z-Repeatability	$<0.3 \mu\text{m}$	N/A
Light Source	RGBW Multi-spectral	

field of view (114×84 mm) with a lateral accuracy of 0.028 mm/pixel. Detailed hardware specifications are summarized in Table 2.

Stage 3: Data Collection and Annotation Refinement. The data collection and annotation process for Real-IAD Variety adheres to the rigorous standards established by Real-IAD [17]. This stage involves meticulous pixel-level manual annotation, algorithmic cross-validation, and iterative refinement. The annotation process iterates until the model’s predicted Average Precision (AP) scores exhibit negligible variation below a predetermined threshold, ensuring annotation consistency and quality.

3.2. Comparison with Existing IAD Datasets

Table 1 presents a comprehensive comparison between Real-IAD Variety and mainstream IAD datasets. Real-IAD Variety is the first IAD dataset to scale object categories to the hundreds, reaching 160 classes, thereby opening new avenues for large-scale multi-class anomaly detection research analogous to the impact of ImageNet [44] in object recognition. Compared to existing benchmarks, Real-IAD Variety offers several distinctive advantages: (1) Unprecedented scale: With 198,950 images including 159,045 anomalous images, it provides substan-

tially larger training and evaluation data compared to prior datasets; (2) Comprehensive annotations: Pixel-level defect masks with rigorous quality control ensure high annotation fidelity; (3) Multi-view coverage: Multiple viewpoints per sample enable robust evaluation of view-invariant detection methods. **Note that pixel-level masks are precisely annotated only in the specific views where defects are visible, as shown in Fig. 1;** (4) Diverse defect taxonomy: **27 defect types** represent the most comprehensive defect categorization among existing IAD datasets; (5) Extended resolution range: Image resolutions spanning, *i.e.*, 260~5,328 (Fig. .2), accommodate industrial products of varying scales, with over 90% of images exceeding 2,000 pixels in resolution. **The lower bound of the resolution range (260 pixels) represents a small fraction of images where objects were tightly cropped to ensure focus on the primary region of interest.** This wide resolution range stems from the diverse scales of industrial products, which are cropped such that the primary object occupies the majority (>90%) of the image area.

3.3. Dataset Characteristics and Statistical Analysis

Statistical Distribution Analysis. **To provide a comprehensive understanding of Real-IAD Variety, we perform a multi-dimensional statistical analysis as illustrated in Fig. 4 and Fig. 5:** (a) *Material Diversity*: Real-IAD Variety encompasses 24 material types (Fig. 4c) across 28 industrial domains (Fig. 2), providing comprehensive coverage from a practical application perspective. (b) *Color Spectrum*: The dataset captures a wide color distribution (Fig. 4d), which is critical for evaluating color-based anomaly detection approaches. (c) *Anomalous Region Proportion*: Compared to Real-IAD [17], Real-IAD Variety exhibits a broader and more uniform distribution of anomaly region proportions (Fig. 4a), encompassing defects of varying sizes and thereby increasing detection difficulty. (d) *Defect Aspect Ratio*: The defect aspect ratio distribution in Real-IAD Variety is comparably diverse to Real-IAD (Fig. 4b), presenting substantial challenges for models to discriminate defects of various geometric configurations. (e) *Distribution of Normal and Anomalous Images*: Real-IAD Variety contains a sub-

stantially more anomalous images than normal images, addressing the “data scarcity” bottleneck in real-world industrial environments (first row of Fig. 5). *(f) Defects Spatial and Scale Distribution:* The spatial distribution of defect centers (second row in Fig. 5) reveals that anomalies occur across almost all possible image coordinates, demonstrating high spatial diversity. The defect bounding boxes relative to the original resolution shows a high density near the origin (third row in Fig. 5), indicating a predominance of challenging small-scale defects. *(g) Anomalous Area Ratio Distribution:* The distribution of anomalous regions (fourth row in Fig. 5) approximately follows a normal distribution on a logarithmic scale.

Representative Defect Visualization. Figs. 1 and 3d presents qualitative visualizations of 27 distinct defect types across various object categories. Notably, identical defect types manifest in diverse visual forms across different materials, reflecting the complexity of real-world industrial scenarios. These defect types, derived from authentic production line observations, enhance the practical applicability of Real-IAD Variety.

Advantages and Contributions. Derived from real production environments, Real-IAD Variety offers several distinctive advantages: (1) Material diversity: 24 material types representing the most comprehensive material coverage among existing IAD datasets; (2) Defect types: 27 defect types providing extensive defect variation; (3) Resolution range: Extended image resolutions (260~5,328 pixels) accommodating diverse industrial product scales; (4) Multi-view annotations: Multiple viewpoints per sample enabling view-invariant evaluation; (5) Annotation quality: Highly accurate pixel-level mask annotations with rigorous quality control; Real-IAD Variety represents the first IAD dataset to systematically consider comprehensive coverage of material types and color variations, thereby robustly supporting the development of large-scale IAD models, industrial foundation models, vision-language IAD frameworks, and high-resolution detection paradigms.

4. Experiments

4.1. Dataset Protocol and Evaluation Metrics

Training and Testing Protocol. The Real-IAD Variety dataset is partitioned into training and testing subsets, comprising 3,991 normal samples (19,955 images) for training and 35,799 samples (3,990 normal and 31,809 anomalous, totaling 178,995 images) for testing. Notably, the testing subset exhibits a nearly balanced distribution between normal and anomalous images. To systematically evaluate the impact of category count on model performance, we partition Real-IAD Variety into three subsets: S1, S2, and S3, containing 30, 60, and 100 categories, respectively. The category selection employs randomization to mitigate potential biases in color and material representation.

Evaluation Metrics. Following established anomaly detection protocols, we employ seven standard metrics to quantify performance across both image-level anomaly classification and pixel-level anomaly localization tasks. These seven metrics include: Image-level Area Under the Receiver Operating Characteristic curve (I-ROC), Area Under the Precision-Recall curve (I-PR), and Maximum F1-score (I-F1_{max}); alongside their pixel-level counterparts (P-ROC, P-PR, P-F1_{max}) and Pixel-level Per-Region Overlap (P-PRO) for fine-grained defect localization assessment.

