

LongT2IBench: A Benchmark for Evaluating Long Text-to-Image Generation with Graph-structured Annotations

Zhichao Yang¹, Tianjiao Gu¹, Jianjie Wang¹, Feiyu Lin¹, Xiangfei Sheng¹, Pengfei Chen^{1*}, Leida Li^{1, 2*}

¹School of Artificial Intelligence, Xidian University

²State Key Laboratory of Electromechanical Integrated Manufacturing of High-Performance Electronic Equipments, Xidian (yangzhichao, gutianjiao, wangjianjie, linfeiyu)@stu.xidian.edu.cn, xiangfeisheng@gmail.com, (chenpengfei, ldli)@xidian.edu.cn

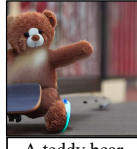
Abstract

The increasing popularity of long Text-to-Image (T2I) generation has created an urgent need for automatic and interpretable models that can evaluate the image-text alignment in long prompt scenarios. However, the existing T2I alignment benchmarks predominantly focus on short prompt scenarios and only provide MOS or Likert scale annotations. This inherent limitation hinders the development of long T2I evaluators, particularly in terms of the interpretability of alignment. In this study, we contribute **LongT2IBench**, which comprises 14K long text-image pairs accompanied by graph-structured human annotations. Given the detail-intensive nature of long prompts, we first design a Generate-Refine-Qualify annotation protocol to convert them into textual graph structures that encompass entities, attributes, and relations. Through this transformation, fine-grained alignment annotations are achieved based on these granular elements. Finally, the graph-structured annotations are converted into alignment scores and interpretations to facilitate the design of T2I evaluation models. Based on LongT2IBench, we further propose **LongT2IExpert**, a LongT2I evaluator that enables multi-modal large language models (LLMs) to provide both quantitative scores and structured interpretations through an instruction-tuning process with Hierarchical Alignment Chain-of-Thought (CoT). Extensive experiments and comparisons demonstrate the superiority of the proposed LongT2IExpert in alignment evaluation and interpretation.

Data, code — <https://github.com/yzc-ipp/LongT2IBench>

Introduction

With the extensive application of Text-to-Image (T2I) generation models in fields such as artistic creation and advertising design, user demand for long T2I generation capabilities has significantly increased (Ko et al. 2023; Zhao et al. 2024; Yang et al. 2024b; Sheng et al. 2025b; Yang et al. 2025). This broadened research efforts from short prompts to increasingly complex and lengthy inputs (Hu et al. 2024; Liu et al. 2024). Long prompts inherently contain more elements and detailed descriptions, presenting significant challenges for alignment evaluation. Given the high costs and

		Benchmarks	Avg. Length
 <p>A teddy bear on a skateboard.</p>	Type of Human annotations:	ImageReward	21.83
	Alignment score	TIFAv1	10.48
	Human Annotation	EvalMuse-40K	10.66
	MOS: 1,2,3,4,5	Q-Eval-100K	35.26
		GenAI-Bench	12.61

Existing T2I Align Benchmarks
 LongT2IBench-(Avg. Length=82.09)

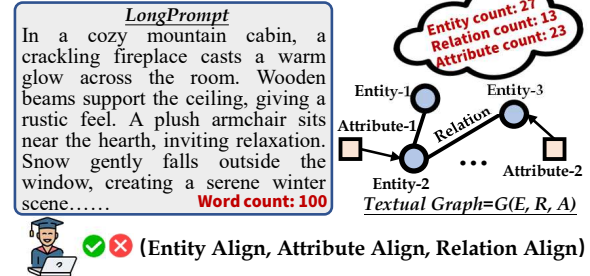


Image	Align Score: 0.5156	Align Interpretation:
	Align: Misalign:	Align: Misalign:
	cabin; blanket;	blanket; steaming
	fireplace; woolen;	woolen; (mug);
	window; mug;	mug; falls
	books; ...	curtain; (snow);
		...

Figure 1: Illustration of the graph-structured annotations of LongT2IBench. Compared to existing benchmarks, we focus on T2I alignment in long prompt scenarios, offering both quantitative scores and fine-grained interpretations.

inefficiencies associated with manual evaluation, developing automatic and interpretable methods for evaluating long T2I alignment has become in urgent need. They can effectively evaluate the performance of existing T2I models and guide the optimization of future models. High-quality, fine-grained annotations of long T2I alignment serve as the essential foundation for developing such evaluation methods.

In the past few years, a proliferation of T2I alignment benchmarks have been proposed, along with corresponding evaluation algorithms. Early benchmarks, such as Pick-a-pic (Kirstain et al. 2023) and HPDv2 (Wu et al. 2023) employed side-by-side comparisons to evaluate the alignment between generated image and text. Other benchmarks like ImageReward (Xu et al. 2023) and EvalAlign (Tan et al. 2024) uti-

*Corresponding authors.

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

lized Likert scale annotation, with human annotators providing holistic scores based on general text-image correspondence. Recently, Benchmarks like T2I-CompBench (Huang et al. 2023) and GenAI-Bench (Li et al. 2024a) introduced curated prompts designed to evaluate specific capabilities of generative models, including counting, spatial relationships, etc. For fine-grained evaluation, benchmarks like TIFA (Hu et al. 2023) and EvalMuse-40K (Han et al. 2024) decomposed prompts into multiple concise phrases, formulating separate assessments that require human annotators to provide element-wise alignment scores.

