

AD-FM: Multimodal LLMs for Anomaly Detection via Multi-Stage Reasoning and Fine-Grained Reward Optimization

Jingyi Liao^{1,2}, Yongyi Su^{2,3}, Rong-Cheng Tu^{1*}, Zhao Jin¹, Wenhao Sun¹, Yiting Li², Xun Xu^{2*}, Dacheng Tao¹, Xulei Yang²

¹Nanyang Technological University
²Institute for Infocomm Research (I²R), A*STAR
³South China University of Technology

Abstract

While Multimodal Large Language Models (MLLMs) demonstrate remarkable capabilities across diverse domains, their application to specialized anomaly detection (AD) remains constrained by domain adaptation challenges. Existing Group Relative Policy Optimization (GRPO) based approaches suffer from two critical limitations: inadequate training data utilization when models produce uniform responses, and insufficient supervision over reasoning processes that encourage immediate binary decisions without deliberative analysis. We propose a comprehensive framework addressing these limitations through two synergistic innovations. First, we introduce a multi-stage deliberative reasoning process that guides models from region identification to focused examination, generating diverse response patterns essential for GRPO optimization while enabling structured supervision over analytical workflows. Second, we develop a fine-grained reward mechanism incorporating classification accuracy and localization supervision, transforming binary feedback into continuous signals that distinguish genuine analytical insight from spurious correctness. Comprehensive evaluation across multiple industrial datasets shows that our method achieves superior accuracy by enabling general-purpose MLLMs to acquire fine-grained visual discrimination for detecting subtle manufacturing defects.

Introduction

Anomaly detection typically involves modeling the distribution of normal data to identify deviations. Recent advances in multi-modal learning have enabled zero-shot generalization by incorporating language descriptions (Jeong et al. 2023; Zhou et al. 2023). The emergence of multimodal large language models (MLLMs) further expands this capability by allowing natural language reasoning over visual inputs (Jiang et al. 2025; Tu et al. 2025; Gu et al. 2024).

However, recent studies (Jiang et al. 2025; Zhang et al. 2025) show that general-purpose MLLMs struggle with specialized tasks like industrial anomaly detection, where fine-grained visual cues, e.g. subtle surface defects require reasoning beyond what is captured in their pretraining. This challenge motivates a paradigm of adapting general MLLMs towards downstream anomaly detection tasks.

*Corresponding authors

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

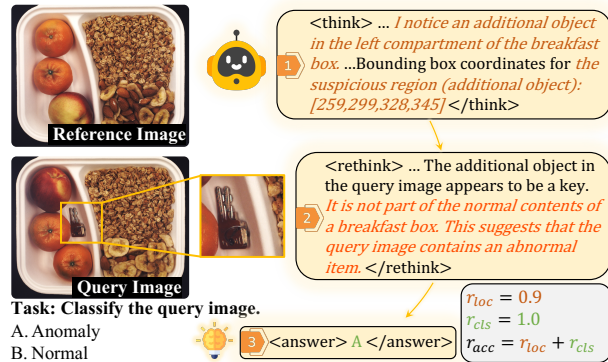


Figure 1: Deliberative inference pipeline of AD-FM. The model performs multi-stage reasoning by identifying suspicious regions (<think>), focused region analysis (<rethink>), and making a final decision (<answer>). The localization reward r_{loc} and classification reward r_{cls} jointly form the fine-grained reward.

Reinforcement fine-tuning (RFT) (Ouyang et al. 2022) offers a promising adaptation path by enabling task-specific reasoning without needing additional supervision. Compared to supervised fine-tuning (SFT), RFT facilitates greater exploration and generalization by optimizing against a reward signal derived from a simple reward function, thereby circumventing the reliance on detailed, token-level reasoning labels. Notably, Group Relative Policy Optimization (GRPO) (Shao et al. 2024) has shown potential when applied to MLLMs for anomaly detection (Chao et al. 2025; Zhao et al. 2025). Yet, existing applications of GRPO to anomaly detection suffer from two critical limitations. First, GRPO’s fundamental reliance on response/reward variance for learning signals leads to severely reduced data utilization efficiency when the model produces uniform outputs. Second, the lack of supervision over the reasoning process allows flawed analytical patterns to persist unchecked.

In GRPO, multiple responses to a prompt are scored and normalized (e.g., via z-score). When responses are uniform, e.g. consistently incorrect/correct, the variance collapses, producing vanishing gradients and halting learning. This is particularly problematic in downstream anomaly-detection tasks under domain shift, where the binary nature of clas-

sification decisions creates an inherently constrained reward space that severely limits response diversity. Moreover, even consistently correct answers may rely on flawed reasoning due to insufficient supervision of the thinking process. For instance, a model may detect an anomaly but reference an irrelevant region, or fail to identify all anomalous regions, which we term as *spurious correctness*.

To address these challenges, we introduce two synergistic improvements: a multi-stage deliberative reasoning framework and fine-grained reward mechanisms.

We first restructure the reasoning process to encourage more careful and methodical visual analysis, which is then utilized for a more fine-grained reward signal. Motivated by recent successes in test-time scaling and chain-of-thought prompting (Madaan et al. 2023; Kojima et al. 2022; Wei et al. 2022), and adding additional reasoning process in GRPO (Fan et al. 2025), we propose a multi-stage deliberative reasoning framework that encourages more cautious decision-making while creating new supervisable intermediate signals. As illustrated in Fig. 1, our method guides models through a structured process: first identifying and localizing regions of interest, then performing focused analysis within these regions before producing final decisions. This approach aligns with human expert cognitive strategies during visual inspection, i.e., rapid global scanning followed by focused examination of suspicious areas, while improving traceability and reliability of the model’s decision process.

With models now producing detailed spatial predictions, we introduce complementary reward mechanisms that leverage this enriched information. These rewards deliver dual benefits: enhanced reward variance for improved GRPO learning and systematic filtering of spurious correctness patterns. Inspired by multi-objective GRPO frameworks (Li et al. 2025a), we design fine-grained reward functions that assess both decision accuracy and localization quality. Specifically, we prompt the model to not only provide a decision and rationale but also localize the anomalies via bounding boxes. Since defects may spread out on the object of interest, our first reward measures count accuracy of predicted boxes against ground truth, while the second computes Generalized Intersection over Union (GIoU) between predicted and ground-truth boxes. These spatially grounded rewards enable fine-grained discrimination between genuine understanding and spurious correctness, allowing models to exploit available data more effectively while developing reliable visual analysis capabilities, thereby enhancing deployment safety and trustworthiness in inspection scenarios.

