

MAU-GPT: Enhancing Multi-type Industrial Anomaly Understanding via Anomaly-aware and Generalist Experts Adaptation

Zhuonan Wang^{1*}, Zhenxuan Fan^{1*}, Siwen Tan^{1*}, Yu Zhong¹, Yuqian Yuan¹, Haoyuan Li^{1,2}, Hao Jiang², Wenqiao Zhang^{1†}, Feifei Shao^{1†}, Hongwei Wang¹, Jun Xiao¹

¹Zhejiang University

²Alibaba Group

{wnzz, zxfan, siwentan, yuzhong, yuqianyuan, lihaoyuan, wenqiaozhang, sff}@zju.edu.cn
aoshu.jh@alibaba-inc.com, hongweiwang@intl.zju.edu.cn, junx@cs.zju.edu.cn

Abstract

As industrial manufacturing scales, automating fine-grained product image analysis has become critical for quality control. However, existing approaches are hindered by limited dataset coverage and poor model generalization across diverse and complex anomaly patterns. To address these challenges, we introduce **MAU-Set**, a comprehensive dataset for **Multi-type industrial Anomaly Understanding**. It spans multiple industrial domains and features a hierarchical task structure, ranging from binary classification to complex reasoning. Alongside this dataset, we establish a rigorous evaluation protocol to facilitate fair and comprehensive model assessment. Building upon this foundation, we further present **MAU-GPT**, a domain-adapted multimodal large model specifically designed for industrial anomaly understanding. It incorporates a novel AMoE-LoRA mechanism that unifies anomaly-aware and generalist experts adaptation, enhancing both understanding and reasoning across diverse defect classes. Extensive experiments show that MAU-GPT consistently outperforms prior state-of-the-art methods across all domains, demonstrating strong potential for scalable and automated industrial inspection.

Introduction

Industrial anomaly analysis plays a vital role in product quality assessment, process monitoring, and defect prevention (Cao et al. 2024; Ren et al. 2022; Bhatt et al. 2021). In practice, quality inspectors conduct detailed quantitative evaluations and generate anomaly reports based on product images, relying on substantial domain expertise to identify subtle or complex defects (Tang et al. 2022; Jiang et al. 2024; Baitieva et al. 2024). However, as production scales, manual inspection becomes increasingly impractical for delivering timely and consistent analysis (Alzarooni et al. 2025), raising a critical question: *Can we automate the in-depth analysis of industrial product images?*

Automating this process requires models capable of both fine-grained visual understanding and the reasoning of pre-

cise, context-aware descriptions. This challenge naturally aligns with methodologies in Visual Question Answering (VQA), where models are designed to perform image-grounded reasoning by interpreting visual content in response to textual queries and producing semantically rich answers (Chen et al. 2021; Lu et al. 2022; Li et al. 2022).

However, directly transferring existing VQA models from general-purpose domains to industrial anomaly analysis poses significant challenges. These challenges include, but are not limited to: **1) Limited Dataset Coverage:** Existing public datasets are often narrowly scoped, focusing on specific anomaly types or product categories and failing to reflect the diversity of real-world industrial defects. For example, MVTEC AD (Bergmann et al. 2019) primarily targets surface-level visual defects, while domain-specific datasets such as DeepPCB (Tang et al. 2019) are restricted to anomalies in printed circuit boards. **2) Poor Model Generalization:** The limited coverage of such datasets constrains model exposure to diverse defect patterns, impairing their ability to retrieve relevant domain knowledge. As a result, existing approaches struggle to differentiate fine-grained anomaly types and fail to accurately reason about their root causes.

To tackle the aforementioned dataset limitation, we construct **MAU-Set**, a comprehensive dataset for **Multi-type industrial Anomaly Understanding**.

MAU-Set defines two Question Answer (QA) styles—**Discriminative QA** and **Open-Ended QA**—covering tasks from basic binary classification to complex reasoning, and further groups them into five distinct tasks with varying difficulty levels. This hierarchical task structure increases both the complexity and depth of the benchmark, promoting progressive knowledge accumulation and equipping models to handle the varied challenges encountered in real-world industrial applications. As illustrated in Figure 1 and Table 1, MAU-Set spans **35 product types and over 100 defect classes across 6 industrial domains**, ensuring both breadth and diversity. In addition, MAU-Set provides broad anomaly coverage, fine-grained supervision, and a unified testbed, laying a solid foundation for the development and evaluation of scalable, automated, and trustworthy industrial anomaly analysis systems.

*These authors contributed equally.

†Corresponding authors.