While existing benchmarks offer valuable insights, they primarily focus on short prompt and only provide Mean Opinion Score (MOS) or Likert-scale human annotations (Figure 1-top). This limitation hinders the development of long T2I evaluators, particularly for alignment interpretability. The detail-intensive nature of long prompts presents two key challenges for obtaining high-quality, fine-grained alignment annotations: (1) **Detail Overload**: it is both challenging and unreliable for annotators to directly provide an overall alignment score. (2) **Alignment Complexity**: element-wise annotations are inadequate in capturing the intricate alignment relationships present in long prompts, such as connections between distant elements. To address these challenges, we contribute **LongT2IBench**, a benchmark specifically designed for long T2I alignment that provides both quantitative alignment scores and fine-grained interpretations by introducing graph-structured annotations (Figure 1-bottom). LongT2IBench consists of 3K long prompts and 14K image-text pairs. To ensure the diversity, long prompts are obtained from three sources: human-generated content, AI-generated content, and long image captions. These prompts are first converted into textual graph structures that encompass entities, attributes, and relations. This conversion follows a Generate-Refine-Qualify protocol to ensure exceptional structural fidelity. Through this transformation, fine-grained alignment annotations are achieved based on these granular elements of textual graphs. We also implement a re-verification process to ensure the reliability of the alignment annotations. Finally, the above annotations are converted into alignment scores and interpretations.

Building upon LongT2IBench, we further propose **LongT2IExpert**, a long T2I evaluator which empowers multi-modal large language models (MLLMs) to provide both quantitative scores and structured interpretations. Inspired by the construction of LongT2IBench, we first introduce a Hierarchical Alignment Chain-of-Thought (HA-CoT) to prompt MLLMs to engage in structured alignment reasoning. HA-CoT features a three-hop reasoning process, progressively guiding models from fine-grained analysis to comprehensive evaluation. Then, human-annotated alignment scores and interpretations are utilized as supervision signals to fine-tune MLLMs in a multi-task manner. With these designs, we achieve a unified scoring and interpreting model that aligns closely with human evaluators.

The contributions of this work are three-fold:

- We contribute LongT2IBench, a benchmark of Long Text-to-Image generation that contains 14K long text-

image pairs with graph-structured annotations. It provides both quantitative scores and fine-grained interpretations, elucidating the intricate alignment relationships between textual elements and visual components.

- We propose LongT2IExpert, a novel long T2I evaluator that equips multimodal large language models (MLLMs) with structured alignment reasoning capabilities, simultaneously performing alignment scoring and interpreting.
- We establish a systematic framework for examining long T2I alignment abilities. It comprehensively assesses both scoring accuracy and interpretation quality, while thoroughly testing performance across varying text lengths and alignment dimensions: entity, attribute, and relation.

Related Works

Text-to-Image Alignment Benchmarks. Many benchmarks have been developed to evaluate the text-to-image alignment of generation models. Early benchmarks such as Pick-a-pic (Kirstain et al. 2023) and HPDv2 (Wu et al. 2023), compare models side-by-side to evaluate generated image quality. Other benchmarks like ImageReward (Xu et al. 2023) and EvalAlign (Tan et al. 2024) utilized Likert scale annotation to collect overall alignment scores. For fine-grained evaluation, benchmarks like TIFA (Hu et al. 2023) and EvalMuse-40K (Han et al. 2024) extract concise phrases from prompts and collect element-wise annotations. With increasing prompt length, benchmarks specifically designed for long T2I alignment remain scarce. In this paper, we contribute LongT2IBench, providing a foundation for examining and developing long T2I alignment evaluation models.

Text-to-Image Alignment Evaluators. Early T2I evaluators directly utilized vision-language models (VLMs) for alignment evaluation, such as CLIPScore (Hessel et al. 2021) and BLIP2Score (Li et al. 2023). With the extensive collection of alignment annotations, PickScore (Kirstain et al. 2023) and ImageReward (Xu et al. 2023) fine-tuned VLMs to align with human ratings. With MLLMs rapidly advanced, VQAScore (Li et al. 2024b) and Q-Eval-Score (Zhang et al. 2025) leverage their powerful vision-language capabilities for T2I alignment. While existing models have achieved remarkable progress, their evaluation performance in long-prompt contexts remains largely unexplored. In this paper, we address this gap and introduce LongT2IExpert, a model specifically designed for long T2I alignment.

Text-to-Image Models for Long Prompt. Traditional CLIP-based T2I models face token length constraints, limiting long prompt processing capabilities. Recent work addresses this by employing large language models (LLMs) as text encoders for enhanced semantic extraction, exemplified by Stable Diffusion 3.5 (Esser et al. 2024). Alternatively, other methods introduced architectural improvements to the CLIP to enable longer inputs. LongCLIP (Zhang et al. 2024) offers a plug-and-play alternative to standard CLIP that supports long-text processing. LongAlign (Liu et al. 2024) decomposes long prompts into individual sentences for separate encoding, enabling long T2I generation. These advances in long-prompt generation has created an urgent need for automatic, interpretable long T2I alignment evaluation models.