4.2. Benchmark on Multi-Class Unsupervised AD

Experimental Setting. The Multi-Class Unsupervised Anomaly Detection (MUAD) paradigm, first introduced in UniAD [4], trains a unified model on all object categories within a dataset simultaneously, employing the same model for both training and inference. This approach eliminates the need for class-specific models, thereby reducing storage costs and potentially enabling the learning of generalizable features across multiple categories. This setting holds significant practical value for real-world industrial deployment.

Unless otherwise specified, MUAD models are trained at a resolution of 256×256 for 100 epochs. For Dinomaly [6], we follow the original configuration,

Table 3: Quantitative MUAD performance comparisons on Real-IAD Variety dataset. Dis., Emb., and Rec. represent discrimination-based, embedding-based, and reconstruction-based methods, respectively. The best results are indicated in **bold**.

Paradigms	Methods	Backbone	I-ROC	I-PR	I-F1 _{max}	P-ROC	P-PR	P-F1 _{max}	P-PRO
Dis.	DRAEM [30]	scratch	54.9	89.3	93.1	53.7	9.4	8.8	21.6
	SimpleNet [11]	WRes-50	53.7	88.5	93.0	62.6	4.8	8.8	29.5
Emb.	CFA [10]	WRes-50	52.9	88.5	93.0	52.9	2.7	5.6	14.1
	CFLOW [36]	WRes-50	63.8	91.7	93.1	85.2	14.8	20.0	59.5
Rec.	UniAD [4]	Eff-B4	67.5	92.4	93.3	87.1	18.0	24.0	62.8
	OneNIP [5]	Eff-B4	73.3	94.2	93.5	88.3	24.0	29.4	64.5
	RD [35]	WRes-50	72.8	93.9	93.6	88.9	22.3	28.1	68.9
	DeSTSeg [37]	Res-18	75.0	95.0	93.2	65.6	37.2	31.2	37.7
	MambaAD [12]	Res-34	81.7	96.2	93.9	91.2	33.1	37.7	73.5
	LGC [38]	WRes-50	81.0	96.0	94.0	91.7	31.9	37.0	74.6
	Dinomaly [6]	ViT-B-14	85.4	97.2	94.5	91.5	42.8	45.8	75.6
	Dinomaly+ (Ours)	ViT-B-14	87.1	97.3	94.0	91.9	49.7	49.2	76.4

utilizing a 392×392 crop from 448×448 input resolution for 50,000 steps (*e.g.*, approximately 160 epochs for Real-IAD Variety). Additionally, we introduce Dinomaly+, an enhanced version that extends the baseline Dinomaly [6] by integrating a refined segmentation head [5]. This model is fine-tuned for 4 epochs on pseudo-anomalous data to facilitate high-precision, coarse-to-fine anomaly localization.

Competitive Methods. We evaluate state-of-the-art unsupervised anomaly detection methods to comprehensively assess their scalability to large-scale datasets. Specifically, we benchmark discrimination-based methods (DRAEM [30], SimpleNet [11]), embedding-based methods (CFA [10], CFLOW-AD [36]), and reconstruction-based methods (UniAD [4], OneNIP [5], RD [35], DesTSeg [37], MambaAD [12], LGC [38], Dinomaly [6], and Dinomaly+) on Real-IAD Variety. Considering computational resource constraints, we focus on representative and strong baselines, SimpleNet [11], CFLOW-AD [36], UniAD [4], OneNIP [5], MambaAD [12], LGC [38], Dinomaly [6], and Dinomaly+ for comprehensive MUAD evaluation. Notably, PatchCore [31] is excluded from this comparison due to its memory bank mechanism, which incurs prohibitive memory consumption as category

count increases, leading to potential out-of-memory issues on large-scale datasets like Real-IAD Variety.

Table 4: Comprehensive comparison of MUAD methods on Real-IAD and Real-IAD Variety datasets. Δ represents the absolute difference between the maximum and minimum values across Real-IAD Variety scales (S1, S2, S3, and Full) for each method, highlighting performance sensitivity to category expansion, and the same below.

