

Agentic Retoucher for Text-To-Image Generation

Shaocheng Shen^{1*}, Jianfeng Liang^{1*}, Chunlei Cai¹, Cong Geng², Huiyu Duan¹,
Xiaoyun Zhang¹, Qiang Hu^{1†}, Guangtao Zhai¹

¹Shanghai Jiao Tong University, Shanghai, China ²JIUTIAN Research, Beijing, China



Figure 1. **Left:** Existing VLMs hallucinate and fail to localize distortions in AIGC-images, even with explicit region cues, whereas our method accurately localizes distorted regions and provides reasonable diagnoses. **Right:** Each before-after pair shows the distorted image and the result refined by our Agentic Retoucher, including diverse distortion artifacts across text, hand, face, and interaction.

Abstract

Text-to-image (T2I) diffusion models such as SDXL and FLUX have achieved impressive photorealism, yet small-scale distortions remain pervasive in limbs, face, text and so on. Existing refinement approaches either perform costly iterative re-generation or rely on vision-language models (VLMs) with weak spatial grounding, leading to semantic drift and unreliable local edits. To close this gap, we propose **Agentic Retoucher**, a hierarchical decision-driven framework that reformulates post-generation correction as a human-like perception-reasoning-action loop. Specifically, we design (1) a **perception agent** that learns contextual saliency for fine-grained distortion localization under text-image consistency cues, (2) a **reasoning agent** that performs human-aligned inferential diagnosis via progressive preference alignment, and (3) an **action agent** that adaptively plans localized inpainting guided by user preference. This design integrates perceptual evidence, linguistic reasoning, and controllable correction into a unified,

self-corrective decision process. To enable fine-grained supervision and quantitative evaluation, we further construct **GenBlemish-27K**, a dataset of 6K T2I images with 27K annotated artifact regions across 12 categories. Extensive experiments demonstrate that Agentic Retoucher consistently outperforms state-of-the-art methods in perceptual quality, distortion localization and human preference alignment, establishing a new paradigm for self-corrective and perceptually reliable T2I generation.

1. Introduction

Text-to-image (T2I) diffusion models [24, 25, 43] such as Imagen [61], Stable Diffusion [55, 59], FLUX [40], and Qwen-Image [1] have revolutionized image synthesis, enabling photorealistic and creative generation from natural language prompts. They are now widely adopted in design, film, and entertainment pipelines, as well as in downstream tasks like editing [6, 20, 39, 49, 74] and video generation [14, 63, 75]. However, even the most advanced models frequently produce small-scale distortions, including misaligned limbs, asymmetric faces, unreadable text, and in-

*: These authors contribute equally to this work.

†: Corresponding author.

consistent object interactions. These flaws typically occur locally within otherwise high-quality outputs, making them difficult to detect and expensive to correct through full-image regeneration. As a result, T2I systems still lack autonomous perceptual reliability, a key barrier to real-world creative and industrial use.

Recent research has explored three main directions to improve generative fidelity: prompt enhancement [31, 67, 71], reinforcement learning-based optimization [5], and fine-grained noise-space alignment [37, 70, 78]. Although these approaches effectively enhance overall realism, they lack explicit spatial reasoning and cannot interpret or correct localized failures. Post-hoc editing pipelines such as Imagic [39], Bagel [20], and Step1x-Edit [49] enable local refinement, but rely on manually crafted masks or heuristic textual hints, preventing autonomous identification of regions requiring correction.

Vision-language models (VLMs) [41, 54] show promise as automated critics due to their semantic reasoning capabilities. However, as shown in Fig. 1 (Left), even state-of-the-art VLMs struggle to reliably localize distorted regions. Explicit queries often yield inconsistent or incorrect assessments, with clearly abnormal regions being misjudged as normal. This stems from two key issues: VLMs are optimized for high-level semantic alignment rather than pixel-level verification, leading to weak spatial grounding and missed fine-scale artifacts. Furthermore, their extensive knowledge priors can override visual evidence, causing hallucinated judgments. For example, a portrait with six fingers is deemed plausible despite explicit highlighting of the defective hand, demonstrating that current VLMs are unreliable for fine-grained artifact detection in AI-generated images.

To address these limitations, we present **Agentic Retoucher**, a hierarchical decision-driven framework that reformulates post-generation correction as a structured *perception-reasoning-action* loop. Agentic Retoucher comprises three collaborative agents that execute a unified self-refinement cycle. The perception agent predicts context-aware distortion saliency by integrating visual evidence with prompt semantics, generating reliable region proposals for fine-scale anomalies. The reasoning agent performs human-aligned diagnostic inference, including identifying distortion types, detailing the appearance of distortions and assessing their inconsistency with global images through progressive preference alignment. The action agent then adaptively selects and executes targeted retouching operations from a modular tool library, supporting both mask-guided and instruction-driven editing under user or environment constraints. Through iterative verification, these components fuse perceptual cues, semantic reasoning, and controllable tool-based correction into a coherent self-corrective process, enabling the proposed Agentic Re-

toucher automatically refine distortion artifacts across text, hand, face, and interaction. (Fig. 1, Right)

To enable fine-grained and region-aware supervision, we construct **GenBlemish-27K**, a dataset of 6K T2I images with 27K pixel-level annotated distortion regions spanning 12 representative artifact categories. This dataset provides both spatial grounding and semantic diagnostic cues, allowing our system to reliably map localized distortions to interpretable region-level feedback and convert them into targeted retouching actions. Beyond supporting our framework, GenBlemish-27K also improves VLM robustness for evaluating AIGC imagery, enhancing region-grounded assessment and steering adaptation toward human-aligned distortion reasoning. Extensive experiments show that Agentic Retoucher significantly boosts local perceptual fidelity across diverse diffusion backbones while preserving global coherence, outperforming state-of-the-art post-editing methods on both objective metrics (plausibility score increasing from 44.21 to **47.10**) and human preference studies (**83.2%** preferred over unretouched outputs).

Our main contributions are summarized as follows:

- We propose **Agentic Retoucher**, a novel paradigm that reformulates post-generation editing as a *perception-reasoning-action* loop, enabling diffusion models to autonomously diagnose and refine their artifacts.
- We design a collaborative three-agent system, where a perception agent performs context-aware distortion localization, a reasoning agent conducts human-aligned fine-grained diagnosis, and an action agent performs adaptive local retouching with user-guided tools.
- We construct the **GenBlemish-27K** with pixel-level masks and textual annotations across 12 artifact types, providing a dataset for fine-grained artifact perception and correction.
- Extensive experiments demonstrate that our framework achieves state-of-the-art performance in perceptual quality enhancement, artifact localization and textual description accuracy across diverse diffusion backbones.

