

TIIF-Bench: How Does Your T2I Model Follow Your Instructions?

Xinyu Wei^{1,3}, Jinrui Zhang^{1,3}, Zeqing Wang^{1,3}, Hongyang Wei^{2,3}, Zhen Guo^{1,3}, Bairui Li^{1,3}, and Lei Zhang^{1,3}

¹ The Hong Kong Polytechnic University, Hong Kong SAR, China
`allen.wei@connect.polyu.hk`, `john.zhang@polyu.edu.hk`

² Tsinghua University, Shenzhen, China

³ OPPO Y-Lab, Shenzhen, China

Abstract. The rapid advancements of Text-to-Image (T2I) models have ushered in a new phase of AI-generated content, marked by their growing ability to interpret and follow user instructions. However, existing T2I model evaluation benchmarks fall short in limited prompt diversity and complexity, as well as coarse evaluation metrics, making it difficult to evaluate the fine-grained alignment performance between textual instructions and generated images. In this paper, we present **TIIF-Bench** (Text-to-Image Instruction Following Benchmark), aiming to systematically assess T2I models’ ability in interpreting and following intricate textual instructions. TIIF-Bench comprises 5,000 prompts organized along multiple dimensions and categorized into three levels of difficulty and complexity. To rigorously evaluate robustness to prompt length, each prompt is provided in both short and long versions with identical core semantics. We further propose a novel Global Normalized Edit Distance (GNED) metric for text rendering and provide aspect-ratio-diverse reference images for each prompt to assess style control. In addition, we collect 100 high-quality designer-level prompts covering diverse scenarios for comprehensive evaluation. To enable scalable and fine-grained evaluation, we explore the best paradigm for leveraging the world knowledge encoded in large Vision-Language Models (VLMs) as automated binary evaluators. Through extensive ablations, we develop a fully reproducible evaluator that provides interpretable reasoning and reliable verification, enabling our benchmark to discern subtle variations in T2I model outputs. Through comprehensive benchmarking of mainstream T2I models on TIIF-Bench, we analyze the strengths and weaknesses of current T2I systems and reveal the limitations of existing evaluation benchmarks.

Keywords: T2I Generation · Benchmarking · Reward Signal

1 Introduction

Text-to-Image (T2I) generation has emerged as a cornerstone of multimodal AI, enabling the translation of abstract textual concepts into detailed visual




Fig. 1: Prompt-length sensitivity example. Short prompt: “The birds are more numerous than the fish.” Long prompt: “The birds, with their feathers catching the gentle light of dawn, vastly outnumber their aquatic counterparts, the fish, which glide silently beneath the rippling surface of the water, their sleek forms moving like shadows in the depths below.” PixArt-Sigma and SD 3.5 show clear sensitivity to prompt length, producing different results for semantically equivalent prompts. In contrast, DALL-E 3 and GPT-Image-1 remain robust, maintaining consistent instruction-following performance.

content, advancing applications from digital art to scientific visualization. Recent T2I models can be categorized into 3 main paradigms. Diffusion-based methods—exemplified by Stable Diffusion [17, 48], PixArt [7, 10], FLUX [30], SANA [9, 67], Qwen-Image [63], and others [4, 14, 28, 37, 40, 55, 78]—leverage U-Net or Diffusion-Transformer backbones to iteratively denoise Gaussian noise into photorealistic images, achieving strong visual fidelity and diversity. Autoregressive (AR) approaches, such as LlamaGen [52], Infinity [22], VAR [54], and other influential open-source models, treat images as token sequences and synthesize them through next-token prediction or scale-progressive generation. Unified-model (UM) approaches, including the Show-o series [68, 69], Janus [64] and Janus-Pro [12], Lumina-mGPT series [42, 70], Bagel [16] and the EMU series [15], unify generation and understanding within a single transformer backbone, aiming to mutually enhance multimodal reasoning and image synthesis. Very recently, commercial T2I systems such as the Nano-Banana series [19–21], GPT-Image-1 [27], Imagen 3 [3], and MidJourney [53] have further propelled T2I generation to a new level. In particular, GPT-Image-1 [27] demonstrates strong instruction-following capability, enabling it to interpret highly intricate prompts and produce visually precise, stylistically coherent images within a single conversational interaction. Meanwhile, the Nano-Banana series [19–21] pushes the boundaries of generative reasoning and world knowledge integration in T2I models.

With the rapid development of T2I techniques, how to comprehensively evaluate the performance, especially the instruction-following capability, of modern T2I models has become an important issue. Existing research on the evaluation of T2I models generally falls into two complementary categories. **Preference-alignment approaches**, such as VQAScore [41], HPSv2 [65], HPSv3 [43], VisionReward [71],




Fig. 2: Semantic uniqueness after de-duplication with a cosine similarity threshold of 0.85. Many existing benchmarks contain substantial semantic redundancy, with fewer than 60% semantically unique prompts. While Gecko [62] and ConceptMix [66] achieve higher semantic uniqueness, they mainly focus on compositional evaluation. In contrast, TIIF-Bench achieves the highest semantic uniqueness while covering multiple dimensions, including composition, reasoning, text rendering, and style control.

UnifiedReward [60], and UnifiedReward-CoT [59], leverage learned reward models to align evaluation with human preferences. **Benchmark-driven approaches**, exemplified by TIFA [25], DSG-Bench [13], DPG-Bench [24], Gecko [62], CompBench++ [26], GenEval [18], and GenAI Bench [32], construct structured prompt sets across compositional dimensions (*e.g.*, object attributes, relations, and numeracy) and employ CLIP-based metrics for automatic evaluation. With the emergence of more capable models such as GPT-Image-1 [27], the overall capability of T2I systems has significantly improved, leading to a new generation of more challenging and comprehensive benchmarks. Recent works—including OneIG-Bench [6], LongT2I-Bench [72], and GenEval2 [29], among others [73]—introduce more complex evaluation settings, while some studies further incorporate world knowledge and reasoning into the evaluation process [11, 35, 36, 44, 51].

While existing benchmarks have advanced T2I evaluation, they still exhibit several limitations. **First**, although T2I models often show different performance when given prompts with identical semantics but different lengths, most benchmarks provide only a single prompt version, preventing systematic study of prompt-length robustness (Fig. 1). **Second**, many benchmarks contain substantial semantic redundancy, limiting their ability to evaluate model generalization across diverse concepts (Fig. 2). **Finally**, current evaluation protocols remain inadequate for fine-grained instruction following. Traditional expert scorers such as CLIP or BLIP capture only coarse image-text alignment and often disagree with human judgments (Fig. 3). Some recent works employ large vision-language models (VLMs) as evaluators [26], but their evaluation queries are typically too coarse to fully exploit VLM reasoning ability. In addition, several benchmarks rely on closed-source APIs [36, 66, 73], making evaluation difficult to reproduce. Directly applying general-purpose open-source VLMs (*e.g.*, Qwen2.5-VL [46]) also yields unreliable judgments without task-specific adaptation.

To address these challenges, we introduce TIIF-Bench (**T**ext-**T**o-**I**mage **I**nstruction **F**ollowing **B**enchmark), a benchmark designed for fine-grained evaluation of

Prompt: a boy on the bottom of a bee															
GPT-Image-1				DALL E 3				Flux.1 Pro				MidJourney v6			
CLIP	BLIP	UNiDet	VQAS	CLIP	BLIP	UNiDet	VQAS	CLIP	BLIP	UNiDet	VQAS	CLIP	BLIP	UNiDet	VQAS
32.08	36.23	69.06	63.81	31.08	71.67	0	83.96	30.25	20.70	0	67.27	33.96	54.36	0	88.77
SANA 1.5				SANA Sprint				Pixart Sigma				SD 3.5			
CLIP	BLIP	UNiDet	VQAS	CLIP	BLIP	UNiDet	VQAS	CLIP	BLIP	UNiDet	VQAS	CLIP	BLIP	UNiDet	VQAS
33.15	48.29	0	71.75	31.79	45.73	0	75.93	31.42	70.44	0	86.20	31.71	78.89	0	88.77
Human: Image-1 CLIP: MidJ v6 BLIP: SD 3.5 VQA Score: Midj v6 UNiDet: failed in most cases!															

