

MCIE: Multimodal LLM-Driven Complex Instruction Image Editing with Spatial Guidance

Xuehai Bai¹, Xiaoling Gu^{1*}, Akide Liu², Hangjie Yuan³, YiFan Zhang⁴, Jack Ma⁵

¹Hangzhou Dianzi University,

²Monash University,

³Zhejiang University,

⁴University of Chinese Academy of Sciences,

⁵Tsinghua University

Abstract

Recent advances in instruction-based image editing have shown remarkable progress. However, existing methods remain limited to relatively simple editing operations, hindering real-world applications that require complex and compositional instructions. In this work, we address these limitations from the perspectives of *architectural design*, *data*, and *evaluation protocols*. Specifically, we identify two key challenges in current models: insufficient instruction compliance and background inconsistency. To this end, we propose **MCIE-E1**, a Multimodal Large Language Model-Driven Complex Instruction Image Editing method that integrates two key modules: a *spatial-aware cross-attention* module and a *background-consistent cross-attention* module. The former enhances instruction-following capability by explicitly aligning semantic instructions with spatial regions through spatial guidance during the denoising process, while the latter preserves features in unedited regions to maintain background consistency. To enable effective training, we construct a dedicated data pipeline to mitigate the scarcity of complex instruction-based image editing datasets, combining fine-grained automatic filtering via a powerful MLLM with rigorous human validation. Finally, to comprehensively evaluate complex instruction-based image editing, we introduce **CIE-Bench**, a new benchmark with two new evaluation metrics. Experimental results on CIE-Bench demonstrate that MCIE-E1 consistently outperforms previous state-of-the-art methods in both quantitative and qualitative assessments, achieving a 23.96% improvement in instruction compliance.

Code — <https://github.com/bxhsort/MCIE>

1 Introduction

Diffusion models have achieved notable progress in instruction-based image editing, attracting substantial attention from both academia and industry. These tasks range from simple edits (Brooks, Holynski, and Efros 2023; Yu et al. 2025; Sheynin et al. 2024) that modify a single element within an image to complex manipulations (Guo and Lin 2024) requiring the simultaneous editing of multiple regions. Although existing diffusion-based models (Zhao et al.

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

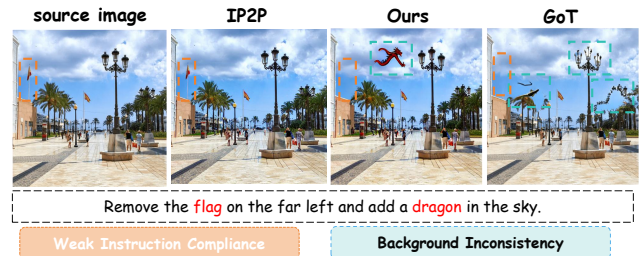


Figure 1: We propose MCIE-E1 to address the challenges of weak instruction compliance and background consistency in complex instruction-based image editing.

2024; Yu et al. 2025) perform well on simple editing instructions, they often struggle with detailed and compositional modifications. This gap between simple and complex editing capabilities in current models raises a fundamental question: how can we advance instruction-based image editing toward reliable and controllable complex editing?

Prior work (Yu et al. 2025; Hertz et al. 2022) has explored instruction-based image editing from different perspectives. One line of research leverages CLIP (Ramesh et al. 2022) to extract intended modifications from user instructions. Another line of research (Huang et al. 2024; Fang et al. 2025) integrates multimodal large language models (MLLMs) into diffusion models to enhance the understanding of fine-grained editing instructions. As shown in Fig. 1, when users provide complex instructions, existing approaches (Brooks, Holynski, and Efros 2023; Fang et al. 2025) often either overlook sub-instructions, leading to suboptimal instruction-following, or introduce unintended changes, compromising background consistency.

Current methods still face challenges in complex instruction-based image editing, primarily for the following reasons: (1) **Lack of high-quality data for complex instruction editing**. Existing data creation pipelines (Zhang et al. 2023a; Brooks, Holynski, and Efros 2023; Yu et al. 2025; Zhao et al. 2024; Bai et al. 2024) mainly focus on simple instructions and low-resolution images, and often lack fine-grained, multi-stage post-processing, limiting the model’s ability to perform instruction-based image editing.



Figure 2: Visual results of MCIE-E1. Our method effectively performs complex instruction-based image editing with accurate and consistent outputs.

(2) **Inadequate instruction-region alignment.** Although recent methods (Brooks, Holynski, and Efros 2023; Fu et al. 2023; Yu et al. 2025; Huang et al. 2024; Fang et al. 2025) have made notable advances, they still struggle to precisely localize editing regions for complex instructions, leading to incomplete or inaccurate edits. This misalignment also affects the model’s ability to preserve background consistency and accurately follow instructions, further reducing output reliability. (3) **Limited benchmarking protocols.** Current evaluation frameworks (Ma et al. 2024; Sheynin et al. 2024; Wang et al. 2023; Basu et al. 2023; Kawar et al. 2023) often fall short in comprehensively assessing a model’s ability to follow complex instructions and maintain background consistency, which hinders the effective evaluation of complex instruction-based image editing.

