

# 000 001 002 003 004 005 FAKEXPLAIN: AI-GENERATED IMAGES DETECTION 006 VIA HUMAN-ALIGNED GROUNDED REASONING 007 008 009

010 **Anonymous authors**  
011 Paper under double-blind review  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026  
027

## ABSTRACT

028 The rapid rise of image generation calls for detection methods that are both in-  
029 terpretable and reliable. Existing approaches, though accurate, act as black boxes  
030 and fail to generalize to out-of-distribution data, while multi-modal large language  
031 models (MLLMs) provide reasoning ability but often hallucinate. To address these  
032 issues, we construct **FakeExplained** dataset of AI-generated images annotated with  
033 bounding boxes and descriptive captions that highlight synthesis artifacts, forming  
034 the basis for human-aligned, visually grounded reasoning. Leveraging **FakeExplained**,  
035 we develop **FakeExplainer** which fine-tunes MLLMs with a progressive  
036 training pipeline, enabling accurate detection, artifact localization, and coherent  
037 textual explanations. Extensive experiments show that **FakeExplainer** not only  
038 sets a new state-of-the-art in detection and localization accuracy (98.2% accuracy,  
039 36.0% IoU), but also demonstrates strong robustness and out-of-distribution gen-  
040 eralization, uniquely delivering spatially grounded, human-aligned rationales.  
041  
042  
043  
044  
045  
046  
047  
048  
049  
050  
051

## 1 INTRODUCTION

052 The past decade has witnessed rapid progress in text-to-image generation, evolving from Genera-  
053 tive Adversarial Networks to Diffusion Models, which are now capable of producing images nearly  
054 indistinguishable from real photographs (Goodfellow et al., 2014; Peebles & Xie, 2023). These ad-  
055 vances have led to an explosion of highly realistic AI-generated content, raising pressing concerns  
056 about misinformation, authenticity, and trust in digital media. Most existing detection methods cast  
057 this task as a binary classification problem, leveraging convolutional neural networks or vision trans-  
058 formers (Wang et al., 2020; Ojha et al., 2023; Park & Owens, 2024). However, binary labels offer  
059 limited insight into *why* an image is classified as AI-generated. In real-world applications, especially  
060 those involving legal, journalistic, or ethical implications, explainable detection is essential. An ef-  
061 fective detection system should not only identify whether an image is AI-generated but also pinpoint  
062 the specific visual cues or logical inconsistencies that betray its synthetic origin. Such explainability  
063 promotes user trust, supports verification workflows, and enables more informed decision-making.  
064  
065

066 The rise of Multi-modal Large Language Models (MLLMs) has enabled cross-modal inference,  
067 allowing models to generate human-readable explanations about AI-generated images. Recent ef-  
068 forts (Li et al., 2024; Zhang et al., 2024; Zhou et al., 2025; Gao et al., 2025; Xu et al., 2024; Liu  
069 et al., 2024; Ji et al., 2025) have advanced interpretable textual reasoning using MLLMs. However,  
070 these methods either depend heavily on prompt engineering, model-generated explanations or plug-  
071 in segmentation modules (Kirillov et al., 2023) to delineate manipulated regions. As illustrated in  
072 Figure 1, existing MLLM-based detectors may hallucinate false claims or provide reasons without  
073 spatial grounding, since their explanations are not validated by human annotations. Without proper  
074 visual grounding or human-aligned supervision, it remains unclear whether the generated rationales  
075 truly reflect the image content or derive from model hallucinations. To improve human alignment,  
076 fine-grained multimodal supervision, such as region-level annotations and captions, is essential. Yet,  
077 the lack of such high-quality datasets poses a major challenge to building reliable and interpretable  
078 MLLM-based detection systems.  
079  
080

081 In this paper, we present **FakeExplained**, a dataset of high-quality AI-generated images with fine-  
082 grained, human-grounded annotations, together with **FakeExplainer**, an RL fine-tuning pipeline for  
083 MLLMs that achieves state-of-the-art detection accuracy and grounding performance. As shown in  
084  
085

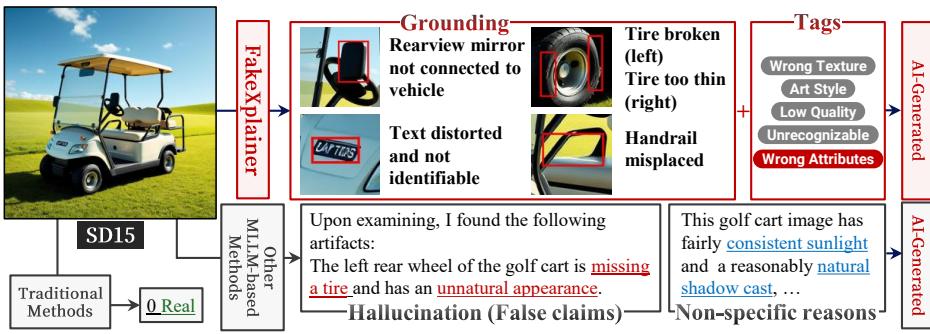


Figure 1: Comparison of our method (**FakeExplainer**) with traditional classification-based detectors (without explanations) and other MLLM-based methods (with hallucinated or non-specific reasons). **FakeExplainer** is trained to localize flawed regions and explain why the image appears AI-generated.

Figure 1, training on FakeExplained enables FakeExplainer to provide comprehensive rationales for fake image detection, performing on par with human experts. Our major contributions are threefold:

- **The FakeExplained dataset** is a curated dataset of 8,772 AI-generated images from diverse state-of-the-art generative models, annotated with bounding boxes and concise captions that highlight visual anomalies and illogical details.
- **The FakeExplainer detector:** By fine-tuning MLLMs on **FakeExplained**, we build an end-to-end system that not only detects AI-generated images but also explains them. Fine-tuning on FakeExplained enables **FakeExplainer** to perform fine-grained visual reasoning and articulate clear, human-aligned observations.
- **Robust performance with explainability:** **FakeExplainer** answers “*where and why does this image look fake?*” with reliable, spatially grounded explanations. Extensive experiments show that it achieves state-of-the-art detection accuracy, generalizes well to out-of-distribution images, and remains robust under perturbations while providing human-aligned, interpretable reasoning.

## 2 RELATED WORKS

**Detection of AI-generated and manipulated images.** Detecting AI-generated images has gained prominence with the improving fidelity of synthetic images from GANs (Goodfellow et al., 2014; Esser et al., 2021), autoregressive transformers (Van Den Oord et al., 2017), diffusion-based models (Le et al., 2025; Ye et al., 2025; Wang et al., 2025b; Li et al., 2025a; Chadebec et al., 2025; Song et al., 2020; Ho et al., 2020) and DiTs (Peebles & Xie, 2023; Chen et al., 2023). Deep learning methods such as ResNets and Vision Transformers trained on real and synthetic data (Wang et al., 2020; Tan & Le, 2019; Park & Owens, 2024; Chang et al., 2023; Yan et al., 2024) leverage strong feature extraction to learn discriminative patterns. However, generalization to unseen models remains challenging (Bi et al., 2023; Yan et al., 2024). As generation techniques evolve, artifact-based cues alone become increasingly unreliable. A complementary research direction focuses on model explainability, as most detectors offer only binary classification without indicating how or where synthetic cues are found. Recent efforts towards fine-grained or localized detection include using multi-branch systems for multi-level labels (Bi et al., 2023), computing local intrinsic dimensionalities (Lorenz et al., 2023), or using gradient visualizations (Selvaraju et al., 2017).

Despite these advances, existing MLLM-based methods still face limitations in providing human-aligned, grounded explanations and in generalizing across rapidly evolving generative techniques. While several recent datasets leverage MLLMs to generate rationales or global labels (Zhang et al., 2024; Zhou et al., 2025; Huang et al., 2025b; Wen et al., 2025; Li et al., 2025b), their reliance on automatic annotations raises concerns about hallucination and weak visual grounding. Although FakeBench (Li et al., 2024) incorporates human refinement, its explanations are initially generated by GPT-4V and remain purely textual, without any region-level grounding; moreover, the full annotations are not publicly released. LOKI covers multiple modalities but relies on GPT-4o for fine-grained labels, which weakens its alignment with humans. Other datasets such as So-Fake-

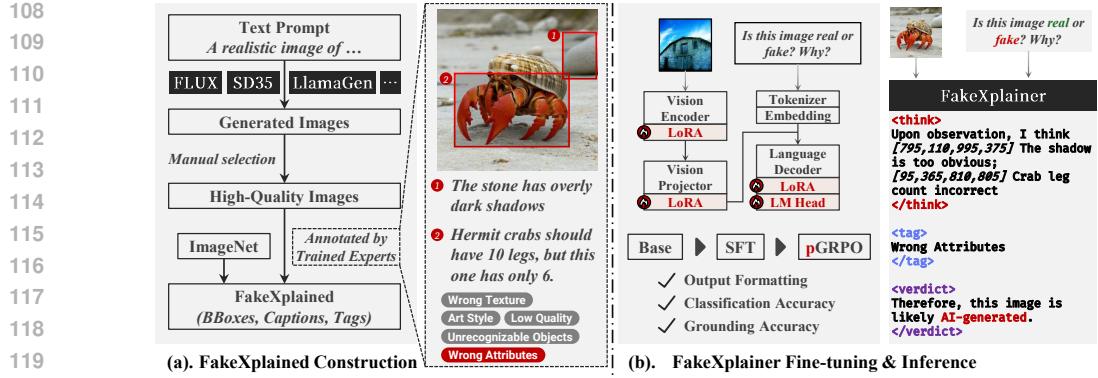


Figure 2: Overview of: (a). **FakeXplained**: Construction of dataset with human-aligned bounding boxes and captions, and (b). **FakeXplainer**: Progressive fine-tuning pipeline of MLLMs, which integrates SFT and GRPO to achieve accurate detection, grounding, and interpretable reasoning.

Set (Huang et al., 2025b) focus on social-media scenarios and similarly lack region-level annotations, limiting their suitability for training grounded, human-aligned detectors. Methodologically, recent detectors such as AIGI-Holmes (Zhou et al., 2025) combine NLP with MLLMs to improve interpretability, whereas localization-focused approaches like FakeShield (Xu et al., 2024) and LEGION (Kang et al., 2025) depend on external segmentation modules (e.g., SAM (Kirillov et al., 2023)), leaving the intrinsic grounding capabilities of MLLMs underutilized. Meanwhile, methods without additional modules, such as So-Fake (Huang et al., 2025b) and FakeVLM (Wen et al., 2025), still cannot generate spatially localized explanations. The absence of high-quality datasets remains a key obstacle to building reliable and interpretable MLLM-based detection systems. Without proper grounding or human-aligned supervision, it is unclear whether generated rationales truly reflect image content or are derived from hallucinations.

**Training & fine-tuning reasoning-capable MLLMs.** Enhancing the reasoning capabilities of MLLMs is crucial for tasks requiring nuanced understanding (Wu et al., 2025a;b; Yang et al., 2025a; Fang et al., 2025; Chen et al., 2024c). Initial strategies involved converting images into formalized textual representations to enable structured, language-driven reasoning (Yang et al., 2025b). Subsequent research has focused on instilling deeper cognitive abilities, including self-verification, self-correction, developing “slow thinking” capabilities (Wang et al., 2025a), and managing reasoning depth to address phenomena like “overthinking” (Xiao et al., 2025). Efforts also explore constructing high-quality multi-modal Chain-of-Thought (CoT) datasets (Huang et al., 2025a) to guide reasoning processes.