The synergy between multi-stage reasoning and fine-grained rewards creates a comprehensive framework that simultaneously addresses GRPO’s variance dependency and the challenge of supervision for reasoning process, transforming both the learning efficiency and reasoning quality of GRPO-based anomaly detection systems. In summary, we propose a framework for **Anomaly Detection** via **Fine-grained reward optimization** and **Multi-stage reasoning (AD-FM)**. Our key contributions are as follows:

- We introduce a multi-stage deliberative reasoning framework that creates supervisory signals for reasoning and

enables systematic visual inspection for AD.

- We design fine-grained rewards that suppress spurious correctness and maximize data utilization via joint optimization of classification accuracy and spatial understanding.
- We demonstrate substantial performance improvements across multiple datasets in both in-distribution and out-of-distribution settings, effectively bridging the gap between general-purpose MLLMs and specialized visual inspection requirements.

Related Works

Anomaly Detection. Anomaly Detection (AD) involves identifying instances that deviate from normal patterns. Modern AD techniques are predominantly unsupervised and aim to model the distribution of normal data (Zhang et al. 2025). Key approaches include: one-class classification (Liu et al. 2023; Liao et al. 2024), reconstruction-based methods that identify anomalies via reconstruction error (He et al. 2024; Zhang, Xu, and Zhou 2024; Li et al. 2025b), feature distillation which compares original and distilled features (Deng and Li 2022), and embedding-based strategies that measure feature divergence between normal and target samples (Roth et al. 2022; Lei et al. 2023).

Zero-shot Anomaly Detection. Zero-shot AD leverages textual prompts to detect anomalies in unseen distributions, enabling deployment without retraining. Early work used CLIP-based models (Li et al. 2025c; Gu et al. 2024) to align visual features with anomaly-related text, while (Cao et al. 2025) incorporated SAM (Kirillov et al. 2023) for improved localization. The rise of Multimodal LLMs (MLLMs) (OpenAI 2023; Hurst et al. 2024) expanded their use in vision tasks, including anomaly detection (Cao et al. 2023; Deng et al. 2025). Although early attempts (Li et al. 2025c; Gu et al. 2024) explored few-shot prompting, they lacked binary decision capability. MMAD (Jiang et al. 2025) introduced the first MLLM-based binary AD benchmark, but both MMAD and (Chen et al. 2025b) reported limited zero/one-shot performance despite improved explainability.

Reinforcement Fine-tuning. Reinforcement fine-tuning (RFT), first applied in O1 (Jaech et al. 2024) and DeepSeek R1 (Guo et al. 2025), has proven effective for structured reasoning tasks. Multimodal RFT (Liu et al. 2025; Chen et al. 2025a) extends this to vision with task-specific reward functions, while (Fan et al. 2025) introduces a dual-thinking process but lacks supervision over intermediate steps. However, applying RFT to AD remains nascent. AnomalyR1 (Chao et al. 2025) used GRPO directly without addressing AD-specific issues such as sparse binary rewards and low feedback diversity. OmniAD (Zhao et al. 2025) relied mainly on supervised tuning, underutilizing GRPO’s exploratory strengths. These limitations highlight the need for GRPO adaptations tailored to AD’s unique constraints.

Methodology

Revisiting GRPO

Group Relative Policy Optimization (GRPO) is a post-training algorithm that aligns large language models (LLMs)

with downstream tasks using reinforcement learning from preference feedback (RLPF). Given a prompt x , the policy model π_θ generates multiple responses $\{y_i\}_{i=1}^G \sim \pi_\theta(\cdot | x)$. Each response is assigned a scalar reward $r(x, y_i) \in \mathbb{R}$, reflecting its quality under human or learned preferences. The model parameters θ are updated via policy gradient to maximize estimated advantages (i.e., the normalized rewards):

$$\begin{aligned} \nabla_\theta J(\theta) &= \mathbb{E}_{x, y_i} \left[\hat{A}(x, y_i) \nabla_\theta \log \pi_\theta(y_i | x) \right], \\ \text{s.t. } \hat{A}(x, y_i) &= \tilde{r} = \frac{r(x, y_i) - \text{mean}(r(x))}{\text{std}(r(x))}. \end{aligned} \quad (1)$$

This objective encourages the model to assign higher probabilities to high-reward responses while down-weighting less favorable ones.

Reward Composition in GRPO. GRPO typically uses a reward function composed of multiple task-specific components. In AD, two main factors are considered:

$$r(x, y) = r_{\text{acc}}(x, y) + r_{\text{format}}(x, y), \quad (2)$$

where r_{acc} measures response correctness against ground truth, and r_{format} assesses compliance with expected output structures (e.g., $\langle \text{think} \rangle \langle \text{answer} \rangle$).

GRPO with KL Regularization. To stabilize learning, GRPO augments the advantage-weighted log-likelihood with a KL divergence term between the current and reference policies:

$$\begin{aligned} \mathcal{L}_{\text{rew}} &= -\mathbb{E}_x \frac{1}{G} \sum_{i=1}^G \frac{1}{|y_i|} \sum_{t=1}^{|y_i|} \left[\hat{A}(x, y_i) \log \pi_\theta(y_{i,t} | x) \right], \\ \mathcal{L}_{\text{reg}} &= \mathbb{D}_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}), \quad \mathcal{L}(\theta) = \mathcal{L}_{\text{rew}} + \beta \cdot \mathcal{L}_{\text{reg}} \end{aligned} \quad (3)$$

where π_{ref} is a fixed baseline policy, e.g. the base model before finetuning, and β balances exploration and stability. The KL term prevents overly aggressive policy shifts.