To address these challenges, we propose a unified framework that combines: (1) **A large-scale and high-quality dataset.** We introduce the MCIE dataset, the first dataset tailored for complex instruction image editing. First, we aggregate multi-round edits into complex instructions, which may potentially introduce semantic conflicts. As illustrated in Fig. 3(b), we use Qwen2.5-VL-72B (Bai et al. 2025) to detect conflicts in the augmented dataset. Additionally, we generate bounding boxes to provide spatial guidance for complex instructions. Finally, human experts meticulously filter the data along three dimensions: instruction consistency, image quality, and editing scenario complexity. (2) **A novel image editing model for complex instruction.** We propose MCIE-E1, a model for complex instruction-based image editing. It leverages MLLMs to disentangle complex instructions into sub-instructions with their corresponding editing regions. Specifically, we introduce spatial-aware cross-attention, which processes sub-instructions independently and injects spatial guidance to enhance instruction following. To preserve background consistency, we propose background-consistent cross-attention, which utilizes pixel-level visual features to maintain coherence in unedited regions. (3) **A general benchmark for evaluating complex instruction-based image editing.** We construct CIE-Bench, which includes 400 carefully crafted evaluation sets, and we introduce two metrics to assess the model’s ability to follow instructions and maintain background consistency.

By leveraging MLLMs to decompose complex instruc-

tions and incorporating a semantically and spatially aware diffusion model, MCIE-E1 achieves superior instruction compliance and background consistency, as shown in Fig. 2. Extensive experimental results demonstrate that our method outperforms existing approaches on complex instruction-based image editing.

Our main contributions are summarized as follows:

- We tackle the novel task of complex instruction-based image editing and present the MCIE dataset, which comprises 90k samples carefully filtered by human experts along multiple dimensions.
- To address the challenge of complex instruction-based image editing, we propose MCIE-E1, a novel model with two key components: a spatial-aware cross-attention module for enhancing instruction following and a background-consistent cross-attention module for preserving unedited areas.
- To validate the effectiveness of our method, we construct CIE-Bench and introduce two evaluation metrics, Instruction Compliance and Background Consistency, on which MCIE-E1 demonstrates clear superiority.

2 Related Work

2.1 Instruction-Based Image Editing

IP2P (Brooks, Holynski, and Efros 2023) is a pioneering work that employs instructions to edit images by fine-tuning a diffusion model on a large-scale dataset. InstructDiffusion (Geng et al. 2024) further generalizes instruction-based image editing to conventional vision tasks. AnyEdit (Yu et al. 2025) introduces a Mixture-of-Experts mechanism to address diverse editing scenarios. Despite these advancements, most existing methods rely on CLIP (Ramesh et al. 2022) or T5 (Raffel et al. 2020), which limits their ability to interpret fine-grained instructions. To mitigate this limitation, SmartEdit (Huang et al. 2024) and MGIE (Fu et al. 2023) integrate MLLMs into diffusion-based frameworks to enhance instruction comprehension. However, these approaches still exhibit limited performance in handling complex instructions. The most relevant work to our study is FOI (Guo and Lin 2024), which leverages the cross-attention mechanism in U-Net to localize editing regions. Nevertheless, due

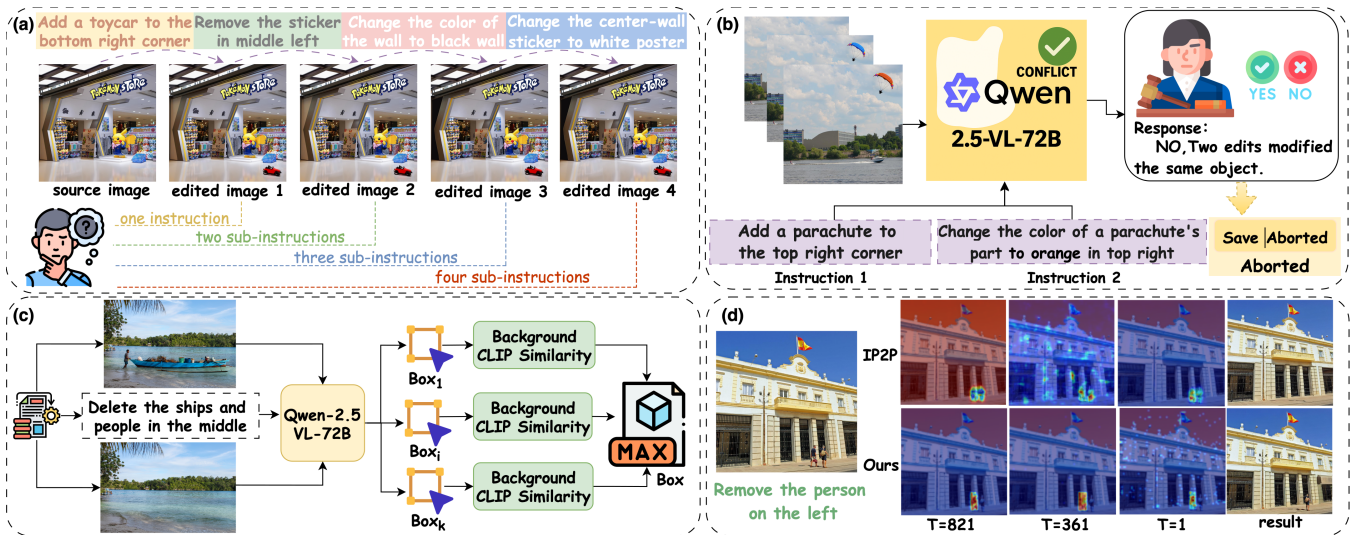


Figure 3: (a) shows how multi-turn editing sequences are expanded into multiple complex instruction editing instances. (b) shows the use of Qwen2.5-VL-72B to detect instruction conflicts. (c) illustrates the generation and selection of bounding boxes. (d) compares attention maps and editing results for IP2P and our method during the denoising process.

to the limited comprehension capability of CLIP, FOI often fails to generate accurate masks for the edited areas. Our key insight is that general diffusion models struggle to execute complex instructions, making it essential to decompose them into sub-instructions with corresponding spatial regions. Furthermore, we incorporate spatial information for each sub-instruction to improve instruction-following performance and integrate fine-grained visual features to maintain background consistency.