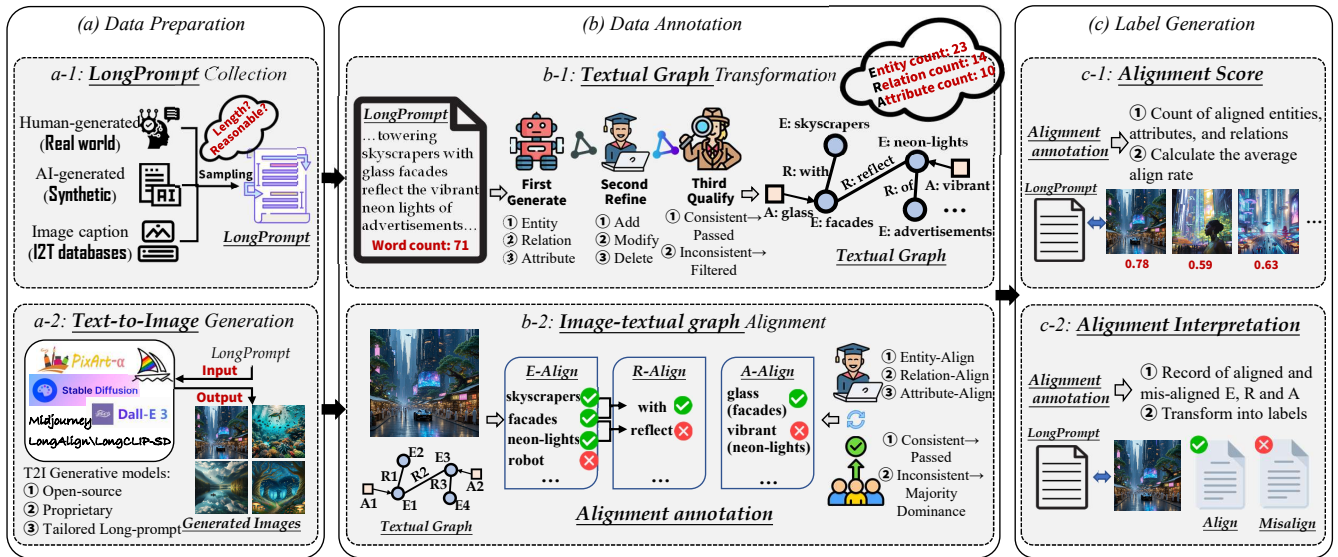


Figure 2: Overview of the construction pipeline for **LongT2IBench**. The pipeline consists of three stages: (a) **Data Preparation**. Long prompts are collected from three sources and input into various T2I models to generate images. (b) **Data Annotation**. Long prompts are converted into textual graph structures, and fine-grained image-textual graph alignment annotations are achieved. (c) **Label Generation**. Two categories of labels: quantitative alignment scores and alignment interpretations are produced based on graph-structured human annotations.

LongT2IBench

In this section, we detail LongT2IBench, a benchmark for evaluating long text-to-image alignment by introducing graph-structured annotations. The overview of construction pipeline for LongT2IBench is illustrated in Figure 2. Throughout the benchmark, we employ multiple strategies to ensure both diversity and reliability. Next, we will elaborate on the construction process from three aspects: data preparation, data annotation, and label generation, followed by a statistical analysis.

Data Preparation

Long Prompt Collection. To ensure prompt diversity and comprehensive evaluation of long T2I alignment capabilities, we collect long prompts from three sources and controlled length distribution across different word-count intervals. Specifically, the long prompts are sampled from human-generated content (Human-Gen), AI-generated content (AI-Gen), and long image captions (Img-Cap). Human-Gen prompts are sourced from DiffusionDB (Wang et al. 2023), which contains numerous prompts created by real users. AI-Gen prompts are synthesized using GPT-4 (Achiam et al. 2023) following specific templates. Img-Cap prompts are derived from DOCCI (Onoe et al. 2024), where each image is accompanied by a long description with fine-grained visual details. Meanwhile, we ensured balanced sampling across five word-count intervals (30-50, 50-70, 70-90, 90-110, 110+) to systematically assess T2I evaluators’ alignment capabilities at varying text lengths.

Text-to-Image Generation. To ensure the diversity of the generated images, we select six distinct types of text-to-image generative models. These models are categorized

into three groups: (1) basic open-source popular T2I models, namely Stable Diffusion v3.5 (Esser et al. 2024) and PixArt- α (Chen et al. 2024); (2) powerful proprietary T2I models, including DALL-E 3 (OpenAI 2023) and Midjourney v6 (Holz 2023); (3) tailored long T2I models, including LongCLIP-SD (Zhang et al. 2024) and LongSD (Liu et al. 2024). For each long prompt, we utilize the default settings of these models to generate images, ensuring a comprehensive dataset for annotation.

Data Annotation

Textual Graph Transformation. Graph representation has proven a powerful paradigm for modeling complex relationships and structural information in various domains (Hamilton 2020; Yang et al. 2024a). Given the detail-intensive nature of long prompts, we perform textual graph transformation on the collected prompts. This establishes the foundation for more objective and accurate text-image alignment annotations. Additionally, the graph structure enables precise localization of aligned and misaligned elements, allowing us to better evaluate long T2I alignment capabilities.

To ensure exceptional graph structural fidelity, we adopt a systematic three-phase transformation protocol. First, we design specific input prompts and output formats to guide GPT-4 to **generate** textual graphs, converting long prompts into graph-structures comprising entities, attributes, and relations. Second, trained human annotators **refine** these generated graph structures to ensure accuracy. This refinement process involves removing inappropriate conversions and making detailed adjustments through additions, deletions, and modifications to coarse transformations. Third, we perform a double-check to **qualify** the refined transformations

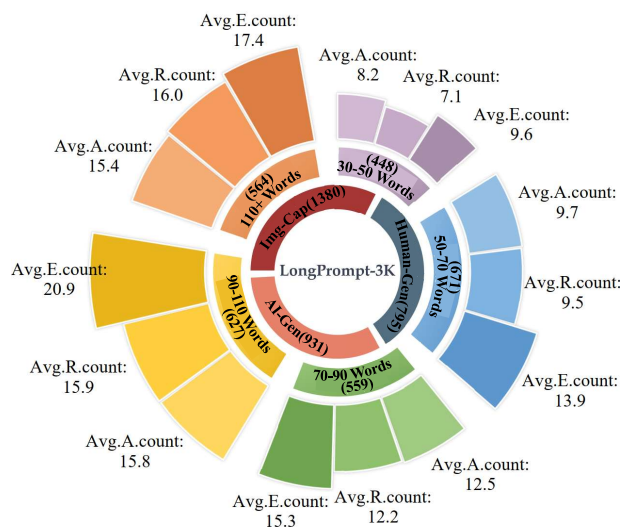


Figure 3: The Source Distribution Map of **LongPrompt-3K**. Long prompts are sampled from three sources: Human-Gen, AI-Gen and Img-Cap. They are evenly distributed across different word count ranges, with entity, attribute, and relation counts in each range also being statistically presented.

from the previous stage, ensuring reliability. Only conversion graphs that receive consistent agreement are retained in the final dataset. This rigorous multi-stage process ensures high-quality structural graph transformation that accurately capture the semantic relationships of long prompts. Throughout this process, 3K precise conversions were retained from an initial collection of 4.5K long prompts. A statistical analysis on the number of entities, attributes, and relations within the transformed graph structures in **LongPrompt-3K** are illustrated in Figure 3. The instructions for graph generation using GPT-4, the operational interface for Refine and Qualify stages, along with annotations training guidelines are detailed in the Appendix.