Discussions in Anomaly Detection Task. The effectiveness of GRPO relies on two factors: the MLLM’s ability to generate diverse, high-quality responses, and an effective scoring mechanism. The diversity in generated responses refers to the variance of the rewards assigned to each response. However, in anomaly detection, the task’s simplicity, i.e. binary classification with only two possible outcomes, limits response diversity. To address this, we propose a Multi-Stage Reasoning architecture, which compels the MLLM to follow a structured reasoning process before generating an answer, thus avoiding making hasty, superficial answers and introducing more fine-grained reward signals. Furthermore, conventional, coarse-grained rewards, e.g., binary accuracy, can cause the learning signal to vanish when all responses share the same final answer, i.e. $\mathcal{L}_{\text{rew}} = 0$. This not only limits sample utilization but also prevents distinguishing between genuinely and spuriously correct responses. To counter this, we analyze the reward variance and propose a Fine-Grained Reward method, which assesses both decision accuracy and localization quality within the reasoning process.

Multi-Stage Reasoning

To enhance the model’s reasoning capability and mitigate hasty decision-making in complex anomaly scenarios, we propose a structured, two-stage reasoning framework inspired by the test-time scaling paradigm and chain-of-thought prompting (Madaan et al. 2023; Kojima et al. 2022; Wei et al. 2022). This approach shifts from single-pass inference to a more deliberate, resource-adaptive process. Recent progress in test-time scaling (Madaan et al. 2023) shows that allocating additional computation at inference can significantly improve performance in challenging cases. Following these insights, our framework is designed as:

- **Stage 1: Initial Spatial Analysis (Think).** The model conducts a rapid, coarse-grained scan to localize potentially anomalous regions, generating region proposals for further scrutiny. This step also supplies the localization cues used by our reward mechanism.
- **Stage 2: Focused Region Examination (Rethink).** For each proposed region, the model allocates more reasoning capacity, enabling fine-grained analysis of suspicious areas to better distinguish true anomalies from benign variations.
- **Stage 3: Integrated Decision Making (Answer).** The model combines global and local information to make a final decision.

This procedure aligns with the cognitive strategies employed by human experts in visual inspection tasks, 1) performing rapid global scanning to identify candidate regions of interest, 2) engaging in deliberate, focused analysis to distinguish genuine anomalies from normal variations.

Additionally, from the GRPO training perspective, this deliberative, multi-stage reasoning framework provides a richer set of supervision signals, which facilitates more fine-grained reward assignment (as detailed in the next section) and offers a crucial mechanism for identifying responses that are spuriously correct. Furthermore, this framework enhances the diversity of the response space, thereby enriching the samples for more effective policy gradient estimation.

Fine-Grained Reward

To address the limitations of binary rewards in GRPO-based anomaly detection, we propose a novel fine-grained reward mechanism. This mechanism fundamentally shifts the reward signal from a discrete, binary value to a more continuous, multi-task score by assessing both final decision accuracy and the explicit localization of anomalous regions.




Figure 2: Conceptual illustration of loss landscape. (a) Our proposed fine-grained reward; (b) existing discrete reward.

As shown in Fig. 2, traditional binary reward schemes (b) create flat, discontinuous optimization surfaces where all correct responses receive identical rewards, resulting in zero gradients i.e. $\nabla_{\theta} \mathcal{L}_{rew} = 0$ and model degradation toward the reference policy θ_{ref} . In contrast, our localization-aware rewards (a) generate a rich, continuous landscape with meaningful gradients that provide strong directional signals: positive reinforcement for high-value responses y_1, y_2 with better spatial reasoning, and negative feedback for spurious correct responses y_3 that achieve classification accuracy through flawed inference.

To this end, two variants of fine-grained rewards oriented towards localization are introduced. Specifically, given a generated response y , we first extract the predicted bounding boxes $\mathcal{B}(y) = \{\hat{b}_i\}_{i=1}^m$ of anomalous regions from the response. The corresponding ground-truth bounding boxes in this sample denotes $\mathcal{B}(x) = \{b_j\}_{j=1}^n$. To assess the quality of the predicted anomalous regions, we use the average normalized GIoU over the matched box pairs, together with a count-based metric, as fine-grained reward signals. We define a cost matrix $C_{ij} = 1 - \text{GIoU}(\hat{b}_i, b_j)$ and apply the Hungarian algorithm to find optimal matches $\mathcal{M}(\mathcal{B}(y), \mathcal{B}(x)) \subseteq \{(i, j)\}_{i=1\dots m, j=1\dots n}$. The classification ground truth of this sample is indicated as $l(x) \in \{0, 1\}$ where 0 denotes a normal sample while 1 denotes an abnormal sample. We compute the following quantities,

$$r_{\text{loc}}(x, y) = \begin{cases} \frac{1}{|\mathcal{M}|} \sum_{(i,j) \in \mathcal{M}} \text{GIoU}(\hat{b}_i, b_j) + \alpha \cdot r_{\text{count}}, & l(x) = 1, \\ \beta \cdot r_{\text{focus}}(m), & \text{s.t. } m = |\mathcal{B}(y)|, \quad l(x) = 0, \end{cases} \quad (4)$$

with counting reward:

$$r_{\text{count}}(x, y) = \begin{cases} 1, & |m - n| = 0 \\ 0.5, & |m - n| = 1, \text{ s.t. } \begin{matrix} m = |\mathcal{B}(y)| \\ n = |\mathcal{B}(x)| \end{matrix} \\ 0, & |m - n| \geq 2 \end{cases} \quad (5)$$

For normal samples, i.e. $l = 0$, where ground-truth boxes are absent, we still encourage the model to engage in focused analysis by incentivizing bounding box predictions via:

$$r_{\text{focus}}(m) = \begin{cases} 0, & m = 0 \\ 0.5, & m = 1 \\ -0.1, & m \geq 2 \end{cases} \quad (6)$$

This gently guides the model toward identifying a single region of interest, akin to human attention. To ensure basic decision correctness, we include a classification reward, $r_{\text{cls}}(x, y) = \mathbb{1}(l(y) = l(x))$, where $l(y)$ indicates the final classification answer. The overall accuracy reward integrates both localization and classification signals:

$$r_{\text{acc}}(x, y) = \omega \cdot r_{\text{loc}}(x, y) + r_{\text{cls}}(x, y). \quad (7)$$

By providing this richer and more differentiated fine-grained reward signal, our method can effectively improve sample utilization and mitigate the model degradation situation while facilitate more accurate visual understanding.