2.2 Datasets for Instruction Editing

In Tab. 1, we compare several popular instruction-based image editing datasets (Brooks, Holynski, and Efros 2023; Zhang et al. 2023a; Yu et al. 2025; Wei et al. 2024). For instance, IP2P uses P2P (Hertz et al. 2022) to generate suboptimal image pairs, which limits its applicability in real-world editing scenarios. MagicBrush (Zhang et al. 2023a) improves real-world applicability through 10,000 high-quality human annotations. Both OmniEdit (Wei et al. 2024) and AnyEdit (Yu et al. 2025) leverage expert models to generate diverse image editing tasks. UltraEdit (Zhao et al. 2024) scales up data generation by incorporating both free-form and region-based samples. SEED-Data-Edit (Ge et al. 2024) provides single-turn and multi-turn editing data, covering a broader range of real-world images. However, these datasets still suffer from limited data quality and overly simple instructions. In contrast, our MCIE dataset is specifically tailored for complex instruction-based image editing.

3 Method

We present a high-fidelity dataset and a state-of-the-art model for complex instruction-based image editing. Section 3.1 details the data creation pipeline, including data augmentation, data filtering, spatial information generation, and post-processing. Section 3.2 introduces MCIE-E1,

an advanced image editing model trained on a subset of OmniEdit-GoT (Fang et al. 2025) and the MCIE dataset.

3.1 MCIE Dataset

Existing datasets (Zhang et al. 2023a; Yu et al. 2025) primarily support simple instructions, limiting their effectiveness for complex instruction-based image editing. Moreover, complex instructions often require auxiliary inputs such as bounding boxes or masks to provide spatial guidance. To address these limitations, we construct a large-scale, high-quality dataset by leveraging the understanding and grounding capabilities of advanced MLLMs. Our MCIE dataset includes meticulously designed semantic-spatial information for complex instruction-based image editing, with each sample containing an instruction, bounding boxes, and corresponding images. Finally, to ensure data quality, experts evaluate the dataset based on three criteria: *instruction consistency*, *image quality*, and *editing scenario complexity*.

Data Preparation and Augmentation. As the foundation of our study, we adopt 20K multi-turn editing instances from the SEED-Data-Edit (Ge et al. 2024) dataset as our primary corpus, as it offers more diverse instructions and more accurate editing results than other datasets. However, the dataset contains only 21K multi-round samples, limiting the generalization capability of models. To address this, we propose an efficient data augmentation approach. As illustrated in Fig. 3(a), a four-turn editing instance can be expanded into **three** two-instruction, **two** three-instruction, and **one** four-instruction editing samples. Nevertheless, this approach inherits the limitations of multi-turn editing, which may involve modifying the same object multiple times.

Data Filtering. As illustrated in Fig. 3(b), we provide the original image, the intermediate results after each edit, and the corresponding instruction for each turn to facilitate conflict detection. The MLLM must comprehend the

Dataset	Size	Res. (px) \uparrow	DeQA-Score \uparrow	Q-Insight \uparrow	Q-align \uparrow	CE	SI	HF
InstructPix2Pix (Brooks, Holynski, and Efros 2023)	313K	512	3.38	3.61	3.67	\times	\times	\times
MagicBrush (Zhang et al. 2023a)	10K	500	3.31	3.32	3.35	\times	\checkmark	\checkmark
OmniEdit (Wei et al. 2024)	1.2M	≥ 768	3.44	3.69	3.83	\times	\times	\times
AnyEdit (Yu et al. 2025)	2.5M	512	3.65	3.70	3.92	\times	\times	\times
UltraEdit (Zhao et al. 2024)	4M	512	3.65	3.67	3.75	\times	\checkmark	\times
SEED-Data-Edit (Ge et al. 2024)	3.7M	≥ 768	3.59	3.73	3.80	\times	\times	\times
MCIE dataset (ours)	90K	≥ 1024	3.90	3.95	4.27	\checkmark	\checkmark	\checkmark

Table 1: Comparison of existing datasets and the proposed MCIE dataset. **CE** denotes Complex Editing, **SI** denotes Spatial Information, and **HF** denotes Human Filter. **DeQA-Score** (You et al. 2025), **Q-Insight** (Li et al. 2025), and **Q-align** (Wu et al. 2023) are used to assess image quality.

editing instructions for each turn along with their associated regions, enabling it to reason about whether the same object undergoes multiple modifications. Leveraging the strong multi-image understanding capabilities of Qwen2.5-VL-72B (Bai et al. 2025), we employ carefully designed in-context prompts to evaluate whether converting multi-turn edits into complex instruction edits introduces any conflicts.

Post Processing. Since coarse filtering has already been performed during data preparation using MLLMs, we further employ 20 human experts to conduct a fine-grained evaluation across three key aspects: instruction consistency, image quality, and editing scenario complexity. Instruction consistency is assessed with two options, i.e., “Yes” or “No”, ensuring that the editing instructions are conflict-free. For image quality and editing scenario complexity, we adopt a 1–5 scoring system with detailed criteria, retaining only samples with scores above 3.