Methods	Datasets	Classes	I-ROC	I-PR	I-F1 _{max}	P-ROC	P-PR	P-F1 _{max}	P-PRO
SimpleNet [11]	Real-IAD	30	54.9	50.6	61.5	76.1	1.9	4.9	42.4
	Real-IAD Variety S1	30	75.4	95.2	93.7	83.4	22.1	26.6	52.7
	Real-IAD Variety S2	60	67.9	93.2	93.2	80.4	14.4	19.2	46.0
	Real-IAD Variety S3	100	61.0	91.0	93.1	78.0	9.2	13.8	41.8
	Real-IAD Variety	160	53.7	88.5	93.0	62.6	4.8	8.8	29.5
	$\Delta \downarrow$		21.7	6.7	0.7	20.8	17.3	17.8	23.2
CFLOW [36]	Real-IAD	30	77.0	75.8	69.9	94.8	17.6	21.7	80.4
	Real-IAD Variety S1	30	77.9	95.4	93.8	87.5	27.6	26.6	63.2
	Real-IAD Variety S2	60	72.5	94.1	93.3	86.7	23.6	24.5	61.3
	Real-IAD Variety S3	100	72.5	94.1	93.3	86.7	23.6	24.5	61.3
	Real-IAD Variety	160	63.8	91.7	93.1	85.2	14.8	20.0	59.5
	$\Delta \downarrow$		14.1	3.7	0.7	2.3	12.8	6.6	3.7
UniAD [4]	Real-IAD	30	83.1	81.2	74.5	97.4	23.3	30.9	87.1
	Real-IAD Variety S1	30	78.5	95.4	94.2	88.9	29.5	34.1	67.1
	Real-IAD Variety S2	60	71.6	93.5	93.7	87.3	21.7	27.6	64.8
	Real-IAD Variety S3	100	70.2	93.1	93.5	87.6	19.0	25.0	64.3
	Real-IAD Variety	160	67.5	92.4	93.3	87.1	18.0	24.0	62.8
	$\Delta \downarrow$		11.0	3.0	0.9	1.8	11.5	10.1	4.3
OneNIP [5]	Real-IAD	30	86.3	84.9	77.2	98.2	37.0	41.5	88.6
	Real-IAD Variety S1	30	82.8	96.7	94.4	90.1	36.7	39.1	69.0
	Real-IAD Variety S2	60	78.7	95.6	94.0	89.2	30.7	34.9	67.2
	Real-IAD Variety S3	100	76.5	94.9	93.7	89.2	27.0	32.2	66.5
	Real-IAD Variety	160	73.3	94.2	93.5	88.3	24.0	29.4	64.5
	$\Delta \downarrow$		9.5	2.5	0.9	1.8	12.7	9.7	4.5
MambaAD [12]	Real-IAD	30	87.0	85.3	77.6	98.6	32.4	38.1	91.2
	Real-IAD Variety S1	30	87.4	97.5	95.0	91.5	40.3	43.2	75.0
	Real-IAD Variety S2	60	76.2	94.9	93.8	89.0	25.9	31.5	68.1
	Real-IAD Variety S3	100	72.5	93.7	93.6	88.5	21.5	27.5	66.5
	Real-IAD Variety	160	81.7	96.2	93.9	91.2	33.1	37.7	73.5
	$\Delta \downarrow$		14.9	3.8	1.4	3.0	18.8	15.7	8.5
LGC [38]	Real-IAD	30	85.9	83.0	77.1	98.1	25.3	33.3	90.7
	Real-IAD Variety S1	30	84.1	96.7	94.8	91.0	36.1	39.9	73.7
	Real-IAD Variety S2	60	82.7	96.4	94.1	91.9	34.5	39.9	76.6
	Real-IAD Variety S3	100	82.4	96.2	94.3	91.8	32.8	37.9	75.0
	Real-IAD Variety	160	81.0	96.0	94.0	91.7	31.9	37.0	74.6
	$\Delta \downarrow$		3.1	0.7	0.8	0.9	4.2	2.9	2.9
Dinomally [6]	Real-IAD	30	89.3	86.5	80.1	98.9	42.8	47.1	93.9
	Real-IAD Variety S1	30	91.4	98.4	95.8	92.7	56.9	56.1	79.8
	Real-IAD Variety S2	60	89.7	98.1	95.2	92.6	53.1	53.7	79.7
	Real-IAD Variety S3	100	87.9	97.7	94.9	92.3	48.3	50.1	78.2
	Real-IAD Variety	160	85.4	97.2	94.5	91.5	42.8	45.8	75.6
	$\Delta \downarrow$		6.0	1.2	1.4	1.2	14.1	10.3	4.2
Dinomally+	Real-IAD	30	89.5	86.3	80.7	98.8	45.6	48.6	93.1
	Real-IAD Variety S1	30	92.4	98.6	96.0	92.7	59.2	56.1	79.3
	Real-IAD Variety S2	60	90.8	98.3	95.3	92.9	56.3	54.4	79.6
	Real-IAD Variety S3	100	89.4	97.9	94.7	92.4	53.6	52.5	78.4
	Real-IAD Variety	160	87.1	97.3	94.0	91.9	49.7	49.2	76.4
	$\Delta \downarrow$		5.3	1.3	2.0	6.9	9.5	6.9	3.2

Results and Analysis under MUAD. Table 3 presents quantitative results of different methods on Real-IAD Variety.

Several critical observations emerge that differ from findings on small-scale datasets. (1) All evaluated methods experience substantial and non-saturating performance degradation on Real-IAD Variety, indicating the dataset’s signif-

icant challenge for future research and confirming successful avoidance of the metric saturation phenomenon frequently observed in smaller-scale IAD benchmarks; (2) Discrimination-based approaches (DRAEM [30], SimpleNet [11]) and the embedding-based method CFA [10] exhibit near-failure performance, with I-ROC scores approaching the 50% random guessing threshold. This demonstrates their severe lack of generalization capability when confronted with large-scale, high-variance normal data distributions; (3) Reconstruction-based methods maintain relatively satisfactory performance despite the challenging conditions. Notably, Dinomaly+ (building upon Dinomaly [6] and OneNIP [5]) achieves the best overall results across most metrics (*e.g.*, 87.1% I-ROC and 49.7% P-PR), underscoring its superior generalization and robustness across large-scale multi-class scenarios. In contrast, UniAD [4], an early MUAD method, exhibits one of the most severe performance drops among reconstruction-based approaches.

Performance Trends with Increasing Categories. To evaluate performance trends as category count increases, we selected eight representative methods across different paradigms. Table 4 presents results on Real-IAD and Real-IAD Variety with incrementally increasing categories: S1 (30 classes), S2 (60 classes), S3 (100 classes), and the full 160 classes (visualized in Fig. 6).

Several consistent conclusions emerge under the MUAD setting. (1) *Real-IAD Variety poses a more significant challenge than the original Real-IAD. Even at the same category scale (30 classes), the increased intra-class variance and higher resolution complexity in Real-IAD Variety lead to consistent performance drops across most baselines.* (2) A fundamental scalability challenge of MUAD is revealed: consistent performance degradation occurs across all evaluated methods as category count increases. Crucially, the magnitude of degradation is method-dependent, highlighting that different architectures possess varying levels of category generalization capability and robustness to scale-up. (3) Methods leveraging powerful backbones (*e.g.*, DINOv2-R [45]) and advanced architectures (*e.g.*, Transformer decoders), such as Dinomaly and Dinomaly+, consistently achieve substantial performance superiority across all




Figure 6: Performance trends of I-AUROC and P-AUPR metrics with increasing categories for MUAD, ZSAD and FSAD methods.

datasets. This advantage is especially prominent on high-scale benchmarks, underscoring their effectiveness over normalizing flow-based (*e.g.*, CFLOW-AD) or simpler backbone-based models (*e.g.*, SimpleNet, UniAD, MambaAD).

The performance degradation of MUAD methods when scaling to a large number of categories primarily stems from two factors. First, as the number of categories increases, models often suffer from a capacity-diversity conflict, converging toward a “generic” feature space that loses category-specific nuances essential for spotting subtle anomalies. To address this, recent methods such as OneNIP [5] and LGC [38] introduce category-level information via explicit visual prompts or class-aware contrastive learning to resolve feature reconstruction and preserve class-specific details. Second, reconstruction-based multi-class models are susceptible to identity mapping. In these cases, the model’s excessive generalization capability allows it to reconstruct the anomalies themselves rather than restoring them to a normal state, thereby failing to produce a significant residual for detection. The state-of-the-art Dinomaly [6] partially mitigates this by leveraging a more powerful vision encoder (*e.g.*, DINOv2-R) and simple dropout to achieve more robust performance.

Table 5: Quantitative comparison of ZSAD and FSAD methods on Real-IAD and Real-IAD Variety datasets. Results for FSAD methods represent the average and standard deviation over three runs with random seeds.