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

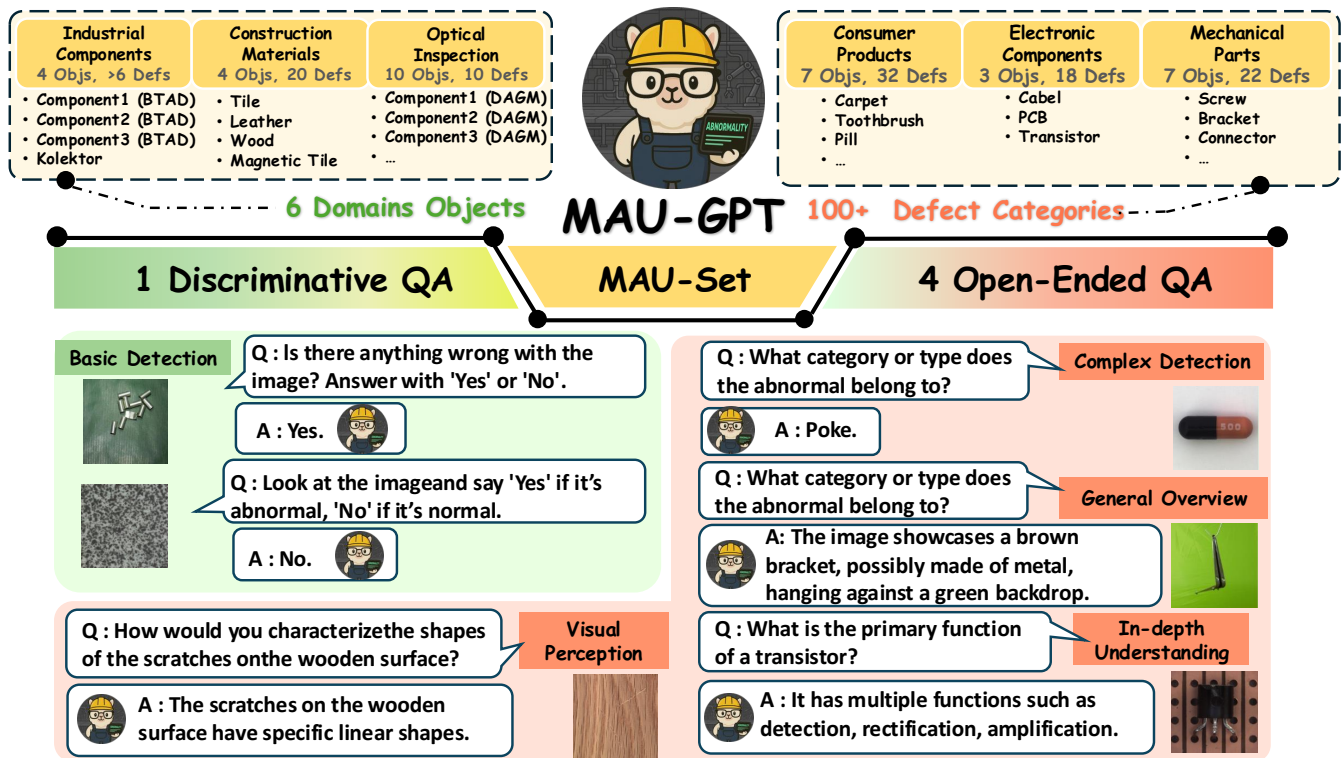


Figure 1: Overview of the MAU-GPT model and MAU-Set dataset. The dataset spans 35 product types and over 100 defect categories across 6 major industrial domains, supporting 5 tasks ranging from discriminative question answering (QA) to open-ended visual reasoning. *Objs* refer to object types, while *Defcs* denote defect categories.

To address the challenges of limited generalization and restricted reasoning capacity, we propose MAU-GPT, a multimodal large model specifically tailored for industrial anomaly understanding. Built upon a pre-trained vision-language foundation, MAU-GPT introduces a novel Adaptive Mixture of Experts with Low-Rank Adaptation (AMoE-LoRA) mechanism that unifies anomaly-aware and generalist experts adaptation, enhancing both generalization and reasoning across diverse anomaly classes. Concretely, for general samples, MAU-GPT draws on the Mixture-of-Experts (MoE) paradigm by dynamically activating multiple general LoRA experts via input-dependent routing. This flexible expert selection alleviates the capacity limitations of single-LoRA modules and enables fine-grained, sample-aware representation learning. For anomalous samples, we integrate an anomaly-aware adaptation strategy in which a hypernetwork generates LoRA parameters on-the-fly, conditioned on the specific anomaly category. This eliminates the need for pre-trained category-specific LoRA modules and allows the model to adapt to previously unseen defects. Crucially, this hybrid adaptation framework facilitates seamless model generalization across diverse domains while enhancing sensitivity to rare and novel anomalies, enabling accurate detection in dynamic and complex industrial environments.

The contributions of this work are summarized as follows:

- **A Comprehensive Dataset.** We present MAU-Set, a broad-coverage, fine-grained dataset that introduces two

QA styles—*Discriminative QA* and *Open-Ended QA*—designed to evaluate both low-level perception and high-level visual reasoning across diverse industrial domains.

- **A Domain-Adaptive Model.** We propose MAU-GPT, a multimodal large model equipped with a novel AMoE-LoRA mechanism, enabling accurate and robust anomaly understanding across diverse domains. Extensive experiments show its strong potential and generalization capabilities for scalable and automated industrial inspection.
- **Open-Source Work.** We will open-source the dataset, model checkpoints, and codebase to facilitate further research and development in industrial anomaly analysis.

Related Work

- **Industrial Anomaly Detection.** Industrial anomaly detection methods are typically divided into unsupervised and supervised approaches. Unsupervised methods do not rely on labeled data, and are mainly based on feature embedding or reconstruction. Feature embedding methods, such as SVDD (Tax and Duin 2004) and PatchSVDD (Yi and Yoon 2020), distinguish normal and abnormal data via hypersphere-based or probabilistic modeling. Reconstruction-based approaches like Autoencoders (Bergmann et al. 2018) and GANs (Yan et al. 2021) use reconstruction loss to identify anomalies. Recently, transformer-based models (Mishra et al. 2021) and diffusion methods (Wyatt et al. 2022) have

Dataset	Product Type	Defect Class	Image Number			Annotation Level	QA Type	QA Pairs Constructed
			Normal	Abnormal	All			
BTAD (Mishra et al. 2021)	3	3	2,250	580	2,830	Segmentation mask	✗	16,863
DAGM (Wieler and Hahn 2007)	10	10	10,000	1,500	11,500	Segmentation mask	✗	109,715
DeepPCB (Tang et al. 2019)	1	6	1,501	1,500	3,001	Bounding box	✗	21,090
KolektorSDD2 (Božič, Tabernik, and Skočaj 2021)	1	3+	2,979	356	3,335	Segmentation mask	✗	22,274
Magnetic Tile (Huang, Qiu, and Yuan 2020)	1	5	952	383	1,335	Segmentation mask	✗	9,002
MPDD (Jezek et al. 2021)	4	8	1,064	282	1,346	Segmentation mask	✗	8,817
MVTec AD (Bergmann et al. 2019)	15	73	4,096	1,258	5,354	Segmentation mask	✗	36,580
MAU-Set (Ours)	35	108+	22,842	5,859	28,701	Instance-level QA	Open/Close	224,341