Spatial Information Generation. Spatial information is critical in image editing, significantly affecting instruction-following performance. We decompose multi-turn editing into single-turn instructions. As depicted in Fig. 3(c), leveraging the strong grounding capabilities of Qwen2.5-VL-72B, we generate k candidate bounding boxes for each instruction and select the one that best aligns with the edited region. Because the bounding box primarily covers the areas to be edited, the remaining regions of the source and target images should remain consistent. Accordingly, we compute the CLIP score between the source and target images outside the bounding box to select the most accurate box.

$$i^* = \arg \max_{i \in \{1, 2, \dots, k\}} \text{CLIP}(I_{\text{src}} \cdot (1 - M_i), I_{\text{tar}} \cdot (1 - M_i)) \quad (1)$$

where I_{src} and I_{tar} denote the source and target images, respectively. M_i denote the mask associated with the i -th candidate bounding box.

The creation of the entire dataset requires substantial computational resources, specifically utilizing 32 H20 GPUs over a week. This robust computational setup ensures the precision and reliability of instructions and bounding boxes, thereby providing a solid and essential foundation for training models on complex instruction image editing.

3.2 Complex Instruction Image Editing Design

Given a source image and a complex instruction $I = \{I_1, I_2, \dots, I_m\}$, where each sub-instruction I_i is associated




Figure 4: Comparison of editing results for sub-instruction bounding boxes with and without the Fourier transform.

with operation type $op_i \in \{\text{ADD}, \text{REMOVE}, \text{CHANGE}\}$. As shown in Fig. 5, MCIE-E1 integrates an MLLM to decompose complex instructions and a diffusion model comprising two key components: spatial-aware cross-attention for enhancing instruction following and background-consistent cross-attention for maintaining background consistency.

Instruction Decomposition. Recently, MLLMs have shown remarkable capabilities in understanding, reasoning, and few-shot generalization. We adopt an in-context learning strategy to enable accurate instruction decomposition and region generation. Since diffusion models inherently struggle with counting, we further decompose quantity-related instructions to enhance the model’s instruction-following capability. The decomposition process is as follows:

$$D : I \rightarrow \{(I_1, B_1), (I_2, B_2), \dots, (I_m, B_m)\} \quad (2)$$

where B_i denotes the bounding box corresponding to the editing region specified by sub-instruction I_i .

The prompt consists of three integrated components: a carefully designed guideline specifying the assistant’s role and decomposition format, a representative example illustrating the decomposition of a complex instruction into sub-instructions, and a test instance containing a complex instruction and a source image.

Spatial-Aware Cross-Attention. Previous studies (Brooks, Holynski, and Efros 2023; Huang et al. 2024; Fang et al. 2025) typically encode entire instructions using a CLIP (Radford et al. 2021) model, which causes unintended interference among sub-instructions. To mitigate this, we introduce an instruction-wise encoding strategy, where each sub-instruction is encoded independently and in parallel. The instruction-wise encoding strategy is defined as:

$$\text{Ins}_i = \text{CLIP}(I_i) \quad (3)$$

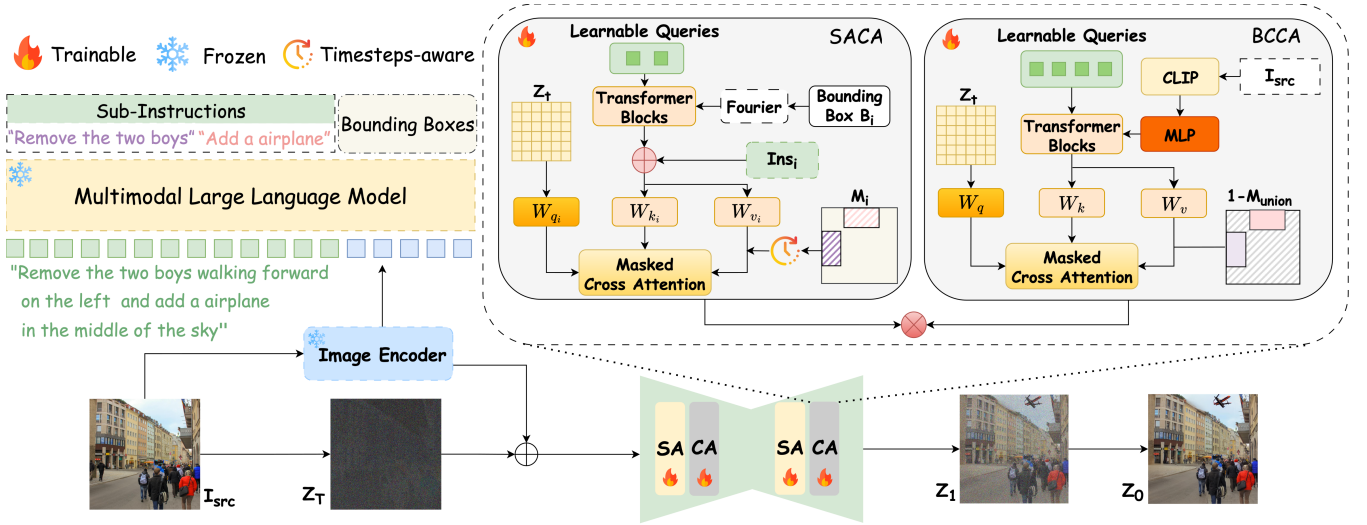


Figure 5: The overall framework of MCIE-E1. It employs an MLLM for instruction decomposition and guiding the diffusion model through two key modules: SACA for enhancing instruction following and BCCA for preserving non-edited regions.

$$\mathbf{T} = [\text{Ins}_1, \dots, \text{Ins}_m] \quad (4)$$

where $[\cdot]$ denotes concatenation.