Image-textual graph Alignment. Based on transformed textual graphs, we implement an initial annotation followed by a re-verification process to obtain reliable image-textual graph alignment annotations. Trained annotators perform initial annotations, making binary decisions on the alignment between transformed entities, attributes, and relations with the corresponding images, namely E-Align, A-Align, and R-Align respectively. To improve efficiency, we incorporate hierarchical annotation logic into our platform design, where annotators first evaluate entities, and the platform automatically filters out attributes and relationships associated with misaligned entities. Furthermore, we engage three independent annotators to review the initial annotation results, removing instances with significant disagreement while retaining alignment annotations that achieve majority consensus. After annotation, 14K image-text pairs are retained from the original 18K samples for the final dataset. Discarded data includes images containing NSFW content, as well as instances with severe generation distortions.

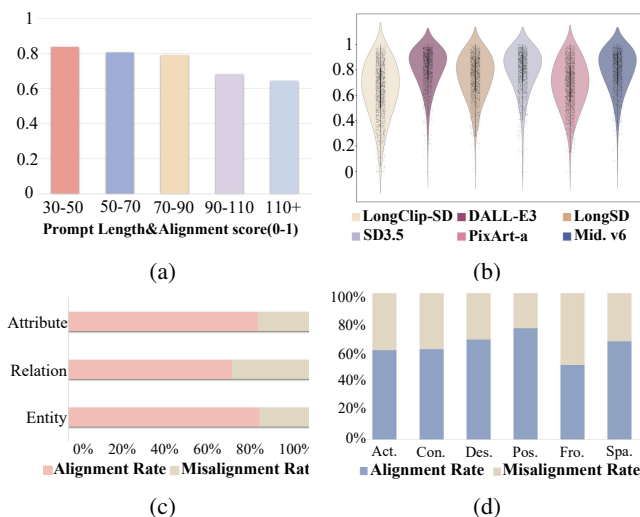


Figure 4: Statistical Analysis. (a) Average alignment scores across five word-count intervals. (b) Distribution of annotated alignment scores for six T2I generative models. (c) Alignment and misalignment rates across entities, attributes, and relations. (d) Alignment percentages among six relation categories (Action, Connection, Description, Possession, From/to and Spatial relation).

Label Generation

Alignment Score. We envision future text-to-image generative models capable of aligning with all textual details. Therefore, we treat all entities, attributes, and relations in the textual graph with equal importance. The alignment score is defined as the ratio of aligned elements to the total number. We analyzed the average alignment scores across five word-count intervals, as illustrated in Figure 4(a). The statistical results demonstrate that alignment scores gradually decrease as prompt length increases, underscoring the importance of evaluating text-image alignment in long-prompt scenarios. Additionally, Figure 4(b) shows the distribution of alignment annotations across six generative models, revealing performance disparities in long-text alignment capabilities, with proprietary models (DALL-E 3 and Midjourney v6) demonstrating relative advantages.

Alignment Interpretation. In long-prompt scenarios, the need for alignment interpretation is more critical, particularly in localizing misaligned elements that can guide optimization for generation models. Leveraging textual graph transformation and image-textual graph alignment annotations, we can derive structured alignment interpretations for each image-text pair, represented as aligned and misaligned entities, attributes, and relations. We analyzed the proportions of alignment and misalignment across these three categories, as illustrated in Figure 4(c). The results reveal that relation misalignment constitutes a particularly significant proportion. During the annotation process, we instructed annotators to classify relations into distinct categories, with Figure 4(d) presenting the alignment and misalignment proportions across these relation types.

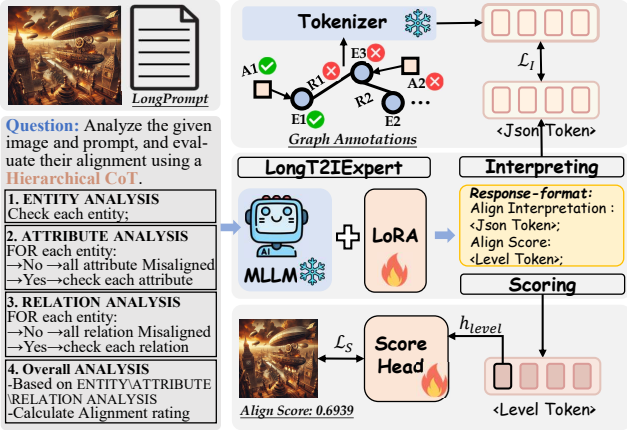


Figure 5: Overall pipeline of the proposed **LongT2IExpert**. A Hierarchical Alignment Chain-of-Thought (CoT) is designed to instruct MLLMs for structured alignment reasoning. Numerical alignment scores and graph-structured interpretations are utilized to train MLLMs for alignment scoring and interpreting in a multi-task manner.

LongT2IExpert

Building on the proposed LongT2IBench, we propose **LongT2IExpert**, a long T2I evaluator that enables multimodal large language models (MLLMs) to provide both quantitative scores and structured interpretations through an instruction-tuning process with Hierarchical Alignment Chain-of-Thought (HA-CoT). The overall pipeline is illustrated in Figure 5.

Instructing MLLMs for Alignment Reasoning

The evaluation of Long T2I alignment requires that models possess both the capability to analyze long texts and the ability to achieve multimodal alignment. This imposes high demands on the foundational capabilities of the model, leading us to select a multimodal large language model as our backbone, i.e. Qwen2.5-VL-7B-Instruct (Bai et al. 2023). However, even with this advanced architecture, the characteristics of Detail Overload and Alignment Complexity inherent in this task render alignment evaluation a challenging endeavor. Inspired by the construction process of LongT2IBench, we introduce a Hierarchical Alignment Chain-of-Thought (HA-CoT) to guide MLLMs in mimicking human-like structured alignment reasoning.