Methods	Datasets	Classes	shot	I-ROC	I-PR	I-F1 _{max}	P-ROC	P-PR	P-F1 _{max}	P-PRO
AnomalyCLIP [7]	Real-IAD	30	0	69.1	63.5	65.4	95.2	26.8	34.7	84.9
	Real-IAD Variety S1	30	0	68.6	88.9	88.6	88.7	36.5	40.3	71.7
	Real-IAD Variety S2	60	0	68.8	88.6	88.3	88.5	36.1	40.0	72.4
	Real-IAD Variety S3	100	0	69.5	88.9	88.4	89.5	36.2	40.2	74.2
	Real-IAD Variety	160	0	69.3	88.9	88.3	89.3	35.3	39.6	74.6
	$\Delta \downarrow$			0.9	0.3	0.3	1.0	1.2	0.7	2.9
AdaCLIP [42]	Real-IAD	30	0	71.1	67.2	67.0	95.6	33.6	38.1	90.2
	Real-IAD Variety S1	30	0	69.7	89.3	88.7	87.9	30.8	35.4	70.0
	Real-IAD Variety S2	60	0	70.1	88.9	88.7	88.0	32.0	36.6	71.3
	Real-IAD Variety S3	100	0	71.3	89.4	88.6	88.9	32.8	37.0	72.7
	Real-IAD Variety	160	0	71.1	89.3	88.6	89.1	32.6	36.8	73.1
	$\Delta \downarrow$			1.6	0.5	0.1	1.2	2.0	1.6	3.1
VCPCLIP [43]	Real-IAD	30	0	72.5	70.9	68.2	95.9	28.8	34.5	83.7
	Real-IAD Variety S1	30	0	71.4	90.0	88.7	87.3	34.9	38.1	67.2
	Real-IAD Variety S2	60	0	73.0	90.4	88.7	87.9	37.5	40.3	69.5
	Real-IAD Variety S3	100	0	72.9	90.5	88.6	89.1	37.6	40.4	71.1
	Real-IAD Variety	160	0	72.6	90.2	88.6	89.3	36.9	39.9	70.9
	$\Delta \downarrow$			1.6	0.5	0.1	2.0	2.7	2.3	3.9
AdaptCLIP [8]	Real-IAD	30	0	74.2	70.8	68.9	94.9	27.9	35.4	83.1
	Real-IAD Variety S1	30	0	73.2	94.7	93.7	88.0	36.5	40.0	68.3
	Real-IAD Variety S2	60	0	72.4	94.2	93.5	88.1	36.3	39.9	69.2
	Real-IAD Variety S3	100	0	72.6	94.2	93.4	89.0	36.6	40.1	71.1
	Real-IAD Variety	160	0	72.9	94.3	93.4	89.1	36.1	39.9	70.8
	$\Delta \downarrow$			0.4	0.5	0.3	1.1	0.5	0.2	2.8
MetaUAS [9]	Real-IAD	30	1	80.0±0.4	77.9±0.4	72.4±0.4	95.6±0.2	36.6±1.1	39.7±1.0	83.5±0.7
	Real-IAD Variety S1	30	1	80.0±0.2	96.1±0.1	94.3±0.0	90.5±0.2	49.0±0.9	48.3±0.7	73.3±0.1
	Real-IAD Variety S2	60	1	80.3±0.3	96.0±0.1	94.1±0.1	91.4±0.1	46.7±0.2	47.1±0.2	74.3±0.0
	Real-IAD Variety S3	100	1	81.7±0.3	96.1±0.0	94.2±0.1	91.9±0.0	48.1±0.5	48.2±0.4	76.0±0.2
	Real-IAD Variety	160	1	81.9±0.1	96.3±0.1	94.1±0.0	92.0±0.1	48.2±0.4	48.3±0.3	76.5±0.1
	$\Delta \downarrow$			1.9	0.3	0.2	1.5	2.3	1.2	3.2
AdaptCLIP [8]	Real-IAD	30	1	81.7±0.3	80.2±0.2	73.4±0.1	97.1±0.1	36.2±0.3	42.2±0.3	88.1±0.2
	Real-IAD Variety S1	30	1	83.3±0.3	96.9±0.0	94.4±0.0	91.1±0.1	49.4±0.1	50.0±0.1	74.2±0.1
	Real-IAD Variety S2	60	1	82.9±0.6	96.6±0.1	94.2±0.1	91.4±0.0	48.2±0.6	49.0±0.5	75.6±0.1
	Real-IAD Variety S3	100	1	83.8±0.3	96.7±0.1	94.2±0.0	92.2±0.1	48.8±0.6	49.8±0.6	77.6±0.2
	Real-IAD Variety	160	1	84.3±0.1	96.9±0.0	94.2±0.0	92.5±0.1	48.9±0.4	49.9±0.4	77.8±0.2
	$\Delta \downarrow$			1.4	0.3	0.2	1.4	1.2	1.0	3.6
AdaptCLIP [8]	Real-IAD	30	4	84.0±0.2	82.6±0.2	75.2±0.2	97.4±0.0	38.8±0.1	44.4±0.1	89.1±0.1
	Real-IAD Variety S1	30	4	87.4±0.2	97.6±0.1	94.8±0.1	91.7±0.0	52.8±0.2	52.8±0.1	76.4±0.1
	Real-IAD Variety S2	60	4	87.3±0.1	97.5±0.0	94.6±0.0	92.0±0.0	52.1±0.2	52.2±0.2	77.7±0.0
	Real-IAD Variety S3	100	4	88.0±0.1	97.6±0.0	94.7±0.0	92.8±0.0	52.6±0.2	52.9±0.2	79.4±0.0
	Real-IAD Variety	160	4	88.1±0.1	97.7±0.0	94.7±0.0	93.0±0.0	52.5±0.2	52.9±0.2	79.4±0.1
	$\Delta \downarrow$			0.8	0.2	0.2	1.3	0.7	0.7	3.0

4.3. Benchmark on Zero-Shot and Few-Shot AD

Experimental Setting. Zero-Shot Anomaly Detection (ZSAD) requires models trained on auxiliary datasets to identify anomalies in novel object categories without class-specific training data. This setting is particularly valuable because obtaining comprehensive coverage of all object classes and anomaly types is impractical in real-world scenarios. ZSAD aims to enhance model robustness and adaptability by facilitating knowledge transfer from known to unknown categories, thereby enabling effective detection of novel anomalies across diverse and dynamic environments. To further evaluate model adaptability in open-world scenarios, we introduce a Few-Shot Anomaly Detection (FSAD) setting along-

side ZSAD evaluation. In FSAD, a limited number of normal images ($k \in \{1, 4\}$) from the target class serve as visual prompts during inference, with their detection results integrated with ZSAD outputs to enhance final anomaly detection performance. **To ensure statistical significance, performance for all FSAD models is averaged across three runs with different random seeds, reporting both mean and standard deviation.**