For complex image editing, spatial information is crucial for guiding the model to execute modifications accurately. As shown in Fig. 3(d), IP2P (Brooks, Holynski, and Efros 2023) often produces imprecise edits. In contrast, our model incorporates rich spatial cues, enabling more controllable and precise editing. We address this by employing masked cross-attention to associate each sub-instruction with its corresponding region explicitly. However, this approach may still encounter failures when processing semantically similar instructions. As illustrated in Fig. 4, when the instruction is to add three apples on the table, it is decomposed into three semantically similar sub-instructions. Due to their similarity, the sub-instructions focus on the same region, causing the model to execute only part of the intended edits. To this end, we propose a Spatial-Aware Cross-Attention (SACA) module that encodes bounding boxes via Fourier mapping and extracts features using learnable queries, enabling the model to learn distinct spatial embeddings for each sub-instruction. This process is formalized as follows:

$$\mathcal{F}(\mathbf{B}_i) = [\sin(2\pi\mathbf{B}_i\mathbf{f}_k), \cos(2\pi\mathbf{B}_i\mathbf{f}_k)] \quad (5)$$

$$\mathbf{C}_i = [\mathbf{T}_i, Q_\beta(\mathcal{F}(\mathbf{B}_i))] \quad (6)$$

$$\text{FG}_t = \text{Softmax}\left(\frac{QK_t^\top \odot M_i}{\sqrt{d}}\right)V_i \quad (7)$$

where \mathbf{f}_k are frequency bands, Q_β denotes L_1 Transformer blocks, \odot represents element-wise multiplication, FG_t are foreground features, $Q = z_t W_q$, $K_i = C_i W_{k_i}$, $V_i = C_i W_{v_i}$, z_t is the noisy latent, and d is the feature dimension.

Prior studies (Hu et al. 2024; Zhang et al. 2023b) have shown that diffusion models in text-to-image generation primarily attend to spatial layouts and global structures in the early stages of denoising, gradually shifting focus to fine-grained features. We observe that instruction-based image

editing exhibits a similar progression, focusing on different types of information throughout the denoising process. Inspired by this observation, we adopt a timestep-aware control schedule to modulate the strength of spatial guidance. Specifically, we define a timestep-aware mask:

$$M'_i = M_i * \exp\left(-\frac{t}{T}\right) \quad (8)$$

where t denotes the current denoising timestep, and T is the total number of denoising steps.

Background-Consistent Cross-Attention. Complex editing instructions often involve multiple region-specific operations, increasing the risk of background inconsistency and potentially compromising the visual coherence of the edited image. To address this, we introduce the Background-Consistent Cross-Attention (BCCA) module. As shown in Fig. 5, we first extract pixel-level visual features from the source image using the CLIP (Radford et al. 2021) model. These features are then refined through an MLP followed by L_2 Transformer blocks, with a set of learnable queries capturing enhanced representations. Finally, the learnable queries interact via masked cross-attention, ensuring high background consistency throughout the denoising process.

$$f = Q_\beta(\text{MLP}(\text{CLIP}(I_{\text{src}}))) \quad (9)$$

$$\text{BG}_t = \text{Softmax}\left(\frac{QK_t^\top \odot (1 - M_{\text{union}})}{\sqrt{d}}\right)V_t \quad (10)$$

where $K_t = W_k f$, $V_t = W_v f$, $M_{\text{union}} = \bigcup_{i=1}^m M_i$, BG_t are background features.

Following IP-Adapter (Ye et al. 2023), we decouple the cross-attention layers into two key modules: SACA, which enhances interactions between foreground features and complex instructions, and BCCA, which preserves the consistency of unedited background regions. Specifically, given

Method	Automatic Metrics						User Study	
	CLIP-I \uparrow	DINO-I \uparrow	L1 \downarrow	L2 \downarrow	IC \uparrow	BC \uparrow	Ins Following \uparrow	BG Consist. \uparrow
IP2P	0.8784	0.9691	0.0710	0.1032	0.4922	0.6270	1.00%	2.00%
MagicBrush	0.8870	0.9813	0.0476	0.0992	0.5246	0.7608	0.67%	1.33%
UltraEdit	0.7820	0.9580	0.0779	0.1209	0.5495	0.6485	6.67%	6.00%
AnyEdit	0.8847	0.9662	0.0701	0.1168	0.5335	0.7659	0.67%	2.33%
InsDiff	0.8815	0.9634	0.0810	0.1283	0.4812	0.7465	1.00%	1.67%
MGIE	0.7679	0.9185	0.1407	0.2005	0.3988	0.6205	0.33%	0.33%
SmartEdit	0.8836	0.9752	0.0629	0.1042	0.5416	0.7419	3.00%	4.33%
GoT	0.8856	0.9799	0.0549	0.1078	0.6491	0.7389	6.67%	9.67%
MCIE-E1 \dagger	0.9200	0.9863	0.0456	0.0864	0.6597	0.7784	–	–
MCIE-E1	0.8899	0.9833	0.0489	0.1029	0.8046	0.7831	80.00%	72.33%

Table 2: Quantitative results on CIE-Bench. \uparrow indicates higher is better, and \downarrow indicates lower is better. The best result in each column is highlighted in **red**, and the second best in **blue**. **MCIE-E1 \dagger** denotes a variant trained only in the first stage.

the foreground features FG_t and background features BG_t , they are fused as:

$$z'_t = \lambda \cdot FG_t + (1 - \lambda) \cdot BG_t \quad (11)$$

where $\lambda \in [0, 1]$ is a weighting factor.