Fig. 3: Limitations of expert scorers widely used in classic T2I benchmarks. CLIP [47], BLIP [34], and VQAScore [41] capture only coarse image-text alignment and struggle to evaluate fine-grained instruction following. UniDet [77] often fails to detect objects in complex generated scenes, leading to scores inconsistent with human judgments.

instruction-following capability in T2I models. We construct prompts using an attribute-composition strategy, combining ten concept pools with six compositional dimensions to ensure semantic uniqueness and diverse sentence structures. Two additional dimensions—text rendering and style control—are introduced to evaluate capabilities beyond traditional compositional reasoning. The benchmark also includes 100 designer-level prompts capturing real-world human aesthetic priors. Each prompt is provided in both short and extended versions to evaluate robustness to prompt length, resulting in a total of 5,000 prompts. In our evaluation framework, each prompt is paired with a set of attribute-specific yes/no questions that serve as a checklist, enabling vision-language models (VLMs) to perform fine-grained verification rather than producing a single coarse score. We further investigate best paradigms for employing VLMs as automated evaluators and develop our own TIIF-Evaluator (see Section 3.4). In addition, text rendering accuracy is measured using the proposed Global Normalized Edit Distance (GNED) metric, while style control is evaluated through similarity with reference images using a weighted combination of CSD [50], DINOv3 [49], and SigLIP2 [56].

We benchmark popular closed-source models, including Nano-Banan series [19, 20], GPT-Image-1 [27], DALL E-3 [5], MidJourney V7 [53], alongside leading open-source diffusion-based models [7–9, 17, 30, 31, 45, 48, 63, 67], and UM pipelines [12, 15, 16, 42, 69, 70]. Our analysis reveals several notable findings. Models that achieve

higher overall scores on TIIF-Bench tend to exhibit stronger robustness to prompt length variations, whereas lower-performing models are significantly more sensitive to such changes. This observation suggests a positive correlation between a model’s ability to comprehend textual instructions and the quality of its generated images. Furthermore, although UM approaches generally produce images with lower visual fidelity compared to sota diffusion models, their instruction-following performance is often comparable. This result highlights a key advantage of unified architectures, where integrating understanding and generation within a single transformer backbone allows multimodal reasoning to directly support image synthesis. The contributions of this paper are summarized as follows:

(i) Prompt-design methodology. We identify critical limitations in existing T2I benchmarks, including fixed-length prompts, high semantic redundancy, and limited syntactic diversity. To address these issues, we propose a novel attribute-composition strategy that improves semantic coverage and diversity, and introduce additional evaluation dimensions together with prompt-length variations to systematically assess model robustness.

(ii) Fine-grained evaluation protocol. We investigate effective paradigms for using VLMs as automated evaluators and develop a dedicated TIIF-Evaluator. Leveraging the world knowledge of VLMs, the evaluator performs attribute-specific yes/no verification with interpretable chain-of-thought reasoning, enabling more fine-grained and reliable alignment assessment than previous metrics.

(iii) Empirical insights for T2I evaluation. Through extensive experiments on TIIF-Bench, we uncover consistent patterns in how T2I models with different architectures follow instructions during image generation, providing practical insights for future research on instruction-following image synthesis.

2 Related Work

T2I models have achieved remarkable progress in generating high-quality images. However, how to evaluate their performance, particularly in terms of the instruction-following capability, remains a challenging task. Existing work in T2I evaluation generally falls into two streams: (i) scoring models aligned with human visual preferences, and (ii) benchmarks for structured evaluation.

Scoring Methods Aligned with Human Preferences. CLIPScore [23] is a widely adopted metric for evaluating image-text alignment, but it often fails on complex prompts due to CLIP’s inherent bag-of-words behavior. To address this limitation, VQAScore [41] fine-tunes a CLIP-FlanT5 specifically for T2I image quality evaluation. HPDv2 [65] mitigates biases in human preference data. When used to fine-tune CLIP, it results in HPSv2, which demonstrates improved alignment with human judgments. However, these approaches remain fundamentally constrained by CLIP-based representations: they struggle to capture fine-grained semantic correspondence. VisionReward [71] attempts to address this limitation by learning reward model through hierarchical visual assessment and interpretable weighting. Nevertheless, it is primarily designed for video quality evaluation, which limits its applicability to static T2I tasks. More recently, the UnifiedReward series [59, 60] employs vision-language models

(VLMs) to evaluate overall image–text alignment. However, these approaches typically produce holistic scores for entire images and cannot provide fine-grained verification at the level of specific semantic attributes or binary questions.

Benchmark-Driven Evaluation Frameworks. Early benchmark-driven evaluation frameworks, such as TIFA [25], DSG-Bench [13], DPG-Bench [24], Gecko [62], CompBench++ [26], GenEval [18], and GenAI Bench [32], construct structured prompt sets across compositional dimensions (e.g., object attributes, relations, and numeracy) and typically rely on CLIP-based metrics or relatively weak VQA models, such as mPLUG [33], for automatic evaluation. As T2I models have rapidly improved—particularly with the emergence of systems such as GPT-Image-1 [27]—recent benchmarks have become increasingly comprehensive and challenging. Works such as OneIG-Bench [6], LongT2I-Bench [72], and GenEval2 [29], among others [1, 61, 73], introduce more complex evaluation settings. In addition, several studies incorporate world knowledge and reasoning into T2I evaluation, including WISE [44], T2I-Bench [11], and T2IReasonBench [51], as well as other benchmarks emphasizing reasoning capabilities [35, 36].

3 Construction of TIIF-Bench

3.1 Limitations of Current T2I Evaluation Benchmarks

As T2I models continue to evolve, accurately evaluating their ability to follow natural language instructions has become increasingly important. However, existing benchmarks still suffer from several limitations, which can be broadly categorized into **prompt-design flaws** and **evaluation-protocol flaws**.

Prompt Design Flaws. We identify two major issues in the prompt design of existing benchmarks. **First**, many benchmarks rely on templated prompts with fixed or narrowly distributed lengths. For example, in COMPBENCH++, the 2D and 3D evaluation dimensions contain a total of 2,000 prompts, yet their lengths are restricted to only four values (5, 6, 7, and 8 words). Such limited variation prevents evaluation of model robustness to prompt length, despite empirical evidence that T2I models can be highly sensitive to prompt length variations, as shown in Fig. 1. **Second**, existing benchmarks often exhibit substantial semantic redundancy. To quantify this issue, we extract CLIP [47] text embeddings for all prompts and compute pairwise cosine similarity. As shown in Fig. 2, many benchmarks contain a large proportion of semantically redundant prompts.

Evaluation Protocol Flaws. We further identify two limitations in current evaluation protocols. **First, reliance on weak expert models.** Most benchmarks rely on CLIP [47] to measure image–text alignment. However, CLIP is known to behave similarly to a “bag-of-words” model [41], often failing to capture fine-grained semantic differences in instruction prompts [38, 39, 74]. Other expert models, such as BLIP [34], offer limited improvements for complex alignment, while object detectors [57, 58] like UniDet [77] frequently struggle with complex scenes generated by modern T2I models (see Fig. 3). **Second, coarse utilization of VLM evaluators.** Some recent benchmarks employ strong VLMs to assess image–text alignment. However, these evaluations typically rely on a single




Fig. 4: Prompt construction and evaluation pipeline of TIIF-Bench. In addition to systematically reusing prompts from existing benchmarks, we incorporate text-rendering prompts from Lex-Art [76], as well as curated style-control prompts (each paired with reference images of multiple aspect ratios) and real-world designer-level prompts. All prompts are further expanded to produce long-form counterparts.

generic query, such as “Evaluate how well the image aligns with the text prompt: $\{xxx\}$.” [26]. Such coarse queries fail to decompose the rich multi-attribute semantics present in modern T2I prompts. Moreover, explicitly including the full T2I prompt (*i.e.*, the caption) in the evaluation question can induce hallucinations in VLM responses, as illustrated in Fig. A1 in the **Appendix**. In addition, some benchmarks rely on closed-source APIs as evaluators, making their evaluation pipelines difficult to reproduce. Others directly adopt general-purpose VLMs (*e.g.*, Qwen2.5-VL [46]) as off-the-shelf evaluators. However, our experiments indicate that without task-specific adaptation, such models struggle to assess the complex prompts in TIIF-Bench.

