

# WeEdit: A Dataset, Benchmark and Glyph-Guided Framework for Text-centric Image Editing

Hui Zhang<sup>1,2,3</sup> Juntao Liu<sup>3</sup> Zongkai Liu<sup>3</sup> Liqiang Niu<sup>3</sup> Fandong Meng<sup>3,†</sup>  
Zuxuan Wu<sup>1,2,†</sup> Yu-Gang Jiang<sup>1,2,†</sup>

<sup>1</sup>Institute of Trustworthy Embodied AI, Fudan University

<sup>2</sup>Shanghai Key Laboratory of Multimodal Embodied AI <sup>3</sup>Weixin AI, Tencent

<https://huizhang0812.github.io/WeEdit/>



**Fig. 1: Left:** WeEdit achieves precise manipulation of textual content within images across diverse editing operations (edited regions are highlighted with blue bounding boxes). **Right:** WeEdit achieves the best performance among all open-source models on both bilingual and multilingual benchmarks, surpassing most proprietary models and ranking second only to Nano Banana Pro.

**Abstract.** Instruction-based image editing aims to modify specific content within existing images according to user-provided instructions while preserving non-target regions. Beyond traditional object- and style-centric manipulation, text-centric image editing focuses on modifying, translating, or rearranging textual elements embedded within images. However, existing leading models often struggle to execute complex text editing precisely, frequently producing blurry or hallucinated characters. We attribute these failures primarily to the lack of specialized training paradigms tailored for text-centric editing, as well as the absence of large-scale datasets and standardized benchmarks necessary for a closed-loop training and evaluation system. To address these limitations, we present WeEdit, a systematic solution encompassing a scalable data construction pipeline, two benchmarks, and a tailored two-stage training strategy. Specifically, we propose a novel HTML-based automatic editing pipeline, which generates 330K training pairs covering diverse editing operations and 15 languages, accompanied by standardized bilingual and multilingual benchmarks for comprehensive evaluation. On the algorithmic side, we employ glyph-guided supervised fine-tuning to inject explicit spatial and content priors, followed by a multi-objective reinforcement learning stage to align generation with instruction adherence, text clarity, and background preservation.

Extensive experiments demonstrate that WeEdit outperforms previous open-source models by a clear margin across diverse editing operations.

**Keywords:** Image Editing · Text-centric Image Editing · Diffusion Models

## 1 Introduction

In recent years, diffusion models [25, 31, 62] have made significant progress in generative tasks, particularly in image generation and editing. Text-to-image models [26, 38, 57, 59] aim to generate high-quality, visually appealing images from textual prompts. Upon this, image editing models [14, 30, 87] have been developed to modify specific content within existing images based on editing instructions, while preserving non-target regions.

Currently, leading proprietary models (e.g. Gemini-3-Pro-Image [12] and GPT-Image-1.5 [7]) and open-source models (e.g. FLUX.2-dev [39] and Qwen-Image-Edit [72]) excel in general image editing, demonstrating precise adherence to user instructions for object manipulation and style transfer. However, a critical yet fundamentally distinct dimension—text-centric image editing, which involves modifying, translating, or rearranging textual elements embedded within images—remains largely underexplored. Such capabilities are important in a wide range of real-world applications, including infographic updating, poster modification, multilingual localization of user interfaces, and document editing. Unlike previous editing tasks, this task requires models to possess accurate text recognition, precise layout planning, and clear text generation abilities, all while strictly preserving the background context (as shown in Fig. 1).

These practical requirements expose the significant limitations of current paradigms. Existing models often struggle to follow complex editing instructions and tend to produce blurry or hallucinated characters, with performance deteriorating further when handling non-Latin scripts (e.g. Arabic, Thai, and Hindi). We attribute these failures to critical gaps in current algorithms, training data, and evaluation methodologies: **I) Algorithmic limitations:** Conventional image editing methods lack specialized training paradigms explicitly designed to handle the unique text content modification within images. **II) Data scarcity:** There is no large-scale, high-quality datasets specifically curated for diverse text-centric editing operations, particularly for multilingual settings. **III) Evaluation gap:** The absence of a comprehensive, standardized benchmark tailored to this specific task hinders systematic model comparison, slowing down progress in this field.

To this end, we propose **WeEdit**, a systematic solution for text-centric image editing that jointly addresses the aforementioned challenges of model capability, training strategy, data scarcity, and evaluation standardization: **I) Glyph-Guided Supervised Fine-Tuning:** To tackle the core difficulty of precise text placement and character-level accuracy, we introduce a text-aware fine-tuning approach. WeEdit first predicts the approximate position and scale

---

<sup>†</sup> Corresponding author.



of the target text, then renders a glyph layout as an explicit spatial prior to condition the diffusion process, granting direct spatial control and significantly improving glyph fidelity. **II) Multi-Objective Reinforcement Learning:** To bridge the gap between pixel-level supervision and human-centric quality goals such as legibility and contextual coherence, we introduce a reinforcement learning stage with a reward function that jointly balances instruction adherence, text readability, and preservation of non-edited regions, leading to higher editing quality. **III) Scalable Data Construction Pipeline:** We develop a scalable, HTML-based data construction pipeline that automatically synthesizes diverse editing pairs. The pipeline naturally extends to multilingual settings via a “translate-then-edit” workflow, yielding a large-scale training dataset covering a broad spectrum of editing operations and languages. **IV) Standardized Multilingual Benchmark:** We establish a comprehensive benchmark encompassing diverse editing operations with both Chinese–English bilingual data and multilingual evaluation across 15 widely used languages. We further design a multi-faceted evaluation scheme to ensure practical and reliable comparisons between the models. Finally, as illustrated on the right side of Fig. 1, our model outperforms existing open-source models and most proprietary counterparts, achieving SOTA performance that is second only to Gemini-3-Pro-Image (i.e. Nano Banana Pro [12]).

## 2 Related Work

### 2.1 General Image Generation and Editing

Text-to-image generation [15, 20, 26–28, 38, 44, 47, 56, 57, 59, 61, 76, 94] aims to generate visual content from textual descriptions. Image editing [14, 24, 40, 46, 73, 74, 78, 84, 87, 88, 90] extends this capability by modifying targeted regions of an existing image according to user instructions while preserving non-targeted content. Recently, leading proprietary [7, 12] and open-source models [16, 23, 39, 45, 49, 65, 66, 72], have demonstrated strong performance on object modification and style transfer. However, text-centric image editing—an equally critical yet fundamentally distinct dimension that involves modifying, translating, or rearranging textual elements embedded within images—remains largely under-explored. When confronted with such tasks, existing models fail to follow editing instructions accurately and produce blurred or misspelled characters. In this paper, we aim to systematically benchmark the text-centric editing capabilities of current models and propose effective methods to bridge this gap.

### 2.2 Text-aware Image Generation and Editing

To improve text rendering quality in generated images, recent approaches incorporate additional information such as character-level linguistic features [51, 52, 92], explicit character bounding boxes [19, 85, 86, 95–97], and auxiliary visual conditions like rendered glyph images [18, 54, 67, 81, 89]. While these approaches have improved text fidelity in image generation, few have focused on the editing task. In this work, we propose to introduce glyph images as explicit spatial priors

into text-centric image editing, and leveraging the understanding and planning capabilities of Vision-Language Models to automatically generate glyph images of target texts.

### 2.3 Datasets and Benchmarks for Image Editing

Existing image editing datasets [21, 22, 34, 55, 58, 69, 70, 82, 87, 90] and benchmarks [17, 29, 49, 68, 83, 91] are primarily designed for general-purpose editing scenarios. Some works involve text-related editing tasks [55, 66], but they are limited in terms of operation diversity, language coverage, and evaluation granularity. Thus, we propose a novel automated data construction pipeline as well as the first comprehensive benchmark tailored for text-centric image editing.

### 2.4 Post-training for Diffusion Models

Parameter-efficient supervised fine-tuning (SFT), such as LoRA [32], has become a mainstream approach for adapting pre-trained diffusion models to downstream tasks or for incorporating auxiliary control signals [33, 60, 64, 77]. Reinforcement learning (RL)-based [63] post-training further aligns generation with human preferences [37, 43, 48, 75, 79, 80, 93]. In this paper, we present a dedicated post-training framework for text-centric image editing, which consists of glyph-guided SFT to inject explicit visual priors, followed by a tailored reward design in the RL stage to ensure editing accuracy.

## 3 Method

In this section, we first review the preliminaries of diffusion and flow-matching models as well as the formulation of image editing (Sec. 3.1). We then present our two-stage training framework: a glyph-guided supervised fine-tuning stage that injects explicit spatial and content priors for accurate text rendering (Sec. 3.2), followed by a reinforcement learning stage that optimizes human-centric objectives such as instruction adherence, text clarity, and background preservation (Sec. 3.3).

### 3.1 Preliminary


**Diffusion and Flow-Matching Models.** Diffusion models generate data by reversing a forward noising process that gradually corrupts clean data  $\mathbf{x}_0 \sim p_{\text{data}}$  with Gaussian noise. The noisy sample at timestep  $t$  is obtained via reparameterization:

$$\mathbf{x}_t = \alpha_t \mathbf{x}_0 + \sigma_t \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (1)$$

where  $\alpha_t, \sigma_t$  define the noise schedule. Under the velocity parameterization, a neural network  $\mathbf{v}_\theta(\mathbf{x}_t, t)$  is trained to predict the trajectory tangent by minimizing:

$$\mathcal{L}_{\text{flow}} = \mathbb{E}_{t, \mathbf{x}_0, \boldsymbol{\epsilon}} [\|\mathbf{v}_\theta(\mathbf{x}_t, t) - \mathbf{v}_t\|_2^2], \quad (2)$$

where  $\mathbf{v}_t = \frac{d\mathbf{x}_t}{dt} = \frac{d\alpha_t}{dt} \mathbf{x}_0 + \frac{d\sigma_t}{dt} \boldsymbol{\epsilon}$ . Samples are then generated by solving the ODE  $d\mathbf{x}_t/dt = \mathbf{v}_\theta(\mathbf{x}_t, t)$  from  $t=1$  (noise) to  $t=0$  (data). Rectified flow [50] sets  $\alpha_t = 1 - t$  and  $\sigma_t = t$ , yielding straight-line interpolation paths with the simplified target  $\mathbf{v} = \boldsymbol{\epsilon} - \mathbf{x}_0$ . This formulation forms the backbone of recent image editing models [39, 72]. Our framework builds upon the flow-based MM-DiT, i.e. Qwen-Image-Edit [72].