Reinforcement Learning (RL) has become pivotal in these advancements, with many sophisticated reasoning developments relying on RL methodologies. Fine-tuning methods have spurred significant interest in RL-based multi-modal reasoning (Chen et al., 2025). RL, particularly when combined with structured reward functions, e.g., using Intersection over Union (IoU) for tasks involving image grounding (Shen et al., 2025) - markedly improves multi-modal alignment, visual reasoning, and human interpretable decision making, demonstrating RL’s capability of advancing model performance in complex vision-language tasks.

### 3 THE FAKEXPLAINED DATASET

Our objective is trustworthy and interpretable detection of AI-generated images. This requires detectors that generalize to unseen generative models, remain robust to perturbations, and provide human-understandable explanations. Conventional detectors often lack interpretability, while MLLMs, though promising, tend to produce unreliable explanations with frequent hallucinations when used without training (see Table 4). To address this, we require models that not only detect AI-generated images but also explain their decisions in natural language for reliability. Achieving this demands a dataset that supports both visual grounding and textual reasoning. Therefore, we introduce the **FakeXplained** dataset, as illustrated in Figure 2(a), to train an MLLM to produce trustworthy ex-

162 explanations. It consists of high-quality synthetic images paired with fine-grained, human annotations  
 163 that indicate the underlying flaws and artifacts responsible for detection as fake.  
 164

### 165 3.1 AI-GENERATED IMAGES SELECTION 166

167 To ensure diversity in image sources, we generated images across the ImageNet-1K classes using 28  
 168 text-to-image generation models (details shown in Appendix A.1). All generated images underwent  
 169 manual quality screening, removing low-quality samples from the dataset. After screening, 8,772  
 170 AI-generated images were selected for subsequent annotation.

### 171 3.2 IMAGE ANNOTATION 172

173 To support interpretable reasoning and help models understand what constitutes an AI-generated  
 174 image, we provide detailed annotations for synthetic images. Real images are not annotated because  
 175 they lack synthesis flaws.  
 176

177 **Flaw regions and explanations.** Precise regional annotations and corresponding textual descriptions  
 178 are essential for visual grounding and interpretability. We recruited 23 trained annotators to  
 179 label the high-quality AI-generated images selected from the previous stage. Their primary task was  
 180 to identify and describe all regions within each image that exhibited signs of being fake (detailed  
 181 guidelines provided in Appendix A.2). Prior to annotation, all participants underwent standardized  
 182 training focused on identifying visual cues of AI-generated content. The training emphasized the  
 183 identification of *fake regions*, which are defined as areas within an image that either violate common  
 184 sense or exhibit noticeable AI-generated artifacts. Examples of common sense violations include  
 185 anomalies such as “a flamingo with three legs” or “bird feathers with a metallic appearance”. Common  
 186 AIGC artifacts include “repetitive patterns on a blanket” or “blurred or illegible text”.

187 Annotators were also introduced to a structured annotation rubric to ensure consistency and alignment  
 188 with the dataset’s objectives. Each annotation consists of one or more fake regions, where  
 189 each region is represented by a tuple  $(R_i, T_i)$ , where  $R_i$  denotes a rectangular bounding box en-  
 190 capsulating the region, and  $T_i$  provides a textual description of the identified anomaly or artifact.  
 191 On average, an annotated image in the dataset contains **5.42** such  $(R_i, T_i)$  pairs, which serve as the  
 192 foundation for grounding and reasoning in downstream model training.

193 **Image-level tagging.** In addition to region-level annotation, annotators were asked to tag images  
 194 based on broader perceptual attributes. These attributes include texture quality, overall realism, cor-  
 195 rectness of attributes, recognizability of objects, and the presence of other significant defects not  
 196 explicitly listed (*e.g.*, the occurrence of multiple sub-images within a single image). These tags  $C_i$   
 197 are mutually independent, allowing annotators to assign zero or multiple tags to each image as ap-  
 198 propriate. This tagging framework allows the dataset to capture holistic image quality assessments,  
 199 particularly in cases where visually realistic AI-generated images may lack distinct localized flaws.

### 200 3.3 QUALITY CONTROL 201

202 To ensure the reliability of the annotations, we implemented a quality control protocol involving  
 203 both manual inspection and algorithmic validation. A subset of annotations was compared against a  
 204 reference set of fake region annotations curated by the research team. Given the inherently subjective  
 205 nature of visual interpretation, we adopted a tolerant validation criterion to accommodate diverse  
 206 perspectives among annotators. Specifically, a minimum Intersection over Union (IoU) threshold of  
 207 20.0% was applied for bounding box overlap, and an accuracy threshold of  $\frac{1}{3}$  was used for image-  
 208 level taggings. These metrics were assessed on a validation set comprising 5% of the annotated  
 209 images. The IoU metric is used to assess the spatial agreement between annotated and reference  
 210 bounding boxes. Let  $R_v$  represent the rectangular bounding box annotated by a volunteer, and  $R_r$   
 211 represent the corresponding reference bounding box from the reference set. The Intersection over  
 212 Union (IoU) is computed as:

$$213 \text{IoU}(R_v, R_r) = \frac{|R_v \cap R_r|}{|R_v \cup R_r|},$$

214 where  $|R_v \cap R_r|$  denotes the area of the intersection between  $R_v$  and  $R_r$ , and  $|R_v \cup R_r|$  denotes  
 215 the area of their union. The IoU value ranges from 0 to 1, with higher values indicating stronger  
 216 alignment. This quality control procedure ensures a baseline level of annotation fidelity while pre-  
 217 serving the diversity of human interpretations. The resulting dataset, enriched with both region-level

216 annotations  $(R, T)$  and image-level tags  $C$ , offers a robust foundation for analyzing the semantic in-  
 217 consistencies and perceptual flaws of AI-generated images.  
 218

## 219 4 METHODOLOGY: FAKEEXPLAINER

220 We propose a training methodology named **FakeExplainer** for MLLMs designed to detect AI-  
 221 generated imagery, localize relevant artifacts, and articulate the rationale for their predictions. The  
 222 training and inference pipeline is shown in Figure 2(b). Inspired by DeepSeek-Math (Shao et al.,  
 223 2024), the training pipeline begins with Supervised Fine-Tuning (SFT) (Ouyang et al., 2022) to  
 224 provide basic reasoning ability and ensure structured outputs. This initial phase is succeeded by  
 225 Reinforcement Learning from Human Feedback (RLHF), which is implemented using *progressive*  
 226 GRPO (pGRPO), leveraging our proposed FakeExplained dataset.  
 227

228 Before training, each image’s annotations are reformatted into a dialogue between a user and an  
 229 assistant, using a prompt structure designed for localization-aware fine-tuning. Region-level anno-  
 230 tations  $(R_i, T_i)$  are enclosed within `<think>` markers, image-level tags  $C_i$  within `<tag>` markers,  
 231 and the final verdict is wrapped in `<verdict>` markers.  
 232

### 233 4.1 COLD START WITH SUPERVISED FINE-TUNING

234 The cold start phase of **FakeExplainer** uses SFT to establish a stable foundation before proceeding  
 235 to RL. During this phase, all linear layers of the vision encoder, projector, and language model  
 236 components in the MLLM are fine-tuned based on the supervision signals from the data. This initial  
 237 fine-tuning is crucial for stabilizing the model prior to full-scale reinforcement learning training,  
 238 preventing instabilities that might arise from pure RL-based updates (Guo et al., 2025).  
 239

240 The SFT process focuses on teaching the model to produce coherent reasoning patterns with a  
 241 clear structure. The training emphasizes the consistent use of the designated marker format with  
 242 `<think>`, `<tag>`, and `<verdict>` fields, ensuring format clarity in the model’s reasoning out-  
 243 puts. This structured Chain-of-Thought (CoT) format reduces errors and improves explainability,  
 244 providing a solid foundation for subsequent GRPO stages that will refine the model’s performance  
 245 on specific metrics.  
 246

### 247 4.2 DESIGN OF REWARD FUNCTIONS

248 Reward design is a critical component of RLHF, guiding the MLLMs to learn not only how to detect  
 249 fake images, but also how to localize relevant regions and provide coherent reasoning. We define  
 250 three core reward functions for this purpose.  
 251

252 **Classification accuracy (Label).** To ensure the model produces the correct verdict, we extract the  
 253 classification decision from within the `verdict` marker and compare it with the ground-truth label.  
 254 Let  $o$  denote the textual output of the MLLM, we have:  
 255

$$\mathcal{R}_C(o) = \begin{cases} 1, & \text{if } V(o) = y, \\ 0, & \text{o.w.} \end{cases}$$

258 where  $V(o)$  is a regex match for the verdict, and  $y$  is the ground-truth label of whether the image is  
 259 real or generated.  
 260

**Grounding accuracy (IoU).** To reward alignment between model-predicted and human-annotated  
 261 regions, we use a relaxed version of the Intersection over Union (IoU):  
 262

$$\mathcal{R}_G(o) = \text{IoU}^{\times\eta} = \min(1, \eta \text{IOU}(R(o), R_y))$$

264 where  $R(o)$  is the region extraction function that parses textual output  $o$  to bounding boxes,  $R_y$   
 265 is the annotated region, and  $\eta$  is a relaxing constant. The relaxation is based on the observation  
 266 that human annotators have slight discrepancies regarding the border of annotated regions. This  
 267 relaxation reward ensures full credit to the model when the regions annotated by models are in good  
 268 correlation with human-annotated ones.  
 269

**Output format validity (Format).** To ensure the model understands the structural requirements of  
 the task, we introduce a format reward that encourages outputs conforming to the expected syntax.  
 270

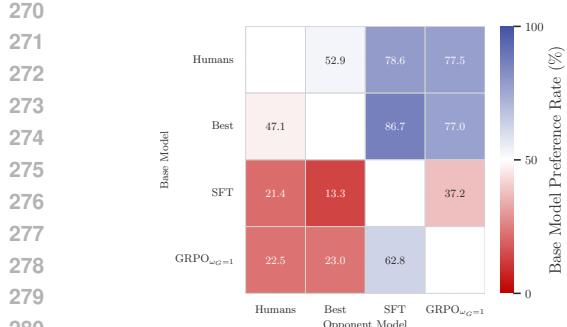


Figure 3: Human preference matrix.

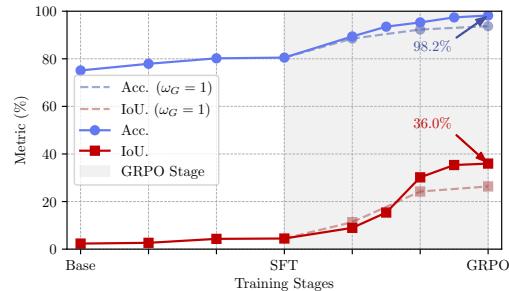


Figure 4: Accuracy and IoU curve of the FakeXplainer during the training process.

A valid output must include correctly structured `<think>`, `<tag>` and `<verdict>` markers, as well as bounding boxes and captions that are syntactically well-formed and can be parsed using regular expressions. Formally, the reward is defined as:

$$\mathcal{R}_F(o) = \begin{cases} 1, & \text{if } \{V, R, T, C\}(o) \text{ are parsable} \\ 0, & \text{o.w.} \end{cases}$$

where  $T(o)$  and  $C(o)$  extract the regional captions and image-level tags from  $o$ , respectively.

### 4.3 RLHF WITH GROUP RELATIVE POLICY OPTIMIZATION

Following SFT, we employ Group Relative Policy Optimization (GRPO) Shao et al. (2024) to progressively align the MLLM with our objectives of interpretable and reliable fake image detection. GRPO combines structured supervision from the dataset with targeted reward signals through a carefully designed training process. The reward function is formulated as:

$$\mathcal{R} = \omega_G(t)\mathcal{R}_G + \omega_C\mathcal{R}_C + \omega_F\mathcal{R}_F, \quad (1)$$

where weights  $\omega_C = \omega_F = 1.0$  remain constant throughout training, while  $\omega_G(t)$  increases linearly from 0.5 to 1.0 over the training process.