3.2 Prompts of TIIF-BENCH

To address the limitations of existing T2I evaluation benchmarks, we build upon the hierarchical prompt taxonomies of prior work by introducing a two-stage process: **concept pool construction** followed by **attribute composition**.

Table 1: Based on difficulty, we define **36 attribute-pool combinations**, grouped into three levels: **Basic Following**, **Advanced Following**, and **Designer Level Following**. **Basic Following** prompts are constructed by combining attributes/relations and objects from a single attribute group, while **Advanced Following** prompts involve cross-group composition, as well as two specialized dimensions—*text rendering* and *style control*—to evaluate text rendering and aesthetic coherence. Finally, 100 real-world **Designer Level Following** prompts are manually curated. The rightmost column reports the number of prompts per subclass, all of which are further expanded into long-form variants.

Level	Dimension	Combination Policy	Count
Basic Following	Attribute	Shape+Color, Color+Texture, Texture+Color, Shape+Texture	{100,50,50,100}
	Relation	2D Spatial, 3D Perspective, Action+2D, Action+3D	{75,75,75,75}
	Reasoning	Numeracy, Negation, Differentiation, Comparison	{75,75,75,75}
Advanced Following	Attribute + Relation	Action+Color, Action+Texture, Color+2D, Color+3D, Shape+2D, Shape+3D, Texture+2D, Texture+3D	{50, 50, 33, 33, 33, 33, 33, 33}
	Attribute + Reasoning	Numeracy+{Color, Texture}, Comparison+{Color, Texture}, Differentiation+{Color, Texture}, Negation+{Color, Texture}	{40, 40, 40, 40, 40, 40, 40, 40}
	Relation + Reasoning	Numeracy+{2D, 3D}, Comparison+{2D, 3D}, Differentiation+{2D, 3D}, Negation+{2D, 3D}	{40, 40, 40, 40, 40, 40, 40, 40}
	Text Generation	—	{150}
	Style Control	—	{150}
Designer Level	Complex	—	{100}

To more comprehensively assess instruction-following performance, we further introduce three new evaluation dimensions: text rendering, style control, and real-world designer-level prompts. Each prompt is processed by a length augmentation module that generates both concise and verbose versions, enabling evaluation across varying linguistic lengths. The data generation pipeline is shown in Fig. 4.

Concept Pool Construction. We first group the prompts in existing benchmarks based on their semantics and leverage GPT-4o to extract the underlying *object-attribute/relation pairs*, forming a set of dimension-specific concept pools. In total, we construct **10 concept pools** from existing benchmarks, categorized them into **three groups**, as summarized in Tab. A1 in **Appendix**.

Attribute Composition. Building upon concept pools, we generate prompts by randomly combining attributes from each pool, leveraging GPT-4o to compose them into natural instructions. As illustrated in Tab. 1, we define **36 distinct combinations**, each paired with a dedicated meta-prompt to guide GPT-4o to assemble instructions. This sampling strategy ensures that the resulting prompts are both semantically unique and compositionally diverse. Prompts that combine elements drawn from a *single* concept-pool group are classified as **Basic Following**. In contrast, **Advanced Following** prompts intertwine elements taken from *different* concept-pool groups, yielding more intricate compositions.

New Dimensions. To extend evaluation beyond conventional capabilities, we introduce three additional dimensions: *text rendering*, *style control*, and *designer-level prompts*. (i) **Text rendering** evaluates a model’s ability to reproduce typographic elements, using prompts sourced from the *Lex-Art* corpus [76]. (ii)

Style control assesses the model’s capacity to follow high-level artistic directives, with prompts curated from leading AIGC communities. (iii) **Designer-level prompts** involve complex instructions that incorporate practical constraints and domain-specific knowledge, collected through manual annotation. The text rendering and style control dimensions are grouped into the **Advanced Following** set, while the designer-level prompts form the **Designer-Level Following** set.

Length Augmentation. Finally, for each generated prompt, we construct a long-form variant by expanding the content through natural language paraphrasing and elaboration, while faithfully preserving its original semantics.

TestMini Subset. Following the above pipeline, the full bench contains 5,000 prompts and 12,074 binary questions. Detailed statistics on question-list lengths and entity densities are provided in the **Appendix**. For efficient evaluation, we further construct a subset of **554 prompts** and **1446 binary questions**. Carefully sampled to maintain balanced category proportions, it achieves higher semantic uniqueness than the full benchmark as illustrated in Fig. 2.

3.3 Evaluation Method of TIIF-Bench

To overcome the limitations of traditional expert scorers (e.g., CLIP and open-set object detectors) in evaluating generated images, we propose an attribute-specific, fine-grained evaluation protocol. Our method leverages the world knowledge encoded in large vision-language models (VLMs) to assess the alignment between textual instructions and generated images. Specifically, we employ a dedicated VLM-based evaluator, termed **TIIF Evaluator**, to perform attribute-level verification. The design and training of this evaluator are described in Sec. 3.4.

As illustrated in Fig. 4, given an input prompt, we first extract its N core concepts $C = \{c_i\}_{i=1}^N$ from the constructed concept pools. For each concept c_i , we employ an closed-source sota LLM to generate a corresponding binary question q_i , forming a set of evaluation questions $Q = \{q_i\}_{i=1}^N$ with ground-truth answers $A = \{\hat{a}_i\}_{i=1}^N$. The generated image, together with the questions Q , is then input to the **TIIF Evaluator**, which produces N predicted answers. The final evaluation score s for the generated image is computed as

$$s = \frac{1}{N} \sum_{i=1}^N \mathbb{I}[a_i = \hat{a}_i],$$

where $\mathbb{I}[\cdot]$ denotes the indicator function and a_i represents the evaluator’s prediction for question q_i . For **designer-level** prompts, which often encode complex real-world constraints and human priors, we manually construct a long list of tailored evaluation questions to ensure reliable assessment.

For **style control** dimension, we provide reference images covering common aspect ratios (16:9, 5:4, 1:1, 4:5, and 9:16). Reference images in the TestMini subset are generated using Nano-Banana-Pro, while those in the full benchmark are generated using Qwen-Image, followed by manual verification. During evaluation, the generated image is compared with the reference image of the closest aspect ratio using the average similarity of CSD [50] and DINOv3 [49].

For **text rendering** evaluation, we propose a Global Normalized Edit Distance (GNED) metric. Let $P = \{p_1, \dots, p_m\}$ denote the set of text tokens specified in the prompt and $G = \{g_1, \dots, g_n\}$ the set of OCR-extracted tokens from the generated image. GNED computes the minimal character-level normalized edit distance (NED) between words in P and G using the Hungarian algorithm to obtain an optimal bipartite matching $\mathcal{M}$. A penalty term $|m - n|$ is added to account for unmatched words (e.g., missing or hallucinated text), and the total cost is normalized by $\max(m, n)$:

$$\text{GNED}(P, G) = \left(\sum_{(i,j) \in \mathcal{M}} \text{NED}(p_i, g_j) + |m - n| \right) / \max(m, n).$$

GNED ranges from 0 to 1, where 0 indicates perfect alignment and 1 indicates maximal deviation. Compared with existing metrics such as OCR Recall and PNED [76], GNED robustly penalizes both over-generation and omission, enabling consistent evaluation across samples with varying text lengths. Illustrative examples of GNED are provided in the **Appendix**.