**Fig. 2:** Overview of the glyph-guided supervised fine-tuning stage. A VLM first predicts the content and layout of the target text to render a glyph image. The original image, instruction, and glyph image are then jointly processed by the MM-DiT block to generate the target image.

**Text-centric Image Editing.** Given a source image  $\mathbf{I}_{src}$  and a user-provided editing instruction  $\mathbf{p}$ , the goal of instruction-based image editing is to generate a target image  $\mathbf{I}_{tgt}$  that semantically aligns with  $\mathbf{p}$  while preserving the non-target regions of  $\mathbf{I}_{src}$ . Distinct from general object or style editing, text-centric editing requires the model to manipulate textual elements—such as modifying, translating, deleting, or rearranging text—within an image. The generated text must accurately reflect the intent of the instruction, be correctly spelled, and maintain a consistent style with the original image, while strictly preserving the integrity of the background. In this work, we focus on this practically significant dimension of image editing tasks that has not yet been fully explored.

### 3.2 Supervised Fine-Tuning (SFT)

While recent diffusion models have demonstrated impressive general editing capabilities, they often struggle with complex text-centric editing scenarios such as long-text replacement or multilingual translation, frequently generating hallucinated or misplaced characters. To address this, we introduce a glyph-guided supervised fine-tuning stage to adapt a pre-trained image editing diffusion model for accurate text-centric editing. As illustrated in Fig. 2, our SFT pipeline comprises two key components: a glyph image prediction mechanism and a LoRA-based glyph-conditioned fine-tuning procedure.

**Glyph Image Prediction.** Given a source image  $\mathbf{I}_{src}$  and an editing instruction  $\mathbf{p}$ , we employ Qwen3-VL-235B-A22B-Instruct [13] to perform a two-step *detect-and-plan* procedure. In the detection step, the VLM identifies all text regions of interest specified by the editing instruction in the source image and outputs their bounding boxes and content as structured tuples  $\{(\mathbf{b}_i^{\text{orig}}, t_i^{\text{orig}})\}_{i=1}^m$ , where  $\mathbf{b}_i^{\text{orig}}$  denotes the normalized coordinates of the top-left and bottom-right corners, and  $t_i^{\text{orig}}$  denotes the text string. In the planning step, the VLM interprets the editing instruction  $\mathbf{p}$  and determines the target text content and spatial placement for each region, producing the target tuples  $\{(\mathbf{b}_j^{\text{tgt}}, t_j^{\text{tgt}})\}_{j=1}^n$ . Once the layout of target texts is determined, we render the glyph image  $\mathbf{x}_{\text{glyph}}$  by drawing each target text



**Fig. 3: Overview of the RL stage.** The model generates multiple candidate images, which are evaluated by four separate reward models targeting four dimensions. Each reward model leverages a Vision-Language Model to produce logit distributions over discrete scores, which are then converted to continuous expected values.

string  $t_j^{\text{tgt}}$  within its corresponding bounding box  $\mathbf{b}_j^{\text{tgt}}$  on a blank canvas that matches the source image dimensions. To ensure high-quality and standardized glyph images, we use Python Pillow package [3] and the multilingual Arial [1] font. For simplicity, all glyph images feature white text on a black background, providing a clear representation of the target texts. This glyph image explicitly includes the character content, spatial positions, and relative scales of all target texts, serving as a spatial prior for the subsequent editing process.

**LoRA-based Glyph-guided Fine-tuning.** We adopt a parameter-efficient fine-tuning strategy using Low-Rank Adaptation (LoRA) [32] to train the editing model. The model takes three inputs: the editing instruction  $\mathbf{p}$ , the source image  $\mathbf{I}_{\text{src}}$ , and the glyph image  $\mathbf{I}_{\text{glyph}}$ . Specifically, the editing instruction, along with the original image and the glyph image, is first fed into the VLM to obtain semantically enriched text tokens. Meanwhile,  $\mathbf{I}_{\text{src}}$  and  $\mathbf{I}_{\text{glyph}}$  are encoded into latent representations via a VAE [36] and concatenated with the noisy latent  $\mathbf{x}_t$  along the token dimension. We keep the original weights of the MM-DiT blocks frozen, and introduce LoRA modules only into the linear layers of the multi-modal attention mechanisms. This approach enables the model to effectively incorporate information from the glyph image, maintaining low training overhead.

### 3.3 Reinforcement Learning (RL)

While the glyph-guided SFT stage equips the model with fundamental text editing capabilities, pixel-level losses cannot capture higher-level perceptual qualities such as legibility, instruction fidelity, and background coherence. To address this limitation, we introduce an RL stage that optimizes the diffusion model against a composite reward function tailored to the unique demands of text-centric image editing, as shown in Fig. 3.

**DiffusionNFT-based Policy Optimization.** We adopt DiffusionNFT [93] as our policy optimization framework, which performs online reinforcement learning

directly on the forward diffusion process via the flow matching objective. For each conditioning input  $\mathbf{c}$  (comprising  $\mathbf{I}_{src}$ ,  $\mathbf{I}_{glyph}$ , and instruction  $\mathbf{p}$ ), we sample  $K$  candidate images  $\{\mathbf{x}_0^{1:K}\}$  from the current policy  $\pi^{\text{old}}$  using an efficient ODE solver (e.g., DPM-Solver [53]). Each candidate is then scored by the reward model to yield an optimality probability  $r \in [0, 1]$ , which serves as a soft indicator of sample quality. The training objective jointly optimizes implicitly parameterized positive and negative policies:

$$\mathcal{L}_{\text{RL}}(\theta) = \mathbb{E}_{\mathbf{c}, \pi^{\text{old}}(\mathbf{x}_0|\mathbf{c}), t} \left[ r \|\mathbf{v}_\theta^+(\mathbf{x}_t, \mathbf{c}, t) - \mathbf{v}\|_2^2 + (1 - r) \|\mathbf{v}_\theta^-(\mathbf{x}_t, \mathbf{c}, t) - \mathbf{v}\|_2^2 \right], \quad (3)$$

where  $\mathbf{v}$  denotes the target velocity, and the implicit positive and negative policies are:

$$\mathbf{v}_\theta^+(\mathbf{x}_t, \mathbf{c}, t) := (1 - \beta) \mathbf{v}^{\text{old}}(\mathbf{x}_t, \mathbf{c}, t) + \beta \mathbf{v}_\theta(\mathbf{x}_t, \mathbf{c}, t), \quad (4)$$

$$\mathbf{v}_\theta^-(\mathbf{x}_t, \mathbf{c}, t) := (1 + \beta) \mathbf{v}^{\text{old}}(\mathbf{x}_t, \mathbf{c}, t) - \beta \mathbf{v}_\theta(\mathbf{x}_t, \mathbf{c}, t). \quad (5)$$

This formulation establishes a contrastive improvement direction: the positive branch pulls the  $\mathbf{v}_\theta$  toward high-reward generations, while the negative branch pushes it away from low-reward ones, with the hyperparameter  $\beta$  controlling the guidance strength.

**Task-Specific Multi-Dimensional Reward.** We design a reward function specifically tailored to the multi-faceted requirements of text-centric image editing. Rather than relying on a single holistic quality score, we decompose the reward into four complementary dimensions, each targeting a critical and distinct aspect of editing quality:

- (1) **Instruction Adherence** ( $R_{\text{task}}^{\text{Adherence}}$ ) evaluates whether the edited image faithfully executes the specified editing operation (e.g. whether the correct text has been replaced, translated, or rearranged at the instructed locations). We prompt Qwen3-VL-235B-A22B-Instruct [13] as a reward model (same below), which takes the original image, the edited image, and the editing instruction as input, and assesses the degree of semantic alignment between the instruction intent and the actual edit.
- (2) **Text Clarity** ( $R_{\text{task}}^{\text{Clarity}}$ ) measures the legibility and typographic quality of the rendered text. The VLM evaluates whether the generated characters are crisp, correctly spelled, and free from common artifacts such as blurriness, stroke distortion, or character merging—issues that are prevalent in text-centric editing tasks involving fine-grained glyph details.
- (3) **Background Preservation** ( $R_{\text{task}}^{\text{Preservation}}$ ) assesses whether non-target regions of the image remain unaltered after editing. Unintended modifications—such as color shifts, structural distortions, or texture degradation—are penalized to ensure that edits are strictly confined to the designated text regions.
- (4) **Relative Quality** ( $R_{\text{task}}^{\text{Quality}}$ ) compares the overall visual quality of the edited image against a reference image, which can be either a ground-truth image or an output produced by a leading editing model. By grounding the evaluation in an explicit comparison target, this metric provides the VLM with a sharper perception of what constitutes a high-quality edit result, effectively establishing a quality anchor. As a result, the editing model is encouraged not only to produce

outputs that merely appear satisfactory, but to generate edits that truly match or surpass the reference.