Competitive Methods. We evaluate AnomalyCLIP [7], AdaCLIP [42], VCP-CLIP [43], and AdaptCLIP [8] on Real-IAD Variety in a zero-shot manner, conducting comprehensive comparisons with Real-IAD [17]. Given that Real-IAD Variety comprises 160 novel categories, it serves as an exceptionally challenging benchmark for assessing ZSAD model robustness and cross-category generalization in complex real-world scenarios. Notably, AdaptCLIP [8] supports both zero- and few-shot generalization across domains. Therefore, we employ it to evaluate FSAD performance on Real-IAD Variety. In addition, we also evaluate the one-shot pure-vision MetaUAS [9].

Results and Analysis under ZSAD and FSAD. Table 5 presents ZSAD and FSAD results on Real-IAD and Real-IAD Variety, with performance trends visualized in Fig. 6.

Some key observations are as follows: (1) ZSAD models demonstrate remarkable stability and robustness, while traditional MUAD methods exhibit significant performance degradation as category complexity increases. This fundamental difference is quantitatively highlighted by low I-ROC Δ values (*e.g.*, 0.9 for AnomalyCLIP, 0.4 for AdaptCLIP) across increasing category counts (30 to 160). This suggests that VLM-based ZSAD approaches, which leverage external knowledge from vast pre-training data rather than modeling intrinsic training set distributions, are inherently less sensitive to data distribution shifts induced by category scale-up. (2) FSAD yields substantial performance gains by **leveraging class-specific prompt information**. Specifically, using only 4 normal images (4-Shot), AdaptCLIP’s performance on Real-IAD Variety (160 classes) improves significantly, with I-AUROC rising from 72.9% (ZSAD) to 88.1% (4-Shot). This improvement confirms the high value of incorporating

even few-shot normal visual prompts to rapidly refine generic VLM knowledge for precise, target-specific anomaly discrimination. The consistent improvement from 1-Shot to 4-Shot further validates the effectiveness of this adaptation mechanism. (3) Few-shot AdaptCLIP (4-Shot) achieves the best overall performance across all benchmarks. On the most challenging 160-class Real-IAD Variety, 4-Shot AdaptCLIP reaches I-ROC of 88.1%, which is highly competitive with the best multi-class Dinomaly [6] (85.4%) that utilizes full normal training data. Crucially, delta analysis reveals that while few-shot adaptation boosts absolute performance, it incurs slightly higher delta for pixel-level metrics, suggesting that local adaptation mechanisms, while powerful, exhibit greater sensitivity to category complexity than global ZSAD paradigms.

ZSAD and FSAD leverage semantic anchoring via textual prompts or visual representation with normal images prompts. Unlike MUAD, these methods utilize a pre-aligned latent space where linguistic concepts of anomalies (*e.g.*, “damaged”, “defect”) provide a unified representation of common anomalous (normal visual) features. Because they rely on frozen VLMs and semantic reasoning rather than statistical learning on specific datasets, they require no category-specific fine-tuning. Consequently, their generalization capability is fundamentally decoupled from the number of test categories. MUAD models may suffer from feature “pollution” or capacity constraints as more classes are added. In contrast, ZSAD models identify defects based on pre-trained semantic alignment rather than statistical novelty relative to a specific training set. To further unlock the potential of these models, particularly in the complex cross-category setting of Real-IAD Variety, future research should focus on developing enhanced methods that effectively utilize detailed textual and semantic information, bridging the gap between high-level semantic understanding and fine-grained localized visual anomaly detection.

4.4. Benchmark on Multi-View AD

Experimental Setting. Multi-View Anomaly Detection (MVAD) originates from high-precision demands in industrial quality inspection. It aims to achieve

detection capabilities unattainable from single-view images by integrating multiple perspective images of the same sample, thereby detecting defects that may be occluded or invisible in certain viewpoints but remain discernible in others. Additionally, this setting leverages multi-view information and consistency constraints to further enhance model performance.

Table 6: MVAD [39] performance comparison of Real-IAD Variety with increasing categories.

Model/Dataset	I-ROC	I-PR	I-F1 _{max}	P-ROC	P-PR	P-F1 _{max}	P-PRO
Real-IAD-30c	86.6	84.8	77.2	97.9	30.3	36.8	91.2
Real-IAD Variety S1	85.4	97.0	94.8	90.7	36.8	39.7	72.4
Real-IAD Variety S2	81.4	96.2	94.1	90.3	31.8	36.2	71.0
Real-IAD Variety S3	80.3	95.8	94.0	90.5	28.8	33.9	70.5
Real-IAD Variety	77.1	95.1	93.6	89.5	24.5	30.2	67.7

Results and Analysis under MVAD. Table 6 presents quantitative results of MVAD [39] on Real-IAD Variety. As category count increases, model performance gradually decreases, consistent with trends observed in MUAD settings. In practical production line applications, multi-view imaging is essential to reduce false negative rates. Nevertheless, current metric results remain relatively modest, presenting a significant challenge for future research to address scalability and robustness in large-scale multi-view anomaly detection scenarios.

4.5. Potential for Foundational AD Models

Experimental Setting. To demonstrate the generalization of foundational AD models pre-trained large-scale Real-IAD Variety dataset. We employ AD-Pretrain [46] framework (DINOv2-base backbone) and compare the effects of pre-training on the original Real-IAD and our proposed Real-IAD Variety. Then, we evaluate few-shot (1, 2, 4, 8-shot) generalization on the independent MVTec and VisA benchmarks.

Few-shot Generalization. As shown in Table 7, the few-shot performance of ADPretrain model pre-trained on Real-IAD Variety consistently outperforms the baseline across all settings. Notably, the performance gains are significantly

Table 7: Few-shot anomaly detection results on MVTec and VisA. We compare the ADPre-train model (FeatureNorm) pre-trained on Real-IAD versus our Real-IAD Variety.