Experiments

Experiment Settings

Datasets & Protocols. To validate the adaptability of our method to diverse anomaly downstream tasks, we first follow a **multi-domain evaluation protocol** using four datasets that cover distinct notions of anomalies: MVTEC (Bergmann et al. 2019), VisA (Zou et al. 2022), MVTEC-LOCO (Bergmann et al. 2022), and GoodSAD (Zhang et al. 2024). This experimental design follows the benchmark protocol introduced in MMAD (Jiang et al. 2025) and specifically evaluates the model’s ability to learn unified representations across heterogeneous anomaly types. For training, we use 20% of the data, uniformly sampled across object categories, and reserve the remaining 80% for in-distribution evaluation.

To further assess the practical utility of our approach in real-world industrial deployment, we evaluate on **cross-dataset generalization** setting. In this setting, the model is trained exclusively on 20% of the MVTEC dataset and tested on several unseen industrial anomaly detection benchmarks: VisA, MPDD (Jezek et al. 2021), DTD (Cimpoi et al. 2014), and DAGM (Wieler and Hahn 2007). Each evaluation is conducted on the standard test split defined by the respective dataset. This setting is designed to simulate a common real-world constraint, where labeled training data is available only in a limited source domain, and models must generalize to novel types of anomalies in unseen target domains without retraining.

Competing Methods. The open-source models include LLaVA-Next (Liu et al. 2024), InternVL2 (Chen et al. 2024), and InternVL3 (Zhu et al. 2025), while the proprietary baselines comprise GPT-4o, GPT-4o-mini, and GPT-4.1. In the **multi-domain evaluation** experiment, we additionally benchmark our approach against two concurrent MLLM fine-tuning methods: AnomalyR1 (Chao et al. 2025) and OmniAD (Zhao et al. 2025). For the **cross-dataset generalization** task, we include comparisons with state-of-the-art CLIP-based approaches (Jeong et al. 2023; Zhou et al. 2023). As CLIP-based methods naturally produce continuous anomaly scores, we convert these into binary decisions using Youden’s J statistic (Ruopp et al. 2008) to optimize the threshold τ . Details of this method are provided in the supplementary. Additionally, we evaluate CLIP-based methods using a fixed threshold of $\tau = 0.5$, representing a baseline approach that requires no dataset-specific optimization.

Implementation Details. All experiments were conducted on two NVIDIA B200 GPUs. During GRPO optimization, we generate six parallel responses per training sample, and train the model for 15 epochs. We adopt Low-Rank Adaptation (LoRA) (Hu et al. 2022) to fine-tune the Qwen2.5-VL (7B) model (Qwen et al. 2025), with LoRA hyperparameters rank $r = 64$ and alpha $\alpha = 128$. Due to hardware compatibility issues with FlashAttention on the Blackwell architecture, we instead employ standard attention mechanisms for all models. Importantly, only the LoRA parameters are updated during training, while the pre-trained backbone weights are frozen.

Bounding Box Generation. Since the original datasets provide only pixel-level segmentation masks, we construct pseudo bounding boxes for localization supervision. Details are provided in the supplementary material.

		Scale	MVTec	VisA	LOCO	GoodsAD	Avg.
Proprietary	GPT-4o-mini [†]	-	-	-	-	-	64.33
	GPT-4o [†]	-	-	-	-	-	68.63
	GPT-4.1	-	78.83	80.99	69.35	57.21	71.60
Open-source	LLaVA-NeXT	34B	72.44	55.91	63.28	56.06	61.92
	InternVL2	76B	79.42	69.78	54.02	53.20	64.11
	InternVL3	8B	78.53	65.02	59.77	50.20	63.38
	InternVL3	38B	78.12	68.75	63.86	55.49	66.56
	QwenVL2.5	7B	62.03	55.50	45.90	40.73	51.04
Fine-tuned	AnomalyR1 [‡]	3B	-	-	-	-	60.62
	OmniAD [‡]	7B	-	-	-	-	68.80
	AD-FM (Ours)	7B	90.72	79.52	65.62	56.74	73.15

Table 1: Binary classification accuracy (%) for multi-domain anomaly detection. LOCO denotes MVTEC-LOCO. Scale is measured in billions of parameters. [†]Results are from (Jiang et al. 2025). [‡]Results are from respective publications.

Multi-Domain Anomaly Detection

We first present the multi-domain anomaly detection results in Tab. 1. Compared to the original Qwen model, our fine-tuned approach improves performance by 22.11%. Our method also outperforms all evaluated large language models, including the latest commercial model GPT-4.1.

Against concurrent fine-tuning methods AnomalyR1 and OmniAD, our approach shows consistent superiority. Since the official implementations and detailed per-dataset results of both methods are unavailable, we directly report their average results as presented in their original papers. Our average performance exceeds AnomalyR1 by 12.53% and OmniAD by 4.35%. These improvements indicate that our design achieves better visual understanding within the GRPO fine-tuning paradigm.

We further conduct comprehensive evaluation across multiple anomaly analysis subtasks (Jiang et al. 2025) to investigate the ability for anomaly explanation. As demonstrated in Tab. 3, our approach consistently outperforms all baseline methods across the evaluated task categories. Detailed experimental protocols of these multi-task evaluations are provided in the supplementary materials.

Cross-Dataset Generalization in Industrial Anomaly Detection

The cross-dataset generalization results on industrial anomaly detection tasks are presented in Tab. 2, where all reported accuracies are obtained using our model trained exclusively on the MVTEC dataset. Without any additional training on the target datasets, our method achieves a notable average improvement of 8.5% over the base Qwen2.5-VL (7B) model. It also consistently outperforms all evaluated zero-shot MLLM baselines. Compared to CLIP-based methods, our approach surpasses them on three out of four datasets (MPDD, DTD, and DAGM), while avoiding the critical reliance on threshold tuning.