Our approach employs a continuous linear interpolation for the localization weight:

$$\omega_G(t) = 0.5 + 0.5 \cdot (t/T), \quad (2)$$

where  $t$  represents the current training step and  $T$  denotes the total number of training steps.

The linear weighting strategy addresses challenges observed in preliminary experiments and offers three benefits. First, without IoU weight adjustment, models trained with equal weights from the start tend to over-optimize localization rewards, producing many small fragmented bounding boxes that achieve high IoU scores but fail to capture meaningful regions. Down-weighting the localization reward in early training prevents this issue. Second, the progressive scheme enables a natural curriculum learning. With reduced localization weight at the beginning, the model first learns output formatting and classification accuracy. As these skills stabilize, the gradually increasing IoU reward improves localization on top of this foundation. Third, continuous weight adjustment avoids reward spikes and stabilizes training, allowing smooth adaptation of optimization objectives. Experiments show that linear reward weighting outperforms static schemes, confirming the effectiveness of gradual reward shaping for training MLLMs in complex visual understanding tasks.

## 5 EXPERIMENTS

### 5.1 EXPERIMENTAL SETUP

We adopt *Qwen2.5-VL-Instruct* (Qwen Team, 2025) as the base model of our FakeXplainer for its strong pre-trained grounding capabilities. We trained FakeXplainer on eight NVIDIA A100 GPUs. Both the SFT and GRPO stages last for three epochs, with a batch size of 1. We use  $\eta = 1.1$  and

324  
325  
326  
327 Table 1: Experimental result for current AI-generated image detectors and our **FakeXplainer** across  
328 different image generation methods.  
329  
330  
331  
332  
333  
334  
335  
336  
337  
338  
339  
340  
341  
342  
343  
344  
345  
346  
347  
348  
349  
350  
351  
352  
353  
354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377

Generators	FakeXplainer		ObjectFormer		SegFormer		NPR	DMD.	ComFor.	AfPr.	AEROB.	DIRE	FakeVLM
	Acc.	IoU	Acc.	IoU	Acc.	IoU	Acc.	Acc.	Acc.	Acc.	Acc.	Acc.	Acc.
DALL-E 2 2022	<b>0.986</b>	0.360	0.957	0.251	0.942	0.285	0.907	0.934	0.877	0.892	0.823	0.916	0.908
DALL-E 3 2023	<b>0.991</b>	0.365	0.949	0.258	0.950	0.292	0.912	0.942	0.872	0.907	0.821	0.923	0.915
DDIM 2020	<b>0.974</b>	0.345	0.954	0.285	0.945	0.280	0.917	0.928	0.879	0.915	0.839	0.912	0.902
DDPM 2020	<b>0.979</b>	0.350	0.951	0.293	0.947	0.288	0.903	0.931	0.876	0.898	0.836	0.917	0.906
FLUX.1-dev 2024	<b>0.988</b>	0.362	0.958	0.299	0.940	0.295	0.922	0.937	0.874	0.779	0.843	0.919	0.938
FLUX.1-schnell	<b>0.972</b>	0.343	0.953	0.287	0.943	0.283	0.926	0.929	0.882	0.805	0.827	0.913	0.941
GLIDE 2021	<b>0.970</b>	0.340	0.950	0.289	0.946	0.286	0.913	0.935	0.873	0.661	0.822	0.922	0.897
Midjourney v4 2023	<b>0.990</b>	0.364	0.956	0.296	0.949	0.294	0.908	0.939	0.869	0.878	0.814	0.925	0.926
Midjourney v5	<b>0.992</b>	0.366	0.959	0.273	0.941	0.297	0.902	0.943	0.871	0.851	0.718	0.927	0.932
SD 1.4 2022a	0.968	0.338	0.952	0.286	0.944	0.282	0.921	<b>0.970</b>	0.880	0.852	0.951	0.909	0.918
SD 1.5	<b>0.975</b>	0.347	0.955	0.294	0.948	0.290	0.916	0.949	0.875	0.866	0.966	0.915	0.921
SD 2.1 2022b	<b>0.980</b>	0.352	0.951	0.291	0.942	0.287	0.911	0.938	0.872	0.881	0.833	0.918	0.913
SD 3.5 Large 2024	<b>0.991</b>	0.365	0.954	0.294	0.945	0.293	0.904	0.944	0.870	0.934	0.830	0.924	0.929
SD 3.5 Large Turbo	<b>0.993</b>	0.368	0.957	0.312	0.950	0.296	0.906	0.947	0.868	0.927	0.837	0.928	0.935
VQDM 2022	<b>0.973</b>	0.342	0.953	0.288	0.943	0.284	0.927	0.932	0.877	0.938	0.932	0.914	0.909
<i>Diffusion</i>	<b>0.983</b>	0.356	0.954	0.287	0.945	0.290	0.913	0.941	0.874	0.864	0.842	0.920	0.919
BigGAN 2018	<b>0.965</b>	0.335	0.950	0.280	0.941	0.278	0.918	0.892	0.903	0.933	0.861	0.887	0.894
GALIP 2023	0.882	0.353	<b>0.941</b>	0.279	<b>0.941</b>	0.289	0.882	0.882	0.941	0.353	0.706	0.882	0.876
VQGAN 2021	<b>0.967</b>	0.337	0.954	0.282	0.943	0.280	0.907	0.889	0.908	0.921	0.932	0.885	0.858
StyleGAN-XL 2018	0.960	0.330	0.951	0.278	0.940	0.275	0.914	0.884	<b>0.980</b>	0.928	0.939	0.879	0.680
<i>GAN</i>	<b>0.955</b>	0.337	0.950	0.280	0.941	0.279	0.912	0.890	0.916	0.866	0.860	0.885	0.827
PixArtAlpha 2023	<b>0.987</b>	0.357	0.956	0.295	0.947	0.291	0.908	0.912	0.891	0.934	0.927	0.903	0.861
PixArtDelta 2024b	<b>0.984</b>	0.354	0.953	0.292	0.943	0.289	0.921	0.909	0.893	0.939	0.922	0.899	0.896
PixArtSigma 2024a	<b>0.989</b>	0.360	0.957	0.296	0.949	0.293	0.919	0.915	0.889	0.924	0.931	0.905	0.904
DiT 2023	<b>0.978</b>	0.349	0.952	0.290	0.942	0.287	0.913	0.907	0.896	0.928	0.938	0.897	0.892
<i>DiT</i>	<b>0.983</b>	0.354	0.954	0.293	0.945	0.289	0.914	0.910	0.893	0.931	0.931	0.900	0.889
VAR 2024	<b>0.976</b>	0.346	0.954	0.287	0.945	0.283	0.928	0.893	0.901	0.934	0.927	0.889	0.886
Infinity 2024	<b>0.974</b>	0.344	0.951	0.289	0.941	0.286	0.914	0.897	0.899	0.938	0.924	0.883	0.872
MaskGIT 2022	<b>0.972</b>	0.342	0.955	0.288	0.948	0.284	0.909	0.895	0.904	0.923	0.933	0.886	0.854
LlamaGen 2024	<b>0.980</b>	0.351	0.953	0.429	0.944	0.289	0.923	0.899	0.897	0.929	0.938	0.892	0.867
<i>Others</i>	<b>0.978</b>	0.348	0.953	0.369	0.944	0.287	0.920	<b>0.898</b>	0.898	0.931	0.933	0.889	0.870
Real Images 2009	<b>0.985</b>	-	0.956	-	0.946	-	0.918	0.903	0.882	0.934	0.854	0.896	0.763
Overall	<b>0.982</b>	0.360	0.954	0.299	0.945	0.289	0.914	0.928	0.882	0.887	0.873	0.911	0.828

set the number of samples to  $G = 4$  during the GRPO stage. All baseline methods are trained on one NVIDIA A100 GPU. More experimental details are provided in the Appendix B. For baseline comparisons, we use the same training data as the FakeXplainer setup. We train SegFormer (Xie et al., 2021) and ObjectFormer (Wang et al., 2022) under a segmentation + classification setting on the FakeXplained dataset by converting bounding boxes to binary masks. For classification-only methods, including NPR (Ojha et al., 2023), DMD (Corvi et al., 2023), ComFor (Park & Owens, 2024), AfPr (Chang et al., 2023), and DIRE (Wang et al., 2023), only image-level labels are used during training. We additionally evaluated state-of-the-art MLLM-based detection methods, specifically FakeShield (Xu et al., 2024), LEGION (Kang et al., 2025), and **FakeVLM** (Wen et al., 2025). For FakeShield, we utilized the pre-trained weights provided by the authors without further fine-tuning. For LEGION, we adhered to the training protocol specified in the original paper, training the model on the SynthScars dataset with identical hyperparameters and experimental configurations as reported.

## 5.2 OVERALL PERFORMANCE

To ensure robustness and mitigate dataset bias, all models are evaluated using four-fold cross-validation. During training, the detection model is exposed to 75% of images from the FakeXplained dataset along with an equal number of real samples. Evaluation is conducted on the remaining 25% of synthetic images, again paired with the same number of real images. We report both classification accuracy and localization performance using the IoU metric on AI-generated images. Robustness tests against perturbations are provided in Appendix C.

**Comparing to other methods.** Quantitative results are reported in Table 1, comparing **FakeXplainer** with traditional detectors. For MLLMs, Table 2 presents post-finetuning performance of LEGION (Kang et al., 2025) and FakeShield (Xu et al., 2024) across different pre-trained models. Our best-performing model achieves an overall classification accuracy of **98.2%**, demonstrating strong robustness and consistent performance across different image generators. For localization,

378 Table 2: Performance comparison of **FakeXplainer** across different base MLLMs and against other  
 379 MLLM-based methods. Post-finetuning results are underlined.  
 380

Method	FakeXplainer						FakeShield	LEGION
	Backbone	InternVL3-8B	InternVL3-14B	Ovis2.5-9B	MiMo-VL-7B-RL	Qwen-2.5-VL-32B		
Acc.	0.584	<u>0.928</u>	0.568	<u>0.951</u>	0.624	<u>0.909</u>	0.515	0.920
IoU	0.039	<u>0.134</u>	0.043	<u>0.289</u>	-	-	-	<u>0.044</u>
BLEU-2	0.061	<u>0.232</u>	0.098	<u>0.235</u>	0.058	<u>0.203</u>	0.083	<u>0.249</u>
ROUGE-L	0.059	<u>0.225</u>	0.092	<u>0.219</u>	0.050	<u>0.184</u>	0.076	<u>0.239</u>
							<u>0.076</u>	<u>0.251</u>

386  
 387 Table 3: Accuracy on external datasets for out-of-distribution generalization testing.  
 388

Sources	FXP.	ObjFormer.	SegFormer	NPR	DMD.	ComFor.	AfPr.	AEROB.	DIRE	FakeShield	LEGION
FakeClue 2025	<b>0.852</b>	0.462	0.485	0.833	0.734	0.766	0.849	0.239	0.727	0.550	0.172
Chameleon 2024	<b>0.843</b>	0.485	0.508	0.794	0.721	0.757	0.803	0.291	0.752	0.587	0.197
GPT-Image-1 2025	<b>0.801</b>	0.513	0.538	0.790	0.735	0.636	0.597	0.458	0.793	0.752	0.238
FaceForensics++ 2019	<b>0.864</b>	0.598	0.716	0.861	0.562	0.429	0.746	0.681	0.850	0.773	0.395
MMFR-Dataset 2025	<b>0.874</b>	0.653	0.657	0.569	0.619	0.595	0.786	0.685	0.624	0.710	0.193

394  
 395 the model achieves an IoU score of **36.0%**, outperforming all segmentation-based baselines. This  
 396 indicates that **FakeXplainer** identifies fake regions more consistently with human annotations than  
 397 competing approaches.  
 398

399 **Reasoning Quality.** Table 2 shows the BLEU-2 and ROUGE-L (Li et al., 2024) of model responses  
 400 against the FakeXplained dataset. The results verify the training effectiveness of FakeXplainer, con-  
 401 siderably outperforming the base model in explanation generation, indicating that both the regions  
 402 and their reasons are generated accurately.

