

🤖 I2E: From Image Pixels to Actionable Interactive Environments for Text-Guided Image Editing

Jinghan Yu^{◆1,2}, Junhao Xiao^{◆1,2,3}, Chenyu Zhu^{1,2}, Jiaming Li^{1,2}, Jia Li¹, Hanming Deng¹, Xirui Wang¹, Guoli Jia⁴, Jianjun Li¹, Xiang Bai¹, Bowen Zhou^{4,5}, Zhiyuan Ma^{◆1}

¹Huazhong University of Science and Technology, ²Kuaishou Technology, ³Central China Normal University, ⁴Tsinghua University, ⁵Shanghai AI Laboratory

jinghanyu0917@gmail.com, xiaojunhao066@gmail.com, mzyth@hust.edu.cn

Abstract

Existing text-guided image editing methods primarily rely on end-to-end pixel-level inpainting paradigm. Despite its success in simple scenarios, this paradigm still significantly struggles with compositional editing tasks that require precise local control and complex multi-object spatial reasoning. This paradigm is severely limited by 1) *the implicit coupling of planning and execution*, 2) *the lack of object-level control granularity*, and 3) *the reliance on unstructured, pixel-centric modeling*. To address these limitations, we propose I2E, a novel “Decompose-then-Action” paradigm that revisits image editing as an actionable interaction process within a structured environment. I2E utilizes a Decomposer to transform unstructured images into discrete, manipulable object layers and then introduces a physics-aware Vision-Language-Action Agent to parse complex instructions into a series of atomic actions via Chain-of-Thought reasoning. Further, we also construct I2E-BENCH, a benchmark designed for multi-instance spatial reasoning and high-precision editing. Experimental results on I2E-BENCH and multiple public benchmarks demonstrate that I2E significantly outperforms state-of-the-art methods in handling complex compositional instructions, maintaining physical plausibility, and ensuring multi-turn editing stability. Code and dataset: [project page](#).

1 Introduction

When performing image editing, humans rarely reason in terms of direct pixel manipulations. Consider a typical editing request: “Move the books that are pressed by the water cup on the desktop to the right side of the cup”. For humans, such an instruction implicitly involves a sequence of intermediate reasoning steps, including object identifi-

◆: Equal contribution as co-first authors. Work completed during joint internship at HUST and Kuaishou.

◆: Corresponding author: Zhiyuan Ma.

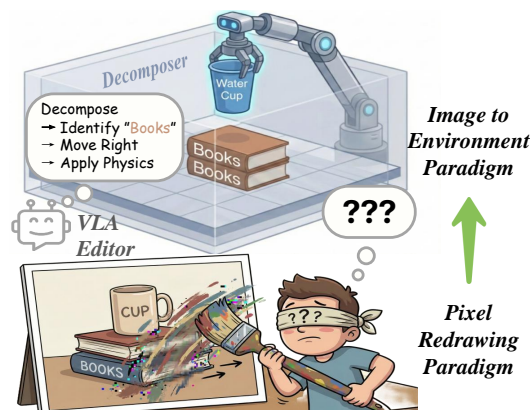


Figure 1: **Paradigm Comparison.** Unlike the *Pixel Redrawing Paradigm* that directly manipulates pixels, I2E transforms images into a structured environment, enabling the VLA Editor to perform spatial and physical reasoning for precise, physically plausible edits.

cation, spatial relationship understanding, physical constraint awareness, and ordered execution.

However, existing text-guided image editing models generally exploit an *End-to-End* paradigm but lack such intermediate representations and *Reasoning-then-Action* process. Specifically, they typically attempt to directly map instructions to final results through one or multiple rounds of pixel-level redrawing. While this end-to-end pixel redrawing paradigm is generally effective in simple editing scenarios, it exposes three major structural limitations when applied to compositional editing tasks that require precise local control and complex multi-object spatial reasoning:

(i) **Tight coupling between semantic reasoning and execution.** Models are required to perform instruction understanding and pixel synthesis within a single generation process, making it difficult to form stable intermediate decision structures and leading to significantly degraded instruction-following performance in complex scenarios.

(ii) **Lack of object-level representations and boundaries.** When editing is performed directly in pixel space, modifications cannot be strictly con-

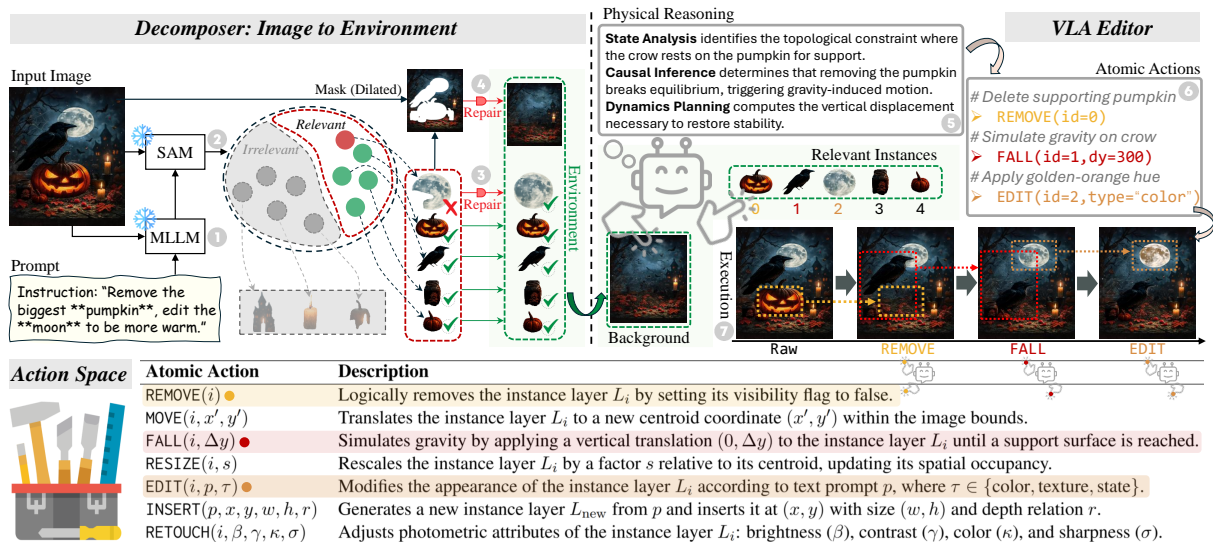


Figure 2: Overview of the I2E. The *Decomposer* transforms unstructured images into a structured environment of actionable physical layers. The physics-aware *VLA Editor* then uses chain-of-thought reasoning to translate instructions into executable atomic actions (see bottom) and executes them sequentially.

fined to target instances and often propagate as global perturbations to non-target regions.

(iii) **Pixel-centric and unstructured modeling.** By treating images as unstructured two-dimensional pixel collections, models struggle to explicitly represent depth relations, support relations, and scale constraints, which frequently results in physically implausible editing outcomes, such as “floating objects” (As described in Figure 4).

These issues are further amplified in multi-round incremental interactive editing. Since each editing step typically redraws the entire image based on the previous output, unstructured pixel-level updates cause errors to accumulate across iterations, leading to severe “feature drift” and making it difficult to achieve fine-grained, controllable continuous editing (see Figure 3). Moreover, repeatedly invoking computationally expensive generation processes substantially degrades interaction efficiency.

To address these challenges, as illustrated in Figure 1, we propose a new image editing paradigm, I2E (Image-to-Environment), which **reformulates image editing as an interactive process within an actionable structured environment**. From this perspective, an image is no longer treated as an indivisible pixel array, but as a composition of entities and background with explicit spatial relationships. Building on this, I2E mainly operates in two stages (Figure 2): 1) *Image-to-Environment Transition*. A *Decomposer* module transforms unstructured pixel representations into environment representations

with explicit spatial structure. This module explicitly recovers the complete appearance of each instance object (e.g., the obscured moon in Figure 2) and their relative physical relationships, encapsulating them as independent and manipulable object-level physical layers, which together with the background layer form an interactive actionable environment. 2) *VLA-based Environment Editing*. On top of this structured environment, we introduce a physics-aware *Vision–Language–Action (VLA) Editor* as the core decision-making component. Rather than directly predicting pixel-level changes, the agent progressively decomposes complex natural language instructions into a sequence of precise atomic actions that satisfy physical constraints through chain-of-thought reasoning.