**Logit-Weighted Continuous Scoring.** Inspired by UniWorld-V2 [45], we adopt a logit-based continuous scoring mechanism to avoid the sparsity in single-integer reward signals. For each evaluation dimension, the VLM [13] receives an input tuple  $\mathbf{X} = (\mathbf{I}_{\text{src}}, \mathbf{I}_{\text{edited}}, T_{\text{prompt}})$ , where  $T_{\text{prompt}}$  is a dimension- and task-specific prompt instructing the VLM to rate the edit on a scale from 0 to 9. Instead of sampling a single token, we compute a softmax over the score token set  $\mathcal{S} = \{0, 1, \dots, 9\}$  at the designated decoding position and obtain the per-dimension reward as the normalized expected score:

$$R_{\text{task}}^{\text{dim}}(\mathbf{X}) = \frac{1}{\max(\mathcal{S})} \sum_{s \in \mathcal{S}} s \cdot \frac{\exp(z_s)}{\sum_{s' \in \mathcal{S}} \exp(z_{s'})}, \quad (6)$$

where  $z_s$  is the logit for score token  $s$ . This soft scoring captures the VLM’s full confidence distribution, yielding a smoother reward score. The composite reward for each candidate is a weighted sum of all four dimensions:

$$R_{\text{task}} = \lambda_{\text{acc}} R_{\text{task}}^{\text{Adherence}} + \lambda_{\text{cla}} R_{\text{task}}^{\text{Clarity}} + \lambda_{\text{pre}} R_{\text{task}}^{\text{Preservation}} + \lambda_{\text{qua}} R_{\text{task}}^{\text{Quality}}, \quad (7)$$

where  $\lambda_{\text{acc}}, \lambda_{\text{cla}}, \lambda_{\text{pre}}, \lambda_{\text{qua}}$  balance the relative importance of each dimension. Notably, the evaluation prompt for each dimension is further customized by editing task type (e.g., replace, translate, rearrange, or delete), enabling task-aware scoring criteria.

To feed  $R_{\text{task}}$  into the policy optimization process, we convert it to an optimality probability  $r \in [0, 1]$  via intra-group normalization:

$$r(\mathbf{x}_0, \mathbf{c}) = \frac{1}{2} + \frac{1}{2} \text{clip} \left[ \frac{R_{\text{task}}(\mathbf{x}_0, \mathbf{c}) - \mu_{\mathbf{c}}}{\sigma_{\mathbf{c}}}, -1, 1 \right], \quad (8)$$

where  $\mu_{\mathbf{c}}$  and  $\sigma_{\mathbf{c}}$  are the mean and standard deviation of  $R_{\text{task}}$  over the  $K$  candidates sampled for conditioning input  $\mathbf{c}$ . This normalization ensures that  $r$  reflects relative quality within each sample group. The value of  $r$  is then used to directly weight the positive and negative branches in  $\mathcal{L}_{\text{RL}}$ .


## 4 Dataset and Benchmark

In this section, we first introduce our data construction pipeline, which generates high-quality source–target image pairs via two complementary routes (Sec. 4.1). We then introduce our evaluation benchmark, covering bilingual (Chinese–English) and multilingual (15 languages) text-centric editing scenarios (Sec. 4.2).

### 4.1 Dataset

Our dataset is designed to cover a broad spectrum of text-centric editing tasks. We define 7 operation types: Add (inserting new text), Replace (substituting existing text with different content), Delete (removing text regions), Rearrange (permuting the spatial layout of text elements), Translate (converting text between languages), Change Style (modifying font, color, or typographic attributes), and Combined (composing multiple atomic operations in a single instruction). Based on the nature of the source images, we identified two fundamentally different data





**Fig. 4:** Overview of our data construction pipelines. **Top:** the structured pipeline converts a source image to HTML, extracts and edits text content via a VLM, and renders both source and target images through a headless browser, yielding pixel-perfect editing pairs. **Bottom:** the unstructured pipeline uses a VLM to propose editing instructions, executes edits with a generative model, and iteratively verifies quality until all acceptance criteria are met.

categories—structured and unstructured—and designed a dedicated construction process for each, as illustrated in Fig. 4.

**Structured Data.** For images with well-organized layouts and relatively uniform text styles—such as web page screenshots, mobile app interfaces, presentation slides, figures, tables, documents, and infographics—we adopt a novel HTML-based construction pipeline (Fig. 4, top). We collected a wide range of source images from various open-source datasets, including Leopard [35], The Cauldron [41], WebSight [42], etc. The pipeline proceeds in four stages:

- (1) **Image to HTML.** Given an original image  $I_{\text{ori}}$ , we employ a VLM [11] to convert it into an HTML representation  $\mathcal{H}_{\text{src}}$  that preserves both the visual layout and the textual content as faithfully as possible, using Tailwind [5] CSS for styling.
- (2) **Content Extraction.** We use BeautifulSoup4 [2] to parse  $\mathcal{H}_{\text{src}}$  and extract a structured JSON record containing all text elements and image elements. For images, we validate external URLs and retrieve semantically similar replacements from the web when the original links are broken, ensuring rendering fidelity.
- (3) **Text Editing Pairs.** Given the  $N$  text entries extracted from  $\mathcal{H}_{\text{src}}$ , we employ a lightweight VLM Qwen3-30B-A3B [13] to generate the corresponding edited texts for each operation type. Specifically, for translation, all  $N$  entries are

translated into the target language; for replacement, 1 to  $N$  entries are randomly selected, and the VLM produces semantically plausible substitutes.

**(4) Content Backfilling.** Once the target texts are determined, we replace the original text entries in  $\mathcal{H}_{\text{src}}$  at their exact positions while keeping all other HTML elements unchanged, yielding the modified HTML code  $\mathcal{H}_{\text{tgt}}$ .

**(5) Rendering.** Both  $\mathcal{H}_{\text{src}}$  and  $\mathcal{H}_{\text{tgt}}$  are rendered into images via Playwright [4]. As the rendering process is fully deterministic and  $\mathcal{H}_{\text{src}}$  and  $\mathcal{H}_{\text{tgt}}$  differ only in the designated text entries, the resulting source image  $\mathbf{I}_{\text{src}}$  and target image  $\mathbf{I}_{\text{tgt}}$  are guaranteed to be pixel-perfect identical in all non-target regions.


Another key advantage of our HTML-based pipeline is its natural extensibility to multilingual data. Before constructing editing pairs, we can first translate all text entries in  $\mathcal{H}_{\text{src}}$  into a target language to obtain a multilingual-based HTML, and then apply the same pair-construction workflow on top of it. In this work, we cover 15 widely used languages: English, Chinese, Hindi, Spanish, French, Arabic, Portuguese, Bengali, Russian, German, Korean, Japanese, Thai, Indonesian, and Vietnamese. This automated pipeline enables efficient construction of large-scale, high-quality, and linguistically diverse text-centric editing pairs.

**Unstructured Data.** For images with complex layouts, diverse typography, and text is tightly entangled with complex visual backgrounds (e.g., street signs, product packaging, posters, and scene photographs), HTML-based reconstruction cannot faithfully recover the intricate visual styles. To address this, we develop an automated edit-verify-and-retry pipeline that operates directly at the image level (Fig. 4, bottom). Given a source image  $\mathbf{I}_{\text{src}}$ , a VLM [11] first analyzes its content and proposes a set of plausible editing instructions covering various task types. Each instruction is then executed by an editing model [12, 39, 72] to produce a candidate target image  $\hat{\mathbf{I}}_{\text{tgt}}$ . A separate VLM-based [11] verifier evaluates the candidate against the source image and instruction on instruction adherence, text legibility, and background preservation. Failed candidates are fed back with verification feedback for re-execution; only candidates passing all checks are retained as valid training pairs. This closed-loop workflow ensures high-quality data for visually challenging scenes, filters out ambiguous or infeasible edits, and provides a clean, diverse complement to the structured data.

**Dataset Statistics.** We constructed a training set containing 330K samples, including approximately 160K unstructured text pairs and 170K structured text pairs. The dataset spans 7 editing operation types and covers 15 languages. The number of editing regions per sample ranges from 1 to over 13, and the edited text length varies widely from under 20 characters to over 1,000 characters, demonstrating the diversity and complexity of our dataset. Detailed statistical breakdowns are provided in the supplementary material.

## 4.2 Benchmark

As illustrated in Fig. 5, we construct a comprehensive benchmark that covers a diverse set of text-centric editing operations, supports multiple languages, and evaluates model performance along three complementary dimensions.



**Fig. 5:** Overview of the proposed WeEdit Benchmark.

**Diverse Editing Operations.** Our benchmark encompasses eight distinct task categories to thoroughly assess model capabilities. Six of these focus on direct textual content manipulation: *Add*, *Replace*, *Delete*, *Rearrange*, *Translate*, and *Combined* (executing multiple atomic operations simultaneously). Furthermore, we include *Change Style* to evaluate the modification of typographic attributes (e.g., font, color). In addition, we introduce a more advanced *Reasoning* task that goes beyond the operation types found in the training set. This task requires the model to first deduce the correct target text based on knowledge or logical context before performing the edit. For instance, replacing “2024 NBA Champions (Boston)” with “2024 Super Bowl Champions” requires the model to recall relevant factual knowledge (e.g., team names, scores, and player details) to obtain the target text content before editing.

**Comprehensive Language Coverage.** To evaluate the robustness of editing models across different linguistic contexts, the benchmark is provided in two distinct versions, each containing 2,000 test cases. The *Bilingual Benchmark* focuses on Chinese and English, representing typical high-frequency scenarios. Meanwhile, the *Multilingual Benchmark* extends the evaluation to 15 widely used languages, testing the models’ generalization ability to diverse character sets and complex glyph structures.

**Evaluation Metrics.** We adopt a VLM-based automatic evaluation protocol, employing Gemini-3-Pro [11] as an impartial judge to score generated images across three core dimensions: *Instruction Adherence* (IA), which measures the degree to which the edited image faithfully fulfills each requirement specified in the editing instruction; *Text Clarity* (TC), which assesses the legibility and spelling correctness of the generated text; and *Background Preservation* (BP), which quantifies the visual integrity of non-target regions. The judge is prompted to first generate a detailed chain-of-thought [71] rationale that examines the edited image against the evaluation criterion, and then assign a score on a 0–9 scale, thereby producing fine-grained quality assessments.