2. Related Works

Visual Quality Assessment. Visual Quality Assessment (VQA) [9, 11, 22, 44, 48, 56] is an important and rapidly evolving field that has made significant contributions to evaluating a wide range of image and video tasks with closer alignment to human subjective perception. For AIGC content assessment, most existing work [65, 73] is limited to applying quantitative metrics at global scales, without explicit localization and assessment of local flaws. RichHF [44] introduces predictors for local structural distortions along with a corresponding scoring procedure. However, these methods focus solely on assessment and have not been integrated into an automated, closed-loop

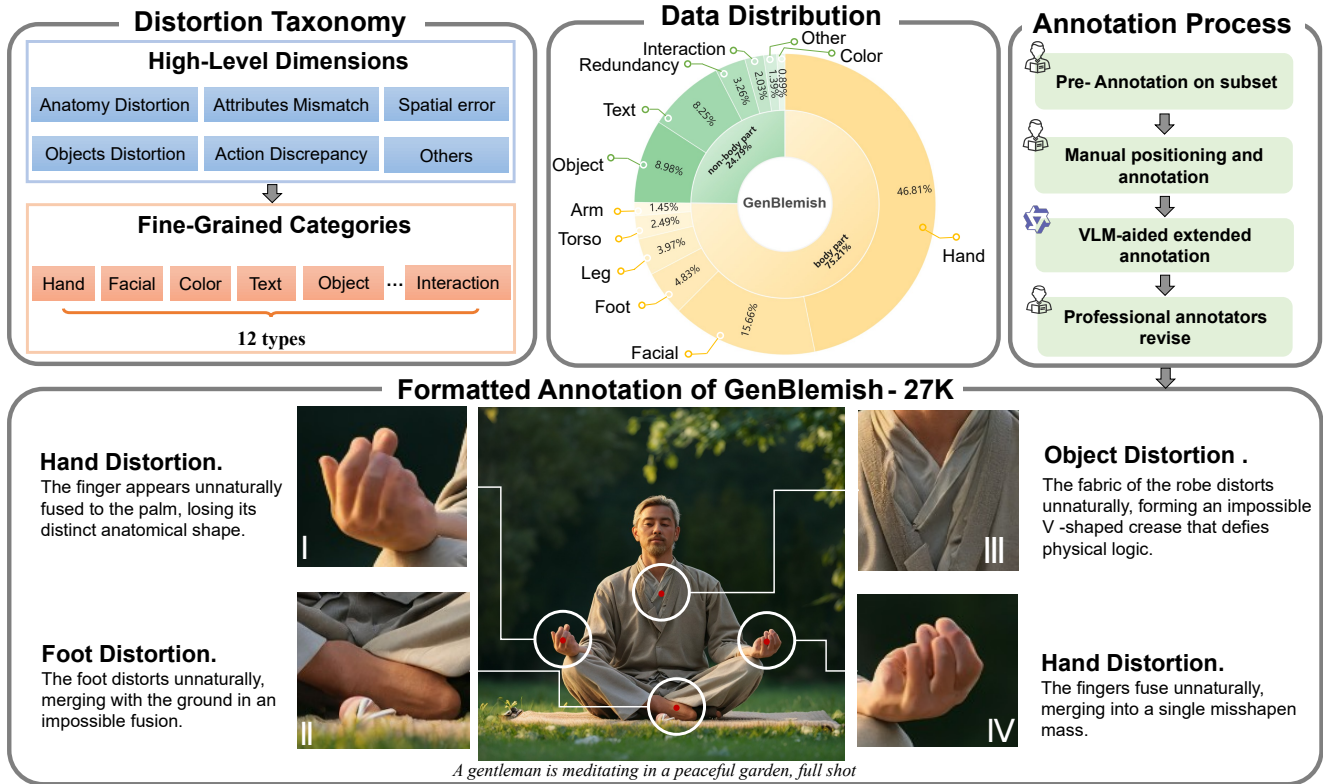


Figure 2. Overview of GenBlemish-27K. The figure illustrates (a) the dual-layer distortion taxonomy with six high-level dimensions and twelve fine-grained categories, (b) the distribution of localized distortion types, (c) the human-AI collaborative annotation pipeline, and (d) representative formatted samples with pixel-level masks and textual descriptions, highlighting how GenBlemish-27K enables fine-grained localization and reasoning over diverse text-to-image distortions.

pipeline for evaluation and refinement.

Vision-Language Model (VLMs). VLMs [51, 62, 72] have become leading drivers of general artificial intelligence, exhibiting remarkable problem-solving and reasoning ability through training on large-scale multimodal data (e.g., GPT-4o [54] and the Qwen [72] family in real-world multimodal interaction). Several works [30, 38, 57] further advance VLMs on image understanding tasks. However, heavy reliance on high-fidelity pretraining data and learned priors often biases VLMs toward prior-based, ungrounded hallucinated responses in the context of text-to-image evaluation.

Agentic System in Vision. Agentic System [13, 18, 36, 42, 60, 64, 77] adopts active, closed-loop perception-decision-action framework, with VLMs increasingly acting as planners due to their strong reasoning. In the 3D domain, VADAR [52] proposes an agentic program synthesis approach, achieving superior performance in 3D spatial reasoning. In image and video restoration, AgenticIR [76] and MoA-VR [47] independently propose VLM-integrated multi-agent repair paradigms. In the realm of artistic creation, JarvisArt [46] enables fine-grained photo retouching via tool invocation based on user instructions.

3. Dataset: GenBlemish-27K

We construct GenBlemish-27K, a large-scale dataset designed for granular distortion diagnosis and reasoning in text-to-image generation. It provides pixel-level annotations and natural-language descriptions for over 27K distorted regions across 12 artifact types, offering comprehensive supervision for perception, reasoning, and localized correction tasks.

3.1. Distortion Taxonomy

Existing T2I evaluation datasets suffer from limited coverage (e.g., HADM [66], Wang et al. [69]), coarse annotation (e.g., RichHF [44]), and insufficient scale (e.g., SynArtifacts-1K [9] with only 1K samples). To address these issues, GenBlemish-27K establishes a hierarchical taxonomy of distortions derived from large-scale inspection of outputs from mainstream T2I models. We define six high-level distortion dimensions including human anatomical distortion, attribute inconsistency, spatial errors, object deformation or redundancy, action and interaction distortion, and miscellaneous cases. These dimensions are further refined into 12 fine-grained categories such as limb defor-