403 **Generalizability of FakeXplainer on other MLLMs.** To assess generalizability beyond Qwen-2.5-  
 404 VL, we further evaluated it on several state-of-the-art MLLMs with diverse architectures and capa-  
 405 bilities in Table 2. The consistent performance gain of FakeXplainer across architectures, whether  
 406 with grounding capabilities (InternVL3 (Zhu et al., 2025)) or without (Ovis2.5 (Lu et al., 2024),  
 407 MiMo (Xiaomi, 2025)), validates the model-agnostic nature of our pipeline.

408 **Out-of-distribution (OoD) evaluation.** We also evaluated the models on five OoD datasets, **Fake-**  
 409 **Clue** (Wen et al., 2025), **Chameleon** (Yan et al., 2024), FaceForensics++ (Rössler et al., 2019),  
 410 images generated by GPT-Image-1 (OpenAI, 2025; Rapidata, 2025) and MMFR-Dataset (eval) pro-  
 411 posed by FakeReasoning (Gao et al., 2025). As shown in Table 3, our model consistently outper-  
 412 forms all other methods across OoD datasets, demonstrating considerable generalization to unseen  
 413 image domains.

414 **Qualitative evaluation.** Figure 5 shows two samples from the FakeXplained test set. We found  
 415 that our model prefers outlining smaller regions than human annotators, demonstrating fine-grained  
 416 localization capability. Compared with LEGION and FakeShield, **FakeXplainer** not only provides  
 417 correct predictions but also delivers reliable, grounded explanations. More qualitative examples will  
 418 be provided in the Appendix A.3.

419 **Human preference evaluation.** While IoU and classification accuracy provide objective metrics for  
 420 detection performance, they do not fully capture the qualitative aspects of region-caption alignment.  
 421 In fact, the model may, in some instances, generate annotations that surpass those of the original  
 422 human annotators. To comprehensively assess the quality and relevance of the generated explana-  
 423 tions, we conducted a human preference study involving an independent group of evaluators. In this  
 424 study, participants were shown pairs of outputs for the same image, each with different bounding  
 425 box annotations and associated captions. With no metadata given, evaluators were asked to choose  
 426 the annotation that demonstrated better alignment between the region and caption, as well as higher  
 427 overall quality. If no clear preference emerged, a neutral option was available.

428 We received 1,525 non-neutral preference votes. In Figure 3, the “Humans” category represents  
 429 annotations from the FakeXplained dataset. When compared with FakeXplainer, human annota-  
 430 tions were preferred in 52.9% of cases, indicating that FakeXplainer achieves near-parity with hu-  
 431 man annotators in producing region-grounded explanations, demonstrating the effectiveness of our  
 framework in generating high-quality visual-textual reasoning.

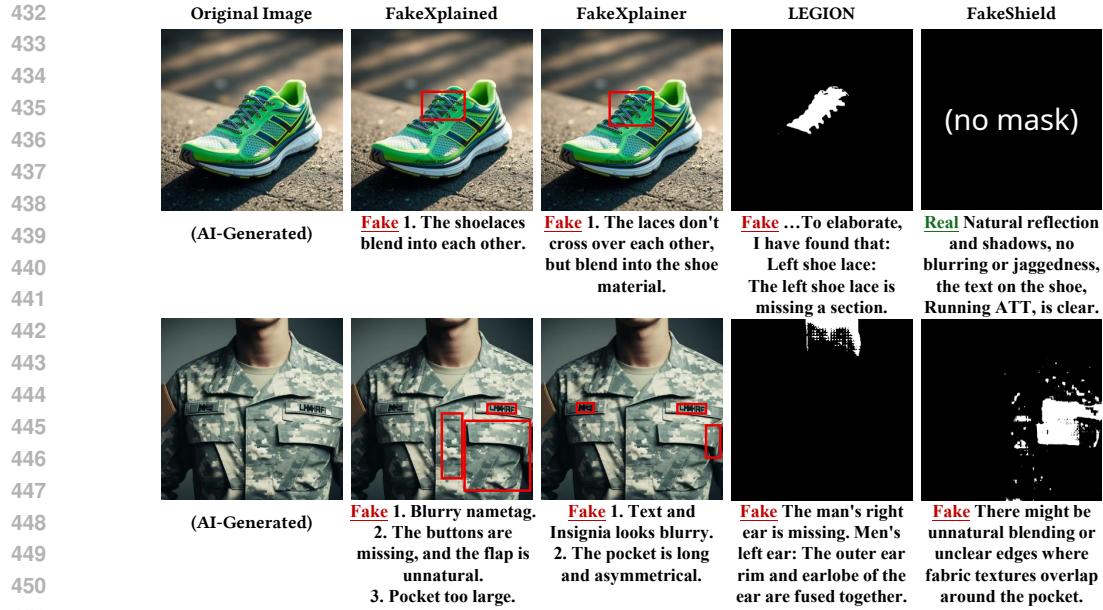


Figure 5: Comparison of responses visualized from the **FakeXplainer** method, the ground truth from the **FakeXplained** dataset, and **LEGION** (Kang et al., 2025) and **FakeShield** (Xu et al., 2024).

Table 4: Performance comparison of FakeXplainer under different training configurations.

Metric	FakeXplainer			No-FT		Training Strategy (32B)			Partial Data (32B)		
	3B	7B	32B	32B	SFT	GRPO $_{\omega_G=1}$	GRPO $_{\omega_G=0.5}$	no-bbox	no-caption	no-tags	label-only
Acc.	0.842	0.958	<b>0.982</b>	0.734	0.893	0.937	0.974	0.952	0.942	0.962	0.937
IoU	0.185	0.255	<b>0.360</b>	0.044	0.043	0.265	0.223	-	0.265	0.358	-
BLEU-2	0.195	0.246	<b>0.267</b>	0.080	0.183	0.257	0.261	0.164	-	0.243	-
ROUGE-L	0.121	0.218	<b>0.251</b>	0.076	0.174	0.239	0.242	0.160	-	0.237	-

### 5.3 ABLATION STUDIES

We ablate each training component, including each data component and each training segment, and report the results in Table 4. Additional ablation results are provided in Appendix D.

**Model size of MLLM.** Model size has a clear impact on the performance of FakeXplainer. With the same training pipeline, the 7B variants can identify the authenticity of the image in 95.8% of the cases, but the 3B variant fails to surpass most traditional methods in detection accuracy and cannot effectively localize fake regions. The 7B variant also outputs a good rationale according to the BLEU and ROUGE-L metrics, making it a good balance between performance and speed.

**Effects of different training stages.** To analyze the impact of each training stage, we report both accuracy and IoU metrics across the entire training process in Figure 4. Without GRPO, SFT alone yields marginal improvements over the base model, especially in localization. The GRPO stage with a constant  $\omega_G = 1$  has a higher IoU than the linear scheme, but struggles to train effectively in later steps, demonstrating the cumulative benefit of the progressive reward design. By the completion of the RLHF stage, the model reaches an accuracy of 98.2% and an IoU of 36.0%.

**Fine-tuning impact.** Without fine-tuning, Qwen-2.5-VL-32B-Instruct achieves only 73.4% accuracy. SFT improves this to 89.3%, and adding GRPO further increases the performance to 98.2%, demonstrating the critical role of our two-stage pipeline and the FakeXplained dataset.

**Data components.** We evaluate three data components: image tags, region annotations (bounding boxes + captions), and binary labels. Using only binary labels yields 93.7% accuracy—the lowest among partial variants but still exceeding DMD’s 92.8%. Removing bounding boxes or captions reduces accuracy by 3.5%, with caption removal severely impacting IoU (-9.5%). While tag removal

486 has a minimal effect on both metrics. These results confirm that structured reasoning information,  
 487 particularly region-level annotations, substantially improves detection performance.  
 488

489 **Training strategies.** Fixed reward weighting ( $\text{GRPO}_{\omega_G=1}$  and  $\text{GRPO}_{\omega_G=0.5}$ ) underperforms our  
 490 progressive GRPO approach across all metrics. Notably, the localization-prioritized  $\text{GRPO}_{\omega_G=1}$   
 491 also shows inferior IoU, validating the necessity of textual explanations and dynamic reward weight-  
 492 ing for the step-by-step acquisition of classification, localization skills, and overall interpretability.  
 493

## 494 6 CONCLUSION

495 Our research presents FakeXplainer, an explainable AI-generated image detection approach util-  
 496 izing MLLMs that provides grounded, human-interpretable explanations alongside detection re-  
 497 sults. The system achieves strong performance metrics (98.2% accuracy, 36.0% IoU) through a  
 498 progressive training pipeline, establishing a foundation for transparent visual media authenti-  
 499 cation. [Although grounded explainability improves FakeXplainer’s generalization, its dependence](#)  
 500 [on human-perceptible artifacts reflects an inherent limitation shared by all current explainable AIGI](#)  
 501 [detectors.](#) Fully realistic synthetic images without semantically describable flaws represent a fun-  
 502 damentally different detection problem and fall outside the scope of this work. At the same time,  
 503 our results show that grounded, human-aligned reasoning provides clear advantages over black-box  
 504 classifiers: models that output only a binary label exhibit substantial degradation under distribu-  
 505 tion shift, whereas FakeXplainer maintains strong performance across multiple out-of-distribution  
 506 benchmarks. These findings suggest that explicit localization and human-aligned supervision offer  
 507 a more robust and verifiable signal than classification alone.

## 508 510 REFERENCES

511 Xiuli Bi, Bo Liu, Fan Yang, Bin Xiao, Weisheng Li, Gao Huang, and Pamela C. Cosman. Detect-  
 512 ing generated images by real images only, 2023. URL <https://arxiv.org/abs/2311.00962>.

514 Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural  
 515 image synthesis. In *International Conference on Learning Representations*, 2018.

517 Clement Chadebec, Onur Tasar, Eyal Benaroche, and Benjamin Aubin. Flash diffusion: Accelerat-  
 518 ing any conditional diffusion model for few steps image generation. In *Proceedings of the AAAI*  
 519 *Conference on Artificial Intelligence*, 2025.

520 Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative  
 521 image transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*  
 522 *recognition*, pp. 11315–11325, 2022.

524 You-Ming Chang, Chen Yeh, Wei-Chen Chiu, and Ning Yu. Antifakeprompt: Prompt-tuned vision-  
 525 language models are fake image detectors. *arXiv preprint arXiv:2310.17419*, 2023.

526 Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Yue Wu, Zhongdao Wang, James  
 527 Kwok, Ping Luo, Huchuan Lu, et al. Pixart- $\alpha$ : Fast training of diffusion transformer for photore-  
 528 alistic text-to-image synthesis. *arXiv preprint arXiv:2310.00426*, 2023.

530 Junsong Chen, Chongjian Ge, Enze Xie, Yue Wu, Lewei Yao, Xiaozhe Ren, Zhongdao Wang, Ping  
 531 Luo, Huchuan Lu, and Zhenguo Li. Pixart- $\sigma$ : Weak-to-strong training of diffusion transformer for  
 532 4k text-to-image generation. In *European Conference on Computer Vision*, pp. 74–91. Springer,  
 533 2024a.