**Table 1: Quantitative Results on the Bilingual Benchmark.** We compare WeEdit with **proprietary** and **open-source** models across 8 editing operations. IA, TC, and BP denote Instruction Adherence, Text Clarity, and Background Preservation. The best open-source results are in **bold**, with **top-3** highlighted. Our method significantly improves text editing over the **base model**, outperforms all open-source models, and surpasses most proprietary systems.

| Model                       | Add               |                  |                   | Replace           |                   |                   | Delete           |                   |                   | Rearrange   |             |             | Translate   |             |             | Style       |             |             | Combined    |             |             | Reasoning   |             |             | Overall     |             |             |  |
|-----------------------------|-------------------|------------------|-------------------|-------------------|-------------------|-------------------|------------------|-------------------|-------------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|--|
|                             | IA                | TC               | BP                | IA                | TC                | BP                | IA               | TC                | BP                | IA          | TC          | BP          | IA          | TC          | BP          | IA          | TC          | BP          | IA          | TC          | BP          | IA          | TC          | BP          | IA          | TC          | BP          |  |
| Gemini-3-Pro-Image [12]     | 9.38              | 9.66             | 9.43              | 9.06              | 8.80              | 8.05              | 8.33             | 8.56              | 8.04              | 8.05        | 8.49        | 8.26        | 8.08        | 9.07        | 9.55        | 9.76        | 9.51        | 9.50        | 9.31        | 9.72        | 9.45        | 4.91        | 9.47        | 9.42        | 8.58        | 9.10        | 8.85        |  |
| Gemini-2.5-Flash-Image [12] | 5.88              | 8.06             | 7.58              | 3.35              | 6.05              | 7.50              | 6.01             | 8.17              | 6.05              | 1.77        | 6.86        | 7.11        | 1.75        | 6.06        | 8.83        | 7.95        | 9.54        | 9.58        | 4.08        | 7.57        | 8.83        | 1.70        | 4.82        | 8.76        | 3.92        | 7.14        | 7.80        |  |
| GPT-Image-1.5 [7]           | 8.02              | 9.32             | 6.80              | 7.16              | 7.58              | 5.55              | 6.43             | 6.49              | 3.82              | 5.44        | 6.91        | 5.35        | 4.09        | 6.97        | 7.71        | 9.30        | 9.44        | 8.31        | 8.12        | 9.12        | 6.34        | 3.37        | 7.02        | 7.93        | 6.52        | 7.78        | 6.15        |  |
| Sedream-1.5 [11]            | 6.86              | 8.21             | 6.40              | 7.23              | 7.93              | 5.91              | 6.60             | 7.37              | 4.70              | 5.29        | 7.08        | 5.36        | 4.97        | 6.93        | 7.15        | 8.76        | 9.08        | 8.37        | 7.33        | 8.62        | 7.66        | 2.17        | 5.69        | 7.67        | 6.29        | 7.66        | 6.38        |  |
| OmniGen2 [23]               | 1.72              | 3.11             | 3.86              | 1.26              | 3.30              | 3.61              | 1.53             | 3.54              | 2.52              | 1.07        | 3.34        | 3.32        | 1.00        | 3.35        | 3.85        | 3.67        | 5.29        | 4.21        | 1.21        | 3.41        | 3.92        | 1.01        | 4.15        | 6.24        | 1.40        | 3.48        | 3.68        |  |
| BAGEL [24]                  | 2.32              | 3.03             | 5.99              | 1.48              | 2.41              | 4.30              | 3.03             | 6.46              | 2.92              | 1.07        | 3.53        | 5.66        | 1.03        | 4.88        | 7.68        | 7.36        | 7.76        | 8.98        | 1.44        | 2.33        | 5.83        | 1.02        | 4.42        | 8.89        | 1.97        | 4.01        | 5.75        |  |
| Emu3.5 [23]                 | 4.89              | 6.60             | 3.72              | 3.26              | 4.60              | 2.81              | 4.76             | 5.20              | 2.29              | 1.46        | 3.94        | 3.36        | 1.27        | 2.73        | 5.93        | 8.27        | 8.66        | 8.41        | 3.96        | 5.98        | 4.13        | 1.23        | 3.34        | 6.32        | 3.42        | 4.96        | 4.07        |  |
| UniWorld-v2 [13]            | 5.15              | 7.16             | 5.41              | 3.27              | 5.06              | 4.29              | 4.23             | 7.13              | 2.88              | 1.47        | 4.33        | 4.39        | 1.32        | 2.36        | 8.50        | 8.78        | 8.98        | 8.55        | 3.52        | 6.22        | 5.08        | 1.22        | 4.10        | 7.97        | 3.34        | 5.49        | 5.41        |  |
| FLUX.2-dev [10]             | 5.40              | <b>7.59</b>      | 4.92              | 3.59              | 6.09              | 4.95              | 3.07             | 8.14              | 5.27              | 1.44        | <b>5.73</b> | 5.84        | 1.24        | 3.59        | <b>8.78</b> | 8.04        | 8.88        | 8.31        | 4.60        | <b>7.31</b> | 5.52        | <b>1.26</b> | <b>6.33</b> | <b>9.55</b> | 3.37        | 6.53        | 6.19        |  |
| Qwen-Image-Edit-2511 [8]    | 4.80              | 5.91             | 8.73              | 3.36              | 3.73              | 3.77              | 3.89             | 4.46              | 3.34              | 1.42        | 2.67        | 3.95        | 1.28        | 1.60        | 6.97        | 8.50        | 8.88        | 8.48        | 3.27        | 4.21        | 4.23        | 1.09        | 1.97        | 6.04        | 3.18        | 3.93        | 4.63        |  |
| LongCat-Image-Edit [10]     | 5.23              | 6.10             | 6.80              | 3.03              | 4.51              | 5.16              | 4.65             | 7.18              | 5.84              | 1.33        | 4.91        | 5.92        | 1.29        | 4.42        | 8.29        | 8.40        | 8.95        | 9.27        | 3.89        | 5.47        | 6.87        | 1.17        | 5.05        | 8.38        | 3.39        | 5.59        | 6.71        |  |
| Step1X-Edit-v1.2 [10]       | 2.73              | 3.98             | 7.38              | 2.44              | 3.22              | <b>6.75</b>       | 5.68             | 7.01              | <b>6.45</b>       | 1.26        | 3.46        | 4.76        | 1.34        | 2.84        | 7.64        | 7.78        | 8.54        | 8.67        | 2.12        | 3.79        | <b>7.89</b> | 1.09        | <b>5.81</b> | <b>9.37</b> | 2.78        | 4.36        | <b>7.03</b> |  |
| HV-image-3-instruct [10]    | <b>6.04</b>       | 7.55             | <b>7.68</b>       | 3.91              | 5.42              | 5.79              | 6.32             | 7.85              | 4.69              | 1.73        | 4.26        | <b>6.59</b> | 1.45        | 3.52        | 8.67        | <b>9.43</b> | <b>9.35</b> | <b>9.52</b> | 4.73        | 6.62        | 7.43        | 1.12        | 4.67        | 8.65        | <b>4.16</b> | 5.99        | 7.03        |  |
| FreeRed-Image-Edit [10]     | 5.44              | 7.54             | 7.67              | <b>4.58</b>       | <b>6.30</b>       | 6.26              | <b>6.80</b>      | <b>8.81</b>       | 5.97              | <b>3.76</b> | 4.78        | 5.52        | 1.39        | 2.67        | 8.44        | <b>9.30</b> | <b>9.23</b> | <b>9.78</b> | 4.19        | 7.23        | 7.48        | 1.25        | <b>5.45</b> | 8.90        | 4.15        | 6.33        | 7.14        |  |
| Qwen-Image-Edit-2509 [8]    | 4.96              | 7.05             | 6.55              | 3.44              | 5.17              | 5.90              | 4.97             | 7.84              | 5.07              | 1.35        | 4.90        | 5.91        | 1.26        | 2.72        | 8.77        | 8.97        | 9.15        | 9.11        | 3.90        | 6.48        | 7.24        | 1.22        | 5.17        | 8.51        | 3.49        | 5.84        | 6.80        |  |
| WeEdit-SFT (ours)           | <b>7.83</b>       | <b>8.51</b>      | <b>9.18</b>       | <b>6.36</b>       | <b>6.75</b>       | <b>7.76</b>       | <b>9.28</b>      | <b>9.12</b>       | <b>8.43</b>       | <b>7.06</b> | <b>7.18</b> | <b>8.50</b> | <b>5.11</b> | <b>5.21</b> | <b>8.72</b> | <b>8.87</b> | <b>8.54</b> | <b>8.94</b> | <b>7.28</b> | <b>8.35</b> | <b>9.14</b> | <b>2.08</b> | <b>2.48</b> | <b>8.56</b> | <b>6.99</b> | <b>7.33</b> | <b>8.65</b> |  |
| WeEdit-RL (ours)            | <b>8.12</b>       | <b>8.99</b>      | <b>9.40</b>       | <b>6.90</b>       | <b>7.77</b>       | <b>8.63</b>       | <b>8.92</b>      | <b>9.49</b>       | <b>8.09</b>       | <b>7.82</b> | <b>7.94</b> | <b>9.00</b> | <b>6.36</b> | <b>7.11</b> | <b>9.44</b> | <b>9.41</b> | <b>9.24</b> | <b>9.47</b> | <b>7.50</b> | <b>8.89</b> | <b>9.38</b> | <b>2.57</b> | <b>5.21</b> | <b>8.99</b> | <b>7.47</b> | <b>8.19</b> | <b>9.01</b> |  |
| vs. Baseline                | +3.16 +1.94 +2.85 | +3.46 +2.6 +2.73 | +3.95 +1.65 +3.02 | +6.47 +3.04 +3.09 | +5.10 +4.39 +0.67 | +0.44 +0.09 +0.36 | +3.6 +2.41 +2.14 | +1.35 +0.07 +0.48 | +3.98 +2.35 +2.21 |             |             |             |             |             |             |             |             |             |             |             |             |             |             |             |             |             |             |  |

## 5 Experiment

### 5.1 Experimental Setup

**Dataset and Evaluation Metrics.** We train our models on the proposed WeEdit dataset. For evaluation, we utilize the WeEdit Benchmark, consisting of a Bilingual Benchmark (Chinese and English) and a Multilingual Benchmark, each containing 2,000 meticulously curated test cases to thoroughly assess model capabilities. Following Sec. 4.2, we employ Gemini-3-Pro [11] to assess the edited images from three dimensions.