Training Procedure. We adopt a two-phase training strategy. We first train the model for 20,000 steps on a subset of the OmniEdit-GoT (Fang et al. 2025) dataset. Subsequently, we fine-tune the model for an additional 10,000 steps on the MCIE dataset to better adapt to complex and fine-grained editing instructions. The optimization process for ϵ_θ can be expressed in two stages as follows:

$$\mathcal{L} = \mathbb{E}_{\mathcal{E}(y), \mathcal{E}(x), c, \epsilon, t} \|\epsilon - \epsilon_\theta(\text{Cat}(z_t, \mathcal{E}(x)), B, y, t)\|_2^2 \quad (12)$$

where x is the source image, y is the target image, $\mathcal{E}(x)$ is the encoded image latent, z_t is the noisy latent, ϵ_θ is the denoising network and Cat denotes concatenation along the channel dimension.

4 Experiment

4.1 CIE-Bench

We introduce CIE-Bench, a new benchmark to evaluate a model’s ability to perform complex instruction-based image editing. Specifically, we select 400 high-quality images from the SEED-data-Edit dataset, covering a variety of realistic scenarios. Additionally, we employ GPT-4o (Hurst et al. 2024) to generate complex instructions, each containing 1 to 4 sub-instructions involving any combination of three operation types: *Add*, *Change*, and *Remove*. To ensure accurate evaluation, human experts manually annotate the corresponding editing regions based on the instructions, and each sample undergoes manual filtering.

Evaluation Metrics. To quantitatively assess editing performance, we use CLIP (Radford et al. 2021), DINOv2 (Oquab et al. 2023) image similarity, L1 distance, and L2 distance.

Additionally, we introduce two new evaluation metrics:

- **Instruction Compliance (IC):** GPT-4o is employed to assess whether the specified modifications align with the given instructions, rated on a scale of 1–10.
- **Background Consistency (BC):** GPT-4o scores the consistency of the unedited regions on a scale of 1–5.

The instruction compliance and background consistency scores are then normalized to the range $[0, 1]$.

4.2 Experimental Settings

Baseline Models. We evaluate our MCIE-E1 model against state-of-the-art image editing methods, including IP2P (Brooks, Holynski, and Efros 2023), MagicBrush (Brooks, Holynski, and Efros 2023), InstructDiffusion (InsDiff) (Geng et al. 2024), AnyEdit (Yu et al. 2025), UltraEdit (Zhao et al. 2024), MGIE (Fu et al. 2023), SmartEdit (Huang et al. 2024) and GoT (Fang et al. 2025).

Implementation Details. We implement MCIE-E1 on Stable Diffusion 1.5 (Rombach et al. 2022), with parameters initialized from IP2P (Brooks, Holynski, and Efros 2023). Low-Rank Adaptation (LoRA) (Hu et al. 2022) is employed to efficiently update only the self-attention layers, while the SACA and BCCA modules are fully optimized within the diffusion model. The SACA module consists of two Transformer blocks and two learnable queries, whereas the BCCA module comprises four Transformer blocks and 16 learnable queries. All experiments use the Euler Ancestral sampler with 20 denoising steps.

4.3 Main Results

Quantitative Evaluation. Tab. 2 presents the quantitative comparison of our method with baseline methods on CIE-Bench. Our method consistently outperforms all baselines across all automatic evaluation metrics, demonstrating its effectiveness in complex instruction-based image editing. Tab. 3 shows the quantitative comparison on MagicBrush for simple instruction-based image editing. Similarly, our method achieves superior performance across all metrics, indicating its advantage over baselines in simpler editing tasks as well.

Qualitative Evaluation. Fig. 6 presents qualitative comparisons on CIE-Bench. As observed, baseline methods often struggle to follow specific editing instructions. For instance, IP2P (Brooks, Holynski, and Efros 2023) and MagicBrush (Zhang et al. 2023a) fail to add a black car on the road, while SmartEdit (Huang et al. 2024) cannot change the weather to sunny. Moreover, they struggle to maintain background consistency. For example, GoT (Fang et al. 2025)