Test Set	Train Set	Shot	I-ROC	I-PR	I-F1 _{max}	P-ROC	P-PR	P-F1 _{max}	P-PRO
MVTec	Real-IAD	1	87.9	93.8	90.2	96.7	48.1	53.4	91.4
		2	95.6	97.0	95.3	97.3	50.2	55.2	92.3
		4	97.4	98.3	96.6	97.4	50.6	55.6	92.5
		8	97.8	98.6	96.8	97.5	51.5	56.0	92.8
	Real-IAD Variety	1	90.7	93.8	93.1	97.3	54.1	56.6	92.0
		2	96.0	97.5	95.8	97.7	56.1	58.1	92.9
		4	97.7	99.0	97.2	98.0	56.9	58.6	93.6
		8	98.1	99.1	97.5	98.1	57.4	59.1	93.7
ViSA	Real-IAD	1	61.4	65.7	73.2	96.4	19.2	28.0	81.0
		2	68.8	75.1	74.1	97.0	25.1	33.5	83.7
		4	75.2	81.5	77.2	97.6	28.8	37.3	85.5
		8	82.6	84.1	80.4	98.0	31.7	40.9	87.1
	Real-IAD Variety	1	58.6	62.4	73.1	96.8	20.8	30.7	83.1
		2	70.4	74.6	75.1	97.6	30.0	37.2	86.5
		4	88.1	90.0	84.8	98.0	34.4	40.5	88.5
		8	93.4	94.4	89.6	98.3	36.1	42.1	89.7

more pronounced on the VisA dataset, which is widely recognized as a more challenging benchmark due to its complex object structures and subtle defects. For instance, in the 4-shot setting on VisA, pre-training on Real-IAD Variety improves I-ROC from 75.2% to 88.1% (+12.9%). We observe a slight fluctuation in the 1-shot ViSA setting, which may be attributed to the high sensitivity of extremely low-shot sampling in certain complex categories. These results empirically demonstrate that the increased category diversity and broader industrial scope of Real-IAD Variety are crucial for enhancing generalization on foundational AD models. By providing a superior foundation for feature alignment, our dataset helps foundational models learn robust visual primitives that generalize far better to challenging unseen domains.

5. Future Work

Foundation Models for Anomaly Detection: Scaling Toward a New Paradigm. Foundation models have fundamentally transformed how diverse tasks are addressed through unified architectures, ushering in a new era of model design [47, 48, 49]. A critical enabler of this transformation is the availability of large-scale, high-quality datasets, which are essential for training versatile and generalizable models across various domains. In the context of Industrial Anomaly Detection (IAD), the development of foundation models is heavily dependent on comprehensive vision-language annotations. These annotations are indispensable not only for accurate anomaly detection but also for providing contextual understanding that enables models to reason about anomalies, thereby enhancing their generalization capabilities across diverse industrial scenarios.

Multimodal Anomaly Detection with Large-Scale Diverse Data. The introduction of Real-IAD Variety opens several promising research directions for advancing anomaly detection. First, the development of robust generalized anomaly detection models represents a critical avenue. Future research should explore training strategies that effectively integrate these textual insights with visual features, *e.g.*, MMAD [50], thereby enhancing the robustness and adaptability of anomaly detection systems. Second, addressing logical anomalies in image-based anomaly detection remains an important challenge. Current methods often struggle with logical anomalies, *where anomalies are defined by misplaced components or incorrect counts rather than distinctive surface textures.* Our ZSAD/FSAD results suggest that while VLMs excel at semantic alignment, *bridging the gap between high-level reasoning and fine-grained spatial logic remains a frontier for foundational IAD research.* Our preliminary ZSAD and FSAD experiments have revealed performance limitations in this regard. To address these constraints, we propose enabling text encoder fine-tuning to enhance model adaptability and contextual reasoning capabilities. Future research should focus on integrating textual descriptions into existing anomaly detection frameworks, enabling models to reason about logical relationships and contex-

tual cues beyond purely visual patterns.

6. Conclusion

Industrial Anomaly Detection (IAD) is undergoing a fundamental transformation with the emergence of large-scale models capable of operating in unified and zero-/few-shot settings. This work makes a substantial contribution to the IAD dataset landscape by introducing Real-IAD Variety, a large-scale benchmark that not only expands the scope of available data but also establishes rigorous evaluation protocols for assessing state-of-the-art (SOTA) IAD methodologies. Real-IAD Variety, characterized by its unprecedented scale and diversity, redefines the benchmark standards for IAD datasets. It encompasses 160 categories spanning 28 industries, 24 material types, 22 color variations and 27 defect types with 198,950 images, effectively overcoming the limitations of previous datasets that were constrained by narrow category ranges and limited scenario representation.

Our comprehensive experimental analysis on Real-IAD Variety reveals several critical findings: (1) Most multi-class unsupervised anomaly detection methods experience significant performance degradation as the category count increases, with I-ROC typically declining by 10%–20% when scaling from 30 to 160 categories; (2) Zero-shot and few-shot approaches demonstrate remarkable resilience to category scale-up, exhibiting minimal performance fluctuation. This suggests their superior scalability and effectiveness in handling diverse industrial scenarios; (3) Vision-language models (VLMs) leveraging external pre-trained knowledge exhibit fundamentally different scalability characteristics compared to traditional unsupervised methods, highlighting the pivotal role of multimodal learning in the next generation of IAD systems. Real-IAD Variety and its associated benchmarks provide a solid foundation for future research and innovation, particularly in multi-class, multi-view and zero-/few-shot anomaly detection paradigms. We anticipate that this benchmark will facilitate the development of models with enhanced generalization capabilities, ultimately improving the

reliability and efficiency of anomaly detection systems in real-world industrial applications.

Acknowledgements

This work was partially supported by the National Natural Science Foundation of China (Grant No. 62171139).