3.4 Ablation and Training of TIIF Evaluator

To construct a diverse training corpus for the evaluator, we collect model outputs from **30 T2I models** spanning different release periods, architectures, generation capabilities, and resolutions. Each model is rolled out twice on all 5,000 prompts in TIIF-Bench, producing approximately 300K generated images and 724K binary questions. To build a reliable test set, we manually annotate 300 images, resulting in 560 balanced questions (280 “yes” and 280 “no”). This set is used exclusively for evaluator validation. On this test set, GPT-4o [27] achieves an accuracy of **90.18%**, indicating that it can serve as a reference evaluator. In contrast, the off-the-shelf Qwen3VL-8B [2] reaches only **84.11%**, suggesting that direct use of a general-purpose VLM may be insufficient for reliable evaluation. We therefore use GPT-4o to annotate the remaining data (excluding the test set), during which chain-of-thought (CoT) reasoning is also generated. After filtering and balancing, we obtain a training set containing 109K positive and 109K negative binary questions. We then investigate three training paradigms for Qwen3VL-8B:

1. **Direct prediction:** the model directly outputs the binary answer.
2. **Representation-based classification:** the final hidden state is extracted and fed into a two-layer MLP for binary classification supervision [75].
3. **Reasoning-based prediction:** the model first generates chain-of-thought reasoning and then outputs the final answer.

Each paradigm is trained for one epoch. On the held-out test set, direct prediction achieves **89.46%** accuracy, while representation-based classification and reasoning-based prediction reach **91.43%** and **91.07%**. Both latter approaches meet the reliability requirement for serving as an automated evaluator. Despite comparable accuracy, the reasoning-based paradigm provides interpretable chain-of-thought explanations that reveal where instruction-following failures occur.

This additional interpretability is valuable for diagnosing model behavior and can potentially serve as reward signals for future model optimization. Therefore, we adopt the reasoning-based paradigm to construct the final **TIIF Evaluator**.

4 Experiments

4.1 Performance of T2I Models on TIIF-Bench

We benchmark TIIF-Bench on a diverse suite of widely used *closed-source* and leading *open-source* text-to-image models. The results are summarized in Tab. 2. For each of the three difficulty levels, *basic following*, *advanced following*, and *designer-level*. We report the average score of their corresponding dimensions under both short and long prompt settings. We further compute the overall average scores across all dimensions and difficulty levels. Note that GNED is a lower-is-better metric, the values reported in the *Text* column are transformed as $1 - \text{GNED}$ so that higher scores consistently indicate better performance. From Tab. 2, we derive several key observations, which are discussed below.

Diffusion-based Open-Source Models (i) Overall performance. From the upper panel of Tab. 2, we observe that Qwen-Image [63] and FLUX.2 Dev [31] exhibit the strongest instruction-following capability among diffusion-based open-source models. Closely following these two models are LuminaImage2 [45] and FLUX.1 Dev [30]. Overall, models built upon the MMDiT architecture continue to dominate the current landscape of high-quality image generation. **(ii) Text Rendering.** Among diffusion-based open-source models, only a subset explicitly supports in-image text rendering, including Qwen-Image, FLUX-1 Dev, FLUX-2 Dev, SANA 1.5 [67] (and SANA Sprint [9]), and the Stable Diffusion family (SD-3 [48], SD-3.5 [17]). While SD-3 and SD-3.5 achieve competitive performance with short prompts, their effectiveness deteriorates when prompts become longer. This behavior is likely due to the increased difficulty of isolating textual elements from the broader descriptive context. **(iii) Style Control.** Prompt length has contrasting effects on different models. In systems such as PixArt-Alpha, the internal vocabulary for style control appears limited; for example, style tokens like *Ghibli* or *Cyberpunk* are weakly grounded in the learned representation. Consequently, short prompts that rely solely on such keywords often fail to produce coherent stylistic outputs. In contrast, longer prompts that provide richer visual descriptions can partially compensate for this limitation and significantly improve generation quality. Conversely, models such as SD-3.5 and PixArt-Sigma [7] are typically trained on datasets where stylistic instructions appear as concise label-like tokens. When these style cues are embedded within longer descriptive prompts, their salience can be diluted, leading to reduced stylistic fidelity. **(iv) Designer-Level Prompts.** Designer-level prompts contain the densest and most diverse set of requirements, providing the most comprehensive evaluation of a model’s instruction-following capability. Consequently, the ranking of models under this setting closely mirrors the overall performance ranking.

Table 2: Performance of closed-source models and state-of-the-art open-source models on T2IF-Bench. Evaluated systems are grouped into (i) **diffusion-based** open-source models, (ii) **unified** open-source models, and (iii) **closed-source** models; within each group, the highest score is shown in bold and color-coded for that group.

Model	Overall		Basic Following								Advanced Following												Designer		
			Avg		Attribute		Relation		Reason		Avg		Attr.+Rela.		Attr.+Rea.		Rela.+Rea.		Style		Text				
	short	long	short	long	short	long	short	long	short	long	short	long	short	long	short	long	short	long	short	long	short	long	short	long	
Diffusion based Open-Source Models																									
Qwen-Image	87.0	87.5	87.1	87.5	87.3	88.0	90.2	90.8	83.9	84.5	86.9	87.6	83.2	83.8	80.4	81.2	83.1	83.7	96.7	97.2	79.0	81.0	89.9	90.6	
FLUX.2 dev	87.2	87.3	86.6	86.1	87.3	87.1	90.0	90.1	82.5	81.9	89.3	88.9	85.5	85.4	90.5	89.8	81.8	81.6	100.0	99.5	71.0	73.0	91.3	90.9	
LuminaImage2	71.2	70.7	79.7	78.9	74.7	74.4	85.0	85.1	79.0	78.4	66.9	66.0	78.6	78.2	75.9	75.7	70.7	70.0	86.7	86.6	20.0	19.0	66.7	66.2	
FLUX.1 dev	72.6	71.0	82.8	80.6	86.9	82.5	88.8	81.5	74.4	73.9	65.1	69.6	68.2	74.4	73.1	72.7	70.3	70.9	67.5	65.9	52.0	55.0	69.8	72.2	
SD 3	66.5	67.0	80.3	76.6	82.6	80.8	81.5	79.7	70.1	75.1	62.7	59.2	62.4	63.1	70.1	69.7	50.2	58.9	65.8	77.6	52.0	28.0	64.2	66.4	
SD 3.5	69.5	68.5	79.4	78.9	79.2	78.4	80.7	84.9	72.2	80.6	67.0	60.6	66.1	63.4	73.1	75.8	59.3	61.0	74.6	62.4	55.0	43.0	66.0	65.6	
SANA Sprint	64.4	57.9	75.7	72.4	74.6	70.4	82.8	73.4	71.7	70.8	58.5	51.0	56.6	54.1	69.4	67.5	61.4	64.3	81.1	58.9	8.4	7.4	66.3	59.4	
SANA 1.5	66.9	67.5	80.6	76.3	81.2	76.8	84.9	84.6	74.7	69.1	62.4	61.7	64.3	55.6	71.5	72.0	64.2	65.1	79.3	81.1	34.0	33.0	72.6	67.8	
PixArt-alpha	45.9	49.6	54.6	62.5	53.1	55.8	65.4	72.9	49.8	53.9	39.5	44.2	37.2	42.4	58.4	51.7	41.3	46.5	51.1	75.4	2.0	0.0	44.8	54.3	
PixArt-sigma	61.0	59.7	69.6	74.2	70.8	77.9	76.1	78.4	66.9	70.9	58.7	48.6	64.6	57.6	66.3	63.2	65.6	55.7	82.0	71.3	1.0	0.0	61.2	53.8	
Lumina-Next	52.4	51.4	63.7	67.6	56.3	60.4	68.9	71.2	70.4	66.1	45.7	46.8	50.7	44.7	52.3	58.6	43.7	55.9	69.1	65.5	1.0	0.0	48.9	48.1	
Hunyuan-DiT	50.4	54.6	68.8	70.6	65.2	71.3	77.1	75.2	63.4	64.6	43.7	44.7	49.3	42.9	58.6	63.4	47.1	52.6	55.4	72.2	0.0	1.0	39.1	45.3	
Open-Source Unified Models																									
Show-o	60.7	58.2	72.5	77.2	73.7	81.1	78.4	79.4	66.9	68.2	54.6	49.7	62.5	56.3	69.9	68.3	65.7	57.5	62.2	67.6	7.0	14.0	56.4	49.8	
JanusPro	65.5	66.6	80.3	77.1	78.7	83.7	77.9	74.6	81.3	78.4	60.6	58.1	65.1	57.7	69.4	72.4	68.1	59.4	61.3	69.2	27.0	32.0	67.4	59.3	
Show-o2	69.2	67.7	76.5	74.7	72.7	71.8	82.6	81.4	73.9	72.9	65.2	63.6	72.4	71.7	67.6	66.2	70.2	69.1	90.0	89.7	22.0	20.0	66.5	65.2	
Lumina-mGPT	66.1	64.6	77.7	76.1	72.7	71.9	83.4	82.3	76.8	75.4	61.5	59.6	65.5	64.8	65.9	64.9	64.7	63.5	100.0	99.5	8.0	5.0	56.6	55.1	
Lumina-mGPT2	58.4	56.8	68.5	67.2	66.0	64.8	75.2	74.0	64.3	63.1	51.8	50.5	61.2	60.0	55.1	53.9	60.4	59.2	73.3	72.4	6.0	5.0	62.2	61.0	
Bagel	71.0	71.3	81.5	80.3	82.8	83.2	82.7	80.2	79.6	76.9	70.0	72.4	74.1	75.2	67.0	70.3	71.8	75.1	86.4	83.6	24.0	32.0	68.0	68.2	
Emu3.5	85.9	84.8	84.8	83.6	84.0	82.9	89.9	88.7	80.5	79.4	87.5	86.3	83.3	82.1	79.7	78.6	81.3	80.1	100.0	99.2	80.0	80.0	85.7	84.5	
Emu3.5-Image	88.4	87.1	86.2	85.1	85.3	84.2	90.2	89.1	83.2	82.1	90.6	89.4	88.6	87.5	84.5	83.4	83.8	82.7	100.0	99.1	82.0	83.0	90.5	89.4	
Closed-Source Models																									
DALL-E 3	72.2	70.2	79.7	77.8	80.6	79.1	82.3	77.9	74.8	78.2	72.4	68.5	74.7	66.6	71.1	72.4	62.5	61.7	88.8	87.5	44.0	52.0	72.0	62.5	
MidJourney v7	67.7	62.7	76.5	77.6	78.9	81.3	83.5	75.8	73.9	68.4	65.6	59.3	66.2	63.9	82.2	70.7	62.1	63.7	84.6	78.7	42.0	26.0	70.3	62.7	
GPT-Image-1	86.1	87.6	89.7	90.6	92.4	86.3	83.7	86.0	95.4	98.7	87.5	89.8	86.3	90.2	86.4	85.0	84.7	84.6	91.2	92.2	79.0	81.0	90.8	92.5	
Nano-Banana	89.2	88.5	89.7	89.3	84.7	84.4	91.0	91.4	93.2	93.0	90.4	90.1	87.7	87.9	87.7	88.0	88.0	88.3	96.7	96.9	82.0	82.0	92.9	93.2	
NB-Pro	93.1	93.3	93.2	93.5	90.7	91.0	94.0	94.3	94.8	94.5	93.6	93.9	91.6	91.9	91.1	91.4	90.9	91.2	96.7	97.0	84.0	85.0	95.2	95.5	