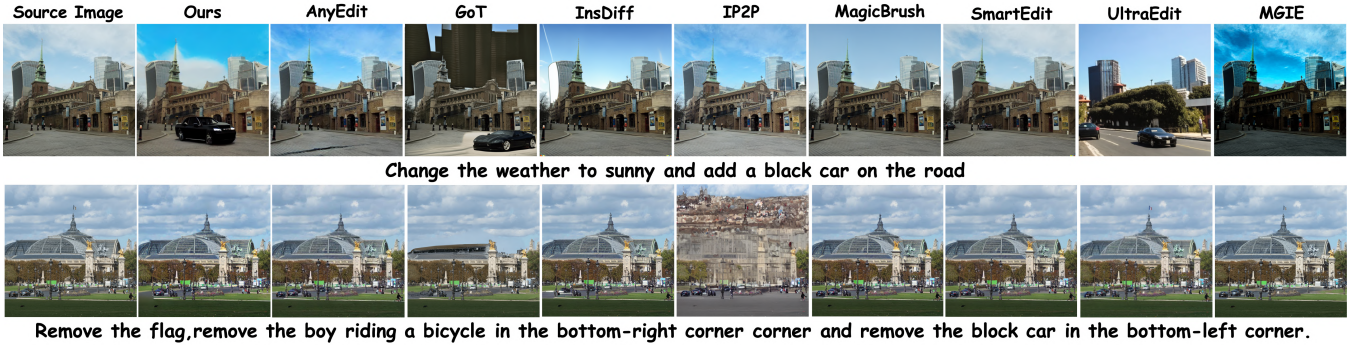


Figure 6: Qualitative results on CIE-bench. Our MCIE-E1 model demonstrates superior performance in terms of instruction compliance and background consistency.

modifies buildings across the entire image when instructed to make the weather sunny, leading to severe distortions. In contrast, our method performs precise edits while preserving background consistency.

Method	L1↓	L2↑	CLIP-I↑	DINO-I↑
IP2P	0.1059	0.0328	0.8674	0.7617
MagicBrush	0.0765	0.0270	0.8951	0.8360
AnyEdit	0.0968	0.0383	0.8689	0.7532
InsDiff	0.0905	0.0363	0.8934	0.8045
MGIE	0.0847	0.0324	0.9033	0.8292
SmartEdit	0.0672	0.0284	0.9017	0.8360
GoT	0.0645	0.0237	0.9047	0.8400
MCIE-E1†	0.0641	0.0228	0.9143	0.8497
MCIE-E1	0.0669	0.0246	0.9122	0.8466

Table 3: Quantitative results on MagicBrush.

User Study. We conducted a user study on complex instruction-based image editing with 50 participants, each evaluating 100 images. The study assessed user preferences across two key aspects: *instruction following* and *background consistency*. Participants were asked to select the method that best adhered to the given instructions and best preserved non-edited regions. As shown in Tab. 2, our method was consistently preferred over baseline methods, with 80% of participants indicating that its instruction-following capability surpassed that of the baselines.

4.4 Ablation Study

We conduct an ablation study to evaluate the contribution of each component in the proposed method on CIE-bench, including the SACA and BCCA modules. In addition, we examine the impact of different MLLMs.

Impact of Spatial-Aware Cross-Attention. As shown in Tab. 4, removing the SACA module markedly degrades the model’s ability to follow sub-instructions, whereas its inclusion substantially enhances instruction-following performance.

Impact of Background-Consistent Cross-Attention. As shown in Tab. 4, integrating the BCCA module yields no-

table gains in CLIP-I, DINO-I, and BC metrics, highlighting the effectiveness of its pixel-level visual encoding in maintaining background consistency.

SACA	BCCA	CLIP-I↑	DINO-I↑	L1↓	L2↓	IC↑	BC↑
-	-	0.8784	0.9691	0.0710	0.1032	0.4922	0.6270
✓	-	0.8566	0.9690	0.0886	0.1320	0.7932	0.6851
✓	✓	0.8899	0.9833	0.0489	0.1029	0.8046	0.7831

Table 4: Quantitative ablation results of SACA and BCCA modules.

Impact of MLLMs. We further perform ablation studies on the MLLMs used to decompose complex instructions into simpler sub-instructions and corresponding regions, covering both open-source and proprietary models. As shown in Tab. 5, proprietary models demonstrate superior decomposition performance compared to open-source alternatives.

Method	CLIP-I↑	DINO-I↑	L1↓	L2↓	IC↑	BC↑
Qwen2.5VL 3B	0.8763	0.9673	0.0715	0.1270	0.7403	0.7139
Qwen2.5VL 7B	0.8789	0.9753	0.0643	0.1111	0.7505	0.7377
Qwen2.5VL 32B	0.8801	0.9768	0.0587	0.1080	0.7773	0.7539
GPT4-o	0.8812	0.9786	0.0565	0.1050	0.7873	0.7705

Table 5: Quantitative ablation results of different MLLMs.

5 Conclusion

In this paper, we present MCIE-E1, a novel model for complex instruction-based image editing that leverages MLLMs to decompose intricate instructions. To enhance instruction adherence and preserve background consistency, we introduce two key modules: spatial-aware cross-attention and background-consistent cross-attention. In addition, we construct MCIE, the first image editing dataset centered on complex instructions, featuring precise spatial annotations and rigorous human filtering. Finally, we propose CIE-Bench and two new evaluation metrics for assessing complex instruction-based image editing. Experimental results on CIE-Bench demonstrate that MCIE-E1 consistently outperforms prior state-of-the-art methods across both quantitative and qualitative evaluations.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grants 62471168, 62422204, and 61802100, 62372147, U21B2040 and in part by the Zhejiang Provincial Natural Science Foundation of China under Grant LDT23F02025F02 and LY21F020019.