HA-CoT features a three-hop reasoning process. In the initial stage, the MLLM acts as an Entity Aligner, analyzing all entities mentioned in the long text and determining their alignment with the image. In the second phase, acting as an Attribute and Relation Aligner, it examines the alignment of attributes and relations associated with each aligned entity. Finally, based on the previous analysis, it synthesizes the overall alignment situation and provides a comprehensive alignment score. Throughout the structured CoT reasoning process, MLLMs perform a multi-faceted alignment evaluation that mimics the deliberation of a human evaluator.

Training MLLMs for Alignment Evaluation

Building on the structured alignment reasoning process, we utilize high-quality alignment scores and interpretations as supervision signals for training MLLMs in alignment evaluation. Inspired by LISA (Lai et al. 2024; Sheng et al. 2025a), which uses special tokens to enable reasoning segmentation capabilities in MLLMs, we expand the original LLM vocabulary with a new token $\langle \text{Level} \rangle$ to signify the request for a numerical score output. Simultaneously, we convert the graph-based human annotations into JSON formats, preserving the original structured information while facilitating easier parsing by MLLMs. These approaches allow us to leverage the powerful understanding and reasoning capabilities of MLLMs for simultaneous alignment scoring and interpreting. Give a long prompt x_{txt} along with the generated image x_{img} and pre-defined CoT, we feed them into the MLLM \mathcal{F} , which in turn outputs a text response \hat{y}_{txt} :

$$\hat{y}_{txt} = \mathcal{F}(x_{txt}, x_{img}, \text{CoT}), \quad (1)$$

where the output \hat{y}_{txt} would include a $\langle \text{Json} \rangle$ token and $\langle \text{Level} \rangle$ token corresponding to the predicted alignment interpretation and alignment score, respectively.

We then extract the MLLM last-layer embedding $\hat{\mathbf{h}}_{Level}$ corresponding to the $\langle \text{Level} \rangle$ token and apply a score head \mathcal{R} to obtain the predicted alignment score \hat{y}_s . The score loss \mathcal{L}_S is calculated using Mean Squared Error (MSE) between the predicted score and the human-annotated score y_s :

$$\hat{y}_s = \mathcal{R}(\hat{\mathbf{h}}_{Level}), \mathcal{L}_S = (\hat{y}_s - y_s)^2. \quad (2)$$

For alignment interpretation, we decode the $\langle \text{Json} \rangle$ token into the graph-structured interpretation $\hat{\mathbf{G}}(E, R, A)$, where E, R, A represent entities, relations, and attributes, respectively. The interpretation loss \mathcal{L}_I is calculated using cross-entropy between the predicted interpretation and the ground-truth interpretation \mathbf{G} :

$$\mathcal{L}_I = \text{CE}(\hat{\mathbf{G}}, \mathbf{G}). \quad (3)$$

We incorporate a parameter-efficient LoRA module (Hu et al. 2022) for fine-tuning the MLLM. The model is trained end-to-end using the score prediction loss \mathcal{L}_S and the interpretations generation loss \mathcal{L}_I . The overall training objective \mathcal{L} is the weighted sum of these losses:

$$\mathcal{L} = \mathcal{L}_I + \lambda \cdot \mathcal{L}_S, \quad (4)$$

where the hyperparameter λ balances the two losses.

Experiments

Experimental Settings

Implementation Details. For LongT2IExpert training, we employ an open-source MLLM, Qwen2.5-VL-7B-Instruct (Bai et al. 2023), and attach LoRA adapters ($r=32$, $\alpha=64$, dropout=0.05) for fine-tuning. The score head consists of three Linear layers. The learning rate is set to $5e^{-5}$ for LoRA parameters and $2e^{-4}$ for the score head. The hyperparameter λ is set to 10, and the model is trained for 3 epochs on an A800 GPU. Details regarding LongT2IBench, including

Model	30-50 words		50-70 words		70-90 words		90-110 words		110+ words		Overall		
	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC	PLCC	Avg.
CLIPScore	0.224	0.235	0.349	0.370	0.271	0.240	0.208	0.202	0.209	0.184	0.269	0.264	0.267
BLIP2Score	0.384	0.231	0.342	0.358	0.190	0.155	0.215	0.191	0.149	0.061	0.266	0.223	0.245
PickScore	0.437	0.550	0.368	0.359	0.313	0.328	0.352	0.373	0.153	0.165	0.327	0.344	0.336
ImageReward	0.283	0.339	0.270	0.319	0.181	0.226	0.187	0.188	-0.023	0.072	0.210	0.244	0.227
HPSv2	0.540	0.534	0.479	0.473	<u>0.394</u>	0.417	0.395	<u>0.408</u>	0.148	0.151	0.381	0.392	0.387
VQAScore	0.504	0.385	0.333	0.360	0.270	0.389	0.195	0.200	0.242	0.225	0.298	0.291	0.295
FGA-BLIP2	0.414	0.399	0.308	0.330	0.193	0.299	0.160	0.171	0.130	0.155	0.230	0.262	0.246
Q-Eval-Score	0.470	0.416	0.460	0.459	0.339	0.382	0.320	0.301	<u>0.422</u>	<u>0.365</u>	0.361	0.355	0.358
PickScore*	0.387	0.510	0.530	0.521	0.393	0.432	0.419	0.372	0.250	0.264	<u>0.438</u>	0.439	0.438
ImageReward*	0.538	<u>0.584</u>	<u>0.546</u>	<u>0.574</u>	0.384	0.398	<u>0.399</u>	0.396	0.167	0.195	<u>0.438</u>	0.439	<u>0.439</u>
FGA-BLIP2*	<u>0.570</u>	0.568	0.514	0.527	0.371	<u>0.475</u>	0.327	0.355	0.126	0.144	0.406	<u>0.444</u>	0.425
LongT2IExpert	0.781	0.789	0.605	0.626	0.548	0.558	0.435	0.443	0.431	0.375	0.558	0.557	0.557

Table 1: Performance comparison between the proposed LongT2IExpert and the state-of-art methods for **Long T2I Alignment Scoring**. Here, ‘*’ denotes the fine-tuned versions of corresponding methods using the training data from LongT2IBench. The SRCC and PLCC across five word-count intervals are reported. Best in **bold**, second underlined.