(v) Robustness to Prompt Length. Top-performing models demonstrate strong robustness to variations in prompt length, producing consistent results for both short and long versions of semantically equivalent prompts. In contrast, weaker models such as SD-3 and PixArt-Alpha exhibit substantial performance discrepancies between the two settings. This observation suggests that the ability to accurately interpret complex instructions is positively correlated with the overall image generation capability of a T2I model.

Unified Open-Source Models (i) Overall Performance. From the middle panel of Tab. 2, we observe that the Emu3.5 series [15] achieves the strongest overall instruction-following performance among unified open-source models. This advantage likely stems from its large-scale pretraining and unified training paradigm that jointly optimizes image generation and multimodal understanding.

(ii) **Text Rendering.** Autoregressive architectures are inherently less effective at rendering text within images. Except for Emu3.5, most unified models exhibit significantly weaker text-rendering performance compared with diffusion-based open-source state-of-the-art models of similar scale and release time. (iii) **Style Control.** Emu3.5 also demonstrates strong capability in style control. Its unified transformer backbone, trained jointly for generation and understanding, appears to facilitate more coherent global visual interpretation, which benefits style-conditioned generation. (iv) **Designer-Level Prompts.** Except for Emu3.5, most unified models struggle to handle complex real-world prompts at the designer level. Nevertheless, the ranking of models under this setting remains broadly consistent with their overall performance. (v) **Robustness to Prompt Length.** Stronger autoregressive models maintain relatively stable performance across both short and long prompt settings, whereas weaker models exhibit substantial performance gaps. This observation further supports the connection between robust instruction understanding and high-quality image generation.

Diffusion-based vs. AR-based Open-Source Models Although AR-based models tend to produce images with lower visual fidelity, their autoregressive architecture, jointly trained on generation and understanding tasks, grants them strong instruction-following capabilities. For instance, Janus-Pro [12] outperforms diffusion-based model PixArt-Sigma [7] on TIIF Bench. We present qualitative examples in Fig. A7 in **Appendix** to illustrate this finding.

Closed-Source Models (i) **Overall Performance.** From the bottom panel of Tab. 2, we observe that the Nano-Banana series [19, 20] and GPT-Image-1 [27] demonstrate substantially stronger instruction-following capability than other models. This advantage likely stems from their unified architectures and superior language understanding. These models perform consistently well across multiple dimensions, including object attributes, spatial relations, and complex logical reasoning. In particular, they achieve exceptionally high success rates in text rendering, with no other model reaching comparable performance. (ii) **Text Rendering.** Benefiting from their strong multimodal comprehension capabilities, the Nano-Banana series [19, 20] and GPT-Image-1 [27] lead by a wide margin in the text rendering dimension, consistently producing accurate and legible in-image text. (iii) **Style Control.** With access to extensive knowledge and strong semantic understanding, the Nano-Banana series and GPT-Image-1 also dominate the style control, reliably generating images that align with the requested artistic styles. (iv) **Designer-Level Prompts.** Closed-source models generally perform well on designer-level prompts, which involve dense and diverse instruction requirements. This advantage likely reflects exposure to large-scale, high-quality training data. In particular, the Nano-Banana series and GPT-Image-1 show clear performance gains over open-source state-of-the-art models under this challenging setting. (v) **Robustness to Prompt Length.** Closed-source models exhibit strong robustness to variations in prompt length. Their performance remains stable across both short and long prompt settings, suggesting that well-trained

multimodal architectures with large-scale data are better equipped to maintain consistent instruction understanding and generation quality.

4.2 TIIF-Bench vs. Existing Benchmarks

To further assess the effectiveness of our TIIF-Bench evaluation methodology, we select four strong open-source models (SD 3.5, SANA 1.5, PixArt-Sigma, and Janus-Pro) and four representative closed-source models (GPT-Image-1, DALL·E 3, NanoBanana, and MidJourney v7), and evaluate them on three benchmarks: CompBench++, GenAI Bench, and our proposed TIIF-BENCH. For each model, we generate images for all prompts. The resulting images are then scored using the evaluation method of each benchmark. Full scoring results are provided in Tab. A2 and Tab. A3 in **Appendix**. To further assess how well each benchmark aligns with human preferences, we conduct a user study as described in **Appendix** A4.3. Specifically, ten volunteers are invited and they are asked to rank the outputs of the eight models for each prompt, without knowing the model identities. We then compute the *Spearman* ρ correlation between the benchmark rankings and human-annotated rankings. As shown in Tab. A4 of the **Appendix**, CompBench++ shows only weak alignment with human judgment. In contrast, TIIF-BENCH exhibits near-perfect correlation with human preferences (see Tab. A5 in **Appendix**), highlighting the effectiveness of our fine-grained evaluation framework and the reliability of the **TIIF Evaluator**.