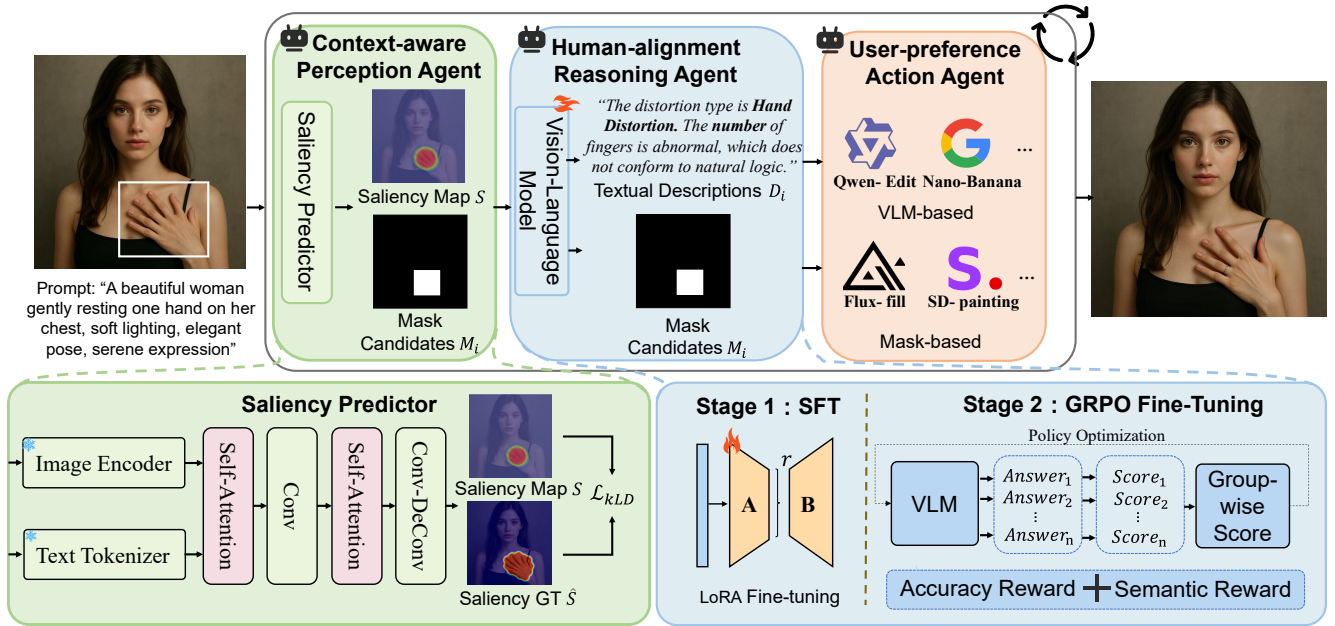


Figure 3. Overview of the proposed Agentic Retoucher. The framework operates as a *perception-reasoning-action* loop for post-generation correction in AIGC. The Perception Agent localizes context-dependent distortions via cross-modal saliency prediction, the Reasoning Agent performs human-aligned diagnosis through iterative reasoning, and the Action Agent executes adaptive localized inpainting guided by reasoning outputs, forming a closed-loop self-corrective process.

mities, face distortion, and text anomalies. This taxonomy captures typical artifacts observed in state-of-the-art diffusion models and enables interpretable reasoning about both *what* and *where* fails (see Fig. 2).

3.2. Data Annotation

We curate 6,025 images from EvalMuse-Structure [28], covering outputs from over 20 T2I models such as Dreamina, Midjourney, Kandinsky [2], and SDXL [55]. A four-stage human-in-the-loop process ensures both semantic richness and annotation consistency. (1) Annotators are first calibrated through a pre-annotation stage. (2) For each distortion region, multiple annotators independently provide the center, category, and a brief textual description; the region radius is $1/20$ of the image height. (3) The textual descriptions are expanded and refined using QwenVL-Max [3]. (4) Final annotations are reconciled through majority voting and expert validation. Each sample includes the generated image, input prompt, distortion mask, category label, and natural-language description, supporting tasks such as saliency prediction, defect classification, and language grounding.

3.3. Dataset Statistics

GenBlemish-27K consists of 6,025 images, 27,507 annotated distortion regions. The agreement rate between majority voting and expert validation exceeds 95%, confirming annotation reliability. Each image contains an average

of 4.6 annotated regions, each paired with an 11.8-word description. As shown in Fig. 2, hand distortions account for 46.8% of all annotations, followed by facial defects at 15.7%. These statistics indicate that fine-grained human generation remains a persistent challenge even for advanced diffusion models. More details will be included in supplementary material.

4. Methodology

4.1. Overview of the Agentic Retoucher

We propose an **Agentic Retoucher** that redefines post-generation image correction as a closed *perception-reasoning-action* loop. Unlike conventional feed-forward editing pipelines that apply static refinement, our framework introduces autonomy, interpretability, and self-correction into the editing process. By framing retouching as a sequential decision process, the model can reason about *what* and *where* distortions occur before performing targeted correction, bridging perceptual evidence, semantic inference and controllable within a unified architecture.

As shown in Fig. 3, the framework consists of three collaborative agents. The Perception Agent detects context-dependent distortions from visual-textual cues and generates a distortion-saliency map. The Reasoning Agent analyzes the detected regions, identifies distortion categories, and produces human-aligned textual descriptions. The Action Agent executes localized retouch guided by reason-

ing outputs, closing the *perception-reasoning-action* cycle through iterative refinement.

Formally, let I_t denote the image to be retouched. At iteration t , the perception agent produces a saliency map S_t highlighting anomalous regions. If the saliency S_t exceeds a threshold τ_s , the reasoning agent infers distortion types and generates region-level descriptions $\{D_i\}$ and masks $\{M_i\}$. The action agent then applies localized refinement to obtain an updated image:

$$I_{t+1} = \Phi_{\text{act}}(I_t, \{M_i \vee D_i\}), \quad t \leftarrow t + 1. \quad (1)$$

This process repeats until all salient distortions are eliminated, producing a perceptually faithful result. Through this iterative loop, the framework transitions post-generation editing from reactive correction to proactive reasoning, integrating perceptual analysis, contextual understanding, and controllable retouching in a single interpretable pipeline.

4.2. Context-Aware Perceptual Distortion Analysis

Text-to-image generations frequently exhibit subtle and context-dependent distortions such as implausible limb, object and text. These artifacts often lack explicit object boundaries, making conventional pixel-wise detection unreliable. To emulate human perceptual sensitivity, inspired by [44], we design a context-aware saliency predictor that estimates a distortion-saliency map $S \in [0, 1]^{H \times W}$ conditioned on both the generated image I and its prompt P . A dual-encoder ViT [21]-T5 [58] backbone encodes image and text representations, which are subsequently concatenated and fused via a self-attention mechanism to capture inherent correspondences between visual structures and textual semantics. A lightweight attention refinement module further aggregates multi-scale contextual cues, improving the detection of distortions whose visibility depends on global images.

The model is optimized using a hybrid loss that balances pixel accuracy and distributional consistency:

$$\mathcal{L}_{\text{sal}} = \alpha \mathcal{L}_{\text{MSE}}(S, \hat{S}) + (1 - \alpha) \mathcal{L}_{\text{KLD}}(S, \hat{S}), \quad (2)$$

where \hat{S} is the ground-truth saliency and α controls the balance between reconstruction precision and perceptual alignment. The KLD term encourages alignment with human fixation distributions, preserving discriminability in ambiguous regions and preventing over-smoothing. The resulting saliency map is binarized and morphologically dilated to form mask candidates $\{M_i\}$ for subsequent reasoning.

Beyond localization, the predicted saliency reflects contextual anomalies, serving as an explicit spatial prior that helps the reasoning agent focus on regions requiring further analysis, ensuring that higher-level diagnosis emerges from low-level visual awareness.

4.3. Human-Aligned Reasoning and Adaptive Action

Given localized regions $\{M_i\}$, the reasoning agent performs inferential diagnosis to generate textual descriptions $\{D_i\}$ that capture distortion types, local characteristics and contextual relationships. This task requires structured reasoning aligned with human perceptual judgment rather than simple classification or captioning. We adopt a progressive preference alignment paradigm consisting of two complementary stages: supervised fine-tuning (SFT) for structural initialization and Group Relative Policy Optimization (GRPO) for human-aligned reinforcement.