This design brings multiple advantages. By decomposing high-level editing intents into executable atomic steps, it substantially improves instruction-following for complex instructions. Action execution grounded in object-level environment states further ensures that edits are strictly localized to target instances, effectively eliminating interference in irrelevant regions. Beyond these benefits, the decoupled atomic actions and the explicit grounding of target entities enable efficient multi-round incremental editing. This transforms the generative editing paradigm from a “one-shot global repainting” process into a “progressive refinement workflow”. When responding to user feedback or during self-correction, the system does not require a reset of the scene, but instead updates

the state by appending corrective actions.

Moreover, we observe that existing benchmarks primarily focus on style transfer or simple single-step instructions, lacking comprehensive evaluation of complex spatial reasoning, multi-instance interaction, and physical constraint consistency. To fill this gap and validate the effectiveness of the proposed approach, we introduce I2E-BENCH, a benchmark designed for multi-instance spatial reasoning and high-precision image editing. Extensive experiments on I2E-BENCH, as well as public benchmarks such as MagicBrush and EmuEdit, demonstrate that I2E significantly outperforms state-of-the-art methods in handling compositional instructions, maintaining physical consistency, and ensuring stability in multi-round interactions.

2 Related Works

2.1 Text-Guided Image Editing and Agentic Approaches

Early text-guided image editing methods (Brooks et al., 2023; Hertz et al., 2022; Meng et al., 2021) primarily rely on end-to-end pixel redrawing, directly mapping textual instructions to global image synthesis. While effective for simple edits, tightly coupling instruction understanding with pixel generation limits their ability to handle compositional commands that require precise local control and multi-object spatial reasoning.

Recent approaches improve semantic interpretation by incorporating multimodal large language models (MLLMs) (Fu et al., 2024; Liu et al., 2025a; Yu et al., 2025) or unifying reasoning and generation within large transformer-based models (Xiao et al., 2025; Betker et al., 2023; Feng et al., 2025; Li et al., 2025). Despite stronger instruction comprehension, edit execution remains bound to global resampling, making these methods prone to unintended changes in non-target regions and shows attribute leakage (Mun et al., 2025).

Agentic editing frameworks (Huang et al., 2024; Hu et al., 2025; Xu et al., 2025) further decompose instructions into sub-tasks, yet most still realize each step through independent image generation, leading to accumulated deviations across interactions. In contrast, I2E performs editing as object-level actions over a structured environment, enabling incremental updates without re-synthesizing the entire image.

2.2 Structured Scene Representation and Amodal Decomposition

Precise local editing requires structured scene representations beyond flat pixel grids. Instance segmentation models (Kirillov et al., 2023; Ravi et al., 2024) provide object localization but operate at the modal level, failing to recover occluded content and often introducing missing regions when objects are manipulated (Yu et al., 2019).

Layered image generation methods (Zhang and Agrawala, 2024) partially improve locality but typically focus on synthesis with fixed layouts rather than interactive editing. Amodal completion approaches (Ozguroglu et al., 2024; Liu et al., 2025b) reconstruct occluded appearances, yet are usually designed as standalone restoration modules and lack a unified representation for downstream manipulation (Ao et al., 2025).

Our work integrates instance-level amodal decomposition with explicit depth-aware ordering to form complete, spatially organized object layers. This representation enables object-specific editing while keeping non-target regions unchanged.

2.3 Physical Reasoning and Vision–Language–Action Models

Maintaining physical plausibility remains challenging in generative image editing. Existing methods often rely on static geometric cues, such as depth maps or edge constraints (Zhang et al., 2023b; Lee and Park, 2022), which cannot capture changes in physical relationships induced by editing actions, leading to implausible outcomes such as unsupported objects (Pan et al., 2023).

Vision–Language–Action (VLA) models (Brohan et al., 2023; Driess et al., 2023; Kim et al., 2024; Chen et al., 2025) demonstrate effective physical reasoning by grounding language instructions in structured environment states and executing actions under explicit constraints. Inspired by this paradigm, we reformulate image editing as task-driven interaction within a structured scene representation (Ha and Schmidhuber, 2018). Rather than directly synthesizing pixels, our framework plans and executes object-level actions conditioned on explicit spatial and relational states, enabling physically consistent editing.

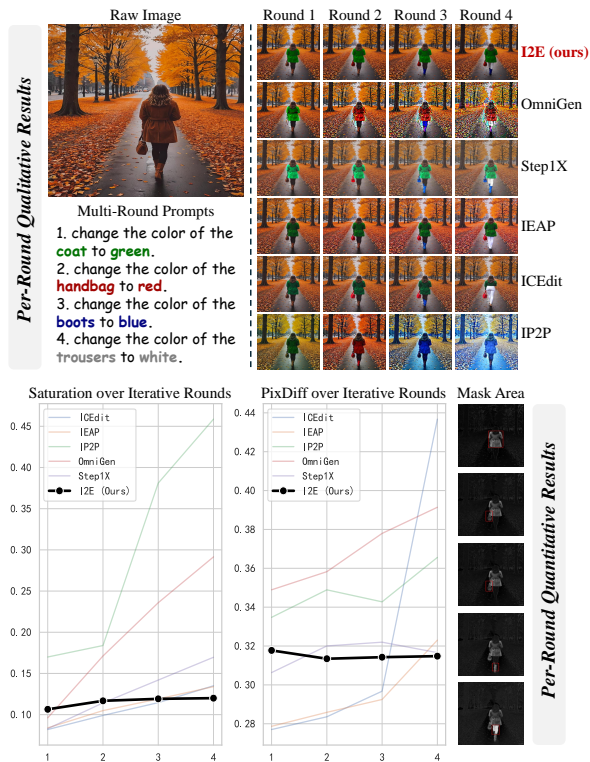


Figure 3: Multi-turn Stability. **Top:** Baselines exhibit severe *error accumulation* (e.g., visual distortion) over 4 rounds, while I2E preserves integrity. **Bottom:** Metrics confirm I2E’s constant consistency in Saturation and Pixel Difference (PixDiff) versus the monotonic degradation of end-to-end models.

3 Motivation: Analysis of End-to-End Editing Bottlenecks

End-to-end pixel inpainting, while effective for simple edits, struggles with compositional tasks that require precise local control and multi-instance spatial reasoning. We analyze these structural bottlenecks to motivate a paradigm shift.

3.1 Instruction Collapse

Figure 4 illustrates that when instructions contain multiple sub-goals, existing models often execute only a subset, ignoring or inconsistently satisfying others (*instruction collapse*). Analysis (Appendix A.2) attributes this to text encoding and conditioning limits: complex instructions are compressed into a single global embedding and injected via cross-attention, causing sub-goal conflicts and unstable execution.

3.2 Inevitability of Global Entanglement

End-to-end models are statistical generators, not deterministic editors, making *lossless local editing* theoretically infeasible. Two structural factors contribute: (i) *VAE bottleneck*: lossy compression-

reconstruction ($x \rightarrow z \rightarrow \hat{x}$) degrades high-frequency details in non-edited regions; (ii) *Self-attention coupling*: local feature changes propagate globally via dense attention. Multi-round simulations (Figure 3) confirm that non-target deviations accumulate with editing steps n , leading to severe visual artifacts.

3.3 Paradigm Shift: From Pixel Resampling to Entity Manipulation

These observations suggest that tightly coupling instruction parsing, planning, and pixel rendering limits edit reliability. We propose a paradigm shift: reformulate image editing by transforming unstructured pixel arrays into interactive structured environments, where a VLA agent executes edits through explicit entity manipulations instead of global pixel resampling.

4 Methodology

As illustrated in Figure 2, we propose the I2E framework, which reformulates image editing as an interaction process within an actionable environment. The framework consists of two cascaded stages: (i) *Environment Construction*, where a Decomposer transforms the input image into an explicit object-level environment representation; and (ii) *Agentic Interaction*, in which a VLA Editor performs physical reasoning and executes edits through atomic entity-level operations.

4.1 Decomposer: Environment Construction

The Decomposer converts an unstructured input image I into an interactive structured environment \mathcal{E} for object-level manipulation. It disentangles and completes relevant instances as independent entities, and organizes them into a physically consistent stacking hierarchy. The resulting object layers with explicit spatial relationships form a manipulable environment for subsequent agentic interaction.