534 Junsong Chen, Yue Wu, Simian Luo, Enze Xie, Sayak Paul, Ping Luo, Hang Zhao, and Zhenguo Li.  
 535 Pixart- $\delta$ : Fast and controllable image generation with latent consistency models. *arXiv preprint*  
 536 *arXiv:2401.05252*, 2024b.

538 Liang Chen, Lei Li, Haozhe Zhao, Yifan Song, and Vinci. R1-v: Reinforcing super generalization  
 539 ability in vision-language models with less than \$3. <https://github.com/Deep-Agent/R1-V>, 2025. Accessed: 2025-02-02.

540 Zhe Chen, Jiannan Wu, Wenhui Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong  
 541 Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning  
 542 for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF conference on computer*  
 543 *vision and pattern recognition*, 2024c.

544 Riccardo Corvi, Davide Cozzolino, Giada Zingarini, Giovanni Poggi, Koki Nagano, and Luisa Ver-  
 545 doliva. On the detection of synthetic images generated by diffusion models. In *IEEE Interna-  
 546 tional Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, 2023. doi:  
 547 10.1109/ICASSP49357.2023.10095167.

548 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, K. Li, and Li Fei-Fei. Imagenet: A large-scale hier-  
 549 archical image database. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp.  
 550 248–255, 2009. URL <https://api.semanticscholar.org/CorpusID:57246310>.

552 Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image  
 553 synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recogni-  
 554 tion*, pp. 12873–12883, 2021.

555 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam  
 556 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion En-  
 557 glish, Kyle Lacey, Alex Goodwin, Yannik Marek, and Robin Rombach. Scaling rectified flow  
 558 transformers for high-resolution image synthesis, 2024. URL <https://arxiv.org/abs/2403.03206>.

559 Wenlong Fang, Qiaofeng Wu, Jing Chen, and Yun Xue. guided mllm reasoning: Enhancing mllm  
 560 with knowledge and visual notes for visual question answering. In *Proceedings of the Computer  
 Vision and Pattern Recognition Conference*, 2025.

561 Yueying Gao, Dongliang Chang, Bingyao Yu, Haotian Qin, Lei Chen, Kongming Liang, and Zhanyu  
 562 Ma. Fakereasoning: Towards generalizable forgery detection and reasoning. *arXiv preprint  
 563 arXiv:2503.21210*, 2025.

564 Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,  
 565 Aaron Courville, and Y. Bengio. Generative adversarial networks. *Advances in Neural Infor-  
 566 mation Processing Systems*, 3, 06 2014. doi: 10.1145/3422622.

567 Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and  
 568 Baining Guo. Vector quantized diffusion model for text-to-image synthesis. In *Proceedings of  
 569 the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10696–10706, 2022.

570 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,  
 571 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms  
 572 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

573 Jian Han, Jinlai Liu, Yi Jiang, Bin Yan, Yuqi Zhang, Zehuan Yuan, Bingyue Peng, and Xiaobing  
 574 Liu. Infinity: Scaling bitwise autoregressive modeling for high-resolution image synthesis. *arXiv  
 575 preprint arXiv:2412.04431*, 2024.

576 Jonathan Ho, Ajay Jain, and P. Abbeel. Denoising diffusion probabilistic models. *ArXiv*,  
 577 abs/2006.11239, 2020. URL <https://api.semanticscholar.org/CorpusID:219955663>.

578 Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and  
 579 Shaohui Lin. Vision-r1: Incentivizing reasoning capability in multimodal large language models.  
 580 *arXiv preprint arXiv:2503.06749*, 2025a.

581 Zhenglin Huang, Tianxiao Li, Xiangtai Li, Haiquan Wen, Yiwei He, Jiangning Zhang, Hao Fei,  
 582 Xi Yang, Xiaowei Huang, Bei Peng, et al. So-fake: Benchmarking and explaining social media  
 583 image forgery detection. *arXiv preprint arXiv:2505.18660*, 2025b.

584 Yikun Ji, Yan Hong, Jiahui Zhan, Haoxing Chen, jun lan, Huijia Zhu, Weiqiang Wang, Liqing  
 585 Zhang, and Jianfu Zhang. Towards explainable fake image detection with multi-modal large  
 586 language models, 2025. URL <https://arxiv.org/abs/2504.14245>.

594 Hengrui Kang, Siwei Wen, Zichen Wen, Junyan Ye, Weijia Li, Peilin Feng, Baichuan Zhou, Bin  
 595 Wang, Dahua Lin, Linfeng Zhang, et al. Legion: Learning to ground and explain for synthetic  
 596 image detection. *arXiv preprint arXiv:2503.15264*, 2025.

597 Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative  
 598 adversarial networks. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recog-  
 599 nition (CVPR)*, pp. 4396–4405, 2018. URL <https://api.semanticscholar.org/CorpusID:54482423>.

600 Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete  
 601 Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceed-  
 602 ings of the IEEE/CVF international conference on computer vision*, pp. 4015–4026, 2023.

603 Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024.

604 Duong H Le, Tuan Pham, Sangho Lee, Christopher Clark, Aniruddha Kembhavi, Stephan Mandt,  
 605 Ranjay Krishna, and Jiasen Lu. One diffusion to generate them all. In *Proceedings of the Com-  
 606 puter Vision and Pattern Recognition Conference*, 2025.

607 Jia Li, Lijie Hu, Jingfeng Zhang, Tianhang Zheng, Hua Zhang, and Di Wang. Fair text-to-image  
 608 diffusion via fair mapping. In *Proceedings of the AAAI Conference on Artificial Intelligence*,  
 609 2025a.

610 Yixuan Li, Xuelin Liu, Xiaoyang Wang, Shiqi Wang, and Weisi Lin. Fakelbench: Uncover the  
 611 achilles' heels of fake images with large multimodal models. *ArXiv*, abs/2404.13306, 2024. URL  
 612 <https://api.semanticscholar.org/CorpusID:269293612>.

613 Yixuan Li, Yu Tian, Yipo Huang, Wei Lu, Shiqi Wang, Weisi Lin, and Anderson Rocha. Fakescope:  
 614 Large multimodal expert model for transparent ai-generated image forensics. *arXiv preprint  
 615 arXiv:2503.24267*, 2025b.

616 Jiawei Liu, Fanrui Zhang, Jiaying Zhu, Esther Sun, Qiang Zhang, and Zheng-Jun Zha. Forgerygpt:  
 617 Multimodal large language model for explainable image forgery detection and localization. *arXiv  
 618 preprint arXiv:2410.10238*, 2024.

619 Peter Lorenz, Ricard Durall, and Janis Keuper. Detecting images generated by deep diffusion models  
 620 using their local intrinsic dimensionality. *2023 IEEE/CVF International Conference on Computer  
 621 Vision Workshops (ICCVW)*, pp. 448–459, 2023. URL <https://api.semanticscholar.org/CorpusID:259342331>.

622 Shiyin Lu, Yang Li, Qing-Guo Chen, Zhao Xu, Weihua Luo, Kaifu Zhang, and Han-Jia Ye. Ovis:  
 623 Structural embedding alignment for multimodal large language model. *arXiv:2405.20797*, 2024.

624 Midjourney. Midjourney, 2023. URL <https://www.midjourney.com/home/>.

625 Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew,  
 626 Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with  
 627 text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021.

628 Utkarsh Ojha, Yuheng Li, and Yong Jae Lee. Towards universal fake image detectors that generalize  
 629 across generative models. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recog-  
 630 nition (CVPR)*, pp. 24480–24489, 2023. URL <https://api.semanticscholar.org/CorpusID:257038440>.

631 OpenAI. Dall-e 3 system card, 2023. URL [https://cdn.openai.com/papers/DALL\\_E\\_3\\_System\\_Card.pdf](https://cdn.openai.com/papers/DALL_E_3_System_Card.pdf).

632 OpenAI. Introducing 4o image generation, Mar 2025. URL <https://openai.com/index/introducing-4o-image-generation/>.

633 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong  
 634 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to fol-  
 635 low instructions with human feedback. *Advances in neural information processing systems*, 35:  
 636 27730–27744, 2022.

648 Jeongsoo Park and Andrew Owens. Community forensics: Using thousands of generators to train  
 649 fake image detectors, 2024. URL <https://arxiv.org/abs/2411.04125>.  
 650

651 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of*  
 652 *the IEEE/CVF international conference on computer vision*, pp. 4195–4205, 2023.

653 Qwen Team. Qwen2.5-vl, January 2025. URL <https://qwenlm.github.io/blog/qwen2.5-vl/>.  
 654

655 Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-  
 656 conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.

657

658 Rapidata. Rapidata openai 4o preference, Mar 2025. URL [https://huggingface.co/datasets/Rapidata/OpenAI-4o\\_t2i\\_human\\_preference](https://huggingface.co/datasets/Rapidata/OpenAI-4o_t2i_human_preference).  
 659

660 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-  
 661 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Con-*  
 662 *ference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10684–10695, June 2022a.

663

664 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-  
 665 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Con-*  
 666 *ference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10684–10695, June 2022b.

667

668 Andreas Rössler, Davide Cozzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias  
 669 Nießner. FaceForensics++: Learning to detect manipulated facial images. In *International Con-*  
 670 *ference on Computer Vision (ICCV)*, 2019.

671

672 Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh,  
 673 and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based local-  
 674 ization. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 618–626, 2017.  
 675 doi: 10.1109/ICCV.2017.74.

676

677 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Jun-Mei Song, Mingchuan Zhang, Y. K. Li,  
 678 Yu Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open  
 679 language models. *ArXiv*, abs/2402.03300, 2024. URL <https://api.semanticscholar.org/CorpusID:267412607>.

680

681 Haozhan Shen, Peng Liu, Jingcheng Li, Chunxin Fang, Yibo Ma, Jiajia Liao, Qiaoli Shen, Zilun  
 682 Zhang, Kangjia Zhao, Qianqian Zhang, Ruochen Xu, and Tiancheng Zhao. Vlm-r1: A stable and  
 683 generalizable r1-style large vision-language model. *arXiv preprint arXiv:2504.07615*, 2025.

684

685 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *ArXiv*,  
 686 abs/2010.02502, 2020. URL <https://api.semanticscholar.org/CorpusID:222140788>.

687

688 Peize Sun, Yi Jiang, Shoufa Chen, Shilong Zhang, Bingyue Peng, Ping Luo, and Zehuan Yuan.  
 689 Autoregressive model beats diffusion: Llama for scalable image generation. *arXiv preprint*  
 690 *arXiv:2406.06525*, 2024.

691

692 Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural  
 693 networks. *ArXiv*, abs/1905.11946, 2019. URL <https://api.semanticscholar.org/CorpusID:167217261>.

694

695 Ming Tao, Bing-Kun Bao, Hao Tang, and Changsheng Xu. Galip: Generative adversarial clips for  
 696 text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*  
 697 *Pattern Recognition*, pp. 14214–14223, 2023.

698

699 Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling:  
 700 Scalable image generation via next-scale prediction. *Advances in neural information processing*  
 701 *systems*, 37:84839–84865, 2024.

702

703 Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in*  
 704 *neural information processing systems*, 30, 2017.

702 Haozhe Wang, Chao Qu, Zuming Huang, Wei Chu, Fangzhen Lin, and Wenhui Chen. Vi-  
 703 rethinker: Incentivizing self-reflection of vision-language models with reinforcement learning.  
 704 *arXiv preprint arXiv:2504.08837*, 2025a.