In the first stage, SFT establishes structured response formats and distortion taxonomy under limited supervision. To reduce computational overhead, we employ Low-Rank Adaptation (LoRA) [34], where the weight update ΔW for a layer $W \in \mathbb{R}^{n \times m}$ is decomposed as $\Delta W = AB$, with $A \in \mathbb{R}^{n \times r}$ and $B \in \mathbb{R}^{r \times m}$. This low-rank decomposition enables efficient specialization of the reasoning model without full-parameter fine-tuning.

In the second stage, GRPO [19] aligns reasoning behavior with human preferences through reinforcement signals:

$$\mathcal{L}_{\text{GRPO}} = \mathbb{E}_{(q,o)} [\min(r_t \hat{A}_t, \text{clip}(r_t, 1 - \varepsilon, 1 + \varepsilon) \hat{A}_t) - \beta D_{\text{KL}}[\pi_\theta || \pi_{\text{ref}}]], \quad (3)$$

where \hat{A}_t denotes the normalized advantage that captures preference consistency. Policy optimization is guided by rewards capturing distortion-type classification accuracy and alignment between textual descriptions and human labels. This stage reduces hallucination, enabling the agent to produce consistent, human-aligned reasoning across diverse distortion patterns.

Building on the reasoning outputs, the Action Agent transforms $\{M_i, D_i\}$ into controllable local editing operations. It determines the spatial extent, tool selection, and inpainting instruction for each region. Depending on computational constraints or user preferences, the agent dynamically chooses between VLM-based or mask-guided inpainting from a modular tool library. The updated image is re-evaluated by the perception agent to close the *perception-reasoning-action* loop. Through iterative perception and reasoning, the framework converges toward high-quality outputs with plausible details.

The proposed framework transitions post-generation editing from reactive correction to proactive reasoning. By integrating perception-driven diagnosis, human-aligned reasoning, and adaptive retouching into a unified loop, the system achieves interpretable and autonomous refinement of generative outputs.

Table 1. Quantitative comparison of Agentic Retoucher with VLM-based and mask-based inpainting baselines on the GenBlemish-27K and SynArtifacts-1K datasets.

Condition Type	Model	GenBlemish-27K				SynArtifacts-1K			
		plausibility↑	aesthetics↑	alignment↑	overall↑	plausibility↑	aesthetics↑	alignment↑	overall↑
VLM-based	Original	44.21	53.69	57.89	47.15	61.53	61.63	60.65	55.35
	Qwen-Edit	44.44	53.71	57.69	47.15	61.45	61.64	60.70	55.33
	Ours w Qwen-Edit	47.10	55.75	59.54	49.27	65.43	64.88	62.61	58.04
	Gemini 2.5 Flash Image	44.41	53.80	57.93	47.27	62.63	63.07	61.21	56.15
	Ours w Gemini 2.5 Flash Image	46.81	55.47	59.22	48.97	65.96	65.27	62.94	58.43
Mask-based	Flux-fill	44.12	53.68	57.91	47.07	61.92	61.78	61.17	55.71
	Ours w Flux-fill	46.18	55.17	59.26	48.66	65.25	64.07	62.74	57.86
	SD-inpainting	45.18	53.85	57.70	47.50	63.60	62.49	60.88	56.26
	Ours w SD-inpainting	46.71	54.71	58.07	48.31	66.66	64.67	62.33	58.27

Table 2. Human evaluation results: preference distribution comparing Agentic Retoucher outputs to original images. Percentages of test cases rated as \gg (significantly better), $>$ (slightly better), \approx (about the same), $<$ (slightly worse), or \ll (significantly worse). Data from 5 participants in a randomized, blind survey.

Preference	\gg	$>$	\approx	$<$	\ll
Baseline	4.2%	22.8%	60.8%	9.2%	3.0%
Ours	48.8%	34.4%	10.2%	5.8%	0.8%

5. Experiments

5.1. Experimental setup

Datasets and Implementation Details. We evaluate our framework on the proposed GenBlemish-27K dataset to assess optimization efficacy and module functionality, and further verify its generalization on SynArtifacts-1K [9]. For the Context-Aware Perception Agent, the learning rate is set to 2×10^{-5} . For the Human-Alignment Reasoning Agent, Stage 1 employs LoRA fine-tuning with rank 64 and $\alpha = 32$. Inference is fully automatic and converges within 2-3 reasoning iterations per image.

Evaluation Metrics. For retouch evaluation, we use the four perceptual metrics from RichHF [44]: plausibility, aesthetics, alignment, and overall, to assess both structural plausibility and perceptual quality, as standard T2I metrics (e.g., FID [32]) fail to capture localized improvements. For perception, following [8], we adopt CC, SIM, KLD, AUC-Judd, and NSS to evaluate distributional and fixation-level consistency. For reasoning, we report distortion-type classification accuracy and semantic alignment using ROUGE [45], METEOR [4], Word2Vec [53], and SimCSE [26]. More setup details are provided in the supplementary material.

5.2. Comparison

Quantitative comparisons. We evaluate Agentic Retoucher using two categories of inpainting tools within our Adaptive Action Toolkit: VLM-based models (Qwen-Edit and Gemini 2.5 Flash Image) and mask-based models

Table 3. Quantitative evaluation of the Context-Aware Perception Agent on distortion-aware saliency prediction. Higher AUC-Judd, NSS, CC, SIM and lower KLD indicate better context perception.

Method/Metric	AUC-Judd ↑	NSS ↑	CC ↑	SIM ↑	KLD ↓
AIM [7]	0.7822	1.1479	0.1667	0.0759	3.0185
GBVS [29]	0.6580	0.5811	0.0080	0.0010	8.5547
SR [33]	0.5336	0.0135	0.0002	0.0524	3.4162
SMVJ [10]	0.8167	0.7121	0.0778	0.0623	3.3722
SWD [23]	0.8170	0.5307	0.0712	0.0761	3.5233
CA [27]	0.4516	0.7633	0.1000	0.0693	3.3286
SALICON [35]	0.9230	1.0774	0.5039	0.2734	1.7171
TranSalNet [50]	0.9029	1.1494	0.4616	0.0989	2.8716
Sal-CFS-GAN [12]	0.7747	0.7810	0.2124	0.1018	2.8589
SAM-VGG [17]	0.8773	1.2072	0.3410	0.1791	2.4094
SAM-ResNet [17]	0.9162	1.0552	0.4040	0.2475	2.0740
MLNet [16]	0.8539	1.0455	0.3535	0.2381	2.2359
InternVL3.5-8B [68]	0.8049	0.7689	0.5104	0.4095	3.9325
Qwen2.5-VL-7B [3]	0.6145	0.4190	0.1710	0.1481	7.4353
GLM-4.1V-9B [62]	0.5461	0.2191	0.0604	0.0902	8.0118
RichHF [44]	0.9211	0.8954	0.4748	0.3309	1.6697
Ours	0.9336	1.2087	0.5568	0.3822	1.4313

(Flux-Fill and SD-inpainting). As shown in Tab. 1, Agentic Retoucher consistently surpasses all baselines across the four perceptual metrics of plausibility, aesthetics, alignment, and overall. On the GenBlemish-27K, the plausibility score improves from 44.21 to **47.10** and the overall score from 47.15 to **49.27**, indicating that our system effectively handles localized distortion while preserving global structure and style. Similar improvements are observed on SynArtifacts-1K (overall score reaching **58.43**), confirming the generalization capability of the proposed framework. In human evaluations, as shown in Tab. 2, the Agentic Retoucher substantially outperforms the baseline model, with **83.2%** of results judged superior to the pre-retouching images, further demonstrating the visual expressiveness of our method.