		Scale	VisA	MPDD	DTD	DAGM	Avg.
CLIP-based	WinClip ($\tau = 0.5$)	-	43.40	72.10	25.90	86.90	57.08
	WinClip*	-	77.70	41.80	90.00	88.90	75.85
	AnomalyClip ($\tau = 0.5$)	-	56.56	56.49	81.71	22.19	54.24
	AnomalyClip*	-	82.17	68.23	88.91	95.36	83.67
Proprietary	GPT4o-mini	-	73.98	67.90	90.56	90.80	80.81
	GPT4o	-	73.77	69.43	90.80	90.17	81.04
	GPT-4.1	-	80.99	66.81	90.18	89.93	81.98
Open-source	LLaVA-NeXT	34B	55.91	61.57	72.62	23.15	53.70
	InternVL2	76B	69.78	66.38	90.58	94.52	80.32
	InternVL3	8B	65.02	70.74	85.05	81.33	75.54
	InternVL3	38B	68.75	69.43	90.79	89.87	79.71
	Qwen2.5-VL	7B	55.50	68.34	85.59	94.83	76.07
Fine-tuned	AD-FM (Ours)	7B	77.43	72.71	92.64	95.51	84.57

Table 2: Cross-dataset generalization performance in industrial anomaly detection. Binary classification accuracy (%) with single normal reference protocol. *denotes threshold optimized using Youden’s J statistic.

On the VisA dataset, where our model shows a modest gap relative to the optimized CLIP baseline, the shortfall primarily arises from the dataset’s subtle defect patterns, which pose significant challenges for current MLLMs. Nonetheless, our method offers substantial practical benefits. CLIP-based methods are highly sensitive to threshold settings, e.g. AnomalyCLIP’s performance drops from 83.67% (optimized) to 54.24% (fixed threshold), and WinCLIP from 75.85% to 57.08%. In contrast, our method provides direct binary predictions without requiring threshold calibration, which is an essential advantage for real-world deployments where full data distributions are not available for tuning.

Qualitative Results

Multi-stage Deliberative Reasoning. To demonstrate our multi-stage deliberative reasoning framework, we present representative examples in Fig. 3 that highlight different aspects of our approach’s reasoning capabilities.

The first two examples demonstrate successful detection of obvious anomalies and accurate classification. In case (a), the model correctly identifies surface irregularities on wood in *think* stage, noting deviations from natural grain patterns with appropriate spatial localization. Case (b) illustrates how our framework cultivates systematic analytical habits even for normal samples while no ground truth bounding boxes are available for normal cases.

Cases (c) and (d) illustrate the critical value of our multi-stage approach in handling initially ambiguous scenarios. In case (c), during *think* stage, the model exhibits uncertainty, noting that “*the slight variation in the surface pattern might be considered within normal variation, and there are no other signs of defects.*”. However, the *rethink* stage enables deeper examination, ultimately leading to correct normal decision. Similarly, case (d) progresses from identifying “*a small irregularity in the grid structure,*” to determining through “*closer inspection*” that the deviation constitutes “*a complete break or significant deformation.*”

This progression from uncertainty to confident demonstrates how our framework prevents premature conclusions and enables recovery from initial analytical uncertainty.

		Scale	Defect Class.	Defect Loc.	Defect Desc.	Defect Anal.	Object Class.	Object Anal.	Avg.
Proprietary	GPT-4o-mini [†]	-	48.58	38.75	63.68	80.40	88.56	79.74	66.62
	GPT-4o [†]	-	65.80	55.62	73.21	83.41	94.98	82.80	75.97
	GPT-4.1	-	65.93	62.64	76.18	84.23	90.95	85.57	77.58
Open-source	LLaVA-NeXT [†]	34B	48.79	52.87	71.34	80.28	81.12	77.80	68.70
	InternVL2 [†]	76B	54.22	56.66	66.30	80.47	86.40	82.92	71.16
	InternVL3	8B	46.35	46.24	66.01	78.28	79.17	84.83	66.81
	InternVL3	38B	58.99	59.63	73.06	83.88	91.62	88.50	75.94
	Qwen2.5-VL	7B	41.07	44.87	61.43	75.99	86.54	79.58	64.91
Fine-tuned	AnomalyR1 [‡]	3B	63.56	70.14	80.47	85.28	92.48	86.15	79.68
	OmniAD [‡]	7B	78.80	75.50	67.20	86.40	96.00	86.40	81.72
	AD-FM (Ours)	7B	73.36	77.53	86.80	86.75	89.98	86.67	83.56

Table 3: Multi-task anomaly analysis performance comparison across six evaluation categories. Results reported as average accuracy (%) across four datasets (MVTec, VisA, MVTec-LOCO, GoodsAD) for defect-related tasks and object-level understanding tasks. [†]Results are from (Jiang et al. 2025). [‡]Results are from respective publications.

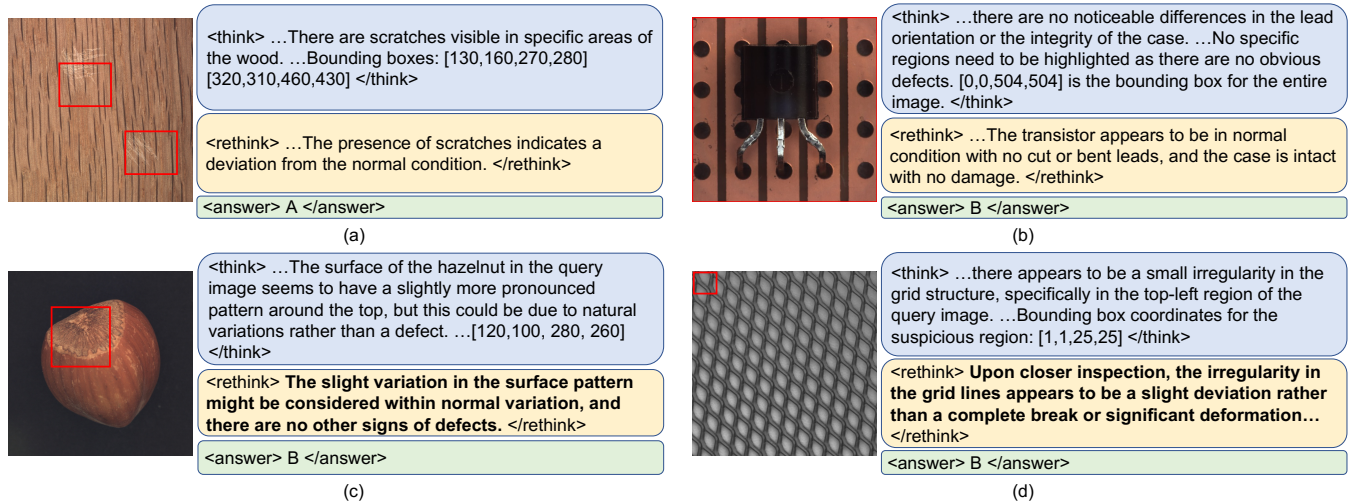


Figure 3: Qualitative examples demonstrating the effectiveness of our multi-stage deliberative reasoning framework.