Instance Disentanglement and Completion. To lift the input image into manipulable layers, we first employ a collaborative perception pipeline to identify and segment high-precision masks m_i for relevant instances (i.e., potential editing targets), while merging irrelevant objects into the background. This stage integrates a Multimodal Large Language Model (MLLM) for semantic reasoning with advanced grounding and segmentation frameworks to ensure mask accuracy. Since the segmented regions are inherently incomplete due

to occlusion, we utilize a generative fill-in mechanism to recover invisible structures. Guided by context-rich prompts from the MLLM, this process synthesizes missing textures and geometry, yielding a set of complete, transparent RGBA layers $\{\tilde{I}_i\}$. Concurrent with foreground processing, an occlusion-aware inpainting module is applied to remove the extracted instances from the original canvas, restoring a clean and cohesive background B . See details in Appendix A.4

Physical Layer Construction. Establishing a globally correct stacking order is a prerequisite for a physically consistent environment. Since explicit occlusion constraints only exist in regions where instances overlap at the pixel level, we propose a DAG-based Spatial Constraint Propagation Algorithm (see Algorithm 1) to infer global layer relationships. Specifically, we construct a directed acyclic graph (DAG) where nodes correspond to instances and directed edges encode occlusion dependencies. We jointly consider two types of constraints: (i) *hard constraints* derived from the pixel-level occlusion matrix predictions, and (ii) *soft constraints* obtained from monocular depth estimation, which refine the relative ordering without violating observed hard occlusions. By computing the transitive closure of the graph and defining the node out-degree as the depth score D_i , we resolve the global topological structure.

The global stacking sequence is then formalized as a permutation $\pi = (\pi_1, \dots, \pi_N)$ that satisfies the monotonicity constraint:

$$D_{\pi_k} \geq D_{\pi_{k+1}}, \quad \forall k \in [1, N-1], \quad (1)$$

where π_1 denotes the front-most instance index. Finally, each instance is encapsulated into an independent physical layer L_i , and combined with the background B to constitute the structured physical environment \mathcal{E} :

$$L_i = \{\tilde{I}_i, m_i, D_i\}, \quad \mathcal{E} = (\{L_i\}_{i=1}^N, B). \quad (2)$$

4.2 VLA Editor: Agentic Interaction

Given the structured environment \mathcal{E} , the VLA Editor serves as the decision and execution core, translating natural language instructions into physics-consistent actions that drive environment evolution.

Physics-Aware CoT Reasoning. We employ an MLLM-based (Yang et al., 2025) agent to perform chain-of-thought (CoT) reasoning under an explicit set of physical constraints \mathcal{C}_{phy} (e.g., gravity and

Algorithm 1 DAG-based Spatial Constraint Propagation

Input: Occlusion matrix $O \in \{0, 1\}^{N \times N}$, Depth soft constraints $O^{\text{soft}} \in \{0, 1\}^{N \times N}$
Output: Depth scores $D \in \mathbb{Z}^N$
1: # Phase 1: Occlusion (hard constraints)
2: **for all** $i, j \in \{1, \dots, N\}$ **do**
3: **if** $O_{ij} = 1$ **then**
4: $G_{ji} \leftarrow 1$
5: **end if**
6: **end for**
7: # Phase 2: Depth (soft constraints)
8: **for all** $i, j \in \{1, \dots, N\}$ **do**
9: **if** $O_{ij}^{\text{soft}} = 1 \wedge G_{ji} = 0$ **then**
10: $G_{ij} \leftarrow 1$
11: **end if**
12: **end for**
13: # Phase 3: Constraint propagation
14: **repeat**
15: $G \leftarrow G \vee (G \cdot G)$
16: **until** G converges
17: # Phase 4: Calculate Scores
18: **for** $i = 1$ to N **do**
19: $D_i \leftarrow \sum_{j \neq i} \mathbb{I}(G_{ij} = 1)$
20: **end for**
21: **return** $D = \{D_1, \dots, D_N\}$

support rules; see Appendix A.4) and a predefined action space \mathcal{A} (illustrated in Figure 2). Given the instruction T and the current scene state, the agent produces structured reasoning outputs, which are compiled into a sequence of parameterized atomic actions $\tilde{\mathcal{A}} = \{a_1, \dots, a_k\}$.

Action Execution. Each atomic action updates \mathcal{E} through object-level operations rather than global pixel resampling. REMOVE hides the target layer, exposing the pre-repaired background B . MOVE and FALL perform rigid transformations on object layers while preserving geometric integrity. RESIZE rescales object layers with fixed aspect ratios. Appearance edits are handled via EDIT and RETOUCH, which modify color, texture, or photometric attributes at the layer level. INSERT synthesizes a new object layer and inserts it into the global stacking order π according to predicted relational constraints, ensuring physically consistent occlusion.

Multi-Round Incremental Refinement. Since the environment state is explicitly maintained, I2E naturally supports incremental editing through action accumulation. User feedback or self-critique is handled by appending corrective actions without resetting the scene. By iteratively executing, evaluating, and revising, the closed loop stabilizes refinement and inhibits the compounding errors that arise with repeated pixel-level re-generation.

5 I2E-Bench

Existing benchmarks mainly target style transfer or simple single-step edits, and therefore inadequately

evaluate complex spatial reasoning, multi-instance interaction, and physical consistency. To fill this gap, we introduce I2E-BENCH, a benchmark for multi-instance spatial reasoning and high-precision image editing. It comprises 200 curated images from diverse open-source platforms, spanning real-world scenes, illustrations, and anime. Each image is paired with 5–10 editing instructions, emphasizing complex multi-action edits that require precise spatial manipulation while preserving stylistic and semantic coherence. Details in Appendix A.1.

6 Comparative Experiments

Baselines. We compare our method with representative state-of-the-art instruction-guided image editing approaches, including IP2P(Brooks et al., 2023), OmniGen(Xiao et al., 2025), Step1X(Liu et al., 2025a), IEAP(Hu et al., 2025), and ICEdit(Zhang et al., 2025b). These methods cover the dominant end-to-end and agent-based paradigms in the literature. For a fair comparison, we restrict all methods to a single refinement round.

Datasets. In addition to our proposed I2E-BENCH, we evaluate all methods on two widely adopted public benchmarks, MagicBrush(Zhang et al., 2023a) and EmuEdit(Sheynin et al., 2024), to ensure comprehensive and objective evaluation.

Metrics. We evaluate performance along three dimensions: (i) **Image fidelity.** We introduce LPIPS-U, a variant of LPIPS (Zhang et al., 2018), to measure perceptual similarity over unedited regions, and employ DINO-ViT (Caron et al., 2021) to assess semantic consistency. (ii) **Constraint adherence.** Spatial and operational constraints are evaluated using Spatial Accuracy (SA) and Constraint Satisfaction Rate (CSR). GroundingDINO (Liu et al., 2023) is used to localize referenced objects and automatically verify spatially constrained operations. (iii) **Instruction completion.** For high-level reasoning evaluation, we adopt Qwen3VL (Yang et al., 2025) to score Physical Consistency (PC) and Instruction Compliance (IC). We further report the Multi-step Score (MS) to quantify overall success in multi-action editing scenarios, which is particularly important for complex benchmarks such as I2E-BENCH. Implementation details are provided in Appendix A.5.

6.1 Quantitative Results

Tables 1 and 2 report quantitative comparisons between our method and state-of-the-art baselines on

Method	LPIPS-U↓	DINO↑	SA↑	CSR↑	PC↑	IC↑	MS↑
IP2P	0.2137	0.9565	0.4654	0.6400	0.7270	0.3747	0.3765
OmniGen	0.2072	0.9507	0.5263	0.6900	0.5620	0.3813	0.4300
Step1X	0.0353	0.9892	0.4861	0.6300	<u>0.8400</u>	<u>0.5898</u>	0.5483
IEAP	0.1532	0.9651	<u>0.6213</u>	<u>0.7700</u>	0.6290	0.5109	<u>0.5583</u>
ICEdit	0.0778	0.9816	0.4825	0.6200	0.7560	0.5125	0.5199
I2E	<u>0.0754</u>	<u>0.9821</u>	0.6923	0.8700	0.9210	0.8645	0.8074

Table 1: Quantitative comparison on I2E-BENCH. **Bold:** best; underline: second best.