Model	Entity			Attribute			Relation			All		
	Align	Mis.	Overall	Align	Mis.	Overall	Align	Mis.	Overall	Align	Mis.	Overall
Gemini-1.5-flash	48.4%	23.2%	43.6%	26.4%	9.2%	23.2%	6.4%	3.2%	5.6%	29.7%	11.8%	25.9%
Gemini-1.5-pro	49.2%	36.1%	46.7%	31.6%	16.6%	28.8%	13.6%	12.8%	13.4%	33.6%	21.9%	31.1%
GPT-4o	44.2%	30.7%	41.6%	27.6%	13.9%	25.0%	10.0%	11.0%	10.2%	29.3%	18.6%	27.0%
GPT-4-Turbo	43.1%	27.1%	40.1%	24.9%	11.9%	22.5%	5.6%	5.0%	5.4%	26.8%	14.6%	24.2%
Grok-3	42.9%	44.0%	43.1%	23.0%	16.4%	21.7%	6.2%	10.7%	7.4%	26.3%	<u>23.8%</u>	25.7%
Qwen2.5-vl-max	<u>49.9%</u>	22.8%	44.8%	28.2%	10.8%	24.9%	9.8%	4.7%	8.4%	31.7%	12.6%	27.7%
LLava-v1.6	33.6%	6.5%	28.4%	11.4%	1.9%	9.6%	1.9%	0.3%	1.5%	17.6%	2.9%	14.5%
LongT2IExpert	80.2%	<u>36.5%</u>	71.9%	53.8%	18.9%	47.3%	40.6%	20.8%	35.2%	60.7%	25.8%	53.2%

Table 2: Comparison between the proposed LongT2IExpert and the advanced multimodal large language models (MLLMs) for **Long T2I Alignment Interpreting**. Aligned and Misaligned accuracy across entity, attribute and relation are reported.

the annotation platforms and MLLM instructions for long prompt synthesis and textual graph transformations, as well as annotator information, are provided in the Appendix.

Evaluation Setting. We employ a standard train-validation-test split on LongT2IBench, ensuring long prompts are evenly distributed across five word-count intervals. For alignment score, we compare LongT2IExpert against two categories of existing state-of-the-art T2I evaluators. The first category comprises VLM-based models: CLIPScore (Hessel et al. 2021), BLIP2Score (Li et al. 2023), PickScore (Li et al. 2023), ImageReward (Xu et al. 2023), and FGA-BLIP2 (Han et al. 2024). The second category encompasses MLLM-based models: HPSv2 (Wu et al. 2023), VQAScore (Li et al. 2024b), and Q-Eval-Score (Zhang et al. 2025). To measure correlation with human annotations, we report both SRCC and PLCC metrics. Few models can simultaneously predict both alignment score and interpretations. To evaluate interpretation quality, we benchmark LongT2IExpert against two categories of MLLMs. Proprietary models include Gemini-1.5-flash/pro (Team et al. 2023), GPT-4o, GPT-4-Turbo (Achiam et al. 2023), and Grok-3 (xAI 2025), while open-source models

include Qwen2.5-vl-max (Bai et al. 2023), and LLava-v1.6 (Liu et al. 2023). We use Accuracy to evaluate interpretation quality in identifying aligned/misaligned elements.

Performance Evaluation

Results on Long T2I Alignment Scoring. Table 1 summarizes the performance comparison between the proposed LongT2IExpert and existing T2I evaluators. To thoroughly examine these models’ scoring capabilities in long-prompt scenarios, we evaluated their performance across different word-count ranges. The results demonstrate that most models show significant performance degradation as prompt length increases. Particularly in the 110+ words evaluation scenario, VLM-based models perform poorly due to text encoder’s token length restrictions, while MLLM-based evaluators such as Q-Eval-Score (Zhang et al. 2025) show notable performance advantages. Additionally, we fine-tuned open-source models using LongT2IBench training data for fair comparison. Our proposed LongT2IExpert demonstrates a clear advantage across all word-count intervals, suggesting its effectiveness in understanding long prompts and evaluating text-image alignment.