Table 1: Comparison of widely used industrial anomaly detection benchmark datasets with our MAU-Set. QA-related columns (QA Type and QA Pairs Constructed) and the annotation level correspond to our newly constructed VQA annotations.

shown promising results in industrial settings. In contrast, supervised methods utilize limited labeled anomalies. Some adopt semi-supervised learning (Chu and Kitani 2020) or loss-based training (Liznerski et al. 2020), while others focus on fine-grained anomaly modeling (Pang et al. 2021; Shao et al. 2024) or disentangling anomaly representations (Ding, Pang, and Shen 2022; Shao et al. 2021). Although these approaches depend on annotations, they often alleviate data imbalance by introducing pseudo-labels or synthesizing anomalies.

• **Large Vision-Language Models.** The rapid advancements in large language models (LLMs) have spurred the development of Large Vision-Language Models (LVLMs), which integrate powerful linguistic understanding with visual perception. These models, often built upon frozen image encoders and large pre-trained LLMs, demonstrate remarkable capabilities in multimodal understanding, reasoning, and generation across diverse tasks such as visual question answering, image captioning, and visual dialogue (Alayrac et al. 2022; Li et al. 2023a,b). By leveraging extensive pre-training on vast amounts of image-text pairs, LVLMs learn rich cross-modal representations, enabling them to comprehend complex visual scenes and respond to intricate textual queries. Recent architectures, such as LLaVA (Liu et al. 2024b), have notably advanced visual instruction tuning, while InternVL (Chen et al. 2024) has demonstrated significant progress in scaling up vision foundation models for large-scale visual document understanding, further refining efficiency and performance, and allowing for more versatile applications. The emergence of these sophisticated LVLMs presents a promising avenue for addressing complex visual analysis challenges in specialized domains, including industrial anomaly detection,

Dataset: MAU-Set

To support fair and comprehensive evaluation in industrial anomaly analysis, we introduce MAU-Set, which features six industrial domains, two progressive QA-style questions, ranging from basic binary classification to complex anomaly reasoning, with detailed collection and annotation provided in Figure 2. This hierarchical task structure increases both the complexity and depth of the benchmark, encouraging models to progressively accumulate knowledge and adapt to

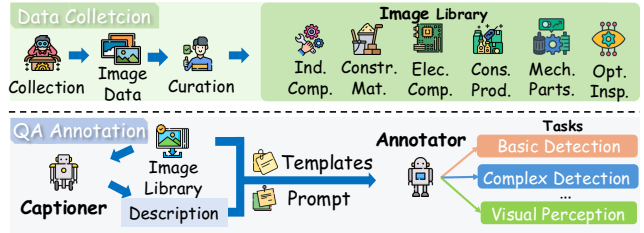


Figure 2: The data collection and annotation of MAU-Set. **Ind. Comp.** and **Constr. Mat.** denote Industrial Components and Construction Materials; **Elec. Comp., Cons. Prod., Mech. Parts** and **Opt. Insp.** are defined analogously.

the diverse challenges encountered in practical scenarios.

Dataset Design

- **Data Collection.** As illustrated in Figure 2, we construct the dataset by curating 7 data sources, incorporating both real and synthetic images to enable robust evaluation and foster progress in industrial anomaly analysis. As shown in Table 1, the dataset comprises 35 product types, encompassing both surface textures and objects themselves.
- **Domain Division.** To enhance field clarity, we further organize these sources into 6 industrial domains. As reported in Figure 3, these domains captures the breadth and complexity of real-world industrial scenarios.
- **QA Style.** To fully utilize existing data and support diverse evaluation, we design two QA styles: (i) Discriminative QA: Following the conventional detection formulation, the model determines whether an image contains an anomaly and produces a binary output: *True* (abnormal) or *False* (normal). (ii) Open-Ended QA: It requires the model to perform fine-grained reasoning by generating free-form responses to image-specific questions. These questions may involve identifying defect types, localizing defects, or inferring causes and implications, etc.
- **Task Division.** To improve the clarity of the task structure, we further decompose the two QA types into 5 distinct tasks. As illustrated in Figure 1, Discriminative QA corresponds to Basic Detection, while Open-Ended QA is subdivided into 4 tasks of increasing difficulty.




Figure 3: The data collection and annotation of MAU-Set.

Data Annotation

We collect plenty of images from 7 sources, however, some sources are designed for (semi-)unsupervised training, resulting in significant class imbalance between normal and abnormal samples in both training and test sets. To address this, we first collect normal and abnormal samples from each source, apply a fixed-ratio split for training and testing, and then merge them to create a more balanced dataset.

Inspired by prior work on the benefits of instruction diversity, we build a question pool with varied prompts for each task. For each image, 1–2 questions per task are sampled, yielding around 7–8 questions across 5 tasks. Grammatically incorrect questions are manually filtered. Unlike other tasks, the “In-depth Understanding” task necessitates domain-specific expertise in industrial products. To meet this requirement, we manually curate a set of high-quality, expert-level questions tailored to the unique characteristics of each product. These questions are further used as prompts for a pre-trained large language model to automatically generate additional professional-grade questions, thereby streamlining the annotation process.