Method	LPIPS-U↓	DINO↑	SA↑	CSR↑	PC↑	IC↑
<i>MagicBrush</i>						
IP2P	0.0865	0.8765	0.6840	0.9700	0.7680	0.4808
OmniGen	<u>0.0456</u>	<u>0.9254</u>	0.6820	0.9700	0.6600	0.6435
Step1X	0.0476	0.9226	<u>0.6920</u>	0.9800	0.7490	0.8617
IEAP	0.1893	0.7085	0.7020	<u>0.9900</u>	0.7640	0.5467
ICEdit	0.0554	0.9251	0.6820	0.9700	0.7240	0.7020
I2E	0.0446	0.9581	0.7120	1.0000	0.8668	<u>0.7924</u>
<i>EmuEdit</i>						
IP2P	0.1392	0.8734	0.6539	0.9300	0.8470	0.5676
OmniGen	0.0923	0.8625	0.6923	<u>0.9700</u>	0.7340	0.6563
Step1X	0.0732	0.9029	<u>0.6969</u>	<u>0.9700</u>	0.8190	<u>0.7672</u>
IEAP	0.1533	0.7592	0.6881	0.9600	0.8230	0.6645
ICEdit	0.0493	<u>0.9261</u>	0.6660	0.9400	<u>0.8570</u>	0.7238
I2E	<u>0.0565</u>	0.9404	0.7183	1.0000	0.8952	0.8107

Table 2: Quantitative comparison on MagicBrush and EmuEdit. **Bold:** best; underline: second best.

I2E-BENCH, MagicBrush, and EmuEdit.

Results on I2E-BENCH As shown in Table 1, our method consistently outperforms all baselines across most metrics. In particular, I2E achieves a substantial improvement on the Multi-step Score (MS), exceeding the second-best method by nearly 0.25. This gain mainly stems from the instance-level disentanglement introduced by the Decomposer, which isolates editing effects and effectively mitigates error accumulation in multi-step interactions. Moreover, equipped with the VLA Editor for explicit physics-aware reasoning and action decomposition, our method also achieves clear advantages in constraint-related metrics, including CSR, IC, and PC. These results demonstrate the effectiveness of reformulating image editing as structured interaction within an explicit physical environment.

We note that LPIPS-U and DINO scores are slightly lower than those of Step1X, which is primarily attributable to the currently adopted background restoration module. As an open framework, I2E can readily incorporate stronger restoration models to further improve perceptual fidelity.

Results on MagicBrush and EmuEdit Table 2 summarizes the results on two widely used public benchmarks. Our method shows consistent superiority across datasets. On MagicBrush, I2E achieves

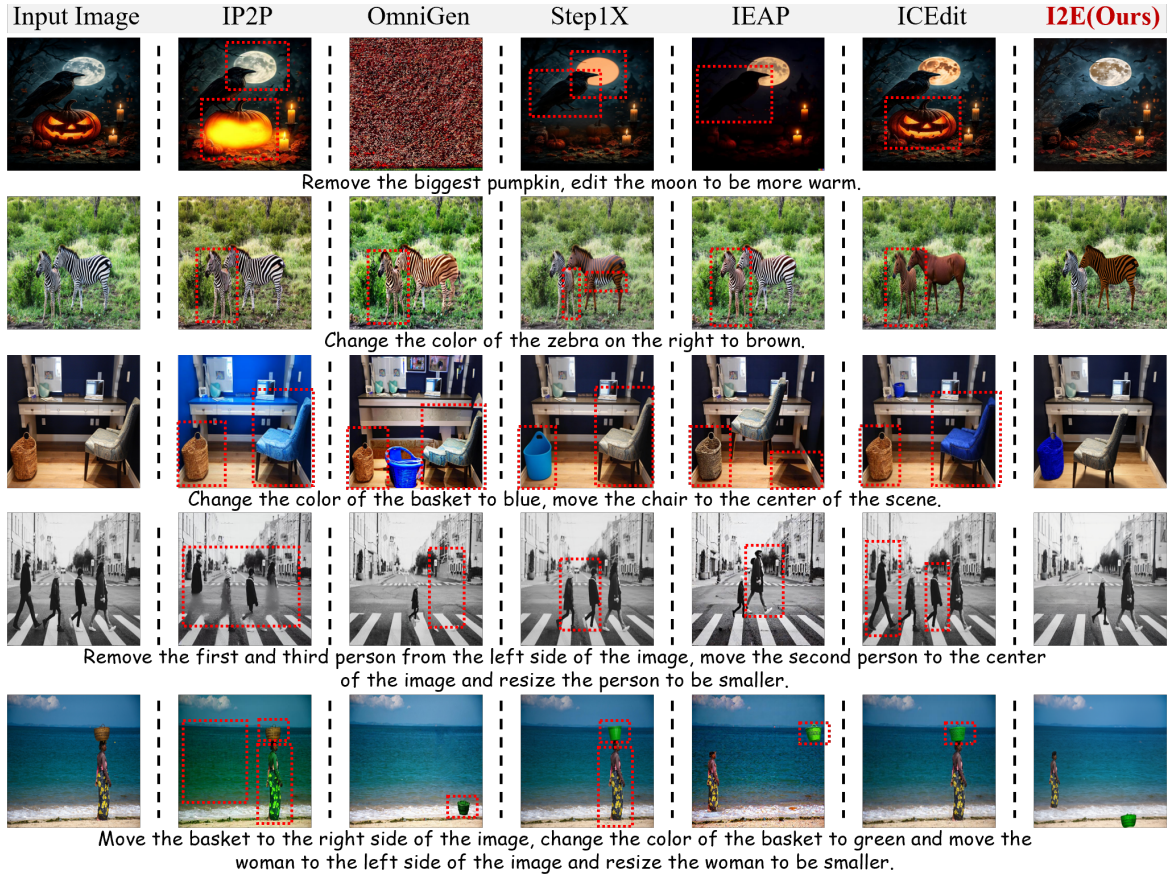


Figure 4: The results of the qualitative comparison on I2E-BENCH. The results are generated randomly within a batch with no re-run to ensure fairness.

Method	Avg. Score \uparrow	Rank	Mean Rank \downarrow	Rank
IP2P	2.24	6	4.76	6
OmniGen	3.39	5	4.08	5
Step1X	4.32	3	3.37	2
IEAP	4.40	2	3.39	3
ICEdit	3.83	4	3.60	4
I2E (Ours)	7.42	1	1.79	1

Table 3: Human evaluation results on I2E-BENCH. Rank columns indicate the ordering among all compared methods. **Bold** and underline denote the best and second-best results, respectively.

the best performance on all reported metrics. On EmuEdit, although LPIPS-U is marginally lower than ICEdit, I2E yields a significant improvement in instruction compliance (IC), with a relative gain of nearly 9%. Notably, I2E attains perfect constraint satisfaction (CSR = 1.0000) and the highest physical consistency (PC) on both datasets, highlighting the effectiveness of explicit spatial planning and physical constraint modeling for complex image editing.

6.2 Qualitative Results

Figure 4 visualizes comparisons on I2E-BENCH, highlighting our framework’s superiority across

three critical dimensions:

Compositionality and Physical Logic. End-to-end models often suffer from semantic coupling in multi-objective tasks. For instance, given the instruction “*remove the pumpkin and adjust the moon*” (Row 1), baselines typically miss sub-goals due to attention conflict. I2E, however, accurately disentangles and executes all constraints. Crucially, I2E enforces physical plausibility. Although baselines like Step1X remove the supporting object (pumpkin), it usually surves a clear physical hallucination like leaving the dependent object (crow) floating. Relatively, I2E detects the support loss and triggers a gravity simulation, naturally landing the crow to ensure logical consistency.

Spatial Precision and Attribute Isolation. I2E excels in maintaining geometric integrity during large-scale manipulations. In tasks requiring significant displacement (e.g., “*move the woman/chair*,” Rows 3&5), baselines frequently distort object structures or misplace targets. In contrast, our 2.5D layered representation enables precise translation and scaling. Furthermore, for local attribute editing



Figure 5: Qualitative results on I2E-Bench compared to commercial models.

(e.g., “recolor the right zebra,” Row 2), our strict layer isolation prevents the attribute leakage (color bleeding to the adjacent zebra) commonly observed in global processing models.

6.3 Human Evaluation

We conduct a blind human evaluation to assess human preference for complex image editing results. Thirty participants each score 10 randomly sampled cases, where six anonymized outputs from I2E and five baselines are shown in random order, using a 0–10 holistic rating. As shown in Table 3, I2E consistently achieves the highest preference.

7 Comparison with Commercial Models

We further present qualitative comparisons between I2E and several recent commercial models, including GPT5.1 (Wang et al., 2025), Seedream4.5 (Seedream et al., 2025), FLUX.1-Kontext-dev (Labs et al., 2025), Qwen3-Max (Team, 2025), and Nano Banana (Team et al., 2025) in Figure 5. The results show that I2E outperforms these closed-source models in compositional instruction-based image editing, particularly for tasks requiring pre-

cise instance-level control and spatially and physically consistent manipulation.