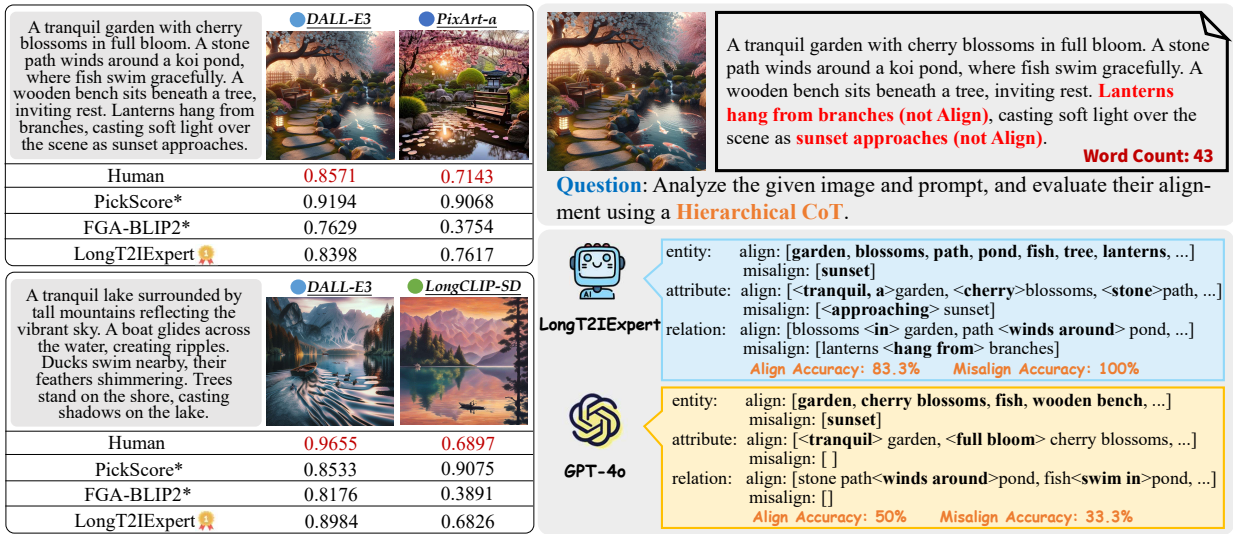


Figure 6: Visualization of long T2I alignment scoring and interpreting. Left: The predicted scores from LongT2IExpert demonstrate stronger correlation with human annotations compared to PickScore and FGA-BLIP2 (fine-tuned versions). Right: LongT2IExpert produces more accurate alignment interpretations than GPT-4o.

Results on Long T2I Alignment Interpreting. Due to the lack of objective alignment interpretation benchmarks and the scarcity of T2I evaluators with both scoring and interpreting capabilities, we compared LongT2IExpert with advanced MLLMs on long T2I alignment interpretation. For fair comparison, we used identical prompting patterns and output formats, including our designed HA-CoT. Table 2 lists the comparative results, where we report the accuracy across entity, attribute and relation. The proposed LongT2IExpert achieves the best performance, with Gemini-1.5-pro (Team et al. 2023) ranking second. The results indicate that models generally find relation alignment more challenging than entity and attribute. In long-prompt scenarios, accurate relation judgment places higher demands on long-context understanding capabilities. We also observe that accuracy for misaligned elements is consistently lower than for aligned elements. Identifying misalignments among numerous aligned elements requires more precise localization and discrimination abilities.

Qualitative Analysis. In Figure 6-left, we present the prediction scores of LongT2IExpert, PickScore and FGA-BLIP2 (fine-tuned versions) on long text-image pairs. LongT2IExpert produces scores that demonstrate greater consistency with human annotations. For alignment interpretation, as shown in Figure 6-right, the proposed method generates high-accuracy alignment interpretations with superior misalignment detection capabilities. (Additional visual examples can be found in the Appendix)

Ablation Study

To validate the effectiveness of each component in LongT2IExpert, we conduct comprehensive ablation experiments. As summarized in Table 3, we systematically evaluate the contributions of different components through various model variants. We first examine the performance when

Setting	Align-S	Align-I
LongT2IExpert (w/o Score)	—	32.8%
LongT2IExpert (w/o Interpretation)	0.474	—
LongT2IExpert (w/o HA-CoT)	0.516	39.9%
LongT2IExpert	0.557	53.2%

Table 3: Ablation study of LongT2IExpert components. The Avg. (SRCC+PLCC)/2 and Overall Acc for long T2I Alignment Scoring and Interpreting respectively, are reported.

LongT2IExpert is trained solely for interpreting (w/o Score) and scoring (w/o Interpretation). The comparative results demonstrate the significant advantages of multi-task training. We then evaluate performance without the Hierarchical Alignment Chain-of-Thought (HA-CoT). The results validate the effectiveness of HA-CoT in long T2I alignment, guiding MLLMs for structured alignment reasoning.

Conclusion

In this paper, we have performed a comprehensive subjective and objective benchmark of long text-to-image (T2I) alignment. We contribute LongT2IBench, the first-of-its-kind benchmark by introducing graph-structured annotations. We have also proposed LongT2IExpert, which fine-tunes a multimodal large language model (MLLM) to output token-level alignment scores and structured alignment interpretations. Extensive experiments demonstrate that LongT2IExpert achieves promising performance. While encouraging results have been obtained in this work, long T2I alignment remains far from solved—a challenge even for human evaluators. We hope our contributions will inspire the community to rethink the problem with more technical approaches and broader perspectives.