705 Junke Wang, Zuxuan Wu, Jingjing Chen, Xintong Han, Abhinav Shrivastava, Ser-Nam Lim, and  
 706 Yu-Gang Jiang. Objectformer for image manipulation detection and localization. In *Proceedings*  
 707 *of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.

708 Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. Cnn-generated  
 709 images are surprisingly easy to spot... for now. In *Proceedings of the IEEE/CVF conference on*  
 710 *computer vision and pattern recognition*, pp. 8695–8704, 2020.

711 Zhendong Wang, Jianmin Bao, Wengang Zhou, Weilun Wang, Hezhen Hu, Hong Chen, and  
 712 Houqiang Li. Dire for diffusion-generated image detection. In *Proceedings of the IEEE/CVF*  
 713 *International Conference on Computer Vision*, pp. 22445–22455, 2023.

714 Zhendong Wang, Jianmin Bao, Shuyang Gu, Dong Chen, Wengang Zhou, and Houqiang Li. Design-  
 715 diffusion: High-quality text-to-design image generation with diffusion models. In *Proceedings of*  
 716 *the Computer Vision and Pattern Recognition Conference*, 2025b.

717 Siwei Wen, Junyan Ye, Peilin Feng, Hengrui Kang, Zichen Wen, Yize Chen, Jiang Wu, Wenjun  
 718 Wu, Conghui He, and Weijia Li. Spot the fake: Large multimodal model-based synthetic image  
 719 detection with artifact explanation. *arXiv preprint arXiv:2503.14905*, 2025.

720 Mengyang Wu, Yuzhi Zhao, Jialun Cao, Mingjie Xu, Zhongming Jiang, Xuehui Wang, Qinbin Li,  
 721 Guangneng Hu, Shengchao Qin, and Chi-Wing Fu. Icm-assistant: Instruction-tuning multimodal  
 722 large language models for rule-based explainable image content moderation. In *Proceedings of*  
 723 *the AAAI Conference on Artificial Intelligence*, 2025a.

724 Shengqiong Wu, Hao Fei, Liangming Pan, William Yang Wang, Shuicheng Yan, and Tat-Seng Chua.  
 725 Combating multimodal llm hallucination via bottom-up holistic reasoning. In *Proceedings of the*  
 726 *AAAI Conference on Artificial Intelligence*, 2025b.

727 Wenyi Xiao, Leilei Gan, Weilong Dai, Wanggui He, Ziwei Huang, Haoyuan Li, Fangxun Shu,  
 728 Zhelun Yu, Peng Zhang, Hao Jiang, et al. Fast-slow thinking for large vision-language model  
 729 reasoning. *arXiv preprint arXiv:2504.18458*, 2025.

730 LLM-Core-Team Xiaomi. Mimo-vl technical report, 2025. URL <https://arxiv.org/abs/2506.03569>.

731 Enze Xie, Wenhui Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Seg-  
 732 former: Simple and efficient design for semantic segmentation with transformers. In *Neural*  
 733 *Information Processing Systems (NeurIPS)*, 2021.

734 Zhipei Xu, Xuanyu Zhang, Runyi Li, Zecheng Tang, Qing Huang, and Jian Zhang. Fakeshield:  
 735 Explainable image forgery detection and localization via multi-modal large language models.  
 736 *arXiv preprint arXiv:2410.02761*, 2024.

737 Shilin Yan, Ouxiang Li, Jiayin Cai, Yanbin Hao, Xiaolong Jiang, Yao Hu, and Weidi Xie. A sanity  
 738 check for ai-generated image detection. *arXiv preprint arXiv:2406.19435*, 2024.

739 Fan Yang, Ru Zhen, Jianing Wang, Yanhao Zhang, Haoxiang Chen, Haonan Lu, Sicheng Zhao, and  
 740 Guiguang Ding. Heie: Mllm-based hierarchical explainable aigc image implausibility evaluator.  
 741 In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025a.

742 Yi Yang, Xiaoxuan He, Hongkun Pan, Xiyan Jiang, Yan Deng, Xingtao Yang, Haoyu Lu, Dacheng  
 743 Yin, Fengyun Rao, Minfeng Zhu, et al. R1-onevision: Advancing generalized multimodal rea-  
 744 soning through cross-modal formalization. *arXiv preprint arXiv:2503.10615*, 2025b.

745 Zilyu Ye, Zhiyang Chen, Tiancheng Li, Zemin Huang, Weijian Luo, and Guo-Jun Qi. Schedule on  
 746 the fly: Diffusion time prediction for faster and better image generation. In *Proceedings of the*  
 747 *Computer Vision and Pattern Recognition Conference*, 2025.

756 Zicheng Zhang, Haoning Wu, Chunyi Li, Yingjie Zhou, Wei Sun, Xiongkuo Min, Zijian Chen,  
 757 Xiaohong Liu, Weisi Lin, and Guangtao Zhai. A-bench: Are Imms masters at evaluating ai-  
 758 generated images? *arXiv preprint arXiv:2406.03070*, 2024.

759 Yuze Zhao, Jintao Huang, Jinghan Hu, Xingjun Wang, Yunlin Mao, Daoze Zhang, Zeyinzi Jiang,  
 760 Zhikai Wu, Baole Ai, Ang Wang, et al. Swift: a scalable lightweight infrastructure for fine-tuning.  
 761 In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 29733–29735,  
 762 2025.

763 Ziyin Zhou, Yunpeng Luo, Yuanchen Wu, Ke Sun, Jiayi Ji, Ke Yan, Shouhong Ding, Xi-  
 764 aoshuai Sun, Yunsheng Wu, and Rongrong Ji. Aigi-holmes: Towards explainable and gen-  
 765 eralizable ai-generated image detection via multimodal large language models. *arXiv preprint*  
 766 *arXiv:2507.02664*, 2025.

767 Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen  
 768 Duan, Weijie Su, Jie Shao, et al. Internvl3: Exploring advanced training and test-time recipes for  
 769 open-source multimodal models. *arXiv preprint arXiv:2504.10479*, 2025.

## 772 A APPENDIX FOR FAKEXPLAINED

773 The FakeXplained dataset contains 8,772 high-quality AI-generated images annotated with fine-  
 774 grained bounding boxes and descriptive captions that highlight synthesis artifacts and logical in-  
 775 consistencies. This dataset addresses the critical gap in explainable and spatially grounded AI-generated  
 776 image detection.

### 779 A.1 SOURCE OF IMAGES

780 To ensure content diversity, we generated images across 1,000 ImageNet categories using 28 differ-  
 781 ent text-to-image generation models, one image per class per model. The prompt used is:

782 “A realistic image of {class\\_name}”

783 The models used are listed as follows:

784 **Diffusion-Based Generators:** Midjourney Midjourney (2023), Stable Diffusion models Rom-  
 785 bach et al. (2022b); Esser et al. (2024), DDIM Song et al. (2020), DDPM Ho et al. (2020), DALLE  
 786 OpenAI (2023), GLIDE Nichol et al. (2021), and VQDM Gu et al. (2022).

787 **GAN-Based Generators:** GALIP Tao et al. (2023), StyleGAN Karras et al. (2018), VQGAN  
 788 Esser et al. (2021), and BigGAN Brock et al. (2018).

789 **DiT-Based Generators:** PixArt Chen et al. (2023; 2024b;a) and DiT Peebles & Xie (2023).

790 **Other Generators:** VAR Tian et al. (2024), Infinity Han et al. (2024), MaskGIT Chang et al.  
 791 (2022), and LlamaGen Sun et al. (2024).

792 Most generated images are at a resolution of  $512 \times 512$ . For methods that do not natively support this  
 793 resolution, the  $1024 \times 1024$  resolution is used, and the images are downsampled to  $512 \times 512$  before  
 794 entering annotation stage. Generated images went through manual screening to remove images  
 795 that can be easily identified as AI-generated. The filter process is conducted manually to filter out  
 796 low-quality images that depict non-identifiable objects or if no part of the image is real.

### 803 A.2 ANNOTATION PROCESS

804 We recruited a team of experts who are trained to identify fake regions accurately. All of them  
 805 have prior experience in photography, have seen AI-generated images before, possess a fundamental  
 806 photographic literacy understanding, and are familiar with related concepts such as “saturation,”  
 807 “shadow,” “perspective,” and “noise.” During the training process, we provide all annotators with a  
 808 detailed instruction handout with examples. The handout contains positive and negative examples  
 809 for each global tag, a detailed bounding box annotation guideline, along with quality control metrics.

810     **Instructions on fake regions.** The rule for annotating fake regions is, if through observation of  
 811     the selected regions of interest, humans should be able to clearly determine that the image is not an  
 812     authentic photograph. Fake regions primarily show objects that do not follow the natural physics  
 813     laws or contradict common sense. Common image generation artifacts are also encouraged to be  
 814     annotated. After selecting a local area in the image, it is necessary to describe the reason for identi-  
 815     fying it as a generated image. The descriptive sentence must start with a noun, followed by one or  
 816     several adjective phrases or short clauses, and must exclusively describe content that appears in the  
 817     region.

818  
 819     **Definition of tags.** We refer to the most prominent depicted object in the non-background portion  
 820     of the image as the *image subject*. There are exactly five different tags that annotators can attach to  
 821     an image. Their definitions are listed as follows:

- 823     • **Perspective errors:** Indicates that the image has an unnatural viewing angle, or errors  
 824     in perspective, vanishing points. Incorrect occlusion and shadow errors do not constitute  
 825     perspective errors, but can be considered as fake regions instead.
- 826     • **Artistic styles:** If the overall image presents any artistic style, including but not limited  
 827     to oil painting, ink painting, or manga style, then select the “Artistic Style” tag. If only a  
 828     certain part of the image contains content in an artistic style, this tag should not be selected.
- 829  
 830     • **Unknown objects:** Indicates that the *subject* of the image does not exist in the world,  
 831     or is obviously unreasonable. There may be unusual insects and furniture with strong  
 832     design elements. Judgment should be based on intuition; unfamiliar or rare subjects do not  
 833     necessarily indicate unreasonable or non-existent objects.
- 834     • **Structure/attribute errors:** Indicates that the **subject** of the image has a structure that is  
 835     inconsistent with common knowledge, or has attributes inconsistent with common knowl-  
 836     edge. Examples include green flower petals, pink elephants, bent iron spoon handles, hu-  
 837     mans with more than two legs, and asymmetrical shapes. For erroneous attributes that only  
 838     occupy a small portion of the image subject, such as an incorrect number of fingers on a  
 839     human hand, fake regions should be marked as well.
- 840     • **Texture errors:** If obvious texture errors appear in the image, this tag needs to be selected.  
 841     For example, the texture of the entire image is blurry, or a portion of an object has a repet-  
 842     itive, tilted, or distorted texture. Unreadable text does not qualify as a texture error and  
 843     should be labeled as a fake region instead. If “Artistic Style” has already been marked, this  
 844     tag is usually omitted.
- 845     • **Other anomalies:** If there are very obvious global errors in the image that do not belong  
 846     to any of the above categories, check this item. This tag can also be marked even if other  
 847     tags have already been chosen.

849  
 850     **Keywords in Annotations** We analyze the captions of the bounding boxes to find the most fre-  
 851     quent phrases. Their occurrences are shown in Figure 6. Since we filter for the highest-quality im-  
 852     ages, it is hard to find a deciding bounding box for some cases. The contours and depth of field are  
 853     more likely to give the image away, leading to a high frequency of related captions. FakeXplainer  
 854     manages to align with most of the FakeXplained traits, with a higher detection rate for abnormal  
 855     object textures.