References

- [1] X. Jiang, J. Liu, J. Wang, Q. Nie, K. Wu, Y. Liu, C. Wang, F. Zheng, Softpatch: Unsupervised anomaly detection with noisy data, NeurIPS 35 (2022).
- [2] C. Wang, X. Jiang, B.-B. Gao, Z. Gan, Y. Liu, F. Zheng, L. Ma, Softpatch+: Fully unsupervised anomaly classification and segmentation, Pattern Recognition 161 (2025) 111295.
- [3] C. Wang, H. Zhu, J. Peng, Y. Wang, R. Yi, Y. Wu, L. Ma, J. Zhang, M3dmnr: Rgb-3d noisy-resistant industrial anomaly detection via multimodal denoising, IEEE Transactions on Pattern Analysis and Machine Intelligence (2025).
- [4] Z. You, L. Cui, Y. Shen, K. Yang, X. Lu, Y. Zheng, X. Le, A unified model for multi-class anomaly detection, NeurIPS 35 (2022).
- [5] B.-B. Gao, Learning to detect multi-class anomalies with just one normal image prompt, in: ECCV, 2024.
- [6] J. Guo, S. Lu, W. Zhang, F. Chen, H. Li, H. Liao, Dinomaly: The less is more philosophy in multi-class unsupervised anomaly detection, in: CVPR, 2025.
- [7] Q. Zhou, G. Pang, Y. Tian, S. He, J. Chen, AnomalyCLIP: Object-agnostic prompt learning for zero-shot anomaly detection, in: ICLR, 2024.

- [8] B.-B. Gao, Y. Zhou, J. Yan, Y. Cai, W. Zhang, M. Wang, J. Liu, Y. Liu, L. Wang, C. Wang, Adaptclip: Adapting clip for universal visual anomaly detection, in: AAAI, 2026.
- [9] B.-B. Gao, Metauas: Universal anomaly segmentation with one-prompt meta-learning, in: NeurIPS, 2024.
- [10] S. Lee, S. Lee, B. C. Song, Cfa: Coupled-hypersphere-based feature adaptation for target-oriented anomaly localization, IEEE Access (2022).
- [11] Z. Liu, Y. Zhou, Y. Xu, Z. Wang, Simplenet: A simple network for image anomaly detection and localization, in: CVPR, 2023.
- [12] H. He, Y. Bai, J. Zhang, Q. He, H. Chen, Z. Gan, C. Wang, X. Li, G. Tian, L. Xie, Mambaad: Exploring state space models for multi-class unsupervised anomaly detection, in: NeurIPS, 2024.
- [13] Z. Gu, B. Zhu, G. Zhu, Y. Chen, M. Tang, J. Wang, Anomalygpt: Detecting industrial anomalies using large vision-language models, in: AAAI, 2024.
- [14] P. Bergmann, M. Fauser, D. Sattlegger, C. Steger, Mvtec ad—a comprehensive real-world dataset for unsupervised anomaly detection, in: CVPR, 2019.
- [15] Y. Zou, J. Jeong, L. Pemula, D. Zhang, O. Dabeer, Spot-the-difference self-supervised pre-training for anomaly detection and segmentation, in: ECCV, 2022.
- [16] Q. Zhou, W. Li, L. Jiang, G. Wang, G. Zhou, S. Zhang, H. Zhao, Pad: A dataset and benchmark for pose-agnostic anomaly detection, NeurIPS 36 (2023).
- [17] C. Wang, W. Zhu, B.-B. Gao, Z. Gan, J. Zhang, Z. Gu, S. Qian, M. Chen, L. Ma, Real-iad: A real-world multi-view dataset for benchmarking versatile industrial anomaly detection, in: CVPR, 2024.

- [18] G. Gui, B.-B. Gao, J. Liu, C. Wang, Y. Wu, Few-shot anomaly-driven generation for anomaly classification and segmentation, in: ECCV, 2024.
- [19] H. Sun, Y. Cao, H. Dong, O. Fink, Unseen visual anomaly generation, in: CVPR, 2025.
- [20] Z. Dai, S. Zeng, H. Liu, X. Li, F. Xue, Y. Zhou, Seas: few-shot industrial anomaly image generation with separation and sharing fine-tuning, in: ICCV, 2025.
- [21] Y. Huang, C. Qiu, K. Yuan, Surface defect saliency of magnetic tile, *The Visual Computer* 36 (2020) 85–96.
- [22] S. Jezek, M. Jonak, R. Burget, P. Dvorak, M. Skotak, Deep learning-based defect detection of metal parts: evaluating current methods in complex conditions, in: ICUMTW, 2021.
- [23] P. Mishra, R. Verk, D. Fornasier, C. Piciarelli, G. L. Foresti, Vt-adl: A vision transformer network for image anomaly detection and localization, in: ISIE, 2021.
- [24] D. Tabernik, S. Šela, J. Skvarč, D. Skočaj, Segmentation-based deep-learning approach for surface-defect detection, *Journal of Intelligent Manufacturing* 31 (3) (2020) 759–776.
- [25] L. Fan, D. Fan, Z. Hu, Y. Ding, D. Di, K. Yi, M. Pagnucco, Y. Song, Manta: A large-scale multi-view and visual-text anomaly detection dataset for tiny objects, in: CVPR, 2025.
- [26] S. Höfer, D. Henning, A. Amiranashvili, D. Morrison, M. Tzes, I. Posner, M. Matvienko, A. Rennola, A. Milan, Kaputt: A large-scale dataset for visual defect detection, in: ICCV, 2025.
- [27] J. Zhang, C. Wang, X. Li, G. Tian, Z. Xue, Y. Liu, G. Pang, D. Tao, Learning feature inversion for multi-class anomaly detection under general-purpose coco-ad benchmark, *arXiv preprint arXiv:2404.10760* (2024).

- [28] H. Ling, J. Guo, Z. Tao, Y. Cao, D. Di, H. Xu, X. Su, Y. Song, L. Fan, Adnet: A large-scale and extensible multi-domain benchmark for anomaly detection across 380 real-world categories, arXiv preprint arXiv:2511.20169 (2025).
- [29] C.-L. Li, K. Sohn, J. Yoon, T. Pfister, Cutpaste: Self-supervised learning for anomaly detection and localization, in: CVPR, 2021.
- [30] V. Zavrtanik, M. Kristan, D. Skočaj, Draem-a discriminatively trained reconstruction embedding for surface anomaly detection, in: ICCV, 2021.
- [31] K. Roth, L. Pemula, J. Zepeda, B. Schölkopf, T. Brox, P. Gehler, Towards total recall in industrial anomaly detection, in: CVPR, 2022.
- [32] M. Rudolph, T. Wehrbein, B. Rosenhahn, B. Wandt, Fully convolutional cross-scale-flows for image-based defect detection, in: WACV, 2022, pp. 1088–1097.
- [33] C. Zhou, R. C. Paffenroth, Anomaly detection with robust deep autoencoders, in: SIGKDD, 2017.
- [34] Y. Liang, J. Zhang, S. Zhao, R. Wu, Y. Liu, S. Pan, Omni-frequency channel-selection representations for unsupervised anomaly detection, IEEE Transactions on Image Processing (2023).
- [35] H. Deng, X. Li, Anomaly detection via reverse distillation from one-class embedding, in: CVPR, 2022.
- [36] D. Gudovskiy, S. Ishizaka, K. Kozuka, Cflow-ad: Real-time unsupervised anomaly detection with localization via conditional normalizing flows, in: WACV, 2022.
- [37] X. Zhang, S. Li, X. Li, P. Huang, J. Shan, T. Chen, Destseg: Segmentation guided denoising student-teacher for anomaly detection, in: CVPR, 2023.