5 Conclusion

We introduce TIIF-BENCH, a comprehensive and hierarchical benchmark for evaluating the instruction-following capabilities of text-to-image (T2I) models. TIIF-Bench covers diverse concept combinations and multiple evaluation dimensions, including text rendering, style control, and designer-level prompts. To enable reliable and fine-grained evaluation, we explore the paradigm of using VLMs as automated evaluators and develop a reproducible reasoning-based evaluator, termed **TIIF Evaluator**. We further propose GNED for evaluating text rendering quality, and provide high-quality aspect-ratio-diverse reference images for style control prompts. Extensive experiments on TIIF-Bench reveal key patterns in how current T2I models follow complex instructions during image generation, offering useful insights for future T2I research.

Limitations. Despite TIIF-Bench’s advantages over existing benchmarks, there remain some limitations. First, the prompts contain mainly common objects. In future work, we will include obscure vocabulary or rare objects to further challenge the instruction following capability of T2I models. Second, currently, all prompts are in English. Incorporating linguistic diversity is an important direction for future work. Last but not least, how could stylistic variations (*e.g.*, formal vs. conversational) impact model performance is not considered yet. An investigation on this problem may offer valuable insight in practical applications.

References

1. An, R., Yang, S., Guo, Z., Dai, W., Shen, Z., Li, H., Zhang, R., Wei, X., Li, G., Wu, W., et al.: Genius: Generative fluid intelligence evaluation suite. arXiv preprint arXiv:2602.11144 (2026)
2. Bai, S., Cai, Y., Chen, R., Chen, K., Chen, X., Cheng, Z., Deng, L., Ding, W., Gao, C., Ge, C., Ge, W., Guo, Z., Huang, Q., Huang, J., Huang, F., Hui, B., Jiang, S., Li, Z., Li, M., Li, M., Li, K., Lin, Z., Lin, J., Liu, X., Liu, J., Liu, C., Liu, Y., Liu, D., Liu, S., Lu, D., Luo, R., Lv, C., Men, R., Meng, L., Ren, X., Ren, X., Song, S., Sun, Y., Tang, J., Tu, J., Wan, J., Wang, P., Wang, P., Wang, Q., Wang, Y., Xie, T., Xu, Y., Xu, H., Xu, J., Yang, Z., Yang, M., Yang, J., Yang, A., Yu, B., Zhang, F., Zhang, H., Zhang, X., Zheng, B., Zhong, H., Zhou, J., Zhou, F., Zhou, J., Zhu, Y., Zhu, K.: Qwen3-vl technical report (2025), <https://arxiv.org/abs/2511.21631>
3. Baldridge, J., Bauer, J., Bhutani, M., Brichtova, N., Bunner, A., Castrejon, L., Chan, K., Chen, Y., Dieleman, S., Du, Y., et al.: Imagen 3. arXiv preprint arXiv:2408.07009 (2024)
4. Betker, J., Goh, G., Jing, L., Brooks, T., Wang, J., Li, L., Ouyang, L., Zhuang, J., Lee, J., Guo, Y., Manassra, W., Dhariwal, P., Chu, C., Jiao, Y., Ramesh, A.: Improving Image Generation with Better Captions
5. Betker, J., Goh, G., Jing, L., Brooks, T., Wang, J., Li, L., Ouyang, L., Zhuang, J., Lee, J., Guo, Y., et al.: Improving image generation with better captions. Computer Science. <https://cdn.openai.com/papers/dall-e-3.pdf> **2**(3), 8 (2023)
6. Chang, J., Fang, Y., Xing, P., Wu, S., Cheng, W., Wang, R., Zeng, X., Yu, G., Chen, H.B.: Oneig-bench: Omni-dimensional nuanced evaluation for image generation (2025), <https://arxiv.org/abs/2506.07977>
7. Chen, J., Ge, C., Xie, E., Wu, Y., Yao, L., Ren, X., Wang, Z., Luo, P., Lu, H., Li, Z.: PixArt- Σ : Weak-to-Strong Training of Diffusion Transformer for 4K Text-to-Image Generation (Mar 2024). <https://doi.org/10.48550/arXiv.2403.04692>
8. Chen, J., Wu, Y., Luo, S., Xie, E., Paul, S., Luo, P., Zhao, H., Li, Z.: PIXART- δ : Fast and Controllable Image Generation with Latent Consistency Models (Jan 2024). <https://doi.org/10.48550/arXiv.2401.05252>
9. Chen, J., Xue, S., Zhao, Y., Yu, J., Paul, S., Chen, J., Cai, H., Xie, E., Han, S.: SANA-Sprint: One-Step Diffusion with Continuous-Time Consistency Distillation (Mar 2025). <https://doi.org/10.48550/arXiv.2503.09641>
10. Chen, J., Yu, J., Ge, C., Yao, L., Xie, E., Wu, Y., Wang, Z., Kwok, J., Luo, P., Lu, H., Li, Z.: PixArt- α : Fast Training of Diffusion Transformer for Photorealistic Text-to-Image Synthesis (Dec 2023). <https://doi.org/10.48550/arXiv.2310.00426>
11. Chen, K., Lin, Z., Xu, Z., Shen, Y., Yao, Y., Rimchala, J., Zhang, J., Huang, L.: R2i-bench: Benchmarking reasoning-driven text-to-image generation (2025), <https://arxiv.org/abs/2505.23493>
12. Chen, X., Wu, Z., Liu, X., Pan, Z., Liu, W., Xie, Z., Yu, X., Ruan, C.: Janus-Pro: Unified Multimodal Understanding and Generation with Data and Model Scaling
13. Cho, J., Hu, Y., Garg, R., Anderson, P., Krishna, R., Baldridge, J., Bansal, M., Pont-Tuset, J., Wang, S.: Davidsonian scene graph: Improving reliability in fine-grained evaluation for text-to-image generation (2024), <https://arxiv.org/abs/2310.18235>
14. Cui, S., Guo, J., An, X., Deng, J., Zhao, Y., Wei, X., Feng, Z.: Idadapter: Learning mixed features for tuning-free personalization of text-to-image models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 950–959 (June 2024)

15. Cui, Y., Chen, H., Deng, H., Huang, X., Li, X., Liu, J., Liu, Y., Luo, Z., Wang, J., Wang, W., Wang, Y., Wang, C., Zhang, F., Zhao, Y., Pan, T., Li, X., Hao, Z., Ma, W., Chen, Z., Ao, Y., Huang, T., Wang, Z., Wang, X.: Emu3.5: Native multimodal models are world learners (2025), <https://arxiv.org/abs/2510.26583>
16. Deng, C., Zhu, D., Li, K., Gou, C., Li, F., Wang, Z., Zhong, S., Yu, W., Nie, X., Song, Z., Shi, G., Fan, H.: Emerging properties in unified multimodal pretraining (2025), <https://arxiv.org/abs/2505.14683>
17. Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., Podell, D., Dockhorn, T., English, Z., Rombach, R.: Scaling Rectified Flow Transformers for High-Resolution Image Synthesis. In: Forty-First International Conference on Machine Learning (Jun 2024)
18. Ghosh, D., Hajishirzi, H., Schmidt, L.: Geneval: An object-focused framework for evaluating text-to-image alignment (2023), <https://arxiv.org/abs/2310.11513>
19. Google: nano-banana. <https://developers.googleblog.com/introducing-gemini-2-5-flash-image/> (2025)
20. Google: nano-banana-pro. <https://deepmind.google/models/gemini-image/pro/> (2025)
21. Google: nano-banana-2. <https://deepmind.google/models/gemini-image/flash/> (2026)
22. Han, J., Liu, J., Jiang, Y., Yan, B., Zhang, Y., Yuan, Z., Peng, B., Liu, X.: Infinity: Scaling bitwise autoregressive modeling for high-resolution image synthesis (2024)
23. Hessel, J., Holtzman, A., Forbes, M., Bras, R.L., Choi, Y.: CLIPScore: A Reference-free Evaluation Metric for Image Captioning (Mar 2022). <https://doi.org/10.48550/arXiv.2104.08718>
24. Hu, X., Wang, R., Fang, Y., Fu, B., Cheng, P., Yu, G.: Ella: Equip diffusion models with llm for enhanced semantic alignment (2024), <https://arxiv.org/abs/2403.05135>
25. Hu, Y., Liu, B., Kasai, J., Wang, Y., Ostendorf, M., Krishna, R., Smith, N.A.: Tifa: Accurate and interpretable text-to-image faithfulness evaluation with question answering (2023), <https://arxiv.org/abs/2303.11897>
26. Huang, K., Duan, C., Sun, K., Xie, E., Li, Z., Liu, X.: T2i-compbench++: An enhanced and comprehensive benchmark for compositional text-to-image generation (2025), <https://arxiv.org/abs/2307.06350>
27. Hurst, A., Lerer, A., Goucher, A.P., Perelman, A., Ramesh, A., Clark, A., Ostrow, A., Welihinda, A., Hayes, A., Radford, A., et al.: Gpt-4o system card. arXiv preprint arXiv:2410.21276 (2024)
28. Imagen-Team-Google, Baldridge, J., Bauer, J., Bhutani, M., Others: Imagen 3 (Dec 2024). <https://doi.org/10.48550/arXiv.2408.07009>
29. Kamath, A., Chang, K.W., Krishna, R., Zettlemoyer, L., Hu, Y., Ghazvininejad, M.: Geneval 2: Addressing benchmark drift in text-to-image evaluation (2025), <https://arxiv.org/abs/2512.16853>
30. Labs, B.F.: Flux. <https://github.com/black-forest-labs/flux> (2024)
31. Labs, B.F.: Flux.2. <https://bfl.ai/models/flux-2> (2026)
32. Li, B., Lin, Z., Pathak, D., Li, J., Fei, Y., Wu, K., Ling, T., Xia, X., Zhang, P., Neubig, G., Ramanan, D.: Genai-bench: Evaluating and improving compositional text-to-visual generation (2024), <https://arxiv.org/abs/2406.13743>
33. Li, C., Xu, H., Tian, J., Wang, W., Yan, M., Bi, B., Ye, J., Chen, H., Xu, G., Cao, Z., Zhang, J., Huang, S., Huang, F., Zhou, J., Si, L.: mplug: Effective and efficient vision-language learning by cross-modal skip-connections (2022), <https://arxiv.org/abs/2205.12005>

34. Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation (2022), <https://arxiv.org/abs/2201.12086>
35. Li, O., Wang, Y., Hu, X., Huang, H., Chen, R., Ou, J., Tao, X., Wan, P., Qi, X., Feng, F.: Easier painting than thinking: Can text-to-image models set the stage, but not direct the play? (2025), <https://arxiv.org/abs/2509.03516>
36. Li, O., Wang, Y., Hu, X., Huang, H., Chen, R., Ou, J., Tao, X., Wan, P., Qi, X., Feng, F.: Easier painting than thinking: Can text-to-image models set the stage, but not direct the play? (2026), <https://arxiv.org/abs/2509.03516>
37. Li, Z., Zhang, J., Lin, Q., Xiong, J., Long, Y., Deng, X., Zhang, Y., Liu, X., Huang, M., Xiao, Z., Chen, D., He, J., Li, J., Li, W., Zhang, C., Quan, R., Lu, J., Huang, J., Yuan, X., Zheng, X., Li, Y., Zhang, J., Zhang, C., Chen, M., Liu, J., Fang, Z., Wang, W., Xue, J., Tao, Y., Zhu, J., Liu, K., Lin, S., Sun, Y., Li, Y., Wang, D., Chen, M., Hu, Z., Xiao, X., Chen, Y., Liu, Y., Liu, W., Wang, D., Yang, Y., Jiang, J., Lu, Q.: Hunyuan-DiT: A Powerful Multi-Resolution Diffusion Transformer with Fine-Grained Chinese Understanding (May 2024). <https://doi.org/10.48550/arXiv.2405.08748>
38. Lin, W., Wei, X., An, R., Peng, G., Zou, B., Luo, Y., Huang, S., Zhang, S., Li, H.: Draw-and-understand: Leveraging visual prompts to enable mllms to comprehend what you want. In: International Conference on Learning Representations. vol. 2025, pp. 46374–46403 (2025)
39. Lin, W., Wei, X., An, R., Ren, T., Chen, T., Zhang, R., Guo, Z., Zhang, W., Zhang, L., Li, H.: Perceive anything: Recognize, explain, caption, and segment anything in images and videos. *Advances in Neural Information Processing Systems* **38**, 96231–96256 (2026)
40. Lin, W., Wei, X., Zhang, R., Zhuo, L., Zhao, S., Huang, S., Xie, J., Peng, G., Li, H.: Pixwizard: Versatile image-to-image visual assistant with open-language instructions. In: International Conference on Learning Representations. vol. 2025, pp. 38680–38711 (2025)
41. Lin, Z., Pathak, D., Li, B., Li, J., Xia, X., Neubig, G., Zhang, P., Ramanan, D.: Evaluating text-to-visual generation with image-to-text generation (2024), <https://arxiv.org/abs/2404.01291>
42. Liu, D., Zhao, S., Zhuo, L., Lin, W., Xin, Y., Li, X., Qin, Q., Qiao, Y., Li, H., Gao, P.: Lumina-mgpt: Illuminate flexible photorealistic text-to-image generation with multimodal generative pretraining (2025), <https://arxiv.org/abs/2408.02657>
43. Ma, Y., Shui, Y., Wu, X., Sun, K., Li, H.: Hpsv3: Towards wide-spectrum human preference score (2025), <https://arxiv.org/abs/2508.03789>
44. Niu, Y., Ning, M., Zheng, M., Jin, W., Lin, B., Jin, P., Liao, J., Feng, C., Ning, K., Zhu, B., Yuan, L.: Wise: A world knowledge-informed semantic evaluation for text-to-image generation (2025), <https://arxiv.org/abs/2503.07265>
45. Qin, Q., Zhuo, L., Xin, Y., Du, R., Li, Z., Fu, B., Lu, Y., Yuan, J., Li, X., Liu, D., Zhu, X., Zhang, M., Beddow, W., Millon, E., Perez, V., Wang, W., He, C., Zhang, B., Liu, X., Li, H., Qiao, Y., Xu, C., Gao, P.: Lumina-image 2.0: A unified and efficient image generative framework (2025), <https://arxiv.org/abs/2503.21758>
46. Qwen, :, Yang, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Li, C., Liu, D., Huang, F., Wei, H., Lin, H., Yang, J., Tu, J., Zhang, J., Yang, J., Yang, J., Zhou, J., Lin, J., Dang, K., Lu, K., Bao, K., Yang, K., Yu, L., Li, M., Xue, M., Zhang, P., Zhu, Q., Men, R., Lin, R., Li, T., Tang, T., Xia, T., Ren, X., Ren, X., Fan, Y., Su, Y., Zhang, Y., Wan, Y., Liu, Y., Cui, Z., Zhang, Z., Qiu, Z.: Qwen2.5 technical report (2025), <https://arxiv.org/abs/2412.15115>

47. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision (2021), <https://arxiv.org/abs/2103.00020>
48. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-Resolution Image Synthesis With Latent Diffusion Models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10684–10695 (2022)
49. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., Massa, F., Haziza, D., Wehrstedt, L., Wang, J., Darcet, T., Moutakanni, T., Sentana, L., Roberts, C., Vedaldi, A., Tolan, J., Brandt, J., Couprie, C., Mairal, J., Jégou, H., Labatut, P., Bojanowski, P.: Dinov3 (2025), <https://arxiv.org/abs/2508.10104>
50. Somepalli, G., Gupta, A., Gupta, K., Palta, S., Goldblum, M., Geiping, J., Shrivastava, A., Goldstein, T.: Measuring style similarity in diffusion models (2024), <https://arxiv.org/abs/2404.01292>
51. Sun, K., Fang, R., Duan, C., Liu, X., Liu, X.: T2i-reasonbench: Benchmarking reasoning-informed text-to-image generation (2025), <https://arxiv.org/abs/2508.17472>
52. Sun, P., Jiang, Y., Chen, S., Zhang, S., Peng, B., Luo, P., Yuan, Z.: Autoregressive Model Beats Diffusion: Llama for Scalable Image Generation (Jun 2024). <https://doi.org/10.48550/arXiv.2406.06525>
53. Team, M.: Midjourney. <https://www.midjourney.com/> (2025)
54. Tian, K., Jiang, Y., Yuan, Z., Peng, B., Wang, L.: Visual autoregressive modeling: Scalable image generation via next-scale prediction (2024)
55. Tong, C., Guo, Z., Zhang, R., Shan, W., Wei, X., Xing, Z., Li, H., Heng, P.A.: Delving into rl for image generation with cot: A study on dpo vs. grpo. *Advances in Neural Information Processing Systems* **38**, 115632–115655 (2026)
56. Tschannen, M., Gritsenko, A., Wang, X., Naeem, M.F., Alabdulmohsin, I., Parthasarathy, N., Evans, T., Beyer, L., Xia, Y., Mustafa, B., Hénaff, O., Harmen, J., Steiner, A., Zhai, X.: Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features (2025), <https://arxiv.org/abs/2502.14786>
57. Wang, G., Wei, X., Liu, J., Zhang, R., Zhang, Y., Zhang, K., Chong, M., Zhang, S.: Mr-mlm: Mutual reinforcement of multimodal comprehension and vision perception. *arXiv preprint arXiv:2406.15768* (2024)
58. Wang, S., Liu, S., Kuang, Y., Wei, X., Liu, Y., Li, Z., Man, Y., Chen, G., Tao, A., Liu, G., et al.: Locateanything: Fast and high-quality vision-language grounding with parallel box decoding. *arXiv preprint arXiv:2605.27365* (2026)
59. Wang, Y., Li, Z., Zang, Y., Wang, C., Lu, Q., Jin, C., Wang, J.: Unified multimodal chain-of-thought reward model through reinforcement fine-tuning (2025), <https://arxiv.org/abs/2505.03318>
60. Wang, Y., Zang, Y., Li, H., Jin, C., Wang, J.: Unified reward model for multimodal understanding and generation (2026), <https://arxiv.org/abs/2503.05236>
61. Wei, X., Cen, K., Wei, H., Guo, Z., Li, B., Wang, Z., Zhang, J., Zhang, L.: Mico-150k: A comprehensive dataset advancing multi-image composition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 29695–29706 (2026)
62. Wiles, O., Zhang, C., Albuquerque, I., Kajić, I., Wang, S., Bugliarello, E., Onoe, Y., Papalampidi, P., Ktena, I., Knutsen, C., Rashtchian, C., Nawalgaria, A., Pont-Tuset, J., Nematzadeh, A.: Revisiting text-to-image evaluation with gecko: On metrics, prompts, and human reinforcements (2025), <https://arxiv.org/abs/2404.16820>

63. Wu, C., Li, J., Zhou, J., Lin, J., Gao, K., Yan, K., ming Yin, S., Bai, S., Xu, X., Chen, Y., Chen, Y., Tang, Z., Zhang, Z., Wang, Z., Yang, A., Yu, B., Cheng, C., Liu, D., Li, D., Zhang, H., Meng, H., Wei, H., Ni, J., Chen, K., Cao, K., Peng, L., Qu, L., Wu, M., Wang, P., Yu, S., Wen, T., Feng, W., Xu, X., Wang, Y., Zhang, Y., Zhu, Y., Wu, Y., Cai, Y., Liu, Z.: Qwen-image technical report (2025), <https://arxiv.org/abs/2508.02324>
64. Wu, C., Chen, X., Wu, Z., Ma, Y., Liu, X., Pan, Z., Liu, W., Xie, Z., Yu, X., Ruan, C., Luo, P.: Janus: Decoupling Visual Encoding for Unified Multimodal Understanding and Generation (Oct 2024). <https://doi.org/10.48550/arXiv.2410.13848>
65. Wu, X., Hao, Y., Sun, K., Chen, Y., Zhu, F., Zhao, R., Li, H.: Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis (2023), <https://arxiv.org/abs/2306.09341>
66. Wu, X., Yu, D., Huang, Y., Russakovsky, O., Arora, S.: Conceptmix: A compositional image generation benchmark with controllable difficulty (2024), <https://arxiv.org/abs/2408.14339>
67. Xie, E., Chen, J., Zhao, Y., Yu, J., Zhu, L., Wu, C., Lin, Y., Zhang, Z., Li, M., Chen, J., Cai, H., Liu, B., Zhou, D., Han, S.: SANA 1.5: Efficient Scaling of Training-Time and Inference-Time Compute in Linear Diffusion Transformer (Mar 2025). <https://doi.org/10.48550/arXiv.2501.18427>
68. Xie, J., Mao, W., Bai, Z., Zhang, D.J., Wang, W., Lin, K.Q., Gu, Y., Chen, Z., Yang, Z., Shou, M.Z.: Show-o: One Single Transformer to Unify Multimodal Understanding and Generation (Oct 2024). <https://doi.org/10.48550/arXiv.2408.12528>
69. Xie, J., Yang, Z., Shou, M.Z.: Show-o2: Improved native unified multimodal models (2025), <https://arxiv.org/abs/2506.15564>
70. Xin, Y., Yan, J., Qin, Q., Li, Z., Liu, D., Li, S., Huang, V.S.J., Zhou, Y., Zhang, R., Zhuo, L., Han, T., Sun, X., Luo, S., Wang, M., Fu, B., Cao, Y., Li, H., Zhai, G., Liu, X., Qiao, Y., Gao, P.: Lumina-mgpt 2.0: Stand-alone autoregressive image modeling (2025), <https://arxiv.org/abs/2507.17801>
71. Xu, J., Huang, Y., Cheng, J., Yang, Y., Xu, J., Wang, Y., Duan, W., Yang, S., Jin, Q., Li, S., Teng, J., Yang, Z., Zheng, W., Liu, X., Ding, M., Zhang, X., Gu, X., Huang, S., Huang, M., Tang, J., Dong, Y.: Visionreward: Fine-grained multi-dimensional human preference learning for image and video generation (2025), <https://arxiv.org/abs/2412.21059>
72. Yang, Z., Gu, T., Wang, J., Lin, F., Sheng, X., Chen, P., Li, L.: Longt2ibench: A benchmark for evaluating long text-to-image generation with graph-structured annotations (2025), <https://arxiv.org/abs/2512.09271>
73. Ye, J., Jiang, D., Wang, Z., Zhu, L., Hu, Z., Huang, Z., He, J., Yan, Z., Yu, J., Li, H., He, C., Li, W.: Echo-4o: Harnessing the power of gpt-4o synthetic images for improved image generation (2025), <https://arxiv.org/abs/2508.09987>
74. Zhang, R., Wei, X., Jiang, D., Guo, Z., Zhang, Y., Tong, C., Liu, J., Zhou, A., Zhang, S., Peng, G., et al.: Mavis: Mathematical visual instruction tuning with an automatic data engine. In: International Conference on Learning Representations. vol. 2025, pp. 87955–87989 (2025)
75. Zhang, Y.F., Yang, H., Zhang, H., Shi, Y., Chen, Z., Tian, H., Fu, C., Wang, H., Wu, K., Cui, B., Wang, X., Pan, J., Wang, H., Zhang, Z., Wang, L.: Basereward: A strong baseline for multimodal reward model (2025), <https://arxiv.org/abs/2509.16127>
76. Zhao, S., Wu, Q., Li, X., Zhang, B., Li, M., Qin, Q., Liu, D., Zhang, K., Li, H., Qiao, Y., Gao, P., Fu, B., Li, Z.: Lex-art: Rethinking text generation via scalable high-quality data synthesis (2025), <https://arxiv.org/abs/2503.21749>
77. Zhou, X., Koltun, V., Krähenbühl, P.: Simple multi-dataset detection (2022), <https://arxiv.org/abs/2102.13086>

78. Zhuo, L., Du, R., Xiao, H., Li, Y., Liu, D., Huang, R., Liu, W., Zhao, L., Wang, F.Y., Ma, Z., Luo, X., Wang, Z., Zhang, K., Zhu, X., Liu, S., Yue, X., Liu, D., Ouyang, W., Liu, Z., Qiao, Y., Li, H., Gao, P.: Lumina-Next: Making Lumina-T2X Stronger and Faster with Next-DiT (Jun 2024). <https://doi.org/10.48550/arXiv.2406.18583>