856  
 857     **Quality Check** To ensure quality, we cross-referenced the annotators’ proposals against control  
 858     samples we marked ourselves. We enforced a rejection policy where an average IoU under 20%  
 859     or a tag agreement lower than  $\frac{1}{3}$  will result in the exclusion of all previous data provided by the  
 860     annotator.

861     Figure 7 qualitatively demonstrates the 20% IoU criteria of the screening process. The criteria  
 862     is selected to accommodate individual understandings over AI-generated images while preventing  
 863     obvious fake regions from being omitted without notice. When the annotation process finishes, the  
 864     annotations and QC samples reached an IoU of 42.35% and the tag-agreement rate was 79.67%.

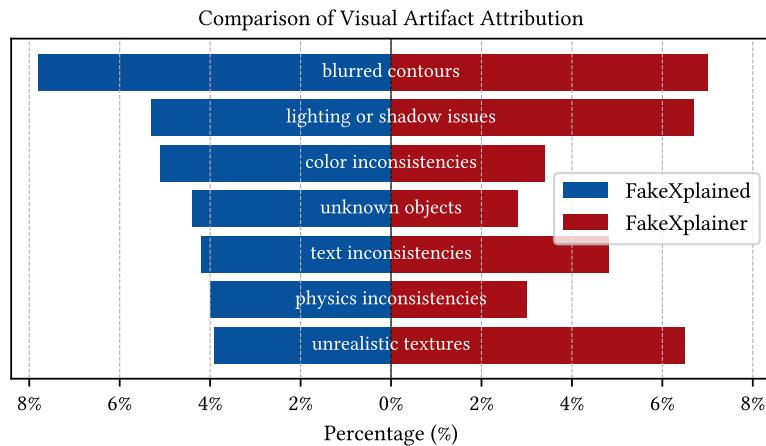


Figure 6: Keyword analysis for FakeXplained and the FakeXplainer responses.



Figure 7: Samples from the quality check process. When our regional annotation has a lower than 20% IoU with the annotation proposal, we reject the annotation.

### A.3 MORE SAMPLES

Figure 9 presents more annotated AI-generated images from FakeXplained. The left column displays the human annotations of FakeXplained. The right column shows the inference results of our best model, FakeXplainer. The center bar indicates the proportion of human preference votes from our user study. Note that since the “neutral” option was allowed, although the third annotated image received 46.2% of the votes, the human annotator is still rated higher than our model response. Our model demonstrates the ability to generate clearer, more descriptive captions for fake regions and reliably identifies content that contradicts common sense. For instance, in the lock-and-keyhole example (row 5), the model successfully detects that the key is not inserted into the correct keyhole. In the volcano example (row 2), in addition to identifying the “broken mountain body” as in the human annotation, the model also detects a subtle issue: the disconnection of the lava flow, highlighting its fine-grained visual reasoning capabilities.

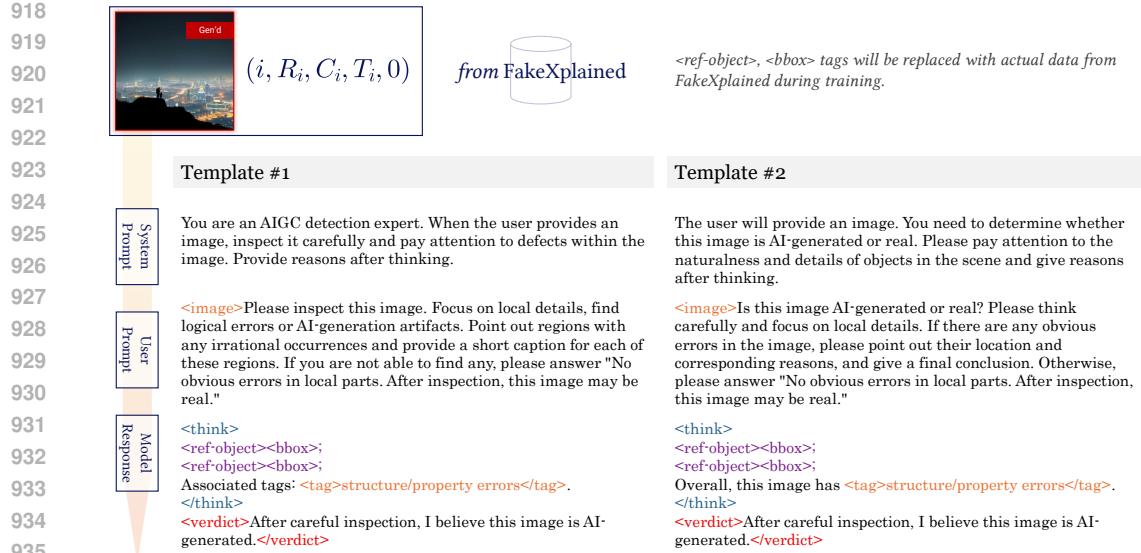


Figure 8: An example showing two different chat templates branched from one annotation entry.

#### A.4 ETHICAL CONSIDERATIONS

All generated images are synthetic with no real individuals. Annotators provided informed consent to this annotation job, allowing us to use the annotated dataset for training. We explicitly ask the annotators not to leave any personal or sensitive information in annotations.

## A.5 KNOWN LIMITATIONS

**Language.** Currently, all annotations are in one language. It is hard to translate the short annotation sentences to other languages without a manual check for language inconsistencies.

**No Real Images.** FakeXplained does not contain real images for the time being, as defining regions for real images can be more subjective than AI-generated images.

## B ADDITIONAL TRAINING DETAILS

## B.1 TWO-STAGE TRAINING

We use ms-swift Zhao et al. (2025) for fine-tuning Qwen-2.5-VL models.

In the LoRA SFT stage, we noticed that freezing either the projector or the vision encoder leads to marginal improvement over the base model without training. To achieve optimal SFT performance, both modules must be fine-tuned jointly.

After the SFT stage, we use GRPO instead of PPO. As noted in Shao et al. (2024), GRPO obviates the need for additional value function approximation as in PPO, and instead uses the average reward of multiple sampled outputs. For each query  $q$ , GRPO samples  $G$  outputs  $\{o_1, o_2, \dots, o_G\}$  from the old policy model  $\pi_{\theta_{old}}$ , and uses the relative *advantage* to optimize the MLLM, making it particularly well-suited for multi-modal reasoning tasks where absolute reward calibration is challenging.

We set the initial learning rate to  $10^{-4}$  for the SFT stage and  $10^{-5}$  for the RLHF stage. Reward signals fluctuated at the beginning of GRPO but quickly converged as the model is generating more human-aligned explanations, confirming the effectiveness of our reward design and training strategy.

## B.2 COMPUTATIONAL RESOURCES

The full training procedure took 41.0 hours on 8x NVIDIA A100 (80G) GPUs.

972  
973  
974  
Table 5: Comparative performance analysis under compression artifacts, spatial transformations,  
and resolution changes.

Degradation & Metric	FakeExplainer	ObjFormer.	SegFormer	FakeShield	LEGION	NPR	DMD.	ComFor.	AfPr.	AEROB.	DIRE
JPEG Compression (80% Quality)	Acc. IoU	<b>0.979</b> <b>0.353</b>	0.940 0.284	0.927 0.231	0.782 0.092	0.544 0.061	0.820 -	0.908 -	0.840 -	0.871 -	0.842 -
JPEG Compression (30% Quality)	Acc. IoU	<b>0.977</b> <b>0.339</b>	0.926 0.267	0.915 0.198	0.735 0.078	0.535 0.059	0.781 -	0.897 -	0.784 -	0.856 -	0.814 0.879
Random Cropping	Acc. IoU	<b>0.962</b> <b>0.314</b>	0.943 0.217	0.934 0.176	0.756 0.067	0.519 0.061	0.903 -	0.915 -	0.829 -	0.879 -	0.858 0.891
Downsampling (0.5x)	Acc. IoU	<b>0.980</b> <b>0.362</b>	0.929 0.259	0.931 0.254	0.748 0.092	0.591 0.076	0.899 -	0.912 -	0.853 -	0.875 -	0.841 0.894
Original Images	Acc. IoU	<b>0.982</b> <b>0.360</b>	0.954 0.299	0.945 0.289	0.801 0.028	0.583 0.098	0.914 -	0.928 -	0.882 -	0.887 -	0.873 0.911

983  
984  
985  
Table 6: Out-of-distribution performance evaluation across different datasets when trained with  
various configurations mentioned in the paper.

Sources	FakeExplainer	No-FT	Partial Data				Training Strategy		
			no-bbox	no-caption	no-tags	label-only	SFT	GRPO <sub>ω_G=1</sub>	GRPO <sub>ω_G=0.5</sub>
GPT-Image-1 Rapidata (2025)	<b>0.801</b>	0.421	0.691	0.760	0.774	0.603	0.591	0.788	0.768
FaceForensics++ Rössler et al. (2019)	<b>0.864</b>	0.519	0.715	0.796	0.817	0.640	0.680	0.826	0.832
MMFR-Dataset Gao et al. (2025)	<b>0.874</b>	0.593	0.859	0.708	0.843	0.612	0.671	0.794	0.773

991  
992  
993  
994  
At inference time, the end-to-end pipeline that takes an image as input to generate the verdict and  
grounding (if the image is deemed AI-generated) takes an average of 7.8 seconds on 2x NVIDIA  
A100 (80G) GPUs.996  
997  
C ROBUSTNESS AGAINST IMAGE PERTURBATIONS999  
1000  
1001  
1002  
To evaluate the practical applicability of our approach, we conduct a comprehensive robustness eval-  
uation under common image degradations that are frequently encountered in real-world scenarios.  
Table 5 presents a comparative performance analysis across three perturbation categories: JPEG  
compression, random cropping, and downsampling.1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
Our method demonstrates exceptional resilience to JPEG compression artifacts, achieving low per-  
formance degradations of merely 0.3% and 0.8% from the uncompressed baseline, significantly  
outperforming current state-of-the-art methods. All of which experience at least a 3% degradation.  
Notably, SegFormer and ObjectFormer show more stability than image-only classification models,  
indicating that grounding enhances robustness, although they still fall short of our method. For  
downsampling, we scaled the input images to 50% of their original width and height. In random  
cropping and downsampling experiments, our approach achieves the accuracy of 96.2 and 98.0, re-  
spectively, indicating robust performance across different resolution scales. Meanwhile, downsam-  
pling does not severely affect the IoU score, which suggests that our grounded reasoning approach  
effectively captures semantic-level artifacts that remain detectable even at reduced resolutions, un-  
like methods that may rely on pixel-level features more susceptible to resolution changes. Since ran-  
dom cropping modifies the overall image layout, this action can remove certain fake regions from an  
image entirely, leading to lower IoU across all methods. Interestingly, we observe a slight increase  
in IoU after downsampling. We hypothesize that this is because our grounding model focuses on  
the dominant artifact region, which remains visible at lower resolutions, while noisy fine details  
are suppressed, leading to more precise and focused localization. Overall, the consistent perfor-  
mance across perturbation types demonstrates that our model captures underlying semantic artifacts  
in AI-generated content, enabling robust detection even under challenging image conditions.1020  
1021  
D ADDITIONAL ABLATION STUDIES1022  
1023  
D.1 OUT-OF-DISTRIBUTION PERFORMANCE1024  
1025  
We evaluate the generalization capabilities of the ablation models in Section 5.2 (Table 2) of the main  
paper on three OoD datasets: images generated by GPT-Image-1 Rapidata (2025), FaceForensics++