- [38] L. Fan, J. Huang, D. Di, A. Su, T. Song, M. Pagnucco, Y. Song, Salvaging the overlooked: Leveraging class-aware contrastive learning for multi-class anomaly detection, in: ICCV, 2025.
- [39] H. He, J. Zhang, G. Tian, C. Wang, L. Xie, Learning multi-view anomaly detection, arXiv preprint arXiv:2407.11935 (2024).
- [40] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sasstry, A. Askell, P. Mishkin, J. Clark, et al., Learning transferable visual models from natural language supervision, in: ICML, 2021.
- [41] J. Jeong, Y. Zou, T. Kim, D. Zhang, A. Ravichandran, O. Dabeer, Winclip: Zero-/few-shot anomaly classification and segmentation, in: CVPR, 2023.
- [42] Y. Cao, J. Zhang, L. Frittoli, Y. Cheng, W. Shen, G. Boracchi, Adacclip: Adapting clip with hybrid learnable prompts for zero-shot anomaly detection, in: ECCV, 2024.
- [43] Z. Qu, X. Tao, M. Prasad, F. Shen, Z. Zhang, X. Gong, G. Ding, Vcp-clip: A visual context prompting model for zero-shot anomaly segmentation, in: ECCV, 2024.
- [44] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, Imagenet: A large-scale hierarchical image database, in: CVPR, 2009.
- [45] T. Darcet, M. Oquab, J. Mairal, P. Bojanowski, Vision transformers need registers, in: ICLR, 2024.
- [46] X. Yao, Y. Luo, Z. Qian, C. Zhang, Adpretrain: Advancing industrial anomaly detection via anomaly representation pretraining, in: NeurIPS, 2025.
- [47] W. X. Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, B. Zhang, J. Zhang, Z. Dong, et al., A survey of large language models, arXiv preprint arXiv:2303.18223 (2023).

- [48] J. Zhang, J. Huang, S. Jin, S. Lu, Vision-language models for vision tasks: A survey, *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2024).
- [49] Y. Du, Z. Liu, J. Li, W. X. Zhao, A survey of vision-language pre-trained models, in: *IJCAI*, 2022.
- [50] X. Jiang, J. Li, H. Deng, Y. Liu, B.-B. Gao, Y. Zhou, J. Li, C. Wang, F. Zheng, Mmad: A comprehensive benchmark for multimodal large language models in industrial anomaly detection, in: *ICLR*, 2025.

The supplementary material presents more detailed industry groups, material names, and color type names.

Industry Groups.

c1: Electrical Manufacturing

- s1: Manufacturing of Power Electronics Components
- s2: Manufacturing of Distribution Switch Control Equipment
- s3: Manufacturing of Wires and Cables
- s4: Manufacturing of Other Power Distribution and Control Equipment
- s5: Battery Manufacturing
- s6: Motor Manufacturing

c2: Transport Manufacturing

- s7: Manufacturing of Distribution Switch Control Equipment
- s8: Plastic Toy Manufacturing

c3: Cultural Products Manufacturing

- s9: Stationery Manufacturing
- s10: Manufacturing of Daily Plastic Products
- s11: Metal Tool Manufacturing

c4: Metal Manufacturing

- s12: Manufacturing of Metal Accessories for Construction and Furniture
- s13: Manufacturing of Decorative Metal Products

c5: General Manufacturing

- s14: Manufacturing of Safety and Fire Protection Metal Products
- s15: Manufacturing of Daily Miscellaneous Products
- s16: General Spare Parts Manufacturing

c6: Electronics Manufacturing

- s17: Manufacturing of Resistors Capacitors and Inductors
- s18: Manufacturing of Sensors
- s19: Manufacturing of Acoustics Devices and Parts
- s20: Manufacturing of Semiconductor Lighting Devices
- s21: Electronic Circuit Manufacturing

c7: Rubber and Plastic Products Manufacturing

- s22: Manufacturing of Plastic Parts and Other Plastic Products
- s23: Rubber Parts Manufacturing
- s24: Daily Plastic Products Manufacturing
- s25: Daily Miscellaneous Products Manufacturing

c8: Other

- s26: Manufacturing of Clothing Supplies
- s27: Manufacturing of Automotive Parts and Accessories
- s28: Clock and Timing Instrument Manufacturing

Material names.

- u1: Metal**
 - v1: Aluminum Alloy
 - v2: Zinc Alloy
 - v3: Iron Oxide
 - v4: Steel
 - v5: Lithium
 - v6: Copper

u2: Plastic	{	v7: ABS
		v8: PC
		v9: PA
		v10: PPS
		v11: PET
		v12: PS
		v13: PP
		v14: POM
		v15: PBT
		v16: PVC
u3: Other	{	v17: Acrylic
		v18: PE
		v19: Silicon
		v20: Carbon
		v21: Rubber
		v22: Ceramics
		v23: Wood
		v24: Glass cup epoxy resin

Color type names.

- a1: Chroma** {
- b1: Gray
 - b2: Pink
 - b3: Green
 - b4: Purple
 - b5: Blue
 - b6: Red
 - b7: Yellow
 - b8: White
 - b9: Gold
 - b10: Brown
- a2: Mixer Color** {
- b11: Green + Silver
 - b12: Black + silver
 - b13: Black + Yellow
 - b14: Yellow + Silver
 - b15: Blue + white
 - b16: Black + Gold
 - b17: Silver + blue
 - b18: Black + white
 - b19: Black + blue
 - b20: Red + silver
 - b21: Silver + gold
 - b22: White + silver
- a3: Silver**
- a4: Black**




Figure .1: The distributions of image height and width in our Real-IAD-Variety.




Figure .2: The class-wise distribution of normal and abnormal images in our Real-IAD Variety dataset.