Dataset Statistics

As shown in Table 1 and Figure 3, MAU-Set comprises 35 product types, over 100 defect classes, and approximately 28K images accompanied by 224K high-quality, diverse QA instruction pairs. The dataset spans 6 industrial domains and includes 2 styles of QA questions, which are further divided into 5 distinct tasks. This design ensures broad coverage of real-world anomaly scenarios and promotes the domain adaptability of detection models. Additionally, its hierarchical task structure—from Discriminative QA (binary classification) to Open-Ended QA (complex reasoning)—facilitates progressive learning, enhancing model generalization and robustness in anomaly detection.

Method: MAU-GPT

Overview

As illustrated in Figure 4, MAU-GPT comprises a frozen vision encoder, a trainable visual projection layer, a word embedding layer, and a stack of LLM blocks. At the core of these blocks lies our proposed **AMoE-LoRA** module, which serves as the primary mechanism for domain adaptation and knowledge specialization.

This design allows MAU-GPT to leverage large models’ pre-trained knowledge and generalization while enabling parameter-efficient domain adaptation and specialization for diverse industrial domains. The AMoE-LoRA architecture and training procedure are described in Sections 4.2 and 4.3, respectively.

Architecture of AMoE-LoRA

AMoE-LoRA is a hybrid expert architecture that integrates both anomaly-aware and generalist experts adaptations. The outputs from these two expert branches are fused into the model by adding them to the original output. In this work, multiple AMoE-LoRA module are inserted at multiple transformer blocks within a multimodal large language model, respectively. For clarity, we will describe the AMoE-LoRA module from the perspective of a single transformer block.

Specifically, the generalist experts component adopts the MoE paradigm, enabling multiple experts (G-LoRA) to collaborate through input-dependent routing that dynamically adjust the contribution weight of each expert. This module comprises N experts and a router. The output of the generalist experts is computed as follows:

$$o_1 = \Delta W x = \sum_{i=1}^N \mathcal{R}(x)_i \cdot E_i(x), \quad (1)$$

where $\mathcal{R}(\cdot) = \text{Softmax}(xW_g)$ denotes the routing function with router weights W_g , x is the embedding of the original model, and $E_i(\cdot)$ represents the i -th expert. Each expert is parameterized by a pair of low-rank matrices B_i and A_i as follows:

$$E_i(x) = B_i A_i x. \quad (2)$$

Both the router weights W_g and the low-rank matrices B_i and A_i are trainable. During forward propagation, the routing function $\mathcal{R}(\cdot)$ produces a contribution weight ω_i for each expert, effectively integrating knowledge from all experts. Thus, the output can be expressed as:

$$o_1 = \frac{\alpha}{r} \sum_{i=1}^N \omega_i \cdot B_i A_i x, \quad (3)$$

where r is the rank of the matrices and α is the scaling factor that modulates the impact of the generalist experts on the overall model.

In addition, the anomaly-aware expert component leverages a trainable hypernetwork to generate adaptive LoRA (A-LoRA) weights, enabling the model to adapt to each individual sample. This eliminates the need for pre-trained category-specific LoRA modules and allows the model to adapt to previously unseen defects. Specifically, the embedding x of the original model is fed into the hypernetwork H to produce the parameters of a pair of low-rank matrices A_0 and B_0 as follows:

$$\begin{aligned} A_0 &= W_a H(x), \\ B_0 &= H(x) W_b, \end{aligned} \quad (4)$$

where W_a and W_b are two learnable weight matrices. The output of the anomaly-aware expert is then defined as:

$$o_2 = \frac{\alpha}{r} \cdot B_0 A_0 x. \quad (5)$$

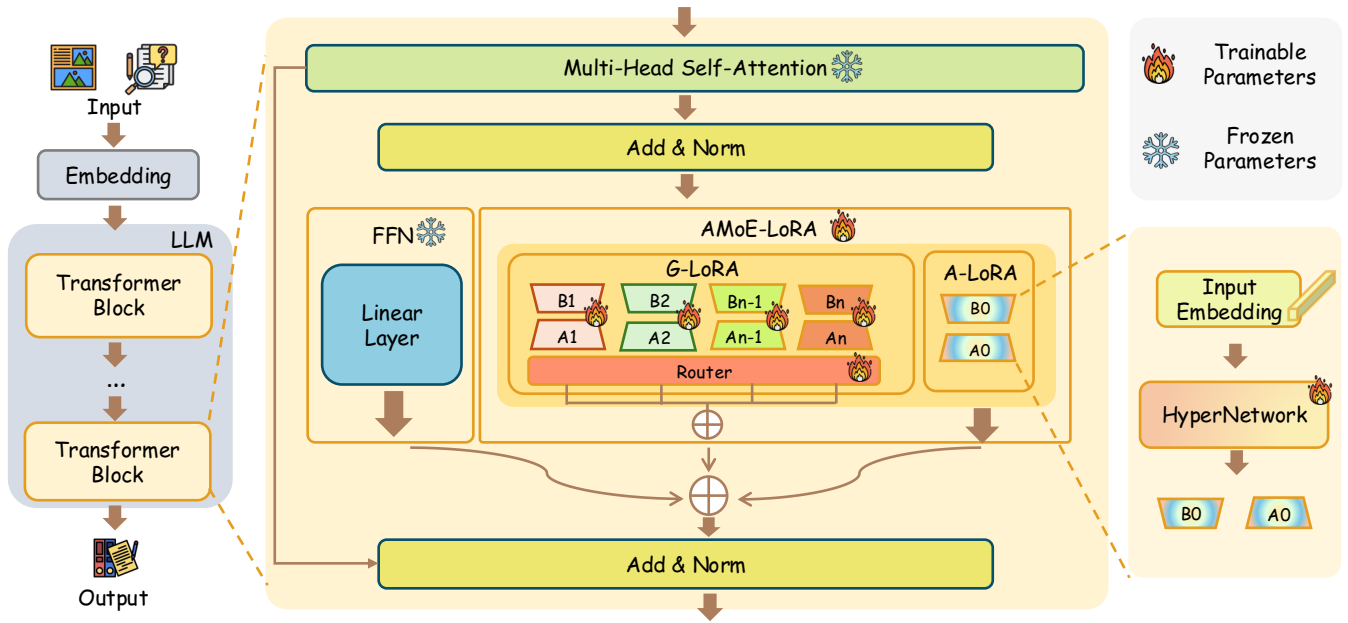


Figure 4: The AMoE-LoRA architecture combines generalist experts with an anomaly-aware expert, enabling the MLLM to incorporate both general knowledge and domain-specific insights for industrial anomaly understanding.

Finally, the outputs of both generalist experts component and anomaly-aware expert component are added to the original model output o_0 :

$$o = o_0 + o_1 + o_2, \quad (6)$$

where o denotes the final output.

Training Pipeline

In this work, we adopt a stage-wise training strategy, as suggested by previous studies (Moryossef, Goldberg, and Dagan 2019; Bolukbasi et al. 2017), to simplify the learning process, enhance robustness, and improve manageability. This strategy facilitates progressive knowledge transfer, efficient parameter tuning, and effective specialization, enabling the model to retain general reasoning capabilities while adapting to complex industrial anomaly scenarios.

- **Stage 1: Multimodal Pre-training.** The vision encoder and the LLM are kept frozen, while only the multimodal projection layer is trained. By Leveraging the high-quality instruction dataset, the model learns to align visual and textual modalities, thereby improving its ability to interpret and reason over image content.

- **Stage 2: Industrial Anomaly Instruction Fine-tuning.** Building on the established vision-language alignment and general reasoning capabilities, we fine-tune the proposed AMoE-LoRA module and the projection layer to specialize the model for industrial anomaly detection.

Experiments

Experimental Setup

- **Data Details.** The first-stage training follows the same procedure as LLaVA, where we pretrain the model on the LLaVA-558K dataset (Liu et al. 2024a). In the second stage,

we specialize the model for industrial anomaly detection by fine-tuning it on our proposed MAU-Set dataset.

- **Model Settings.** To comprehensively evaluate the performance of MAU-GPT, we compare it with a range of state-of-the-art models, including six mainstream open-source multimodal large language models (MMLMs), such as LLaVA-1.5 (with two parameter variants), Unified-IO2 (with two parameter variants), Gemma-3, Yi-VL, and InternVL-2.5. Additionally, we include AnomalyGPT, an MLLM specialized for industrial anomaly detection, which is evaluated in few-shot settings ranging from one-shot to four-shot.

- **Evaluation Metrics.** To better evaluate the model’s performance, we adopt both automatic and GPT-4o-judged metrics. Specifically, we report ROUGE (ROUGE-1/2/L) (Ganesan 2018) and BLEU (Saadany and Orasan 2021) to measure n-gram/sequence overlap between generated outputs and references. In addition, we leverage GPT-4o as an evaluator to assess answer correctness against the annotated ground truth, enabling a finer-grained evaluation.

Main Experiments

- **Performance on Discriminative QA.**

Table 2 reports results on the Discriminative QA task of MAU-Set. Early open-source models (e.g., LLaVA, UIO2) provide moderate baselines, and simply scaling (e.g., LLaVA-1.5) brings limited gains without domain alignment. More recent general-purpose VLMs (e.g., Yi-VL, InternVL-2.5) generally perform better, with Yi-VL being the most consistent. Although AnomalyGPT benefits from more in-context examples, it still lags behind due to prompt-only specialization. Overall, MAU-GPT achieves the best performance, demonstrating the advantage of domain-specific fine-tuning.

Type	Model	Industrial Components	Electronic Components	Construction Materials	Mechanical Parts	Consumer Products	Optical Inspection	Avg.
Generalist Model	LLaVA-1.5(7B)	40.10	30.21	39.46	50.66	45.83	55.40	43.61
	LLaVA-1.5(13B)	11.57	27.31	25.26	32.13	27.76	15.24	23.21
	UIO2-Large(1B)	66.61	53.26	53.28	57.06	61.61	66.40	59.70
	UIO2-XI(3B)	55.55	55.17	64.04	53.28	60.38	67.84	59.38
	Gemma-3(4B)	63.06	58.12	43.75	48.64	51.50	33.86	42.45
	Yi-VL(6B)	62.39	60.05	58.54	47.96	52.91	83.17	73.54
	InternVL-2.5(8B)	46.58	57.77	47.71	49.66	55.31	41.62	45.31
Specialized Model	AnomalyGPT(1-shot)	50.93	66.07	78.70	60.41	79.31	39.76	62.53
	AnomalyGPT(2-shot)	56.20	70.16	79.46	64.38	75.79	41.25	64.54
	AnomalyGPT(3-shot)	57.04	69.74	80.19	67.14	76.62	42.92	65.61
	AnomalyGPT(4-shot)	58.14	69.78	81.00	66.44	77.41	43.39	66.03
	MAU-GPT(4B)	89.37	53.73	73.75	77.56	74.75	86.81	82.12

Table 2: Performance of different models on Discriminative QA.