8 Ablation Study

We conduct an ablation study on I2E-Bench to quantify the contribution of each component, as reported in Table 4 and Figure 6. Specifically, we evaluate the effects of foreground reconstruction (FGR), background reconstruction (BGR), DAG-based spatial constraint propagation (DAG), physical reasoning (PR), and action reasoning (AR).

8.1 Quantitative Analysis.

Table 4 confirms all modules are essential.

Decomposer: Removing *Background Reconstruction* yields artificially better LPIPS-U because the model edits fewer pixels, but causes a sharp drop in IC (to 0.3387), indicating a failure to handle background-dependent instructions. Removing *Foreground Reconstruction* compromises the geometric integrity of occluded objects, significantly degrading PC and MS. Absence of *DAG-based Spatial Constraint Propagation* lowers CSR (0.87

Variant	Metrics (Value / Δ)													
	LPIPS-U \downarrow		DINO \uparrow		SA \uparrow		CSR \uparrow		PC \uparrow		IC \uparrow		MS \uparrow	
Full (Ours)	0.0754	-	0.9821	-	0.6923	-	0.8700	-	0.9210	-	0.8645	-	0.8074	-
<i>Decomposer</i>														
w/o BGR	0.0347	-0.0407	0.9837	+0.0016	0.6562	-0.0361	0.8100	-0.0600	0.6540	-0.2670	0.3387	-0.5258	0.5432	-0.2642
w/o FGR	0.0803	+0.0049	0.9797	-0.0024	0.6507	-0.0416	0.7800	-0.0900	0.7450	-0.1760	0.4943	-0.3702	0.6383	-0.1691
w/o DAG	0.0828	+0.0074	0.9787	-0.0034	0.6253	-0.0670	0.7700	-0.1000	0.8490	-0.0720	0.4792	-0.3853	0.5557	-0.2517
<i>VLA Editor</i>														
w/o PR	0.0822	+0.0068	0.9774	-0.0047	0.6483	-0.0440	0.8000	-0.0700	0.8360	-0.0850	0.4742	-0.3903	0.5142	-0.2932
w/o AR	0.0896	+0.0142	0.9798	-0.0023	0.6177	-0.0746	0.8500	-0.0200	0.8100	-0.1110	0.3151	-0.5494	0.3449	-0.4625

Table 4: Quantitative ablation on I2E-BENCH.




Figure 6: Qualitative ablation on I2E-BENCH.

→ 0.77), confirming the necessity of DAG-based sorting for modeling complex occlusions.

VLA Editor: Without *Action Reasoning*, IC and MS drop to 0.3151 and 0.3449, highlighting the critical role of decomposing high-level instructions into atomic actions for long-horizon stability. Disabling *Physical Reasoning* reduces PC (0.92 → 0.84), verifying that explicit constraints (e.g., gravity and support) are essential to prevent physical anomalies such as floating objects.

8.2 Qualitative Analysis

Figure 6 qualitatively supports the above results.

Decomposer: Removing *background reconstruction* causes visible discontinuities, while disabling *foreground reconstruction* leads to incomplete instances after movement or occlusion changes. Without *DAG-based propagation occlusion*, the ordering in crowded scenes will become inconsistent.

VLA Editor: Ablating *action reasoning*, edits become incomplete or incorrectly ordered. Disabling *physical reasoning* results in floating objects or invalid support relations.

9 Conclusion

We propose I2E, a structured framework for instruction-based image editing via explicit reasoning and execution. By decomposing instructions

into spatially and physically grounded actions, I2E enables more controllable and consistent editing. Future work aims to extend I2E to professional domains like interior design, envisioning the VLA Editor evolving from a passive “Instruction Executor” into an active “Intelligent Designer.”

Limitations

Despite the encouraging results, the current framework still faces several limitations that warrant further investigation. First, the quality of scene decomposition is highly dependent on the performance of the underlying foundation models. Although SAM 2 and Flux provide state-of-the-art results, they may still struggle with extremely complex occlusion patterns or transparent objects such as glass and water, leading to imprecise segmentation or inpainting artifacts that can propagate to subsequent editing stages. Second, at the rendering stage, the system does not yet fully rely on physical simulation to handle complex material properties, such as soft-body deformation or fluid dynamics, nor fine-grained lighting interactions, such as caustics and interreflections. Consequently, edits involving significant lighting changes or object deformation may lack the highest level of photorealism.

Acknowledgments

This paper is supported by the National Natural Science Foundation of China (No. 62406161) and sponsored by CCF-Kuaishou Large Model Explorer Fund (NO. CCF-KuaiShou 2025003).

References

- Jiayang Ao, Yanbei Jiang, Qihong Ke, and Krista A Ehinger. 2025. Open-world amodal appearance completion. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6490–6499.
- James Betker, Gabriel Goh, Li Jing, Tim Brooks, Jianfeng Wang, Linjie Li, Long Ouyang, Juntang Zhuang, Joyce Lee, Yufei Guo, Wesam Manassra, Prafulla Dhariwal, Casey Chu, Yunxin Jiao, and Aditya Ramesh. 2023. Improving image generation with better captions. *Computer Science*. <https://cdn.openai.com/papers/dall-e-3.pdf>, 2(3):8.
- Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, Pete Florence, Chuyuan Fu, Montse Gonzalez Arenas, Keerthana Gopalakrishnan, Kehang Han, Karol Hausman, Alexander Herzog, Jasmine Hsu, Brian Ichter, and 35 others. 2023. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *Conference on Robot Learning*, pages 2165–2183. PMLR.
- Tim Brooks, Aleksander Holynski, and Alexei A Efros. 2023. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18392–18402.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660.
- Ruihan Chen, Qiming Li, Xiaocheng Feng, Xiaoliang Yang, Weihong Zhong, Yuxuan Gu, Zekun Zhou, and Bing Qin. 2025. [Mpr-gui: Benchmarking and enhancing multilingual perception and reasoning in gui agents](#). *Preprint*, arXiv:2512.00756.
- Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, and 3 others. 2023. Palm-e: An embodied multimodal language model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 2969–2977.
- Tsu-Jui Fu, Wenze Hu, Xianzhi Du, William Yang Wang, Yinfei Yang, and Zhe Gan. 2024. Guiding Instruction-based Image Editing via Multimodal Large Language Models. In *International Conference on Learning Representations (ICLR)*.
- David Ha and Jürgen Schmidhuber. 2018. World models. *arXiv preprint arXiv:1803.10122*, 2(3).
- Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. 2022. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*.
- Yujia Hu, Songhua Liu, Zhenxiong Tan, Xingyi Yang, and Xinchao Wang. 2025. Image editing as programs with diffusion models. *arXiv preprint arXiv:2506.04158*.
- Yuzhou Huang, Liangbin Xie, Xintao Wang, Ziyang Yuan, Xiaodong Cun, Yixiao Ge, Jiantao Zhou, Chao Dong, Rui Huang, Ruimao Zhang, and Ying Shan. 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*, pages 8362–8371.
- Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. 2024. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*.
- Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. 2023. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4015–4026.
- Black Forest Labs, Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, Cyril Diagne, Tim Dockhorn, Jack English, Zion English, Patrick Esser, Sumith Kulal, Kyle Lacey, Yam Levi, Cheng Li, Dominik Lorenz, Jonas Müller, Dustin Podell, Robin Rombach, Harry Saini, and 2 others. 2025. [Flux.1 kontext: Flow matching for in-context image generation and editing in latent space](#). *Preprint*, arXiv:2506.15742.
- Hyunmin Lee and Jaesik Park. 2022. Instance-wise Occlusion and Depth Orders in Natural Scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Qiming Li, Xiaocheng Feng, Yixuan Ma, Zekai Ye, Ruihan Chen, Xiachong Feng, and Bing Qin. 2025. Unlocking multilingual reasoning capability of llms and lvlms through representation engineering. *arXiv preprint arXiv:2511.23231*.

- Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, and Lei Zhang. 2023. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*.
- Shiyu Liu, Yucheng Han, Peng Xing, Fukun Yin, Rui Wang, Wei Cheng, Jiaqi Liao, Yingming Wang, Honghao Fu, Chunrui Han, Guopeng Li, Yuang Peng, Quan Sun, Jingwei Wu, Yan Cai, Zheng Ge, Ranchen Ming, Lei Xia, Xianfang Zeng, and 5 others. 2025a. Step1x-edit: A practical framework for general image editing. *arXiv preprint arXiv:2504.17761*.
- Zhaochen Liu, Limeng Qiao, Xiangxiang Chu, Lin Ma, and Tingting Jiang. 2025b. Towards efficient foundation model for zero-shot amodal segmentation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 20254–20264.
- Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. 2021. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*.
- Sunung Mun, Jinhwan Nam, Sunghyun Cho, and Jungseul Ok. 2025. Addressing text embedding leakage in diffusion-based image editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16451–16460.
- Ege Ozguroglu, Ruoshi Liu, Dídac Surís, Dian Chen, Achal Dave, Pavel Tokmakov, and Carl Vondrick. 2024. pix2gestalt: Amodal segmentation by synthesizing wholes. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3931–3940. IEEE Computer Society.
- Xingang Pan, Ayush Tewari, Thomas Leimkühler, Lingjie Liu, Abhimitra Meka, and Christian Theobalt. 2023. Drag your gan: Interactive point-based manipulation on the generative image manifold. In *ACM SIGGRAPH 2023 conference proceedings*, pages 1–11.
- Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryalí, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. 2024. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*.
- Team Seedream, :, Yunpeng Chen, Yu Gao, Lixue Gong, Meng Guo, Qiushan Guo, Zhiyao Guo, Xiaoxia Hou, Weilin Huang, Yixuan Huang, Xiaowen Jian, Huafeng Kuang, Zhichao Lai, Fanshi Li, Liang Li, Xiaochen Lian, Chao Liao, Liyang Liu, and 32 others. 2025. *Seedream 4.0: Toward next-generation multimodal image generation*. *Preprint*, arXiv:2509.20427.
- Shelly Sheynin, Adam Polyak, Uriel Singer, Yuval Kirstain, Amit Zohar, Oron Ashual, Devi Parikh, and Yaniv Taigman. 2024. Emu edit: Precise image editing via recognition and generation tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8871–8879.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, David Silver, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy Lillicrap, Angeliki Lazaridou, and 1332 others. 2025. *Gemini: A family of highly capable multimodal models*. *Preprint*, arXiv:2312.11805.
- Qwen Team. 2025. Qwen3-max: Just scale it.
- Yuhan Wang, Siwei Yang, Bingchen Zhao, Letian Zhang, Qing Liu, Yuyin Zhou, and Cihang Xie. 2025. *Gpt-image-edit-1.5m: A million-scale, gpt-generated image dataset*. *Preprint*, arXiv:2507.21033.
- Runpu Wei, Zijin Yin, Shuo Zhang, Lanxiang Zhou, Xueyi Wang, Chao Ban, Tianwei Cao, Hao Sun, Zhongjiang He, Kongming Liang, and Zhanyu Ma. 2025. *Omnieraser: Remove objects and their effects in images with paired video-frame data*. *arXiv preprint arXiv:2501.07397*.
- Shitao Xiao, Yueze Wang, Junjie Zhou, Huaying Yuan, Xingrun Xing, Ruiran Yan, Chaofan Li, Shuting Wang, Tiejun Huang, and Zheng Liu. 2025. *Omnigen: Unified image generation*. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 13294–13304.
- Hongshen Xu, Zihan Wang, Zichen Zhu, Lei Pan, Xingyu Chen, Shuai Fan, Lu Chen, and Kai Yu. 2025. Alignment for efficient tool calling of large language models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 17787–17803.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chuji Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. 2019. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4471–4480.
- Qifan Yu, Wei Chow, Zhongqi Yue, Kaihang Pan, Yang Wu, Xiaoyang Wan, Juncheng Li, Siliang Tang, Hanwang Zhang, and Yueting Zhuang. 2025. Anyedit: Mastering unified high-quality image editing for any idea. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 26125–26135.

Hao Zhang, Mengsi Lyu, Chenrui He, Yulong Ao, and Yonghua Lin. 2025a. Trimtokenator: Towards adaptive visual token pruning for large multimodal models. *arXiv preprint arXiv:2509.00320*.

Kai Zhang, Lingbo Mo, Wenhui Chen, Huan Sun, and Yu Su. 2023a. Magicbrush: A manually annotated dataset for instruction-guided image editing. *Advances in Neural Information Processing Systems*, 36:31428–31449.

Lvmin Zhang and Maneesh Agrawala. 2024. Transparent image layer diffusion using latent transparency. *arXiv preprint arXiv:2402.17113*.

Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. 2023b. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3836–3847.

Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595.

Zechuan Zhang, Ji Xie, Yu Lu, Zongxin Yang, and Yi Yang. 2025b. In-context edit: Enabling instructional image editing with in-context generation in large scale diffusion transformer. *arXiv preprint arXiv:2504.20690*.

A Appendix

A.1 Details of I2E-Bench

Details of Data Collection The 200 base images for I2E-Bench were manually selected from Pixabay¹. We intentionally chose this source due to its high-resolution content and the diversity of artistic styles it offers. For each image, we curated a diverse set of 5–10 instructions, ranging from straightforward single-action prompts to intricate multi-step tasks. This design is specifically used to provide a rigorous assessment of the model’s capabilities in spatial reasoning and sequential execution. Representative examples are provided in Figure 7.

A.2 Theoretical Analysis: Structural Limitations of End-to-End Editing

Theorem A.1 (Instruction Collapse under Global Conditioning). *Let an end-to-end image editor condition generation on a single text embedding $c \in \mathbb{R}^d$ produced from a composite instruction $T = \{s_1, \dots, s_K\}$, where each s_k denotes a semantically independent sub-instruction. For sufficiently large K , the model cannot guarantee the simultaneous and stable execution of all sub-instructions in a single forward generation.*

Proof Sketch. The editor encodes the full instruction as

$$c = E_{\text{text}}(T), \quad (3)$$

where E_{text} has fixed output dimension d independent of K . Each sub-instruction introduces at least one independent semantic factor, so the intrinsic degrees of freedom required to represent T grow with K . When $K > d$, the mapping E_{text} is necessarily non-injective, i.e.,

$$\exists T \neq T' \quad \text{s.t.} \quad E_{\text{text}}(T) = E_{\text{text}}(T'), \quad (4)$$

implying unavoidable information loss.

Even when $K \leq d$, the single embedding c jointly represents all sub-instructions. Since no constraint enforces disentanglement, the induced sub-instruction representations occupy overlapping directions in \mathbb{R}^d :

$$\langle E_{\text{text}}(s_i), E_{\text{text}}(s_j) \rangle \neq 0, \quad i \neq j. \quad (5)$$

During generation, c is broadcast to all spatial or latent tokens via attention. Without an explicit routing mechanism, sub-instructions compete for the

¹<https://pixabay.com/>

same conditioning channel, leading to unstable or selective execution. This phenomenon is referred to as *instruction collapse*. \square

Proposition A.2 (Elimination of Same-Step Gradient Conflict by Sequential Decomposition). *Consider the joint optimization objective*

$$\mathcal{L}_{\text{joint}}(\theta) = -\log p_{\theta}(x_0 | c_1, \dots, c_K), \quad (6)$$

where $\{c_k\}$ denote sub-instruction conditions. Let \mathbf{g}_k be the gradient contribution induced by c_k . Then the gradient of $\mathcal{L}_{\text{joint}}$ contains cross terms $\langle \mathbf{g}_i, \mathbf{g}_j \rangle$ for $i \neq j$, which may be negative.

By contrast, if the objective is decomposed into a sequence of conditional sub-objectives

$$\mathcal{L}_{\text{seq}}(\theta) = \sum_{k=1}^K \mathcal{L}_k(\theta), \quad (7)$$

$$\mathcal{L}_k = -\log p_{\theta}(x_0^{(k)} | x_0^{(k-1)}, c_k), \quad (8)$$

then no cross-gradient terms arise within the same optimization step.

Proof Sketch. For the joint objective,

$$\nabla_{\theta} \mathcal{L}_{\text{joint}} = \sum_{k=1}^K \mathbf{g}_k, \quad (9)$$

and the squared gradient norm expands as

$$\|\nabla_{\theta} \mathcal{L}_{\text{joint}}\|^2 = \sum_k \|\mathbf{g}_k\|^2 + \sum_{i \neq j} \langle \mathbf{g}_i, \mathbf{g}_j \rangle, \quad (10)$$

where negative inner products correspond to gradient conflict.

Under the sequential formulation, each sub-objective \mathcal{L}_k depends on a single condition c_k . At optimization step k , the gradient is

$$\nabla_{\theta} \mathcal{L}_k = \mathbf{g}_k, \quad (11)$$

and no terms involving \mathbf{g}_j for $j \neq k$ appear in the same update. Thus, while interactions across steps may still exist, the same-step gradient conflict inherent to joint optimization is structurally eliminated. \square

A.3 Theoretical Analysis: Global Degradation in Multi-Round Editing

observation A.3 (Global Degradation under Iterative End-to-End Editing). *Consider an end-to-end image editor that applies a sequence of sub-instructions $\{c_1, \dots, c_N\}$ through repeated generative updates*

$$x^{(t+1)} = F_{\theta}(x^{(t)}, c_t). \quad (12)$$

Even when each sub-instruction c_t is intended to modify only a localized region, the deviation of non-target regions increases with the number of editing rounds N .