Qualitative comparisons. Fig. 4 presents qualitative comparisons across various scenes and prompt conditions. Agentic Retoucher autonomously identifies distortion re-




Figure 4. Qualitative comparison of retouching results across diverse prompts. White bounding boxes indicate zoomed-in fine-grained regions. Agentic Retoucher retouches local distortions and implausibilities while maintaining global visual harmony, outperforming both VLM-based and mask-based baselines. More qualitative results are included in the supplementary material.

gions and performs targeted refinement while preserving global composition. Zoomed-in crops reveal that our method excels at retouching fine-grained geometric details (e.g., faces, fingers, feet), with coherent shading and natural boundaries. In contrast, VLM-based methods fail to localize distortions, and their retouch performance degrades without fine-grained instruction guidance, whereas mask-based models revert to stochastic refinement once explicit masks are removed. These results highlight the effectiveness of our agentic perception-reasoning-action loop in achieving both localized precision and holistic consistency.

5.3. Perception and Reasoning Analysis

Context-Aware Perception. Tab. 3 compares the proposed Context-Aware Perception Agent with conventional saliency detectors, deep saliency networks, and vision-language models. Hand-crafted methods (e.g., AIM [7], GBVS [29]) rely on low-level contrast, while deep saliency networks (e.g., SAM-ResNet [17], TranSalNet [50]) yield only moderate gains; both fail to capture context-aware distortion regions. General-purpose VLMs (e.g., In-

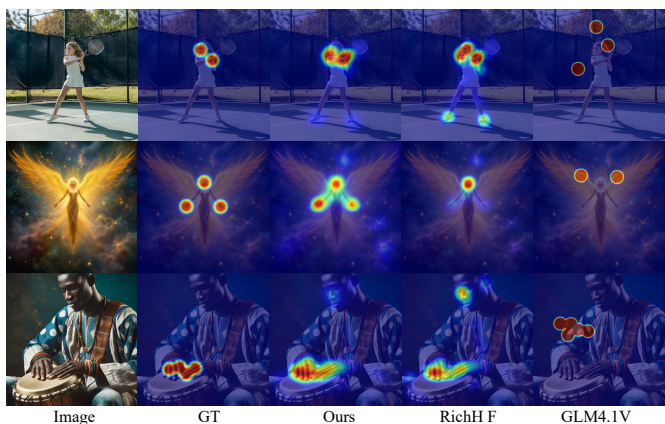


Figure 5. Qualitative visualization of saliency prediction. Our method yields sharper, context-aware localization than RichHF and GLM4.1V.

ternVL3.5 [68], Qwen2.5-VL [3], GLM-4.1V [62]) also perform poorly, lacking task-specific grounding for distortion awareness. In contrast, our perception agent achieves state-of-the-art results across all metrics (AUC-Judd =

Table 4. Quantitative evaluation and ablation of the Human-Alignment Reasoning Agent

Method/Metric	Accuracy \uparrow	SimCSE \uparrow	Word2Vec \uparrow	Meteor \uparrow	ROUGE \uparrow
GPT 5 Zero-Shot	61.31%	0.6928	0.6214	0.1699	0.1131
Gemini-2.5 Pro Zero-Shot [15]	60.28%	0.6856	0.6245	0.1702	0.1121
Qwen2.5-VL-7B [3]	57.76%	0.6658	0.6110	0.1678	0.0733
Qwen2.5-VL-7B + GRPO	58.97%	0.7020	0.6592	0.1741	0.1003
Qwen2.5-VL-7B + SFT	78.34%	0.8405	0.7768	0.4011	0.3515
Ours	80.10%	0.8426	0.7785	0.4037	0.3530
GLM-4.1V-9B [62]	58.25%	0.6723	0.6189	0.1701	0.0966
GLM-4.1V-9B + GRPO	60.15%	0.7182	0.6235	0.1833	0.1226
GLM-4.1V-9B + SFT	77.13%	0.8357	0.7734	0.3970	0.3592
Ours	79.26%	0.8416	0.7811	0.4172	0.3757
Ovis2.5-9B [51]	56.94%	0.6801	0.6056	0.1678	0.1035
Ovis2.5-9B + GRPO	69.67%	0.7264	0.6650	0.1901	0.1563
Ovis2.5-9B + SFT	78.85%	0.8287	0.7616	0.3589	0.3286
Ours	80.62%	0.8392	0.7730	0.3865	0.3521

0.9336, NSS = 1.2087, KLD = 1.4313), demonstrating robust localization of artifact-prone regions.

Fig. 5 shows that, unlike RichHF (overemphasizing facial and limb regions) and GLM4.1V (dispersing attention to irrelevant areas), our saliency maps deliver sharper spatial focus and stronger agreement with ground-truth labels. This precise localization provides a reliable foundation for subsequent reasoning and retouching within the loop.

Human-Alignment Reasoning. Tab. 4 reports quantitative results for the proposed Human-Alignment Reasoning Agent under different training strategies. Within each model family (Qwen2.5-VL, GLM-4.1V, Ovis2.5), our method consistently achieves the highest scores across all metrics, confirming that progressive alignment improves reasoning accuracy and adaptation to human preferences, effectively bridges distortion perception and linguistic reasoning within the *perception-reasoning-action* loop.

Beyond validating the reasoning agent itself, these results also highlight the diagnostic value of our GenBlemish-27K dataset. Across three large-scale backbones, performance consistently improves when trained or evaluated on our dataset, indicating its strong capability to reveal distortion-related reasoning gaps and to guide human-aligned adaptation. Conversely, Zero-Shot settings, including advanced closed-source models such as GPT-5 and Gemini 2.5 Pro, exhibit limited generalization, underscoring the intrinsic difficulty of distortion-type reasoning and further validating the dataset’s discriminative and instructional effectiveness.

5.4. Ablation Studies

We conduct three ablation analyses across the Perception, Reasoning, and Action Agents to evaluate the contribution of each component within our framework. Specifically, we isolate (i) the lightweight attention mechanism and the KLD loss in the Perception Agent, (ii) progressive alignment strategies in the Reasoning Agent, and (iii) adaptive

Table 5. Ablation study of the Context-Aware Perception Agent on attention and KLD loss components.

Method/Metric	AUC-Judd \uparrow	NSS \uparrow	CC \uparrow	SIM \uparrow	KLD \downarrow
Ours w/o attn&KLD_loss	<u>0.9335</u>	1.1957	0.5518	<u>0.3766</u>	1.4436
Ours w/o attn	<u>0.9335</u>	1.2153	0.5544	0.3731	<u>1.4412</u>
Ours w/o KLD_loss	0.9313	1.1892	<u>0.5546</u>	0.3525	1.5008
Ours	0.9336	<u>1.2087</u>	0.5568	0.3822	1.4313

conditioning schemes in the Action Agent.