Response Variance Visualization. To provide deeper insight into how our fine-grained rewards differentiate responses with greater sophistication, we visualize the responses generated by our framework on a representative sample. Fig. 4 presents four different responses (4 responses chosen from all 6 due to space limitations) generated by our model for the same input image containing wood surface defects, alongside the corresponding reward variance analysis.

As seen from Fig. 4 (left), while all responses (a-d) correctly identify anomaly presence, there is a clear progression in analytical quality from (a) to (d). Specifically, responses (a-b) demonstrate incorrect defect enumeration and imprecise localization, response (c) shows improved awareness but lacks precision, and response (d) achieves accurate counting with precise bounding box alignment.

Further examining the reward value plot in Fig. 4 (right), under binary reward schemes, i.e., r_{cls} , all four responses receive identical scores since they all correctly identify anomaly presence, resulting in zero variance and no learning signal for GRPO optimization. In contrast, our fine-grained reward mechanisms inject meaningful variance into

the training signal. The count-based enhancement, i.e., $r_{cls} + r_{count}$, generates moderate variance while failing to differentiate between response (c) and (d), while our localization-aware formulation, i.e., $r_{cls} + r_{loc}$, provides feedback that incentivizes the most accurate response (d) over the progressively less precise responses. This granular differentiation enables GRPO to learn from the full spectrum of analytical quality rather than treating all correct detections as equivalent. This further improves the interpretability and reliability of the reasoning process, contributing to safer and more transparent model behavior.

Ablation Studies

Component Analysis. To validate the effectiveness of our GRPO optimization components, we conduct comprehensive ablation studies across the four multi-domain dataset, systematically examining the contributions of different reward mechanisms and reasoning protocols. The results are demonstrated in Tab. 4. From the results, we make the following key observations:

- The baseline configuration without any training achieves

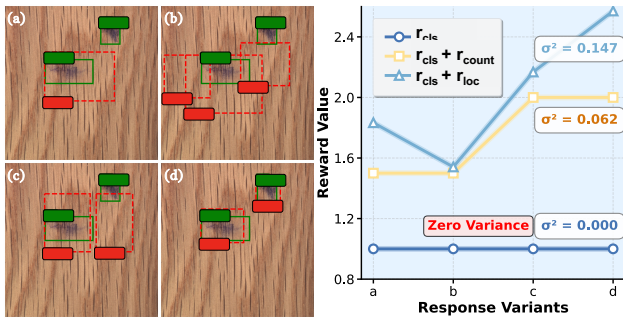


Figure 4: Response variance visualization for the same training sample under different reward mechanisms. Left: Four model responses (a-d) showing varying quality in defect localization and spatial reasoning. Right: Reward variance analysis under different reward schemes.

Reward r_{acc}	Reasoning	MVTec	VisA	LOCO	GoodsAD	Avg.
-	-	62.03	55.50	45.90	40.73	51.04
r_{cls}	one-stage	84.64	73.79	60.08	53.77	68.07
r_{cls}	multi-stage	86.67	75.29	61.42	55.91	69.82
$r_{cls} + r_{random}$	multi-stage	88.41	75.44	60.60	56.44	70.22
$r_{cls} + r_{loc}$	multi-stage	90.72	79.52	65.62	56.74	73.15

Table 4: Ablation study on reward mechanisms and reasoning processes on multi-domain AD experiments. Performance measured as binary classification accuracy (%).

an average accuracy of 51.04%, representing the inherent zero-shot capability of the base Qwen2.5-VL model. The introduction of **binary reward signals with one-stage reasoning** boosts performance to 68.07% (+17.03%), demonstrating effective domain-specific fine-tuning.

- Extending the binary reward framework to incorporate **multi-stage reasoning** processes further enhances performance to 69.82% (+1.75% over one-stage), demonstrating that structured multi-step inference protocols enable more sophisticated anomaly analysis capabilities.
- To isolate the contribution of variance enhancement from structured localization supervision, we evaluate a **random reward** component r_{random} that introduces artificial variance without spatial meaning. This approach builds upon Random Reward Perturbation (RRP) (Ma et al. 2025), where zero-mean Gaussian noise added to reward signals has been shown to enhance exploration through increased policy diversity. This configuration achieves 70.22% with multi-stage reasoning, providing modest improvements (+0.40% over structured reasoning alone), confirming that variance injection per se contributes positively but limited to GRPO effectiveness.
- In contrast, our **fine-grained reward** formulation achieves 73.15% (+3.33% over binary multi-stage), representing a cumulative improvement of +22.11% over the zero-shot baseline. This performance gap between random and ours variance injection validates that our designed rewards deliver dual benefits: variance enhancement that provides essential signal diversity, and fine-

grained quality assessment that selectively reinforces responses demonstrating genuine understanding while filtering out spurious correctness patterns.

Response Variance Analysis. To empirically validate our hypothesis that continuous reward mechanisms enhance GRPO training effectiveness by increasing response variance, we conduct an ablation analysis on MVTEC dataset. We uniformly sample 500 images from MVTEC dataset and generate 6 responses per sample using Qwen2.5-VL model under different reward configurations. For each sample, we compute the variance in reward values across 6 generated responses. We then measure the percentage of samples with zero variance, indicating cases where all responses receive identical rewards and thus contribute no meaningful gradient signal to GRPO optimization. Fig. 5 presents the zero-variance percentages across different reward mechanisms:




Figure 5: Percentage of samples with zero reward variance across 6 generated responses under different reward formulations. Lower percentages indicate enhanced data utilization for GRPO training.

The results demonstrate a significant improvement in data utilization through our reward design. The binary-only baseline yields 41.25% zero-variance samples, which empirically confirms the challenge we identified: traditional reward schemes create sparse learning environments that severely limit gradient flow. In contrast, adding bounding box counting reduces this to 15.25%, and our complete formulation with localization lowers it further to 11.00%, which is a 73% reduction compared to the binary baseline. These results empirically validate our theoretical motivation: continuous reward mechanisms transform GRPO from a sparse-signal optimization into a rich-feedback learning environment.