Acknowledgments

This work is supported by National Natural Science Foundation of China under Grants 62471349, 62171340 and 62301378, Fundamental Research Funds for the Central Universities under Grant QTZX25076 and YJSJ25004, the China Postdoctoral Science Foundation under Grant 2024M762553. We also sincerely thank all annotators for their dedicated efforts and significant contributions to the construction of the benchmark.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; Dang, K.; Deng, X.; Fan, Y.; Ge, W.; Han, Y.; Huang, F.; et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Chen, J.; Jincheng, Y.; Chongjian, G.; Yao, L.; Xie, E.; Wang, Z.; Kwok, J.; Luo, P.; Lu, H.; and Li, Z. 2024. PixArt-alpha: Fast Training of Diffusion Transformer for Photorealistic Text-to-Image Synthesis. In *The Twelfth International Conference on Learning Representations*.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*.
- Hamilton, W. L. 2020. *Graph representation learning*. Morgan & Claypool Publishers.
- Han, S.; Fan, H.; Fu, J.; Li, L.; Li, T.; Cui, J.; Wang, Y.; Tai, Y.; Sun, J.; Guo, C.; et al. 2024. Evalmuse-40k: A reliable and fine-grained benchmark with comprehensive human annotations for text-to-image generation model evaluation. *arXiv preprint arXiv:2412.18150*.
- Hessel, J.; Holtzman, A.; Forbes, M.; Le Bras, R.; and Choi, Y. 2021. CLIPScore: A Reference-free Evaluation Metric for Image Captioning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 7514–7528.
- Holz, D. 2023. Midjourney. <https://www.midjourney>.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Hu, X.; Wang, R.; Fang, Y.; Fu, B.; Cheng, P.; and Yu, G. 2024. Ella: Equip diffusion models with llm for enhanced semantic alignment. *arXiv preprint arXiv:2403.05135*.
- Hu, Y.; Liu, B.; Kasai, J.; Wang, Y.; Ostendorf, M.; Krishna, R.; and Smith, N. A. 2023. Tifa: Accurate and interpretable text-to-image faithfulness evaluation with question answering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 20406–20417.
- Huang, K.; Sun, K.; Xie, E.; Li, Z.; and Liu, X. 2023. T2i-compbench: A comprehensive benchmark for open-world compositional text-to-image generation. *Advances in Neural Information Processing Systems*, 36: 78723–78747.
- Kirstain, Y.; Polyak, A.; Singer, U.; Matiana, S.; Penna, J.; and Levy, O. 2023. Pick-a-pic: An open dataset of user preferences for text-to-image generation. *Advances in neural information processing systems*, 36: 36652–36663.
- Ko, H.-K.; Park, G.; Jeon, H.; Jo, J.; Kim, J.; and Seo, J. 2023. Large-scale text-to-image generation models for visual artists’ creative works. In *Proceedings of the 28th international conference on intelligent user interfaces*, 919–933.
- Lai, X.; Tian, Z.; Chen, Y.; Li, Y.; Yuan, Y.; Liu, S.; and Jia, J. 2024. Lisa: Reasoning segmentation via large language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9579–9589.
- Li, B.; Lin, Z.; Pathak, D.; Li, J.; Fei, Y.; Wu, K.; Ling, T.; Xia, X.; Zhang, P.; Neubig, G.; et al. 2024a. Genai-bench: Evaluating and improving compositional text-to-visual generation. *arXiv preprint arXiv:2406.13743*.
- Li, B.; Lin, Z.; Pathak, D.; Li, J.; Fei, Y.; Wu, K.; Xia, X.; Zhang, P.; Neubig, G.; and Ramanan, D. 2024b. Evaluating and improving compositional text-to-visual generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5290–5301.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, 19730–19742. PMLR.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual instruction tuning. *Advances in neural information processing systems*, 36: 34892–34916.
- Liu, L.; Du, C.; Pang, T.; Wang, Z.; Li, C.; and Xu, D. 2024. Improving long-text alignment for text-to-image diffusion models. *arXiv preprint arXiv:2410.11817*.
- Onoe, Y.; Rane, S.; Berger, Z.; Bitton, Y.; Cho, J.; Garg, R.; Ku, A.; Parekh, Z.; Pont-Tuset, J.; Tanzer, G.; et al. 2024. Docci: Descriptions of connected and contrasting images. In *European Conference on Computer Vision*, 291–309. Springer.
- OpenAI. 2023. Dall-e 3. <https://openai.com/dall-e-3>.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Sheng, X.; Xie, P.; Zou, W.; Chen, P.; Zhu, T.; and Li, L. 2025a. InstructCrop: Teaching Multimodal Large Language Models to Crop Aesthetic Images. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 6830–6839.
- Sheng, X.; Zou, W.; Chen, P.; Cai, L.; He, C.; and Li, L. 2025b. Text-to-Image Diffusion Models are AI-Generated Image Quality Scorers. In *Proceedings of the IEEE International Conference on Multimedia and Expo (ICME)*, 1–6.
- Tan, Z.; Yang, X.; Qin, L.; Yang, M.; Zhang, C.; and Li, H. 2024. EvalAlign: Evaluating Text-to-Image Models through Precision Alignment of Multimodal Large Models with Supervised Fine-Tuning to Human Annotations. *arXiv preprints*, arXiv:2406.

Team, G.; Anil, R.; Borgeaud, S.; Alayrac, J.-B.; Yu, J.; Soricut, R.; Schalkwyk, J.; Dai, A. M.; Hauth, A.; Millican, K.; et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.

Wang, Z. J.; Montoya, E.; Munechika, D.; Yang, H.; Hoover, B.; and Chau, D. H. 2023. DiffusionDB: A Large-scale Prompt Gallery Dataset for Text-to-Image Generative Models. In *The 61st Annual Meeting Of The Association For Computational Linguistics*.

Wu, X.; Hao, Y.; Sun, K.; Chen, Y.; Zhu, F.; Zhao, R.; and Li, H. 2023. Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. *arXiv preprint arXiv:2306.09341*.

xAI. 2025. Grok 3 beta — the age of reasoning agents. <https://x.ai/blog/grok-3>.

Xu, J.; Liu, X.; Wu, Y.; Tong, Y.; Li, Q.; Ding, M.; Tang, J.; and Dong, Y. 2023. Imagereward: Learning and evaluating human preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36: 15903–15935.

Yang, Z.; Li, L.; Chen, P.; Wu, J.; and Dong, W. 2024a. Semantics-Aware Image Aesthetics Assessment using Tag Matching and Contrastive Ranking. In *ACM International Conference on Multimedia*, 2632–2641.

Yang, Z.; Li, L.; Chen, P.; Wu, J.; and Valenzise, G. 2025. Language-Guided Visual Perception Disentanglement for Image Quality Assessment and Conditional Image Generation. *arXiv preprint arXiv:2503.02206*.

Yang, Z.; Li, L.; Yang, Y.; Li, Y.; and Lin, W. 2024b. Multi-Level Transitional Contrast Learning for Personalized Image Aesthetics Assessment. *IEEE Transactions on Multimedia*, 26: 1944–1956.

Zhang, B.; Zhang, P.; Dong, X.; Zang, Y.; and Wang, J. 2024. Long-clip: Unlocking the long-text capability of clip. In *European conference on computer vision*, 310–325. Springer.

Zhang, Z.; Kou, T.; Wang, S.; Li, C.; Sun, W.; Wang, W.; Li, X.; Wang, Z.; Cao, X.; Min, X.; et al. 2025. Q-eval-100k: Evaluating visual quality and alignment level for text-to-vision content. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 10621–10631.

Zhao, K.; Zhao, X.; Jin, Z.; Yang, Y.; Tao, W.; Han, C.; Li, S.; and Liu, L. 2024. Enhancing Baidu Multimodal Advertisement with Chinese Text-to-Image Generation via Bilingual Alignment and Caption Synthesis. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2855–2859.