Perception Agent Ablation. Tab. 5 analyzes the effect of the lightweight attention mechanism and the KLD loss in the Perception Agent. Removing attention (“w/o attn”) results in lower SIM and CC, indicating reduced global structural consistency. Eliminating the KLD loss (“w/o KLD loss”) decreases NSS and AUC-Judd, suggesting less accurate fixation-level localization. These two components are complementary: the attention module maintains coherent contextual structure, while the KLD term sharpens focus on human-attended regions. Their joint optimization yields the best overall balance between local precision and global awareness, validating the perception agent’s contribution to context-aware saliency modeling.

Reasoning Agent Ablation. Tab. 4 provides an ablation on reasoning alignment strategies. Progressive training consistently outperforms single-stage SFT or GRPO-only configurations across all metrics. Notably, applying GRPO at early stages destabilizes response formatting and causes factual drift, whereas progressive alignment enhances both reasoning stability and human-aligned semantic grounding.

Action Agent Ablation. Tab. 1 compares different conditioning schemes for the Action Agent under both VLM-based and mask-based refinement settings. Across all datasets, every tool in our tool library consistently achieves higher scores on all metrics, confirming robustness to diverse distortion types. By dynamically selecting among multiple refinement backbones, the Action Agent ensures locally correction and global coherence, effectively closing the *perception-reasoning-action* loop.

6. Conclusions

We propose Agentic Retoucher, a hierarchical, decision-driven framework that reformulates post-generation editing for T2I diffusion as a human-like *perception-reasoning-action* loop. The perception agent localizes small-scale distortions, the reasoning agent performs human-aligned diagnosis, and the action agent plans localized inpainting guided by user intent. In addition, we introduce GenBlemish-27K for fine-grained supervision and evaluation. Extensive experiments demonstrate consistent improvements over state-of-the-art methods in perceptual quality, distortion localization, and human preference alignment, establishing a self-corrective and perceptually reliable T2I paradigm.

Acknowledgment

This work is supported by National Natural Science Foundation of China (62271308, 62571322, 62431015), STCSM (24ZR1432000, 24511106902, 24511106900, 22DZ2229005), 111 plan (BP0719010), and State Key Laboratory of UHD Video and Audio Production and Presentation.

References

- [1] Qwen-image technical report, 2025. 1
- [2] Vladimir Arkhipkin, Andrei Filatov, Viacheslav Vasilev, Anastasia Maltseva, Said Azizov, Igor Pavlov, Julia Agafonova, Andrey Kuznetsov, and Denis Dimitrov. Kandinsky 3.0 technical report, 2023. 4
- [3] Shuai Bai and Chen et al. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 4, 6, 7, 8
- [4] Satanjeev Banerjee and Alon Lavie. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, 2005. 6
- [5] Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion models with reinforcement learning. 2023. 2
- [6] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*, pages 18392–18402, 2023. 1
- [7] Neil Bruce and John Tsotsos. Saliency based on information maximization. In *NIPS*. MIT Press, 2005. 6, 7
- [8] Zoya Bylinskii, Tilke Judd, Aude Oliva, Antonio Torralba, and Frédo Durand. What do different evaluation metrics tell us about saliency models? *IEEE TPAMI*, 41(3):740–757, 2019. 6
- [9] Bin Cao, Jianhao Yuan, Yexin Liu, Jian Li, Shuyang Sun, Jing Liu, and Bo Zhao. Synartifact: Classifying and alleviating artifacts in synthetic images via vision-language model. *arXiv preprint arXiv:2402.18068*, 2024. 2, 3, 6
- [10] Moran Cerf, Jonathan Harel, Wolfgang Einhaeuser, and Christof Koch. Predicting human gaze using low-level saliency combined with face detection. In *NIPS*. Curran Associates, Inc., 2007. 6
- [11] Jingjing Chang, Yixiao Fang, Peng Xing, Shuhan Wu, Wei Cheng, Rui Wang, Xianfang Zeng, Gang Yu, and Hai-Bao Chen. Oneig-bench: Omni-dimensional nuanced evaluation for image generation. *arXiv preprint arXiv:2506.07977*, 2025. 2
- [12] Zhaohui Che, Ali Borji, Guangtao Zhai, Xionghuo Min, Guodong Guo, and Patrick Le Callet. How is gaze influenced by image transformations? dataset and model. *IEEE TIP*, 29:2287–2300, 2020. 6
- [13] Haoyu Chen, Wenbo Li, Jinjin Gu, Jingjing Ren, Sixiang Chen, Tian Ye, Renjing Pei, Kaiwen Zhou, Fenglong Song, and Lei Zhu. Restoreagent: Autonomous image restoration agent via multimodal large language models, 2024. 3
- [14] Junsong Chen, Yuyang Zhao, Jincheng Yu, Ruihang Chu, Junyu Chen, Shuai Yang, Xianbang Wang, Yicheng Pan, Daquan Zhou, Huan Ling, Haozhe Liu, Hongwei Yi, Hao Zhang, Muyang Li, Yukang Chen, Han Cai, Sanja Fidler, Ping Luo, Song Han, and Enze Xie. Sana-video: Efficient video generation with block linear diffusion transformer. In *The Fourteenth International Conference on Learning Representations*, 2026. 1
- [15] Gheorghe Comanici, Eric Bieber, and Mike Schaekermann et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities, 2025. 8
- [16] Marcella Cornia, Lorenzo Baraldi, Giuseppe Serra, and Rita Cucchiara. A Deep Multi-Level Network for Saliency Prediction. In *ICPR*, 2016. 6
- [17] Marcella Cornia, Lorenzo Baraldi, Giuseppe Serra, and Rita Cucchiara. Predicting Human Eye Fixations via an LSTM-based Saliency Attentive Model. *IEEE Transactions on Image Processing*, 27(10):5142–5154, 2018. 6, 7
- [18] Ekin D. Cubuk, Barret Zoph, Dandelion Man’e, Vijay Vasudevan, and Quoc V. Le. Autoaugment: Learning augmentation policies from data. *2019 IEEE/CVF CVPR*, pages 113–123, 2019. 3
- [19] DeepSeek-AI, Daya Guo, and Dejian Yang et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025. 5
- [20] Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, Guang Shi, and Haoqi Fan. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*, 2025. 1, 2
- [21] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. 2020. 5
- [22] Huiyu Duan, Qiang Hu, Jiarui Wang, Liu Yang, Zitong Xu, Lu Liu, Xionghuo Min, Chunlei Cai, Tianxiao Ye, Xiaoyun Zhang, et al. Finevq: Fine-grained user generated content video quality assessment. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 3206–3217, 2025. 2
- [23] Lijuan Duan, Chunpeng Wu, Jun Miao, Laiyun Qing, and Yu Fu. Visual saliency detection by spatially weighted dissimilarity. In *CVPR 2011*, pages 473–480, 2011. 6
- [24] Yushun Fang, Yuxiang Chen, Shibo Yin, Qiang Hu, Jiangchao Yao, Ya Zhang, Xiaoyun Zhang, and Yanfeng Wang. One-step diffusion transformer for controllable real-world image super-resolution. *arXiv preprint arXiv:2511.17138*, 2025. 1
- [25] Yushun Fang, Lu Liu, Xiang Gao, Qiang Hu, Ning Cao, Jian Cui, Gang Chen, and Xiaoyun Zhang. Robust id-specific face restoration via alignment learning. *ArXiv*, abs/2507.10943, 2025. 1
- [26] Tianyu Gao, Xingcheng Yao, and Danqi Chen. SimCSE: Simple contrastive learning of sentence embeddings. In *EMNLP*, 2021. 6
- [27] Stas Goferman, Lih Zelnik-Manor, and Ayellet Tal. Context-aware saliency detection. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 2376–2383, 2010. 6