1026

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043

1044

1045

1046

1047

1048

1049

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

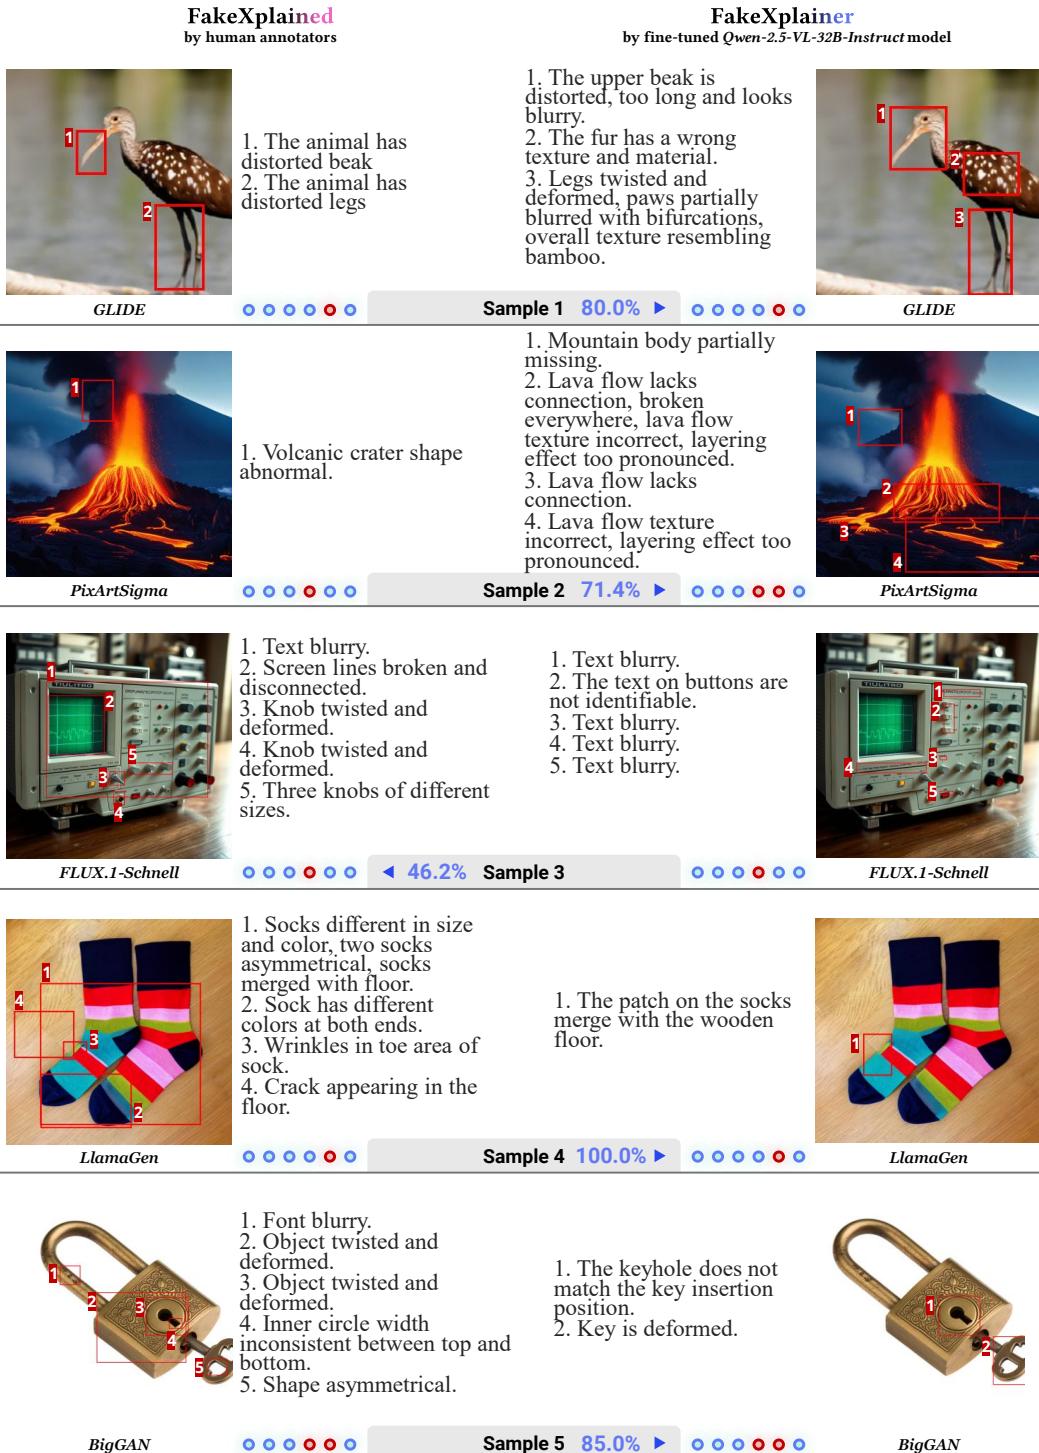
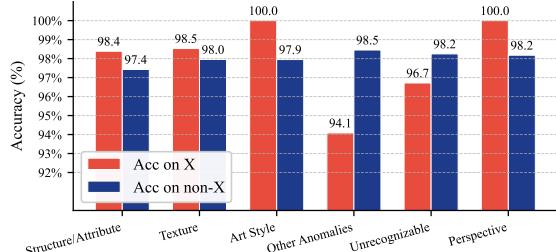


Figure 9: More annotation examples from FakeExplained and model response visualized from FakeExplainer. The ratio in the center shows the human preference score.

1080  
 1081 Table 7: Performance analysis of  
 1082 **FakeExplainer-32B** across different im-  
 1083 age tags. The table shows the percentage  
 1084 of samples containing each tag, and the  
 1085 classification accuracy for images with (Acc  
 1086 on X) and without (Acc on non-X) that  
 1087 specific tag.

Tag (X)	With X	Without X	Samples
Structure/Attribute Error	98.39%	97.43%	80.52%
Wrong Texture	98.53%	97.95%	42.65%
Artistic Style	<b>100.00%</b>	97.95%	12.18%
Other Anomalies	94.08%	98.45%	5.78%
Unrecognizable Objects	96.72%	98.25%	3.12%
Perspective Errors	<b>100.00%</b>	98.18%	1.16%



1093  
 1094  
 1095 Figure 10: Visual representation of detection accuracy  
 1096 across different tag categories.  
 1097

1098 Rössler et al. (2019) and MMFR-Dataset (eval) proposed by FakeReasoning Gao et al. (2025). Ta-  
 1099 ble 6 shows that our complete pipeline achieves accuracies of 80.1, 86.4 and 87.4 respectively,  
 1100 compared to 42.1, 51.9 and 59.3 for the base model without fine-tuning.

1101 Among partial data ablations, the label-only configuration performs the worst among all partial data  
 1102 category entries, yielding near no-finetuning performance. This OoD evaluation further confirms  
 1103 that both spatial grounding and textual reasoning are essential for generalization.

1104 The SFT stage alone yields moderate performance (59.1 on GPT-Image-1, 68.0 on FF++, 67.1 on  
 1105 MMFR). Further into the GRPO training, we see a better overall performance. This result is con-  
 1106 sistent with findings discussed in our main paper, as the RLHF stages give more performance boost  
 1107 than the SFT stage. The consistent improvements across both datasets suggest our approach learns  
 1108 generalizable features for AI-generated content detection rather than dataset-specific patterns.

## 1109 D.2 DISABLING LORA

1110 We employ LoRA during training to reduce computational cost and memory usage. While full-  
 1111 parameter fine-tuning is technically possible, our results show that it does not improve accuracy  
 1112 or IoU (Accuracy:  $98.2\% \rightarrow 97.9\%$ , IoU:  $36.0\% \rightarrow 35.4\%$ ), likely due to the limited amount of  
 1113 annotated data. This suggests that LoRA provides a more efficient and suitable training strategy  
 1114 under current data constraints. With significantly more training data, full fine-tuning may yield  
 1115 better results.

## 1118 E DISCUSSIONS ON TAG-WISE PERFORMANCE

1119 To assess the impact of different artifact types on detection performance, we analyzed the accuracy  
 1120 of **FakeExplainer-32B** on the **FakeExplained** dataset across the six distinct image-level tags defined in  
 1121 Section 3. Table 7 details the prevalence of each tag within the dataset and compares the model’s  
 1122 accuracy on images containing a specific tag versus those without it.

1123 Notably, the model achieves 100% accuracy on images tagged with **Artistic Style** and **Perspective**  
 1124 **Errors**. While Perspective Errors are rare (1.16%), the perfect detection rate implies that vanishing  
 1125 point inconsistencies are distinct features that the MLLM can easily leverage for classification.  
 1126 Meanwhile, the model exhibits a slight performance drop on images labeled with the **Other Anoma-**  
 1127 **lies** tag, which is typically assigned to images with global inconsistencies or subtle defects that are  
 1128 difficult to categorize into specific localized regions. This decrease in performance suggests that  
 1129 localized flaws are powerful cues to draw the conclusion that an image is AI-generated. This can  
 1130 potentially be improved with a global explanation to images in the **FakeExplained** dataset. Similarly,  
 1131 **Unrecognizable Objects** (96.72%) presents a moderate challenge, likely because defining “recog-  
 1132 nizability” can be subjective and occasionally overlaps with abstract artistic intent.

1134 **F LIMITATIONS**  
11351136 Despite promising results, our approach still has limitations. The Qwen-2.5-VL-32B-Instruct model  
1137 incurs substantial computational costs, which may limit its deployment in resource-constrained en-  
1138 vironments. Our evaluation does not sufficiently cover domain-specific or real-world image types,  
1139 such as medical, industrial, or artistic imagery. **As generated media becomes more and more realis-  
1140 tic and cross-modal, FakeExplainer, designed to detect visible, human-interpretable artifacts that can  
1141 be localized and explained, may not be able to produce explainable responses.** This scope is inher-  
1142 ent to all explainable AIGI detectors: grounded, human-aligned reasoning is only feasible when the  
1143 underlying cues are perceptible to humans. However, as long as visible artifacts are still present,  
1144 **FakeExplainer can effectively detect and explain. Nonetheless, as long as perceptible cues remain,  
1145 FakeExplainer provides significantly more robust and verifiable predictions than black-box classi-  
1146 fiers, especially under distribution shift, due to its grounded and human-aligned reasoning process.**  
11471148 **G BROADER IMPACT**  
11491150 While our system improves interpretability in detecting AI-generated content, it may also intro-  
1151 duce risks. The detailed explanations of detection rationale could inadvertently assist malicious  
1152 adversaries in developing more sophisticated evasion techniques, potentially contributing to an ad-  
1153 versarial “arms race.” The deployment of such systems without careful consideration could lead  
1154 to over-censorship of legitimate content, particularly affecting artists and creators who use AI tools  
1155 ethically. To mitigate these risks, we recommend responsible deployment frameworks, ongoing  
1156 monitoring for bias and fairness, and collaborative development with stakeholders to ensure the  
1157 technology serves the public interest while preserving legitimate creative expression.  
11581159 **H THE USE OF LARGE LANGUAGE MODELS**  
11601161 During manuscript preparation, we employed LLMs only for language polishing and grammar re-  
1162 finement. All research ideas, methods, and results were conceived, implemented, and validated  
1163 entirely by the authors. Since our work studies MLLMs in the context of forgery detection, we  
1164 necessarily employed LLMs as research subjects. Specifically, MLLMs were used to generate or  
1165 assist in generating annotations within our dataset and to serve as baseline models in our experi-  
1166 ments. These usages are intrinsic to the research problem itself and should not be interpreted as  
1167 LLMs contributing to the ideation or authorship of this paper.  
1168  
1169  
1170  
1171  
1172  
1173  
1174  
1175  
1176  
1177  
1178  
1179  
1180  
1181  
1182  
1183  
1184  
1185  
1186  
1187