Type	Model	Industrial Components	Electronic Components	Construction Materials	Mechanical Parts	Consumer Products	Optical Inspection	Avg.
Generalist Model	LLaVA-1.5(7B)	16.11	30.32	39.85	30.72	40.48	40.60	33.01
	LLaVA-1.5(13B)	18.93	34.80	32.79	29.94	41.50	41.37	33.22
	UIO2-Large(1B)	14.33	25.06	33.38	29.14	34.45	13.38	24.96
	UIO2-XI(3B)	21.29	26.60	33.92	25.39	34.02	16.90	26.35
	Gemma-3(4B)	31.96	54.31	40.36	39.80	50.59	42.41	43.24
	Yi-VL(6B)	31.47	39.73	34.96	40.89	47.84	36.59	38.58
	InternVL-2.5(8B)	30.52	50.94	50.23	45.35	51.12	61.08	48.21
Specialized Model	AnomalyGPT(1-shot)	11.19	29.27	25.12	17.97	37.08	9.00	21.61
	AnomalyGPT(2-shot)	11.54	31.27	26.73	19.27	36.47	9.05	22.39
	AnomalyGPT(3-shot)	12.00	29.30	27.01	19.83	36.54	9.23	22.32
	AnomalyGPT(4-shot)	12.64	28.16	27.31	19.73	36.71	9.23	22.30
	MAU-GPT(4B)	88.75	81.66	85.64	84.37	86.97	94.39	91.05

Table 3: Performance of different models on GPT-4o judged accuracy for Open-Ended QA.

Model	ROUGE-1	ROUGE-2	ROUGE-L	BLEU-4
LLaVA-1.5(7B)	0.2495	0.0894	0.2008	0.0343
LLaVA-1.5(13B)	0.2459	0.0884	0.1976	0.0332
UIO2-Large(1B)	0.1312	0.0204	0.1284	0.0060
UIO2-XI(3B)	0.1666	0.0348	0.1556	0.0102
Gemma-3(4B)	0.2936	0.0905	0.2425	0.0357
Yi-VL(6B)	0.3044	0.1045	0.2469	0.0414
InternVL-2.5(8B)	0.2982	0.1103	0.2367	0.0440
AnomalyGPT(1-shot)	0.1947	0.0516	0.1730	0.0216
AnomalyGPT(2-shot)	0.1937	0.0514	0.1720	0.0217
AnomalyGPT(3-shot)	0.1943	0.0519	0.1725	0.0216
AnomalyGPT(4-shot)	0.1954	0.0521	0.1734	0.0216
MAU-GPT(4B)	0.7026	0.4537	0.6531	0.4773

Table 4: Performance of different models on ROUGE and BLEU metrics for Open-Ended QA.

• Performance on Open-Ended QA.

Compared to Discriminative questions, Open-Ended QA is inherently more challenging as models must generate coherent, contextually grounded text. Using GPT-4o as a no-reference evaluator alongside ROUGE and BLEU, we observe that newer general-purpose VLMs (e.g., InternVL, Gemma-3) outperform earlier baselines such as LLaVA and UIO2, while AnomalyGPT shows only limited gains with additional in-context examples—highlighting the limitations of prompt-only adaptation. In contrast, MAU-GPT achieves the strongest performance across all evaluation metrics, demonstrating that domain-tailored fine-tuning is

crucial for producing accurate and semantically aligned open-ended responses in industrial anomaly reasoning.

• Performance on MMAD Benchmark.

Table 5 summarizes results on MMAD, spanning seven tasks in anomaly discrimination and defect/object-level understanding. Although large general-purpose models (e.g., GPT-4o, LLaVA-NeXT) achieve higher absolute scores, they require substantially more parameters. Under comparable model sizes, MAU-GPT attains the best overall performance among open-source VLMs and approaches the accuracy of much larger 34B models, highlighting strong parameter efficiency for industrial anomaly understanding.

Ablation and In-depth Study

• **Effect of the Architecture.** To evaluate the effectiveness of the AMoE-LoRA architecture, we compare it with baseline methods such as LoRA (Hu et al. 2022) and LoRAMoE (Dou et al. 2023). The experiments are designed to verify how AMoE-LoRA improves performance by unifying anomaly-aware and generalist experts adaptation. As shown in Table 6, AMoE-LoRA achieve superior results across all metrics, which comprehensively demonstrates the significant advantages of our proposed architecture.

• **Effect of the Model Parameters.** To investigate the impact of parameter configurations on the performance of MAU-GPT, we conduct a series of ablation studies. We keep the total number of parameters constant while varying the

Model	Anomaly Discrimination	Defect				Object		Average
		Classification	Localization	Description	Analysis	Classification	Analysis	
Random Chance	50.00	25.00	25.00	25.00	25.00	25.00	25.00	28.57
Human (expert)	95.24	75.00	92.31	83.33	94.20	86.11	80.37	86.65
GPT-4o-mini	64.33	48.58	38.75	63.68	80.40	88.56	79.74	66.29
GPT-4o	68.63	65.80	55.62	73.21	83.41	94.98	82.80	74.92
Uio2-large-1.5B	49.27	27.52	18.95	25.28	37.85	43.75	47.24	36.08
Uio2-XI-6B	24.52	32.23	36.61	36.99	48.64	57.85	59.21	38.62
AnomalyGPT-7B	61.53	31.39	27.20	30.76	22.24	44.48	30.23	35.41
Gemma3-4B	57.37	37.07	35.71	47.49	76.32	85.11	73.65	58.96
Qwen-VL-Chat-7B	53.65	31.33	28.62	41.66	63.99	74.46	67.94	51.66
SPHINX-7B	53.13	33.93	52.27	50.96	71.23	85.07	73.10	59.96
LLaVA-1.5-13B	49.96	38.78	46.17	58.17	73.09	73.62	70.98	58.68
InternVL2-34B	59.97	43.85	47.91	57.60	78.10	74.18	80.37	63.14
LLaVA-NeXT-34B	57.92	48.79	52.87	71.34	80.28	81.12	77.80	67.16
MAU-GPT-4B	67.50	39.31	41.24	66.37	77.04	67.09	71.31	61.41

Table 5: Accuracy of various models on the MMAD Benchmark across different industrial anomaly understanding tasks.