Explanation. Each update F_{θ} operates on the full image or latent representation. Due to finite latent capacity and stochastic sampling, the reconstruction at each step can be written as

$$x^{(t+1)} = x^{(t)} + \epsilon^{(t)}, \quad (13)$$

where $\epsilon^{(t)}$ denotes the reconstruction error affecting both target and non-target regions.

Since updates are applied recursively, the accumulated deviation after N rounds is

$$x^{(N)} - x^{(0)} = \sum_{t=1}^N \epsilon^{(t)} + \mathcal{O}(\epsilon^2), \quad (14)$$

which implies

$$\mathbb{E}\|x^{(N)} - x^{(0)}\| \geq \sum_{t=1}^N \mathbb{E}\|\epsilon^{(t)}\|. \quad (15)$$

Moreover, modern generative editors rely on dense token mixing mechanisms, causing local modifications to perturb global feature statistics. As a result, errors introduced in early rounds propagate to non-target regions and are amplified over subsequent rounds.

This observation explains the empirically observed background degradation and increasing pixel-wise deviation outside the edited regions as the number of editing rounds grows. \square

A.4 Details of Implementation

Decomposer Our perception pipeline is implemented using Qwen-VL (Yang et al., 2025) as the MLLM, Grounding DINO (Liu et al., 2023) for bounding box localization, and SAM (Kirillov et al., 2023) for pixel-level segmentation. For layer completion, we employ Flux-Fill (Labs et al., 2025)

to perform generative outpainting and content synthesis. Specifically, we use the MLLM to generate detailed descriptive prompts for each occluded instance; these prompts guide Flux-Fill to hallucinate the missing textures and structures behind occlusions, resulting in the final RGBA layers $\{\tilde{I}_i\}$. Simultaneously, to ensure a clean slate for downstream editing, we utilize OmniEraser (Wei et al., 2025) to perform background inpainting. It effectively removes the foreground instances by filling the holes in the original image using surrounding context, thereby yielding a spatially coherent and artifact-free background B .

Physical Reasoning Prompt We utilize a multimodal large language model (MLLM) to perform explicit physical reasoning over the image environment. The model is instructed to simulate physical constraints such as gravity, support, and balance using a structured “mind’s eye” reasoning process. Below we provide the full prompt used for physical reasoning in our framework.

```
You are a physics reasoning expert. Analyze this scene and the user instruction.
Scene Description: {scene_desc}
User Instruction: "{text_prompt}"
=== Physical World Simulation Rules (Mind's Eye) ===
Apply the following explicit physics rules:
**Gravity Rules:** - Rule 1: If Support Object X is removed, Supported Object Y MUST fall -
Rule 2: All objects are affected by gravity unless supported
**Support Rules:** - Rule 3: Object Y is supported by X if and only if Y is above X and in contact -
Rule 4: If X supports Y, removing X causes Y to lose support
**Balance Rules:** - Rule 5: An object remains balanced if its center of mass is within the support base -
Rule 6: If support shifts causing center of mass to move outside base, object loses balance
=== Reasoning Steps ===
1. Identify Target Objects 2. Current State 3. Action Effects 4. Physical Constraints 5. Secondary Effects 6. Final State
=== Output Format ===
Provide structured JSON: { "reasoning": "...", "target_objects": [...], "constraints": [...], "affected_objects": [...], "predicted_outcome": "..." }
```

Atomic Actions planner prompt We utilize a multimodal large language model (MLLM) to perform explicit chain-of-thought (CoT) reasoning over the image environment. The MLLM acts as an action planner that translates the reasoning results

into a sequence of executable atomic actions, which are subsequently applied to the scene to achieve the desired editing goal. Below we provide the full prompt used for the action planner in our framework.

```
You are an action planning expert. Based on the physical reasoning, generate atomic actions.
Scene Description: {scene_desc}
User Instruction: "{text_prompt}"
Physical Reasoning Result: {reasoning_result}
Generate a sequence of atomic actions to achieve the instruction while respecting physical constraints.
=== Image Coordinate System ===
Image Dimensions: {width} × {height} pixels
Reference Points: - Top-Left: (0, 0) - Top-Right: ({width}, 0) - Bottom-Left: (0, {height}) - Bottom-Right: ({width}, {height}) - Image Center: ({width}/2, {height}/2)
Coordinate Rules: - Origin at top-left - X increases left to right - Y increases top to bottom - All coordinates must satisfy image bounds
=== Available Atomic Actions ===
REMOVE(object_id)
MOVE(object_id, x, y)
KEEP(object_id)
FALL(object_id, delta_y)
RESIZE(object_id, scale)
RETOUCH(object_id, brightness, contrast, color, sharpness)
EDIT(object_id, prompt, edit_type)
INSERT(prompt, x, y, width, height, layer_relation)
Guidelines: 1. Respect gravity and support relations 2. Maintain physical plausibility 3. Execute actions in correct order 4. Strictly adhere to image boundaries
=== Output Format ===
{ "action_sequence": [ { "action": "...", "object_id": ..., "reason": "..." } ] }
Generate the action sequence.
```

A.5 Details of Evaluation Metrics

Comprehensive Metrics. Beyond fundamental metrics like redesigned LPIPS-U and DINO-ViT for fidelity and perceptual quality, we specifically design SA and CSR to quantify constraint-following capabilities. To further capture the nuances of instruction completion, we incorporate MLLM-based metrics, including PC, IC, and MS. The specifics of these metrics are detailed in the subsequent discussion.

LPIPS-U (Unedited-Region LPIPS). A key challenge in image editing evaluation is disentangling edit correctness from background preservation. To emphasize preservation, we define

LPIPS-U that measures perceptual distance *only on unedited regions*.

Let I be the original image and \hat{I} be the edited image. We first estimate a binary edit mask $M \in \{0, 1\}^{H \times W}$, where $M(p) = 1$ indicates edited pixels and $M(p) = 0$ indicates unedited pixels; $\bar{M} = 1 - M$ denotes the unedited mask. Let $\phi_\ell(\cdot)$ be the feature map at layer ℓ of a fixed perceptual backbone (as in LPIPS), and \odot denote element-wise masking (broadcast to channels). LPIPS-U is defined as

$$\text{LPIPS-U}(I, \hat{I}) = \sum_{\ell \in \mathcal{L}} w_\ell \cdot \left\| \left(\phi_\ell(I) - \phi_\ell(\hat{I}) \right) \odot \bar{M}_\ell \right\|_2, \quad (16)$$

where \bar{M}_ℓ is the unedited mask resized to match the spatial resolution of $\phi_\ell(\cdot)$, and $\{w_\ell\}$ are layer weights.

Unlike LPIPS, LPIPS-U suppresses the contribution from edited regions, making it more sensitive to *unintended changes* outside the target area (i.e., background/irrelevant-region degradation), which is crucial for instruction-based editing.

Constraint following Metrics (SA and CSR). We quantify spatial adherence using a GroundingDINO-based detector by comparing object presence and location between I and \hat{I} .

For each sample i , suppose the prompt yields a set of spatial constraints \mathcal{C}_i . Each constraint $c \in \mathcal{C}_i$ specifies a target object and an expected spatial relation (e.g., left/right/top/bottom/center) and implies an operation type (e.g., REMOVE, MOVE, INSERT). Let $a_{i,c} \in [0, 1]$ be the per-constraint accuracy computed from detections (object existence and/or normalized center coordinates).

We define **Spatial Accuracy (SA)** as the mean accuracy across all evaluated constraints:

$$\text{SA} = \frac{1}{\sum_i |\mathcal{C}_i|} \sum_i \sum_{c \in \mathcal{C}_i} a_{i,c}. \quad (17)$$

We further define **Constraint Satisfaction Rate (CSR)** as the fraction of constraints that are satisfied under a threshold τ (e.g., $\tau = 0.7$):

$$\text{CSR} = \frac{1}{\sum_i |\mathcal{C}_i|} \sum_i \sum_{c \in \mathcal{C}_i} \mathbb{I}[a_{i,c} \geq \tau], \quad (18)$$

where $\mathbb{I}[\cdot]$ is the indicator function. Intuitively, SA reflects average spatial precision, while CSR measures strict compliance frequency.