**Implementation Details.** During the SFT stage, we fine-tune Qwen-Image-Edit-2509 [8] using LoRA with a rank of 256. The model is optimized with AdamW at a learning rate of  $5e-5$ , trained for 8,000 steps. In the RL stage, we initialize the model from the SFT weights and further fine-tune it with LoRA of rank 256 and AdamW at a learning rate of  $5 \times 10^{-5}$  for 140 epochs. The reward weights are set to  $\lambda_{acc}=0.3$ ,  $\lambda_{cla}=0.3$ ,  $\lambda_{pre}=0.15$ , and  $\lambda_{qua}=0.25$ .

### 5.2 Main Results

**Baseline Methods.** To evaluate WeEdit, we compare it against 15 previous SOTA baselines, including 4 prominent proprietary models [6, 7, 10, 12] and 11 leading open-source models [8, 9, 16, 23, 24, 39, 45, 49, 65, 66, 73]. The evaluation is conducted on both our Bilingual and Multilingual benchmarks, systematically covering 8 diverse editing operations. For each operation, the edited images are rigorously assessed across three complementary criteria: Instruction Adherence (IA), Text Clarity (TC), and Background Preservation (BP).

**Quantitative Comparison.** Tab. 1 and Tab. 2 present results on the Bilingual and Multilingual benchmarks, respectively. We observe that text-centric image editing remains a fundamental weakness of existing proprietary and, in particular,

**Table 2:** Quantitative Results on the Multilingual Benchmark.

| Model                      | Add               |             |             | Replace           |             |             | Delete            |             |             | Rearrange         |             |             | Translate         |             |             | Style             |             |             | Combined          |             |             | Reasoning         |             |             | Overall           |             |             |
|----------------------------|-------------------|-------------|-------------|-------------------|-------------|-------------|-------------------|-------------|-------------|-------------------|-------------|-------------|-------------------|-------------|-------------|-------------------|-------------|-------------|-------------------|-------------|-------------|-------------------|-------------|-------------|-------------------|-------------|-------------|
|                            | IA                | TC          | BP          | IA                | TC          | BP          | IA                | TC          | BP          | IA                | TC          | BP          | IA                | TC          | BP          | IA                | TC          | BP          | IA                | TC          | BP          | IA                | TC          | BP          | IA                | TC          | BP          |
| Gemini-3-Pro-Image [12]    | 9.59              | 9.35        | 8.83        | 8.39              | 8.21        | 6.97        | 8.17              | 8.08        | 6.71        | 7.37              | 7.97        | 7.11        | 6.34              | 7.60        | 8.66        | 8.82              | 8.77        | 7.78        | 8.83              | 8.77        | 8.12        | 6.04              | 8.87        | 9.29        | 8.05              | 8.38        | 7.82        |
| Gemini-2.5-Flash-Image [4] | 6.23              | 8.12        | 8.49        | 3.25              | 6.42        | 7.56        | 6.31              | 8.06        | 6.32        | 1.84              | 6.67        | 7.09        | 1.83              | 5.28        | 8.8         | 6.74              | 8.68        | 8.76        | 5.07              | 7.74        | 8.94        | 1.75              | 4.08        | 8.67        | 4.11              | 6.99        | 7.95        |
| GPT-Image-1.5 [7]          | 7.75              | 8.65        | 3.16        | 5.37              | 6.34        | 3.29        | 5.34              | 4.36        | 1.72        | 4.29              | 5.61        | 3.21        | 2.97              | 5.75        | 5.54        | 6.50              | 7.37        | 3.34        | 7.54              | 7.70        | 3.04        | 2.51              | 4.47        | 5.09        | 5.44              | 6.34        | 3.41        |
| Seedream4.5 [10]           | 6.33              | 7.56        | 5.88        | 5.01              | 6.14        | 5.17        | 6.25              | 6.51        | 4.23        | 3.82              | 5.61        | 4.37        | 3.04              | 5.14        | 6.53        | 7.27              | 8.07        | 6.22        | 6.51              | 7.73        | 6.42        | 1.85              | 4.55        | 7.49        | 5.10              | 6.43        | 5.56        |
| OmniGen2 [73]              | 1.71              | 2.74        | 3.04        | 1.13              | 2.93        | 2.68        | 1.75              | 3.48        | 2.08        | 1.07              | 3.38        | 2.58        | 1.10              | 3.70        | 4.49        | 2.66              | 4.19        | 2.83        | 1.70              | 3.89        | 3.32        | 1.00              | 4.22        | 5.59        | 1.45              | 3.44        | 3.15        |
| BAGEL [34]                 | 2.74              | 3.63        | 8.37        | 1.42              | 2.98        | 4.90        | 3.37              | 5.53        | 3.00        | 1.10              | 3.33        | 4.51        | 1.09              | 3.56        | 6.01        | 5.49              | 6.31        | 7.72        | 3.27              | 4.46        | 6.80        | 1.02              | 4.75        | 7.20        | 2.27              | 4.08        | 5.78        |
| Emu3.5 [23]                | 5.69              | 7.03        | 3.16        | 2.99              | 4.04        | 2.50        | 4.88              | 4.18        | 1.91        | 1.36              | 3.36        | 2.14        | 1.47              | 2.96        | 4.95        | 5.51              | 6.00        | 3.96        | 5.76              | 6.36        | 3.31        | 1.12              | 1.82        | 3.81        | 3.68              | 4.60        | 3.06        |
| UniWorld-v2 [45]           | 5.03              | 6.98        | 5.94        | 2.51              | 4.00        | 3.86        | 4.11              | 5.59        | 2.25        | 1.30              | 3.55        | 2.90        | 1.36              | 2.86        | 6.15        | 5.63              | 6.41        | 5.26        | 4.62              | 6.03        | 4.96        | 1.13              | 3.35        | 7.62        | 3.18              | 4.84        | 4.55        |
| FLUX.2-dev [89]            | 5.86              | 7.66        | 8.38        | 3.53              | 5.79        | 5.27        | 3.43              | 8.01        | 4.68        | 1.42              | 6.03        | 5.58        | 1.40              | 3.97        | 7.88        | 5.90              | 7.68        | 6.26        | 4.82              | 6.46        | 4.75        | 1.34              | <b>5.63</b> | <b>9.12</b> | 3.42              | 6.36        | 6.26        |
| Qwen-Image-Edit-2511 [8]   | 5.51              | 7.41        | 8.05        | 3.11              | 4.75        | 6.17        | 5.90              | 7.44        | 5.65        | 1.54              | 4.32        | 5.70        | 1.56              | 2.92        | 8.28        | 6.51              | 7.44        | 8.02        | 5.24              | 6.92        | 7.50        | 1.19              | 4.61        | <b>8.83</b> | 3.81              | 5.67        | 7.04        |
| LongCat-Image-Edit [67]    | 5.48              | 6.40        | 7.21        | 2.83              | 4.46        | 5.13        | 5.06              | 6.68        | 5.22        | 1.35              | 4.78        | 5.32        | 1.36              | 3.68        | 7.54        | 6.95              | 7.30        | 7.87        | 5.76              | 6.59        | 6.95        | 1.09              | 5.34        | 7.80        | 3.68              | 5.52        | 6.39        |
| StepIX-Edit-v1.2 [16]      | 2.29              | 5.31        | 8.56        | 1.31              | 3.92        | 5.53        | 4.97              | 6.97        | 5.05        | 1.24              | 3.62        | 4.61        | 1.09              | 3.04        | 6.88        | 4.41              | 6.02        | 6.26        | 2.35              | 5.22        | 6.55        | 1.12              | 5.26        | 7.97        | 2.26              | 4.77        | 6.29        |
| HY-Image-3-instruct [16]   | <b>8.19</b>       | <b>8.80</b> | <b>8.72</b> | 4.33              | 5.83        | 6.50        | <b>6.88</b>       | <b>7.92</b> | <b>5.82</b> | 1.64              | 5.35        | 7.06        | 1.75              | 3.88        | <b>8.66</b> | <b>8.05</b>       | <b>8.36</b> | 8.20        | 7.78              | <b>8.36</b> | 8.13        | 1.17              | 4.82        | 8.72        | <b>5.03</b>       | <b>6.67</b> | <b>7.58</b> |
| FireRed-Image-Edit [66]    | 5.52              | 7.42        | 8.18        | 3.25              | 5.02        | 6.53        | 6.57              | <b>8.27</b> | <b>5.52</b> | 1.71              | 4.91        | 5.98        | 1.62              | 3.78        | 8.47        | <b>7.69</b>       | 7.78        | <b>8.31</b> | 5.86              | 7.26        | 8.13        | 1.20              | 5.09        | 9.00        | 4.12              | 6.14        | 7.29        |
| Qwen-Image-Edit-2509 [8]   | 5.35              | 7.09        | 7.03        | 2.88              | 4.83        | 6.08        | 5.46              | 7.12        | 4.85        | 1.39              | 4.77        | 5.39        | 1.44              | 3.70        | 8.26        | 6.96              | 7.73        | 7.45        | 5.01              | 7.02        | 7.49        | 1.16              | 5.13        | 8.60        | 3.63              | 5.82        | 6.67        |
| WeEdit-SFT (ours)          | <b>8.29</b>       | <b>8.50</b> | <b>9.25</b> | 5.52              | 6.08        | 8.33        | <b>8.60</b>       | <b>8.87</b> | <b>8.17</b> | <b>5.72</b>       | <b>6.63</b> | <b>8.07</b> | 3.92              | 4.08        | <b>8.77</b> | 7.98              | 8.05        | <b>8.90</b> | 7.72              | 7.71        | 8.82        | 1.80              | 4.88        | 8.55        | 6.46              | 6.98        | <b>8.58</b> |
| WeEdit-RL (ours)           | <b>8.47</b>       | <b>8.33</b> | <b>8.65</b> | <b>6.12</b>       | <b>6.30</b> | <b>8.45</b> | 8.21              | 8.79        | 7.92        | 5.58              | 6.32        | <b>8.10</b> | <b>5.13</b>       | <b>5.18</b> | 8.50        | 7.14              | 8.15        | 8.62        | <b>8.15</b>       | 7.61        | <b>9.16</b> | <b>2.54</b>       | 5.25        | 8.63        | <b>6.70</b>       | <b>7.10</b> | 8.49        |
| vs. Baseline               | +3.12 +1.24 +1.62 |             |             | -3.24 +1.47 +2.37 |             |             | +2.75 +1.64 +3.07 |             |             | +4.19 +1.55 +2.71 |             |             | +3.69 +1.48 +0.24 |             |             | +0.18 +0.42 +1.17 |             |             | -3.14 +0.59 +1.67 |             |             | +1.38 +0.12 +0.03 |             |             | -3.07 +1.28 +1.82 |             |             |