Figure 5: (a) describes the influence of the number and rank of the generalist experts. (b) shows that our model receives the highest endorsement from human experts. (c) is the 2D representation of hypernetwork-generated parameters across samples.

Architecture	ROUGE-1	ROUGE-2	ROUGE-L	BLEU-4
LoRA	0.6971	0.4446	0.6449	0.4529
LoRAMoE	0.6991	0.4452	0.6468	0.4598
AMoE-LoRA	0.7026	0.4537	0.6531	0.4773

Table 6: Ablation study about architecture design decisions.

number of experts in the MoE module and the rank of each expert’s LoRA matrix. Specifically, we maintain their product at 64. The results, as shown in Figure 5 (a), indicate that the model achieves optimal performance on evaluation metrics such as ROUGE and BLEU when the number of experts is set to 16 and the rank of each expert is 4.

• **Expert Evaluation.** We conduct an expert evaluation to further investigate the performance of our model in MAU-Set. We randomly sampled 500 QA pairs from Open-Ended QA. Subsequently, we recruit 10 domain experts to perform a blind ranking of the responses generated by LLaVA-1.5-7B, UIO2-XI-3B, Gemma-3, Yi-VL, InternVL-2.5, AnomalyGPT and MAU-GPT-4B. As shown in Figure 5 (b), the evaluation results indicate that the responses generated by MAU-GPT are significantly superior to those from other methods in terms of quality and relevance, better meeting real-world application requirements.

• **Adaptive Behavior of the Anomaly-Aware Expert.** To examine the adaptability of our anomaly-aware expert, we analyze the LoRA parameters generated by the hypernetwork across different samples and visualize them with t-SNE in 2D. As shown in Figure 5 (c), the embeddings form clear clusters aligned with object categories, indicating that the hypernetwork produces input-conditioned adapters. Unlike conventional LoRA-based MLLMs with fixed parameters, our model dynamically generates sample-specific LoRA weights, better capturing diverse and subtle industrial anomalies.

Conclusion

In this work, we present MAU-Set, a comprehensive dataset for multi-type industrial anomaly understanding, featuring broad coverage, fine-grained supervision, and a hierarchical QA task structure that spans from binary classification to complex reasoning. To further advance this domain, we introduce MAU-GPT, a domain-adaptive multimodal large model empowered by a novel AMoE-LoRA mechanism that unifies anomaly-aware and generalist experts adaptation. We believe our dataset and model will foster further research in industrial visual understanding and anomaly reasoning.

Acknowledgments

This work was supported by the National Key Research and Development Program of China (2024YFB3312900), the NSFC (No. 62402426, 62272411, 62441617, 62506333, 62276230), the Key Research and Development Projects in Zhejiang Province (No. 2025C01128, 2024C01106, 2025C01030, 2025C02156), Ningbo Yongjiang Talent Introduction Programme (2023A-400-G), Zhejiang University Education Foundation Qizhen Scholar Foundation, Zhejiang Provincial Natural Science Foundation of China (No. LD25F020001, LDT23F02023F02), the Postdoctoral Fellowship Program of CPSF (GZC20251077), Zhejiang Province Postdoctoral Research Excellence Funding Project (ZJ2025065) and Fundamental Research Funds for the Central Universities (226-2025-00057).

References

- Alayrac, J.-B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hasson, Y.; Lenc, K.; Mensch, A.; Millican, K.; Reynolds, M.; et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35: 23716–23736.
- Alzarooni, A.; Iqbal, E.; Khan, S. U.; Javed, S.; Moyo, B.; and Abdulrahman, Y. 2025. Anomaly Detection for Industrial Applications, Its Challenges, Solutions, and Future Directions: A Review. *arXiv preprint arXiv:2501.11310*.
- Baitieva, A.; Hurych, D.; Besnier, V.; and Bernard, O. 2024. Supervised Anomaly Detection for Complex Industrial Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17754–17762.
- 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.
- Bergmann, P.; Löwe, S.; Fauser, M.; Sattlegger, D.; and Steger, C. 2018. Improving unsupervised defect segmentation by applying structural similarity to autoencoders. *arXiv preprint arXiv:1807.02011*.
- Bhatt, P. M.; Malhan, R. K.; Rajendran, P.; Shah, B. C.; Thakar, S.; Yoon, Y. J.; and Gupta, S. K. 2021. Image-based surface defect detection using deep learning: A review. *Journal of Computing and Information Science in Engineering*, 21(4): 040801.
- Bolukbasi, T.; Wang, J.; Dekel, O.; and Saligrama, V. 2017. Adaptive neural networks for efficient inference. In *ICML*.
- Božič, J.; Tabernik, D.; and Skočaj, D. 2021. Mixed supervision for surface-defect detection: From weakly to fully supervised learning. *Computers in Industry*, 129: 103459.
- Cao, Y.; Xu, X.; Zhang, J.; Cheng, Y.; Huang, X.; Pang, G.; and Shen, W. 2024. A survey on visual anomaly detection: Challenge, approach, and prospect. *arXiv preprint arXiv:2401.16402*.
- Chen, Z.; Chen, W.; Smiley, C.; Shah, S.; Borova, I.; Langdon, D.; Moussa, R.; Beane, M.; Huang, T.-H.; Routledge, B.; et al. 2021. Finqa: A dataset of numerical reasoning over financial data. *arXiv preprint arXiv:2109.00122*.
- Chen, Z.; Wang, W.; Cao, Y.; Liu, Y.; Gao, Z.; Cui, E.; Zhu, J.; Ye, S.; Tian, H.; Liu, Z.; et al. 2024. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*.
- Chu, W.-H.; and Kitani, K. M. 2020. Neural batch sampling with reinforcement learning for semi-supervised anomaly detection. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI 16*, 751–766. Springer.
- Ding, C.; Pang, G.; and Shen, C. 2022. Catching both gray and black swans: Open-set supervised anomaly detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7388–7398.
- Dou, S.; Zhou, E.; Liu, Y.; Gao, S.; Zhao, J.; Shen, W.; Zhou, Y.; Xi, Z.; Wang, X.; Fan, X.; et al. 2023. LoRAMoE: Alleviate world knowledge forgetting in large language models via MoE-style plugin. *arXiv preprint arXiv:2312.09979*.
- Ganesan, K. 2018. Rouge 2.0: Updated and improved measures for evaluation of summarization tasks. *arXiv preprint arXiv:1803.01937*.
- 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.
- Huang, Y.; Qiu, C.; and Yuan, K. 2020. Surface defect saliency of magnetic tile. *The Visual Computer*, 36(1): 85–96.
- Jezeck, 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. IEEE.
- Jiang, X.; Li, J.; Deng, H.; Liu, Y.; Gao, B.-B.; Zhou, Y.; Li, J.; Wang, C.; and Zheng, F. 2024. Mmad: The first-ever comprehensive benchmark for multimodal large language models in industrial anomaly detection. *arXiv preprint arXiv:2410.09453*.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, 19730–19742. PMLR.
- Li, L.; Chen, L.; Huang, Y.; Zhang, Z.; Zhang, S.; and Xiao, J. 2022. The devil is in the labels: Noisy label correction for robust scene graph generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18869–18878.
- Li, L.; Xiao, J.; Chen, G.; Shao, J.; Zhuang, Y.; and Chen, L. 2023b. Zero-shot visual relation detection via composite visual cues from large language models. *Advances in Neural Information Processing Systems*, 36: 50105–50116.
- Liu, H.; Li, C.; Li, Y.; and Lee, Y. J. 2024a. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 26296–26306.

- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2024b. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Liznerski, P.; Ruff, L.; Vandermeulen, R. A.; Franks, B. J.; Kloft, M.; and Müller, K.-R. 2020. Explainable deep one-class classification. *arXiv preprint arXiv:2007.01760*.
- Lu, P.; Mishra, S.; Xia, T.; Qiu, L.; Chang, K.-W.; Zhu, S.-C.; Tafjord, O.; Clark, P.; and Kalyan, A. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35: 2507–2521.
- Mishra, P.; Verk, R.; Fornasier, D.; Piciarelli, C.; and Foresti, G. L. 2021. VT-ADL: A vision transformer network for image anomaly detection and localization. In *2021 IEEE 30th International Symposium on Industrial Electronics (ISIE)*, 01–06. IEEE.
- Moryossef, A.; Goldberg, Y.; and Dagan, I. 2019. Step-by-step: Separating planning from realization in neural data-to-text generation. *arXiv preprint arXiv:1904.03396*.
- Pang, G.; Ding, C.; Shen, C.; and Hengel, A. v. d. 2021. Explainable deep few-shot anomaly detection with deviation networks. *arXiv preprint arXiv:2108.00462*.
- Ren, Z.; Fang, F.; Yan, N.; and Wu, Y. 2022. State of the art in defect detection based on machine vision. *International Journal of Precision Engineering and Manufacturing-Green Technology*, 9(2): 661–691.
- Saadany, H.; and Orasan, C. 2021. BLEU, METEOR, BERTScore: evaluation of metrics performance in assessing critical translation errors in sentiment-oriented text. *arXiv preprint arXiv:2109.14250*.
- Shao, F.; Luo, Y.; Gao, F.; Yang, Y.; and Xiao, J. 2024. Knowledge-guided causal intervention for weakly-supervised object localization. *IEEE Transactions on Knowledge and Data Engineering*, 36(11): 6477–6489.
- Shao, F.; Luo, Y.; Zhang, L.; Ye, L.; Tang, S.; Yang, Y.; and Xiao, J. 2021. Improving weakly supervised object localization via causal intervention. In *Proceedings of the 29th ACM International Conference on Multimedia*, 3321–3329.
- Tang, S.; He, F.; Huang, X.; and Yang, J. 2019. Online PCB defect detector on a new PCB defect dataset. *arXiv preprint arXiv:1902.06197*.
- Tang, Y.; Sun, K.; Zhao, D.; Lu, Y.; Jiang, J.; and Chen, H. 2022. Industrial Defect Detection Through Computer Vision: A Survey. In *2022 7th IEEE International Conference on Data Science in Cyberspace (DSC)*, 605–610. IEEE.
- Tax, D. M.; and Duin, R. P. 2004. Support vector data description. *Machine learning*, 54: 45–66.
- Wieler, M.; and Hahn, T. 2007. Weakly supervised learning for industrial optical inspection. In *DAGM symposium in*, volume 6, 11.
- Wyatt, J.; Leach, A.; Schmon, S. M.; and Willcocks, C. G. 2022. Anoddpn: Anomaly detection with denoising diffusion probabilistic models using simplex noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 650–656.
- Yan, X.; Zhang, H.; Xu, X.; Hu, X.; and Heng, P.-A. 2021. Learning semantic context from normal samples for unsupervised anomaly detection. In *Proceedings of the AAAI conference on artificial intelligence*, 3110–3118.
- Yi, J.; and Yoon, S. 2020. Patch svdd: Patch-level svdd for anomaly detection and segmentation. In *Proceedings of the Asian conference on computer vision*.