- [28] Shuhao Han, Haotian Fan, Jiachen Fu, Liang Li, Tao Li, Junhui Cui, Yunqiu Wang, Yang Tai, Jingwei Sun, Chunle Guo, and Chongyi Li. Evalmuse-40k: A reliable and fine-grained benchmark with comprehensive human annotations for text-to-image generation model evaluation, 2024. 4
- [29] Jonathan Harel, Christof Koch, and Pietro Perona. Graph-based visual saliency. In *NIPS*. MIT Press, 2006. 6, 7
- [30] Miran Heo, Min-Hung Chen, De-An Huang, Sifei Liu, Subhashree Radhakrishnan, Seon Joo Kim, Yu-Chiang Frank Wang, and Ryo Hachiuma. Omni-rgpt: Unifying image and video region-level understanding via token marks. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, pages 3919–3930, 2025. 3
- [31] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. 2022. 2
- [32] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*. Curran Associates, Inc., 2017. 6
- [33] Xiaodi Hou and Liqing Zhang. Saliency detection: A spectral residual approach. In *2007 IEEE CVPR*, pages 1–8, 2007. 6
- [34] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *ICLR*. OpenReview.net, 2022. 5
- [35] Xun Huang, Chengyao Shen, Xavier Boix, and Qi Zhao. Salicon: Reducing the semantic gap in saliency prediction by adapting deep neural networks. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 262–270, 2015. 6
- [36] S. Hutchinson, G.D. Hager, and P.I. Corke. A tutorial on visual servo control. *IEEE Transactions on Robotics and Automation*, 12(5):651–670, 1996. 3
- [37] Amir Mohammad Izadi, Seyed Mohammad Hadi Hosseini, Soroush Vafaie Tabar, Ali Abdollahi, Armin Saghafian, and Mahdih Soleymani Baghshah. Fine-grained alignment and noise refinement for compositional text-to-image generation, 2025. 2
- [38] Omri Kaduri, Shai Bagon, and Tali Dekel. What’s in the image? a deep-dive into the vision of vision language models, 2024. 3
- [39] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF CVPR*, pages 6007–6017, 2023. 1, 2
- [40] Black Forest Labs, Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, Cyril Diagne, Tim Dockhorn, Jack English, Zion English, Patrick Esser, Sumith Kulal, Kyle Lacey, Yam Levi, Cheng Li, Dominik Lorenz, Jonas Müller, Dustin Podell, Robin Rombach, Harry Saini, Axel Sauer, and Luke Smith. Flux.1 kontext: Flow matching for in-context image generation and editing in latent space, 2025. 1
- [41] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. Llava-onevision: Easy visual task transfer, 2024. 2
- [42] Xinghang Li, Peiyan Li, Minghuan Liu, Dong Wang, Jirong Liu, Bingyi Kang, Xiao Ma, Tao Kong, Hanbo Zhang, and Huaping Liu. Towards generalist robot policies: What matters in building vision-language-action models. *arXiv preprint arXiv:2412.14058*, 2024. 3
- [43] Jianfeng Liang, Shaocheng Shen, Botao Xu, Qiang Hu, and Xiaoyun Zhang. D²-vr: Degradation-robust and distilled video restoration with synergistic optimization strategy. *arXiv preprint arXiv:2602.08395*, 2026. 1
- [44] Youwei Liang, Junfeng He, Gang Li, Peizhao Li, Arseniy Klimovskiy, Nicholas Carolan, Jiao Sun, Jordi Pont-Tuset, Sarah Young, Feng Yang, Junjie Ke, Krishnamurthy Dj Dvijotham, Katherine M. Collins, Yiwen Luo, Yang Li, Kai J Kohlhoff, Deepak Ramachandran, and Vidhya Navalpakkam. Rich human feedback for text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19401–19411, 2024. 2, 3, 5, 6
- [45] Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain, 2004. Association for Computational Linguistics. 6
- [46] Yunlong Lin, Zixu Lin, Kunjie Lin, Jinbin Bai, Panwang Pan, Chenxin Li, Haoyu Chen, Zhongdao Wang, Xinghao Ding, Wenbo Li, and Shuicheng Yan. Jarvisart: Liberating human artistic creativity via an intelligent photo retouching agent. *arXiv preprint arXiv:2506.17612*, 2025. 3
- [47] Lu Liu, Chunlei Cai, Shaocheng Shen, Jianfeng Liang, Weimin Ouyang, Tianxiao Ye, Jian Mao, Huiyu Duan, Jiangchao Yao, Xiaoyun Zhang, Qiang Hu, and Guangtao Zhai. Moa-vr: A mixture-of-agents system toward all-in-one video restoration. *IEEE Journal of Selected Topics in Signal Processing*, 19(8):1822–1837, 2025. 3
- [48] Lu Liu, Huiyu Duan, Qiang Hu, Liu Yang, Chunlei Cai, Tianxiao Ye, Huayu Liu, Xiaoyun Zhang, and Guangtao Zhai. F-bench: Rethinking human preference evaluation metrics for benchmarking face generation, customization, and restoration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10982–10994, 2025. 2
- [49] Shiyu Liu and Yucheng Han et al. Step1x-edit: A practical framework for general image editing. *arXiv preprint arXiv:2504.17761*, 2025. 1, 2
- [50] Jianxun Lou, Hanhe Lin, David Marshall, Dietmar Saupe, and Hantao Liu. Transalnet: Towards perceptually relevant visual saliency prediction. *Neurocomputing*, 494:455–467, 2022. 6, 7
- [51] Shiyin Lu and Yang Li et.al. Ovis2.5 technical report, 2025. 3, 8
- [52] Damiano Marsili, Rohun Agrawal, Yisong Yue, and Georgia Gkioxari. Visual agentic ai for spatial reasoning with a dynamic api. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, pages 19446–19455, 2025. 3