open-source models. On the Bilingual benchmark, the best-performing open-source baseline achieves only 4.16 overall IA, and the situation deteriorates further on more challenging tasks: on Rearrange and Translate, all open-source baselines score below 1.8 IA, indicating near-complete failure. Even strong proprietary systems (e.g. GPT-Image-1.5 and Seedream4.5) exhibit quality degradation on these tasks, confirming that text-centric editing poses challenges qualitatively different from conventional object or style editing.

WeEdit establishes a new SOTA among open-source models by a clear margin and is surpassed only by the proprietary Gemini-3-Pro-Image. On the Bilingual benchmark, our RL version (WeEdit-RL) achieves 7.47/8.19/9.01 in IA/TC/BP, improving over the base model by +3.98/+2.35/+2.21 and surpassing the previous best open-source results by a clear margin. The gains are more pronounced in tasks such as Translate and Rearrange. Moreover, WeEdit demonstrates top-tier performance across all other editing operations.

The Multilingual benchmark (Tab. 2) further reveals that WeEdit generalizes robustly to 15 diverse languages, including those with complex glyph systems such as Arabic, Thai, and Hindi. Our model achieves 6.70/7.10/8.49 overall, improving over the base model by +3.07/+1.28/+1.82. In contrast, several proprietary and open-source models suffer performance degradation when handling non-Latin scripts. This cross-lingual robustness can be attributed to our HTML-based data construction pipeline, which naturally scales to multilingual training data.

Finally, comparing WeEdit-SFT and WeEdit-RL reveals that the RL stage provides consistent and targeted improvements. On the Bilingual benchmark, the RL stage lifts overall IA from 6.99 to 7.47 and TC from 7.33 to 8.19. This confirms that the multi-dimensional reward design effectively steers the model toward higher editing quality beyond what pixel-level supervision alone can achieve.

**Qualitative Comparison.** Fig. 6 presents qualitative comparisons across all 8 editing operations. Open-source models such as FLUX.2-dev and Qwen-Image-Edit-2509 frequently produce blurry or hallucinated characters, especially when the instructions are complex or the target text is dense. The proprietary GPT-Image-1.5 encounters similar issues under challenging tasks. In contrast, WeEdit generates clear, accurately rendered text while faithfully executing instructions across all editing types, demonstrating robust performance on a wide range






Fig. 6: **Qualitative Results.** Regions with errors are highlighted in red boxes.

of challenging text-centric editing tasks. More detailed qualitative results are provided in the appendix.

**User Study.** To subjectively evaluate the editing quality, we conducted a user study comparing WeEdit against two open-source models and two proprietary systems. As illustrated in Fig. 7, human evaluators assessed the generated results across two dimensions: Instruction Adherence and Text Clarity on both Bilingual and Multilingual benchmarks. The results demonstrate that WeEdit significantly outperforms the open-source baselines Qwen-Image-Edit-2509 and FLUX.2-dev, as well as the proprietary GPT-Image-1.5, while achieving comparable performance with Gemini-3-Pro-Image.

### 5.3 Ablation Study

To systematically evaluate the contribution of each proposed module in WeEdit, we conduct a comprehensive ablation study, as summarized in Fig. 8. As shown in the first row, the base model exhibits poor performance across three dimensions. Applying glyph guidance directly to the baseline model at the inference stage (Row 2) results in negligible performance changes, indicating that the base model cannot effectively incorporate explicit spatial priors without specific training. Comparing the baseline with the SFT-only model (Row 3) reveals a significant performance leap across all three dimensions, which demonstrates the high quality and effectiveness of our constructed dataset. Furthermore, introducing RL directly




Fig. 7: User Study.

| Ablation |       |    |    |     | Performance |             |             |
|----------|-------|----|----|-----|-------------|-------------|-------------|
| SFT      | Glyph | RL | RI | SRM | IA          | TC          | BP          |
|          |       |    |    |     | 3.49        | 5.84        | 6.80        |
|          | ✓     |    |    |     | 3.58        | 4.67        | 6.82        |
| ✓        |       |    |    |     | 5.32        | 6.62        | 8.11        |
| ✓        | ✓     |    | ✓  | ✓   | 5.41        | 6.72        | 8.07        |
| ✓        | ✓     |    |    |     | 6.99        | 7.33        | 8.63        |
| ✓        | ✓     | ✓  | ✓  | ✓   | 7.38        | 7.91        | 8.92        |
| ✓        | ✓     | ✓  | ✓  | ✓   | 7.34        | 8.03        | 8.89        |
| ✓        | ✓     | ✓  | ✓  | ✓   | <b>7.47</b> | <b>8.19</b> | <b>9.01</b> |

Fig. 8: Ablation study on different modules in WeEdit.

on top of this SFT-only baseline (Row 4) struggles to break through to a higher performance ceiling. In contrast, combining SFT with Glyph guidance (Row 5) yields massive improvements, demonstrating a strong synergy where SFT aligns the model with the task while the Glyph prior provides the crucial spatial constraints needed for precise text rendering. Building upon this foundation, applying RL on top of the SFT-with-glyph model leads to further consistent enhancements. Within the RL stage, the design of the reward mechanism proves critical. Removing the Reference Image (RI) deprives the reward model of an explicit comparison anchor (Row 6), resulting in less accurate assessment of edit quality than when the reference is included. Relying on a single reward model to score multiple dimensions simultaneously leads to sub-optimal performance due to metric entanglement (Row 7), whereas using Separate Reward Models (SRM) ensures the independent and accurate assessment of each dimension. The full WeEdit pipeline (Row 8) achieves the best results across all metrics, confirming the effectiveness of our training strategies and data.

## 6 Conclusion

In this paper, we presented WeEdit, a systematic solution for text-centric image editing that jointly addresses the challenges of model capability, data scarcity, and evaluation standardization. We introduced a glyph-guided supervised fine-tuning approach that leverages rendered glyph images as explicit spatial priors to enable precise text placement and character-level fidelity. We further proposed a multi-objective reinforcement learning stage with separate reward models and reference image grounding to optimize for instruction adherence, text readability, and background preservation. To overcome the lack of high-quality training data, we proposed a novel, scalable, HTML-based data construction pipeline that automatically synthesizes diverse editing pairs and naturally extends to multilingual settings. Additionally, we established a comprehensive benchmark covering diverse editing operations with bilingual and multilingual evaluation across 15 languages. Extensive experiments demonstrated that WeEdit surpasses all existing open-source models and most proprietary counterparts.



## Acknowledgments

This work is supported by the National Natural Science Foundation of China (Grant No. 62472098) and the Science and Technology Commission of Shanghai Municipality (No. 25511106100).