References

- Bai, J.; Chow, W.; Yang, L.; Li, X.; Li, J.; Zhang, H.; and Yan, S. 2024. HumanEdit: A High-Quality Human-Rewarded Dataset for Instruction-based Image Editing. *arXiv preprint arXiv:2412.04280*.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Basu, S.; Saberi, M.; Bhardwaj, S.; Chegini, A. M.; Masiceti, D.; Sanjabi, M.; Hu, S. X.; and Feizi, S. 2023. Editval: Benchmarking diffusion based text-guided image editing methods. *arXiv preprint arXiv:2310.02426*.
- Brooks, T.; Holynski, A.; and Efros, A. A. 2023. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 18392–18402.
- Fang, R.; Duan, C.; Wang, K.; Huang, L.; Li, H.; Yan, S.; Tian, H.; Zeng, X.; Zhao, R.; Dai, J.; et al. 2025. Got: Unleashing reasoning capability of multimodal large language model for visual generation and editing. *arXiv preprint arXiv:2503.10639*.
- Fu, T.-J.; Hu, W.; Du, X.; Wang, W. Y.; Yang, Y.; and Gan, Z. 2023. Guiding instruction-based image editing via multimodal large language models. *arXiv preprint arXiv:2309.17102*.
- Ge, Y.; Zhao, S.; Li, C.; Ge, Y.; and Shan, Y. 2024. Seed-data-edit technical report: A hybrid dataset for instructional image editing. *arXiv preprint arXiv:2405.04007*.
- Geng, Z.; Yang, B.; Hang, T.; Li, C.; Gu, S.; Zhang, T.; Bao, J.; Zhang, Z.; Li, H.; Hu, H.; et al. 2024. Instructdiffusion: A generalist modeling interface for vision tasks. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, 12709–12720.
- Guo, Q.; and Lin, T. 2024. Focus on your instruction: Fine-grained and multi-instruction image editing by attention modulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6986–6996.
- Hertz, A.; Mokady, R.; Tenenbaum, J.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2022. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*.
- 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*.
- Huang, Y.; Xie, L.; Wang, X.; Yuan, Z.; Cun, X.; Ge, Y.; Zhou, J.; Dong, C.; Huang, R.; Zhang, R.; et al. 2024. Smartedit: Exploring complex instruction-based image editing with multimodal large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8362–8371.
- Hurst, A.; Lerer, A.; Goucher, A. P.; Perelman, A.; Ramesh, A.; Clark, A.; Ostrow, A.; Welihinda, A.; Hayes, A.; Radford, A.; et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Kawar, B.; Zada, S.; Lang, O.; Tov, O.; Chang, H.; Dekel, T.; Mosseri, I.; and Irani, M. 2023. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6007–6017.
- Li, W.; Zhang, X.; Zhao, S.; Zhang, Y.; Li, J.; Zhang, L.; and Zhang, J. 2025. Q-insight: Understanding image quality via visual reinforcement learning. *arXiv preprint arXiv:2503.22679*.
- Ma, Y.; Ji, J.; Ye, K.; Lin, W.; Wang, Z.; Zheng, Y.; Zhou, Q.; Sun, X.; and Ji, R. 2024. I2EBench: A Comprehensive Benchmark for Instruction-based Image Editing. *arXiv preprint arXiv:2408.14180*.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- 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.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140): 1–67.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2): 3.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Sheynin, S.; Polyak, A.; Singer, U.; Kirstain, Y.; Zohar, A.; Ashual, O.; Parikh, D.; and Taigman, Y. 2024. Emu edit: Precise image editing via recognition and generation tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8871–8879.
- Wang, S.; Saharia, C.; Montgomery, C.; Pont-Tuset, J.; Noy, S.; Pellegrini, S.; Onoe, Y.; Laszlo, S.; Fleet, D. J.; Soricut, R.; et al. 2023. Imagen editor and editbench: Advancing and evaluating text-guided image inpainting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 18359–18369.

Wei, C.; Xiong, Z.; Ren, W.; Du, X.; Zhang, G.; and Chen, W. 2024. Omniedit: Building image editing generalist models through specialist supervision. In *The Thirteenth International Conference on Learning Representations*.

Wu, H.; Zhang, Z.; Zhang, W.; Chen, C.; Liao, L.; Li, C.; Gao, Y.; Wang, A.; Zhang, E.; Sun, W.; et al. 2023. Q-align: Teaching Imms for visual scoring via discrete text-defined levels. *arXiv preprint arXiv:2312.17090*.

Ye, H.; Zhang, J.; Liu, S.; Han, X.; and Yang, W. 2023. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*.

You, Z.; Cai, X.; Gu, J.; Xue, T.; and Dong, C. 2025. Teaching Large Language Models to Regress Accurate Image Quality Scores using Score Distribution. *arXiv preprint arXiv:2501.11561*.

Yu, Q.; Chow, W.; Yue, Z.; Pan, K.; Wu, Y.; Wan, X.; Li, J.; Tang, S.; Zhang, H.; and Zhuang, Y. 2025. Anyedit: Mastering unified high-quality image editing for any idea. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 26125–26135.

Zhang, K.; Mo, L.; Chen, W.; Sun, H.; and Su, Y. 2023a. Magicbrush: A manually annotated dataset for instruction-guided image editing. *Advances in Neural Information Processing Systems*, 36: 31428–31449.

Zhang, Y.; Dong, W.; Tang, F.; Huang, N.; Huang, H.; Ma, C.; Lee, T.-Y.; Deussen, O.; and Xu, C. 2023b. Prospect: Prompt spectrum for attribute-aware personalization of diffusion models. *ACM Transactions on Graphics (TOG)*, 42(6): 1–14.

Zhao, H.; Ma, X. S.; Chen, L.; Si, S.; Wu, R.; An, K.; Yu, P.; Zhang, M.; Li, Q.; and Chang, B. 2024. Ultraedit: Instruction-based fine-grained image editing at scale. *Advances in Neural Information Processing Systems*, 37: 3058–3093.