Conclusion

In this paper, we propose AD-FM, a fine-tuned multi-modal MLLM framework tailored for anomaly detection, addressing two core limitations of existing GRPO-based applications: response variance collapse and under-constrained reasoning procedure. Our multi-stage deliberative reasoning creates supervisory signals and promotes cautious inference, while our fine-grained, spatially grounded reward design introduces dense and discriminative learning signals that suppress spurious correctness. These components jointly enhance model interpretability, training efficiency, and generalization. Extensive results across multi-domain and cross-dataset evaluations demonstrate the effectiveness of our approach in adapting general-purpose MLLMs to anomaly detection tasks.

Acknowledgments

This research is supported by Agency for Science, Technology and Research (A*STAR) under its MTC Programmatic Funds (Grant No. M23L7b0021); and the National Research Foundation, Singapore, and the CyberSG R&D Programme Office (“CRPO”), under the National Cybersecurity R&D Programme (“NCRP”), RIE2025 NCRP Funding Initiative (Award CRPO-GC1-NTU-002).

References

- Bergmann, P.; Batzner, K.; Fauser, M.; Sattlegger, D.; and Steger, C. 2022. Beyond dents and scratches: Logical constraints in unsupervised anomaly detection and localization. *International Journal of Computer Vision*, 130(4): 947–969.
- Bergmann, P.; Fauser, M.; Sattlegger, D.; and Steger, C. 2019. MVTEC AD–A comprehensive real-world dataset for unsupervised anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9592–9600.
- Cao, Y.; Xu, X.; Cheng, Y.; Sun, C.; Du, Z.; Gao, L.; and Shen, W. 2025. Personalizing Vision-Language Models With Hybrid Prompts for Zero-Shot Anomaly Detection. *IEEE Transactions on Cybernetics*, 55(4): 1917–1929.
- Cao, Y.; Xu, X.; Sun, C.; Huang, X.; and Shen, W. 2023. Towards Generic Anomaly Detection and Understanding: Large-scale Visual-linguistic Model (GPT-4V) Takes the Lead. arXiv:2311.02782.
- Chao, Y.; Liu, J.; Tang, J.; and Wu, G. 2025. AnomalyR1: A GRPO-based End-to-end MLLM for Industrial Anomaly Detection. arXiv:2504.11914.
- Chen, L.; Li, L.; Zhao, H.; Song, Y.; and Vinci. 2025a. R1-V: Reinforcing Super Generalization Ability in Vision-Language Models with Less Than \$3. <https://github.com/Deep-Agent/R1-V>. Accessed: 2025-02-02.
- Chen, Z.; Chen, H.; Imani, M.; and Imani, F. 2025b. Can multimodal large language models be guided to improve industrial anomaly detection? In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 89213, V02BT02A051. American Society of Mechanical Engineers.
- Chen, Z.; Wang, W.; Tian, H.; Ye, S.; Gao, Z.; Cui, E.; Tong, W.; Hu, K.; Luo, J.; Ma, Z.; et al. 2024. How Far Are We to GPT-4V? Closing the Gap to Commercial Multimodal Models with Open-Source Suites. *arXiv preprint arXiv:2404.16821*.
- Cimpoi, M.; Maji, S.; Kokkinos, I.; Mohamed, S.; ; and Vedaldi, A. 2014. Describing Textures in the Wild. In *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Deng, H.; and Li, X. 2022. Anomaly detection via reverse distillation from one-class embedding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9737–9746.
- Deng, H.; Luo, H.; Zhai, W.; Guo, Y.; Cao, Y.; and Kang, Y. 2025. Vmad: Visual-enhanced multimodal large language model for zero-shot anomaly detection. *IEEE Transactions on Automation Science and Engineering*.
- Fan, Y.; He, X.; Yang, D.; Zheng, K.; Kuo, C.-C.; Zheng, Y.; Narayanaraju, S. J.; Guan, X.; and Wang, X. E. 2025. GRIT: Teaching MLLMs to Think with Images. arXiv:2505.15879.
- Gu, Z.; Zhu, B.; Zhu, G.; Chen, Y.; Tang, M.; and Wang, J. 2024. Anomalygpt: Detecting industrial anomalies using large vision-language models. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 1932–1940.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- He, H.; Bai, Y.; Zhang, J.; He, Q.; Chen, H.; Gan, Z.; Wang, C.; Li, X.; Tian, G.; and Xie, L. 2024. Mambaad: Exploring state space models for multi-class unsupervised anomaly detection. *Advances in Neural Information Processing Systems*, 37: 71162–71187.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Hurst, A.; Lerer, A.; Goucher, A. P.; Perelman, A.; Ramesh, A.; Clark, A.; Ostrow, A.; Welihinda, A.; Hayes, A.; Radford, A.; et al. 2024. GPT-4o System Card. *arXiv preprint arXiv:2410.21276*.
- Jaech, A.; Kalai, A.; Lerer, A.; Richardson, A.; El-Kishky, A.; Low, A.; Helyar, A.; Madry, A.; Beutel, A.; Carney, A.; et al. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- Jeong, J.; Zou, Y.; Kim, T.; Zhang, D.; Ravichandran, A.; and Dabeer, O. 2023. Winclip: Zero-/few-shot anomaly classification and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19606–19616.
- Jezeq, S.; Jonak, M.; Burget, R.; Dvorak, P.; and Skotak, M. 2021. Deep learning-based defect detection of metal parts: evaluating current methods in complex conditions. In *2021 13th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*, 66–71.
- Jiang, X.; Li, J.; Deng, H.; Liu, Y.; Gao, B.-B.; Zhou, Y.; Li, J.; Wang, C.; and Zheng, F. 2025. MMAD: The Comprehensive Benchmark for Multimodal Large Language Models in Industrial Anomaly Detection. In *The Thirteenth International Conference on Learning Representations*.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4015–4026.
- Kojima, T.; Gu, S. S.; Reid, M.; Matsuo, Y.; and Iwasawa, Y. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*.
- Lei, J.; Hu, X.; Wang, Y.; and Liu, D. 2023. Pyramid-Flow: High-Resolution Defect Contrastive Localization Using Pyramid Normalizing Flow. In *Proceedings of the*

- IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 14143–14152.
- Li, X.; Li, Z.; Kosuga, Y.; and Bian, V. 2025a. Optimizing Safe and Aligned Language Generation: A Multi-Objective GRPO Approach. *arXiv preprint arXiv:2503.21819*.
- Li, Y.; Liu, F.; Liao, J.; Tian, S.; Foo, C.-S.; and Yang, X. 2025b. FIND: Few-Shot Anomaly Inspection with Normal-Only Multi-Modal Data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 23290–23299.
- Li, Y.; Wang, H.; Yuan, S.; Liu, M.; Zhao, D.; Guo, Y.; Xu, C.; Shi, G.; and Zuo, W. 2025c. Myriad: Large Multimodal Model by Applying Vision Experts for Industrial Anomaly Detection. *arXiv:2310.19070*.
- Liao, J.; Xu, X.; Nguyen, M. C.; Goodge, A.; and Foo, C. S. 2024. Coft-ad: Contrastive fine-tuning for few-shot anomaly detection. *IEEE Transactions on Image Processing*, 33: 2090–2103.
- Liu, H.; Li, C.; Li, Y.; Li, B.; Zhang, Y.; Shen, S.; and Lee, Y. J. 2024. LLaVA-NeXT: Improved reasoning, OCR, and world knowledge.
- Liu, Z.; Sun, Z.; Zang, Y.; Dong, X.; Cao, Y.; Duan, H.; Lin, D.; and Wang, J. 2025. Visual-RFT: Visual Reinforcement Fine-Tuning. *arXiv preprint arXiv:2503.01785*.
- Liu, Z.; Zhou, Y.; Xu, Y.; and Wang, Z. 2023. SimpleNet: A Simple Network for Image Anomaly Detection and Localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20402–20411.
- Ma, H.; Fu, G.; Luo, Z.; Wu, J.; and Leong, T.-Y. 2025. Exploration by Random Reward Perturbation. *arXiv preprint arXiv:2506.08737*.
- Madaan, A.; Tandon, N.; Gupta, P.; Hallinan, S.; Gao, L.; Wiegrefe, S.; Alon, U.; Dziri, N.; Prabhumoye, S.; Yang, Y.; et al. 2023. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*.
- OpenAI. 2023. ChatGPT. <https://chat.openai.com/>.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*.
- Qwen; ; Yang, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Li, C.; Liu, D.; Huang, F.; Wei, H.; Lin, H.; Yang, J.; Tu, J.; Zhang, J.; Yang, J.; Yang, J.; Zhou, J.; Lin, J.; Dang, K.; Lu, K.; Bao, K.; Yang, K.; Yu, L.; Li, M.; Xue, M.; Zhang, P.; Zhu, Q.; Men, R.; Lin, R.; Li, T.; Tang, T.; Xia, T.; Ren, X.; Ren, X.; Fan, Y.; Su, Y.; Zhang, Y.; Wan, Y.; Liu, Y.; Cui, Z.; Zhang, Z.; and Qiu, Z. 2025. Qwen2.5 Technical Report. *arXiv:2412.15115*.
- Roth, K.; Pemula, L.; Zepeda, J.; Schölkopf, B.; Brox, T.; and Gehler, P. 2022. Towards total recall in industrial anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14318–14328.
- Ruopp, M. D.; Perkins, N. J.; Whitcomb, B. W.; and Schisterman, E. F. 2008. Youden Index and optimal cut-point estimated from observations affected by a lower limit of detection. *Biometrical Journal: Journal of Mathematical Methods in Biosciences*, 50(3): 419–430.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Tu, R.-C.; Jin, Z.; Liao, J.; Luo, X.; Wang, Y.; Shen, L.; and Tao, D. 2025. MLLM-Guided VLM Fine-Tuning with Joint Inference for Zero-Shot Composed Image Retrieval. *arXiv preprint arXiv:2505.19707*.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*.
- Wieler, M.; and Hahn, T. 2007. Weakly supervised learning for industrial optical inspection. In *DAGM symposium in*, volume 6, 11.
- Zhang, J.; Ding, R.; Ban, M.; and Dai, L. 2024. PKU-GoodsAD: A supermarket goods dataset for unsupervised anomaly detection and segmentation. *IEEE Robotics and Automation Letters*, 9(3): 2008–2015.
- Zhang, J.; He, H.; Gan, Z.; He, Q.; Cai, Y.; Xue, Z.; Wang, Y.; Wang, C.; Xie, L.; and Liu, Y. 2025. A comprehensive library for benchmarking multi-class visual anomaly detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3524–3533.
- Zhang, X.; Xu, M.; and Zhou, X. 2024. Realnet: A feature selection network with realistic synthetic anomaly for anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16699–16708.
- Zhao, S.; Lin, Y.; Han, L.; Zhao, Y.; and Wei, Y. 2025. OmniAD: Detect and Understand Industrial Anomaly via Multimodal Reasoning. *arXiv:2505.22039*.
- Zhou, Q.; Pang, G.; Tian, Y.; He, S.; and Chen, J. 2023. AnomalyCLIP: Object-agnostic Prompt Learning for Zero-shot Anomaly Detection. In *The Twelfth International Conference on Learning Representations*.
- Zhu, J.; Wang, W.; Chen, Z.; Liu, Z.; Ye, S.; Gu, L.; Tian, H.; Duan, Y.; Su, W.; Shao, J.; Gao, Z.; Cui, E.; Wang, X.; Cao, Y.; Liu, Y.; Wei, X.; Zhang, H.; Wang, H.; Xu, W.; Li, H.; Wang, J.; Deng, N.; Li, S.; He, Y.; Jiang, T.; Luo, J.; Wang, Y.; He, C.; Shi, B.; Zhang, X.; Shao, W.; He, J.; Xiong, Y.; Qu, W.; Sun, P.; Jiao, P.; Lv, H.; Wu, L.; Zhang, K.; Deng, H.; Ge, J.; Chen, K.; Wang, L.; Dou, M.; Lu, L.; Zhu, X.; Lu, T.; Lin, D.; Qiao, Y.; Dai, J.; and Wang, W. 2025. InternVL3: Exploring Advanced Training and Test-Time Recipes for Open-Source Multimodal Models. *arXiv:2504.10479*.
- Zou, Y.; Jeong, J.; Pemula, L.; Zhang, D.; and Dabeer, O. 2022. Spot-the-difference self-supervised pre-training for anomaly detection and segmentation. In *European Conference on Computer Vision*, 392–408. Springer.