## References

1. Arial. <https://learn.microsoft.com/en-us/typography/fontlist/arial> 6
2. Beautifulsoup4. <https://pypi.org/project/beautifulsoup4/> 9
3. Pillow. <https://github.com/python-pillow/Pillow> 6
4. Playwright. <https://github.com/microsoft/playwright> 10
5. Tailwind. <https://tailwindcss.com/> 9
6. Gemini-2.5-flash-image. <https://developers.googleblog.com/introducing-gemini-2-5-flash-image/> (2025) 12, 13
7. Gpt-image-1.5. <https://developers.openai.com/api/docs/models/gpt-image-1.5> (2025) 2, 3, 12, 13
8. Qwen-image-edit-2509. <https://huggingface.co/Qwen/Qwen-Image-Edit-2509> (2025) 12, 13
9. Qwen-image-edit-2511. <https://huggingface.co/Qwen/Qwen-Image-Edit-2511> (2025) 12, 13
10. Seedream4.5. [https://seed.bytedance.com/en/seedream4\\_5](https://seed.bytedance.com/en/seedream4_5) (2025) 12, 13
11. Gemini-3.0-pro. <https://blog.google/products-and-platforms/products/gemini/gemini-3/> (202511) 9, 10, 11, 12
12. Gemini-3.0-pro-image. <https://blog.google/innovation-and-ai/technology/developers-tools/gemini-3-pro-image-developers/> (202511) 2, 3, 10, 12, 13
13. Bai, S., Cai, Y., Chen, R., Chen, K., Chen, X., Cheng, Z., Deng, L., Ding, W., Gao, C., Ge, C., et al.: Qwen3-vl technical report. arXiv preprint arXiv:2511.21631 (2025) 5, 7, 8, 9
14. Brooks, T., Holynski, A., Efros, A.A.: Instructpix2pix: Learning to follow image editing instructions. In: CVPR (2023) 2, 3
15. Cai, Q., Chen, J., Chen, Y., Li, Y., Long, F., Pan, Y., Qiu, Z., Zhang, Y., Gao, F., Xu, P., et al.: Hidream-i1: A high-efficient image generative foundation model with sparse diffusion transformer. arXiv preprint arXiv:2505.22705 (2025) 3
16. Cao, S., Chen, H., Chen, P., Cheng, Y., Cui, Y., Deng, X., Dong, Y., Gong, K., Gu, T., Gu, X., et al.: Hunyuanimage 3.0 technical report. arXiv preprint arXiv:2509.23951 (2025) 3, 12, 13
17. Chang, D., Cao, M., Shi, Y., Liu, B., Cai, S., Zhou, S., Huang, W., Wetzstein, G., Soleymani, M., Wang, P.: Bytemorph: Benchmarking instruction-guided image editing with non-rigid motions. arXiv preprint arXiv:2506.03107 (2025) 4
18. Chen, J., Huang, Y., Lv, T., Cui, L., Chen, Q., Wei, F.: Textdiffuser: Diffusion models as text painters. NeurIPS (2023) 3
19. Chen, J., Huang, Y., Lv, T., Cui, L., Chen, Q., Wei, F.: Textdiffuser-2: Unleashing the power of language models for text rendering. In: ECCV (2024) 3
20. Chen, J., Jincheng, Y., Chongjian, G., Yao, L., Xie, E., Wang, Z., Kwok, J., Luo, P., Lu, H., Li, Z.: Pixart- $\alpha$ : Fast training of diffusion transformer for photorealistic text-to-image synthesis. In: ICLR (2024) 3
21. Chen, J., Cai, Z., Chen, P., Chen, S., Ji, K., Wang, X., Yang, Y., Wang, B.: Sharegpt-4o-image: Aligning multimodal models with gpt-4o-level image generation. arXiv preprint arXiv:2506.18095 (2025) 4

22. Chen, Z., Bai, X., Shi, Y., Fu, C., Zhang, H., Wang, H., Sun, X., Zhang, Z., Wang, L., Zhang, Y., et al.: Opengpt-4o-image: A comprehensive dataset for advanced image generation and editing. arXiv preprint arXiv:2509.24900 (2025) [4](#)
23. Cui, Y., Chen, H., Deng, H., Huang, X., Li, X., Liu, J., Liu, Y., Luo, Z., Wang, J., Wang, W., et al.: Emu3. 5: Native multimodal models are world learners. arXiv preprint arXiv:2510.26583 (2025) [3](#), [12](#), [13](#)
24. Deng, C., Zhu, D., Li, K., Gou, C., Li, F., Wang, Z., Zhong, S., Yu, W., Nie, X., Song, Z., et al.: Emerging properties in unified multimodal pretraining. arXiv preprint arXiv:2505.14683 (2025) [3](#), [12](#), [13](#)
25. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. In: NeurIPS (2021) [2](#)
26. Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., et al.: Scaling rectified flow transformers for high-resolution image synthesis. In: ICML (2024) [2](#), [3](#)
27. Gao, P., Zhuo, L., Lin, Z., Liu, C., Chen, J., Du, R., Xie, E., Luo, X., Qiu, L., Zhang, Y., et al.: Lumina-t2x: Transforming text into any modality, resolution, and duration via flow-based large diffusion transformers. arXiv preprint arXiv:2405.05945 (2024) [3](#)
28. Gao, Y., Gong, L., Guo, Q., Hou, X., Lai, Z., Li, F., Li, L., Lian, X., Liao, C., Liu, L., et al.: Seedream 3.0 technical report. arXiv preprint arXiv:2504.11346 (2025) [3](#)
29. Han, F., Wang, Y., Li, C., Liang, Z., Wang, D., Jiao, Y., Wei, Z., Gong, C., Jin, C., Chen, J., et al.: Unireditbench: A unified reasoning-based image editing benchmark. arXiv preprint arXiv:2511.01295 (2025) [4](#)
30. Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., Cohen-Or, D.: Prompt-to-prompt image editing with cross-attention control. In: ICLR (2023) [2](#)
31. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: NeurIPS (2020) [2](#)
32. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. In: ICLR (2022) [4](#), [6](#)
33. Huang, L., Wang, W., Wu, Z.F., Shi, Y., Dou, H., Liang, C., Feng, Y., Liu, Y., Zhou, J.: In-context lora for diffusion transformers. arXiv preprint arXiv:2410.23775 (2024) [4](#)
34. Hui, M., Yang, S., Zhao, B., Shi, Y., Wang, H., Wang, P., Xie, C., Zhou, Y.: Hq-edit: A high-quality dataset for instruction-based image editing. In: ICLR (2024) [4](#)
35. Jia, M., Yu, W., Ma, K., Fang, T., Zhang, Z., Ouyang, S., Zhang, H., Yu, D., Jiang, M.: Leopard: A vision language model for text-rich multi-image tasks. arXiv preprint arXiv:2410.01744 (2024) [9](#)
36. Kingma, D.P.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013) [6](#)
37. Kirstain, Y., Polyak, A., Singer, U., Matiana, S., Penna, J., Levy, O.: Pick-a-pic: An open dataset of user preferences for text-to-image generation. NeurIPS (2023) [4](#)
38. Labs, B.F.: Flux.1-dev. <https://blackforestlabs.ai/announcing-black-forest-labs> (2024) [2](#), [3](#)
39. Labs, B.F.: Flux.2-dev. <https://huggingface.co/black-forest-labs/FLUX.2-dev> (2025) [2](#), [3](#), [4](#), [10](#), [12](#), [13](#)
40. Labs, B.F., Batifol, S., Blattmann, A., Boesel, F., Consul, S., Diagne, C., Dockhorn, T., English, J., English, Z., Esser, P., et al.: Flux. 1 kontext: Flow matching for in-context image generation and editing in latent space. arXiv preprint arXiv:2506.15742 (2025) [3](#)
41. Laurençon, H., Tronchon, L., Cord, M., Sanh, V.: What matters when building vision-language models? NeurIPS (2024) [9](#)

42. Laurençon, H., Tronchon, L., Sanh, V.: Unlocking the conversion of web screenshots into html code with the websight dataset. arXiv preprint arXiv:2403.09029 (2024) [9](#)
43. Li, J., Cui, Y., Huang, T., Ma, Y., Fan, C., Yang, M., Zhong, Z.: Mixgrpo: Unlocking flow-based grpo efficiency with mixed ode-sde. arXiv preprint arXiv:2507.21802 (2025) [4](#)
44. Li, Z., Zhang, J., Lin, Q., Xiong, J., Long, Y., Deng, X., Zhang, Y., Liu, X., Huang, M., Xiao, Z., et al.: Hunyuan-dit: A powerful multi-resolution diffusion transformer with fine-grained chinese understanding. arXiv e-prints (2024) [3](#)
45. Li, Z., Liu, Z., Zhang, Q., Lin, B., Wu, F., Yuan, S., Yan, Z., Ye, Y., Yu, W., Niu, Y., et al.: Uniworld-v2: Reinforce image editing with diffusion negative-aware finetuning and mllm implicit feedback. arXiv preprint arXiv:2510.16888 (2025) [3](#), [8](#), [12](#), [13](#)
46. Lin, B., Li, Z., Cheng, X., Niu, Y., Ye, Y., He, X., Yuan, S., Yu, W., Wang, S., Ge, Y., et al.: Uniworld-v1: High-resolution semantic encoders for unified visual understanding and generation. arXiv preprint arXiv:2506.03147 (2025) [3](#)
47. Liu, B., Akhgari, E., Visseratin, A., Kamko, A., Xu, L., Shrirao, S., Souza, J., Doshi, S., Li, D.: Playground v3: Improving text-to-image alignment with deep-fusion large language models. arXiv preprint arXiv:2409.10695 (2024) [3](#)
48. Liu, J., Liu, G., Liang, J., Li, Y., Liu, J., Wang, X., Wan, P., ZHANG, D., Ouyang, W.: Flow-grpo: Training flow matching models via online rl. In: The Thirty-ninth Annual Conference on Neural Information Processing Systems (2025) [4](#)
49. Liu, S., Han, Y., Xing, P., Yin, F., Wang, R., Cheng, W., Liao, J., Wang, Y., Fu, H., Han, C., et al.: Step1x-edit: A practical framework for general image editing. arXiv preprint arXiv:2504.17761 (2025) [3](#), [4](#), [12](#), [13](#)
50. Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. In: ICLR (2023) [4](#)
51. Liu, Z., Liang, W., Liang, Z., Luo, C., Li, J., Huang, G., Yuan, Y.: Glyph-byt5: A customized text encoder for accurate visual text rendering. In: ECCV (2024) [3](#)
52. Liu, Z., Liang, W., Zhao, Y., Chen, B., Liang, L., Wang, L., Li, J., Yuan, Y.: Glyph-byt5-v2: A strong aesthetic baseline for accurate multilingual visual text rendering. arXiv preprint arXiv:2406.10208 (2024) [3](#)
53. Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. NeurIPS (2022) [7](#)
54. Ma, J., Zhao, M., Chen, C., Wang, R., Niu, D., Lu, H., Lin, X.: Glyphdraw: Seamlessly rendering text with intricate spatial structures in text-to-image generation. arXiv preprint arXiv:2303.17870 (2023) [3](#)
55. Ma, J., Zhu, X., Pan, Z., Peng, Q., Guo, X., Chen, C., Lu, H.: X2edit: Revisiting arbitrary-instruction image editing through self-constructed data and task-aware representation learning. arXiv preprint arXiv:2508.07607 (2025) [4](#)
56. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: CVPR (2023) [3](#)
57. Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: Sdxl: Improving latent diffusion models for high-resolution image synthesis. In: ICLR (2024) [2](#), [3](#)
58. Qian, Y., Bocek-Rivele, E., Song, L., Tong, J., Yang, Y., Lu, J., Hu, W., Gan, Z.: Pico-banana-400k: A large-scale dataset for text-guided image editing. arXiv preprint arXiv:2510.19808 (2025) [4](#)
59. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022) [2](#), [3](#)

60. Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In: CVPR (2023) [4](#)
61. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. In: NeurIPS (2022) [3](#)
62. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: ICLR (2021) [2](#)
63. Sutton, R.S., Barto, A.G.: Reinforcement learning: An introduction. MIT press (2018) [4](#)
64. Tan, Z., Liu, S., Yang, X., Xue, Q., Wang, X.: Ominicontrol: Minimal and universal control for diffusion transformer. In: ICCV. pp. 14940–14950 (2025) [4](#)
65. Team, M.L., Ma, H., Tan, H., Huang, J., Wu, J., He, J.Y., Gao, L., Xiao, S., Wei, X., Ma, X., et al.: Longcat-image technical report. arXiv preprint arXiv:2512.07584 (2025) [3](#), [12](#), [13](#)
66. Team, S.I., Qiao, C., Hui, C., Li, C., Wang, C., Song, D., Zhang, J., Li, J., Xiang, Q., Wang, R., et al.: Fired-image-edit-1.0 technical report. arXiv preprint arXiv:2602.13344 (2026) [3](#), [4](#), [12](#), [13](#)
67. Tuo, Y., Xiang, W., He, J.Y., Geng, Y., Xie, X.: Anytext: Multilingual visual text generation and editing. arXiv preprint arXiv:2311.03054 (2023) [3](#)
68. Wang, J., Wang, J., Duan, H., Kang, J., Zhai, G., Min, X.: I2i-bench: A comprehensive benchmark suite for image-to-image editing models. arXiv preprint arXiv:2512.04660 (2025) [4](#)
69. Wang, Y., Yang, S., Zhao, B., Zhang, L., Liu, Q., Zhou, Y., Xie, C.: Gpt-image-edit-1.5 m: A million-scale, gpt-generated image dataset. arXiv preprint arXiv:2507.21033 (2025) [4](#)
70. Wei, C., Xiong, Z., Ren, W., Du, X., Zhang, G., Chen, W.: Omniedit: Building image editing generalist models through specialist supervision. In: ICLR (2024) [4](#)
71. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q.V., Zhou, D., et al.: Chain-of-thought prompting elicits reasoning in large language models. NeurIPS (2022) [11](#)
72. Wu, C., Li, J., Zhou, J., Lin, J., Gao, K., Yan, K., Yin, S.m., Bai, S., Xu, X., Chen, Y., et al.: Qwen-image technical report. arXiv preprint arXiv:2508.02324 (2025) [2](#), [3](#), [4](#), [10](#)
73. Wu, C., Zheng, P., Yan, R., Xiao, S., Luo, X., Wang, Y., Li, W., Jiang, X., Liu, Y., Zhou, J., et al.: Omnigen2: Exploration to advanced multimodal generation. arXiv preprint arXiv:2506.18871 (2025) [3](#), [12](#), [13](#)
74. Wu, J.Z., Ren, X., Shen, T., Cao, T., He, K., Lu, Y., Gao, R., Xie, E., Lan, S., Alvarez, J.M., et al.: Chronoedit: Towards temporal reasoning for image editing and world simulation. arXiv preprint arXiv:2510.04290 (2025) [3](#)
75. Wu, K., Jiang, S., Ku, M., Nie, P., Liu, M., Chen, W.: Editreward: A human-aligned reward model for instruction-guided image editing. arXiv preprint arXiv:2509.26346 (2025) [4](#)
76. Wu, K., Yang, Z., Zhang, K., Wang, S., Zhu, H., Leng, S., Yang, Z., Wang, Q., Wang, S., Wang, Z., et al.: Visual generation in the new era: An evolution from atomic mapping to agentic world modeling. arXiv preprint arXiv:2604.28185 (2026) [3](#)
77. Wu, S., Huang, M., Wu, W., Cheng, Y., Ding, F., He, Q.: Less-to-more generalization: Unlocking more controllability by in-context generation. In: ICCV (2025) [4](#)

78. Xiao, S., Wang, Y., Zhou, J., Yuan, H., Xing, X., Yan, R., Li, C., Wang, S., Huang, T., Liu, Z.: Omnigen: Unified image generation. In: CVPR (2025) [3](#)
79. Xu, J., Liu, X., Wu, Y., Tong, Y., Li, Q., Ding, M., Tang, J., Dong, Y.: Imagereward: Learning and evaluating human preferences for text-to-image generation. NeurIPS (2023) [4](#)
80. Xue, Z., Wu, J., Gao, Y., Kong, F., Zhu, L., Chen, M., Liu, Z., Liu, W., Guo, Q., Huang, W., et al.: Dancegrpo: Unleashing grpo on visual generation. arXiv preprint arXiv:2505.07818 (2025) [4](#)
81. Yang, Y., Gui, D., Yuan, Y., Liang, W., Ding, H., Hu, H., Chen, K.: Glyphcontrol: Glyph conditional control for visual text generation. NeurIPS (2023) [3](#)
82. Ye, J., Jiang, D., Wang, Z., Zhu, L., Hu, Z., Huang, Z., He, J., Yan, Z., Yu, J., Li, H., et al.: Echo-4o: Harnessing the power of gpt-4o synthetic images for improved image generation. arXiv preprint arXiv:2508.09987 (2025) [4](#)
83. Ye, Y., He, X., Li, Z., Lin, B., Yuan, S., Yan, Z., Hou, B., Yuan, L.: Imgedit: A unified image editing dataset and benchmark. arXiv preprint arXiv:2505.20275 (2025) [4](#)
84. Yu, Q., Chow, W., Yue, Z., Pan, K., Wu, Y., Wan, X., Li, J., Tang, S., Zhang, H., Zhuang, Y.: Anyedit: Mastering unified high-quality image editing for any idea. In: CVPR (2025) [3](#)
85. Zhang, H., Hong, D., Wang, Y., Shao, J., Wu, X., Wu, Z., Jiang, Y.G.: Creatilayout: Siamese multimodal diffusion transformer for creative layout-to-image generation. In: ICCV (2025) [3](#)
86. Zhang, H., Hong, D., Yang, M., Cheng, Y., Zhang, Z., Shao, J., Wu, X., Wu, Z., Jiang, Y.G.: Creatidesign: A unified multi-conditional diffusion transformer for creative graphic design. In: ICLR (2026) [3](#)
87. Zhang, K., Mo, L., Chen, W., Sun, H., Su, Y.: Magicbrush: A manually annotated dataset for instruction-guided image editing. NeurIPS (2023) [2, 3, 4](#)
88. Zhang, Z., Xie, J., Lu, Y., Yang, Z., Yang, Y.: Enabling instructional image editing with in-context generation in large scale diffusion transformer. In: NeurIPS (2025) [3](#)
89. Zhang, Z., Cheng, Y., Hong, D., Yang, M., Shi, G., Ma, L., Zhang, H., Shao, J., Wu, X.: Creatiposter: Towards editable and controllable multi-layer graphic design generation. arXiv preprint arXiv:2506.10890 (2025) [3](#)
90. Zhao, H., Ma, X.S., Chen, L., Si, S., Wu, R., An, K., Yu, P., Zhang, M., Li, Q., Chang, B.: Ultraedit: Instruction-based fine-grained image editing at scale. NeurIPS (2024) [3, 4](#)
91. Zhao, X., Zhang, P., Tang, K., Zhu, X., Li, H., Chai, W., Zhang, Z., Xia, R., Zhai, G., Yan, J., et al.: Envisioning beyond the pixels: Benchmarking reasoning-informed visual editing. arXiv preprint arXiv:2504.02826 (2025) [4](#)
92. Zhao, Y., Lian, Z.: Udifftext: A unified framework for high-quality text synthesis in arbitrary images via character-aware diffusion models. In: ECCV (2024) [3](#)
93. Zheng, K., Chen, H., Ye, H., Wang, H., Zhang, Q., Jiang, K., Su, H., Ermon, S., Zhu, J., Liu, M.Y.: Diffusionnft: Online diffusion reinforcement with forward process. arXiv preprint arXiv:2509.16117 (2025) [4, 6](#)
94. Zheng, W., Teng, J., Yang, Z., Wang, W., Chen, J., Gu, X., Dong, Y., Ding, M., Tang, J.: Cogview3: Finer and faster text-to-image generation via relay diffusion. In: ECCV (2024) [3](#)
95. Zhou, D., Li, M., Yang, Z., Yang, Y.: Dreamrenderer: Taming multi-instance attribute control in large-scale text-to-image models. In: ICCV. pp. 16712–16722 (2025) [3](#)

96. Zhou, D., Li, Y., Ma, F., Yang, Z., Yang, Y.: Migc++: Advanced multi-instance generation controller for image synthesis. TPAMI (2024) [3](#)
97. Zhou, D., Li, Y., Ma, F., Zhang, X., Yang, Y.: Migc: Multi-instance generation controller for text-to-image synthesis. In: CVPR (2024) [3](#)