- [53] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space, 2013. 6
- [54] OpenAI and Aaron Hurst et al. Gpt-4o system card, 2024. 2, 3
- [55] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 1, 4
- [56] Jiaying Qian, Ziheng Jia, Zicheng Zhang, Zeyu Zhang, Guangtao Zhai, and Xiongkuo Min. Towards explainable partial-aigc image quality assessment, 2025. 2
- [57] Liao Qu, Huichao Zhang, Yiheng Liu, Xu Wang, Yi Jiang, Yiming Gao, Hu Ye, Daniel K. Du, Zehuan Yuan, and Xinlong Wu. Tokenflow: Unified image tokenizer for multimodal understanding and generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, pages 2545–2555, 2025. 3
- [58] Adam Roberts and Hyung Won Chung et al. Scaling up models and data with $t5x$ and *seqio*, 2022. 5
- [59] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF CVPR*, pages 10684–10695, 2022. 1
- [60] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. “grabcut”: interactive foreground extraction using iterated graph cuts. *ACM Trans. Graph.*, 23(3):309–314, 2004. 3
- [61] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *NIPS*, 35:36479–36494, 2022. 1
- [62] V Team, Wenyi Hong, Wenmeng Yu, Xiaotao Gu, Guo Wang, Guobing Gan, Haomiao Tang, Jiale Cheng, Ji Qi, Junhui Ji, Lihang Pan, Shuaiqi Duan, Weihang Wang, Yan Wang, Yean Cheng, Zehai He, Zhe Su, Zhen Yang, Ziyang Pan, Aohan Zeng, Baoxu Wang, Bin Chen, Boyan Shi, Changyu Pang, Chenhui Zhang, Da Yin, Fan Yang, Guoqing Chen, Haochen Li, Jiale Zhu, Jiali Chen, Jiaying Xu, Jiazhen Xu, Jing Chen, Jinghao Lin, Jinhao Chen, Jinjiang Wang, Junjie Chen, Leqi Lei, Letian Gong, Leyi Pan, Mingdao Liu, Mingde Xu, Mingzhi Zhang, Qinkai Zheng, Ruiliang Lyu, Shangqin Tu, Sheng Yang, Shengbiao Meng, Shi Zhong, Shiyu Huang, Shuyuan Zhao, Siyan Xue, Tianshu Zhang, Tianwei Luo, Tianxiang Hao, Tianyu Tong, Wei Jia, Wenkai Li, Xiao Liu, Xiaohan Zhang, Xin Lyu, Xinyu Zhang, Xinyue Fan, Xuancheng Huang, Yadong Xue, Yanfeng Wang, Yanling Wang, Yanzi Wang, Yifan An, Yifan Du, Yiheng Huang, Yilin Niu, Yiming Shi, Yu Wang, Yuan Wang, Yuanchang Yue, Yuchen Li, Yusen Liu, Yutao Zhang, Yuting Wang, Yuxuan Zhang, Zhao Xue, Zhengxiao Du, Zhenyu Hou, Zihan Wang, Peng Zhang, Debing Liu, Bin Xu, Juanzi Li, Minlie Huang, Yuxiao Dong, and Jie Tang. Glm-4.5v and glm-4.1v-thinking: Towards versatile multimodal reasoning with scalable reinforcement learning, 2026. 3, 6, 7, 8
- [63] Team Wan, Ang Wang, and Baole Ai et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025. 1
- [64] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *ICLR*, 2021. 3
- [65] Jifang Wang, Xue Yang, Longyue Wang, Zhenran Xu, Yiyu Wang, Yaowei Wang, Weihua Luo, Kaifu Zhang, Baotian Hu, and Min Zhang. A unified agentic framework for evaluating conditional image generation, 2025. 2
- [66] Kaihong Wang, Lingzhi Zhang, and Jianming Zhang. Detecting human artifacts from text-to-image models. *arXiv preprint arXiv:2411.13842*, 2024. 3
- [67] Linqing Wang, Ximing Xing, Yiji Cheng, Zhiyuan Zhao, Donghao Li, Tiankai Hang, Jiale Tao, Qixun Wang, Ruihuang Li, Comi Chen, Xin Li, Mingrui Wu, Xincheng Deng, Shuyang Gu, Chunyu Wang, and Qinglin Lu. Promptenhancer: A simple approach to enhance text-to-image models via chain-of-thought prompt rewriting, 2025. 2
- [68] Weiyun Wang and Zhangwei Gao et al. Internvl3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency, 2025. 6, 7
- [69] Zeqing Wang, Qingyang Ma, Wentao Wan, Haojie Li, Keze Wang, and Yonghong Tian. Is this generated person existed in real-world? fine-grained detecting and calibrating abnormal human-body. In *Proceedings of the IEEE/CVF CVPR*, pages 21226–21237, 2025. 3
- [70] Haoning Wu, Shaocheng Shen, Qiang Hu, Xiaoyun Zhang, Ya Zhang, and Yanfeng Wang. Megafusion: Extend diffusion models towards higher-resolution image generation without further tuning. In *WACV*, 2025. 2
- [71] Mingrui Wu, Lu Wang, Pu Zhao, Fangkai Yang, Jianjin Zhang, Jianfeng Liu, Yuefeng Zhan, Weihao Han, Hao Sun, Jiayi Ji, et al. Reprompt: Reasoning-augmented reprompting for text-to-image generation via reinforcement learning. *arXiv preprint arXiv:2505.17540*, 2025. 2
- [72] An Yang and Anfeng Li et al. Qwen3 technical report, 2025. 3
- [73] Zicheng Zhang, Haoning Wu, Chunyi Li, Yingjie Zhou, Wei Sun, Xiongkuo Min, Zijian Chen, Xiaohong Liu, Weisi Lin, and Guangtao Zhai. A-bench: Are llms masters at evaluating ai-generated images?, 2024. 2
- [74] Zhenghao Zhang, Ziyang Zhang, Junchao Liao, Xiangyu Meng, Qiang Hu, Siyu Zhu, Xiaoyun Zhang, Long Qin, and Weizhi Wang. Lato: Landmark-tokenized diffusion transformer for fine-grained human face editing. *arXiv preprint arXiv:2509.25731*, 2025. 1
- [75] Zangwei Zheng, Xiangyu Peng, Tianji Yang, Chenhui Shen, Shenggui Li, Hongxin Liu, Yukun Zhou, Tianyi Li, and Yang You. Open-sora: Democratizing efficient video production for all. *arXiv preprint arXiv:2412.20404*, 2024. 1
- [76] Kaiwen Zhu, Jinjin Gu, Zhiyuan You, Yu Qiao, and Chao Dong. An intelligent agentic system for complex image restoration problems. In *The Thirteenth ICLR*, 2025. 3
- [77] Xizhou Zhu, Yuntao Chen, Hao Tian, Chenxin Tao, Weijie Su, Chenyu Yang, Gao Huang, Bin Li, Lewei Lu, Xiaogang Wang, Yu Qiao, Zhaoxiang Zhang, and Jifeng Dai. Ghost in

the minecraft: Generally capable agents for open-world environments via large language models with text-based knowledge and memory. *arXiv preprint arXiv:2305.17144*, 2023. [3](#)

- [78] Le Zhuo, Liangbing Zhao, Sayak Paul, Yue Liao, Renrui Zhang, Yi Xin, Peng Gao, Mohamed Elhoseiny, and Hongsheng Li. From reflection to perfection: Scaling inference-time optimization for text-to-image diffusion models via reflection tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15329–15339, 2025. [2](#)