Instruction-completion evaluation Metrics (PC, IC, and MS). We utilize a multimodal learning model (Qwen3-VL) to generate structured assessments that better align with human perception and instruction execution. Scores range from 1 to 10, ultimately mapping to a 0-1 scale.

Physical Consistency (PC). The judge enumerates physical flaws introduced by the edit (e.g., inconsistent shadows/lighting, implausible occlusion, perspective violations, visible artifacts) with severities. We compute a deduction-style score:

$$\text{PCPC} = \text{clip}_{[1,10]} \left(10 - \sum_{k \in \mathcal{E}_{\text{phys}}} d_k \right), \quad (19)$$

where $\mathcal{E}_{\text{phys}}$ is the set of detected physical issues and d_k is the penalty associated with issue k (larger for more severe flaws).

Instruction Following (IC). The judge decomposes the prompt into a set of atomic, visually verifiable requested edits \mathcal{R} , and assigns each request $r \in \mathcal{R}$ a fulfillment score $f_r \in [0, 1]$ (fulfilled/partial/not fulfilled). Additionally, it reports penalties for unrequested changes and non-target-region damage. We compute:

$$\text{IC} = \text{clip}_{[1,10]} \left(10 \cdot \frac{1}{|\mathcal{R}|} \sum_{r \in \mathcal{R}} f_r - \sum_{u \in \mathcal{E}_{\text{unwanted}}} d_u - \sum_{p \in \mathcal{E}_{\text{preserve}}} d_p \right). \quad (20)$$

This formulation emphasizes both completing requested edits and avoiding collateral changes.

Multi-step Score (MS). While IC focuses on localized fidelity, the Multi-step Score (MS) evaluates the model’s ability to execute complex, sequential instructions. For multi-step instructions, the judge first lists the expected action steps \mathcal{S} (restricted to explicit edit actions such as REMOVE, MOVE, EDIT, INSERT, RESIZE), then scores each step $s \in \mathcal{S}$ with success $g_s \in [0, 1]$ based on comparing (I, \hat{I}) . We define:

$$\text{MS} = 10 \cdot \frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} g_s. \quad (21)$$

This metric directly measures step-wise executability and is conservative when steps are missing or not clearly satisfied.

A.6 Details of Evaluation

Evaluation is conducted on three benchmarks: I2E-Bench, MagicBrush, and EmuEdit. For each benchmark, we randomly sample 100 image–instruction pairs from the official evaluation sets. All quantitative metrics are computed using the same evaluation protocol and fixed random seeds to ensure consistency and reproducibility. No additional data filtering or manual selection is performed. The same sampled instances are used consistently across all compared methods and evaluation protocols.

A.7 Details of VLA Editing

We provide more visualization of our VLA process in Figure 8.

A.8 Details of Human Evaluation

Additional visualizations of the human evaluation process are provided in Figure 9 and Figure 10. All evaluations were conducted by unpaid volunteers. The data were collected in a blind manner, where participants were not informed of the methods associated with the evaluated results. The evaluation interface and instructions were presented in English. No personal identifying information was collected, and no assumptions or guarantees are made regarding the demographic attributes (e.g., region or gender) of the participants.

A.9 More qualitative Results

In this section, we present additional qualitative results to further demonstrate the effectiveness of our method across various scenarios in the I2E-Bench. Figures 11, 12, and 13 provide extensive cases demonstrating that our I2E framework consistently outperforms representative baselines in executing multi-step composite instructions, maintaining physical consistency, and ensuring precise spatial localization.

Additionally, we provide qualitative results on real-world product image editing scenarios from Kuaishou Mall in Fig. 14, where our method demonstrates strong practicality and robustness.

A.10 Clarification of Concerns

Data Contains Personally Identifying Info Or Offensive Content All datasets used in this work are collected from publicly available, open-source websites and benchmarks that are explicitly released for non-commercial academic research. We do not intentionally collect or curate any data containing personally identifying information (PII).

The data sources are commonly used in prior research and do not target specific individuals. No additional annotation or manual data collection involving human subjects was conducted in this work. As such, the risk of containing sensitive or identifying personal information is minimal.

The Use of Large Language Models (LLMs)

Large language models (LLMs) were used solely for auxiliary purposes, including grammar checking, formatting refinement, and translation assistance during manuscript preparation. In addition, minor illustrative elements in Figure 1 were generated with the assistance of an LLM-based image generation tool for visualization purposes only (e.g., depicting a human figure). These generated elements are purely illustrative and do not contribute to the core methodology, experimental results, or scientific claims of this work. The overall conceptual design and technical content are entirely human-authored. The use of LLMs does not affect the validity or reproducibility of the experimental findings.

Discuss The License For Artifacts All external artifacts used in this work, including datasets and pretrained models, are obtained from open-source research projects and publicly accessible repositories. These artifacts are used in accordance with their original licenses and terms of use, which permit non-commercial academic research. No proprietary or restricted data or models are used.

Artifact Use Consistent With Intended Use

The usage of existing artifacts in this work is consistent with their intended research purposes as specified by their original authors. All experiments are conducted strictly within a non-commercial, academic research context. Any derived artifacts or outputs produced by our pipeline are intended solely for research and evaluation purposes and are not deployed in real-world or commercial applications.

Model Size And Budget This work adopts a training-free pipeline and does not involve training or fine-tuning large-scale models. Therefore, there is no additional model parameter reporting or large-scale computational budget associated with model training. The experiments are conducted using standard academic computing resources for inference and evaluation. Furthermore, the model can be further lightweighted and deployed efficiently via


mainstream model pruning (Zhang et al., 2025a) and quantization techniques.

Documentation Of Artifacts The datasets and artifacts used in this work are standard benchmarks in the research community and are documented in their original releases, including information about domains and language usage. Our work focuses on instruction-based multimodal inputs, with all textual instructions provided in English. No claims are made regarding demographic representativeness beyond what is specified in the original datasets.

Potential Risks This work is based entirely on publicly available, open-source models, data and datasets that are widely used in prior academic research.

Potential risks are limited to the general misuse of image editing technologies, such as generating misleading or manipulated visual content. However, these risks are not unique to the proposed method and are shared by existing generative and editing models. The intended use of this work is strictly non-commercial academic research, and no deployment in real-world decision-making systems is considered. We do not anticipate any significant ethical, societal, or safety concerns arising from the use of this work under its intended scope.

Original Image



Instructions

"Keep the documents in their current position and insert a new coin next to the calculator."

"Remove the green euro note from the left side of the image."

"Move the spoon to the right of the calculator and resize the potato to be half its current size."

"Remove the green euro note from the left side of the image, causing the scattered coins to shift toward the center."

"Insert a new document under the calculator and move the spoon to the top of the new document."

"Remove the fence and move the horse to the left side of the road."

"Remove the green euro note from the left side of the image."

"Move the sign behind the trees and resize the mountains to be shorter."

"Edit the trees to have more autumn colors and move the horse closer to the sign."

"Change the color of the left horse to white."

"Remove the rock causing the person to fall onto the grass and then move the suitcase next to the bush."

"Move the hat from the person's head to the left side of the image and resize the bush to be smaller."

"Edit the color of the suitcase to navy blue and move the person to sit on the grass near the pole."

"Remove the bush next to the rock and move the rock to the right side of the person."

"Change the color of the hat to bright red."

Figure 7: Examples of I2E-bench.




Original Image	VLA Editing process		I2E Result
	Instruction	Atomic Actions	
	<p>Remove the spoon, move the calculator in front of the potato and change the color of potato to green.</p>	<p>REMOVE(id=0), FALL(id=1, dy=400), FALL(id=4, dy=300), MOVE(id=2, x'=800, y'=700), EDIT(id=1, prompt, r=color)</p>	
	<p>Remove the fence and the right horse, and move the left horse in the road.</p>	<p>REMOVE(id=2), REMOVE(id=1), MOVE(id=0, x'=250, y'=650)</p>	
	<p>Edit the color of the suitcase to navy blue and move the person to sit on the rock near the pole.</p>	<p>EDIT(id=2, prompt, r=color), MOVE(id=0, x'=897, y'=443)</p>	
	<p>Resize the car to be larger and place it closer to the pole with the flag.</p>	<p>RESIZE(id=0, s=1.2), MOVE(id=0, x'=900, y'=683)</p>	

Figure 8: Visualization of VLA process on our proposed I2E-Bench.




Figure 9: Example of human evaluation system.

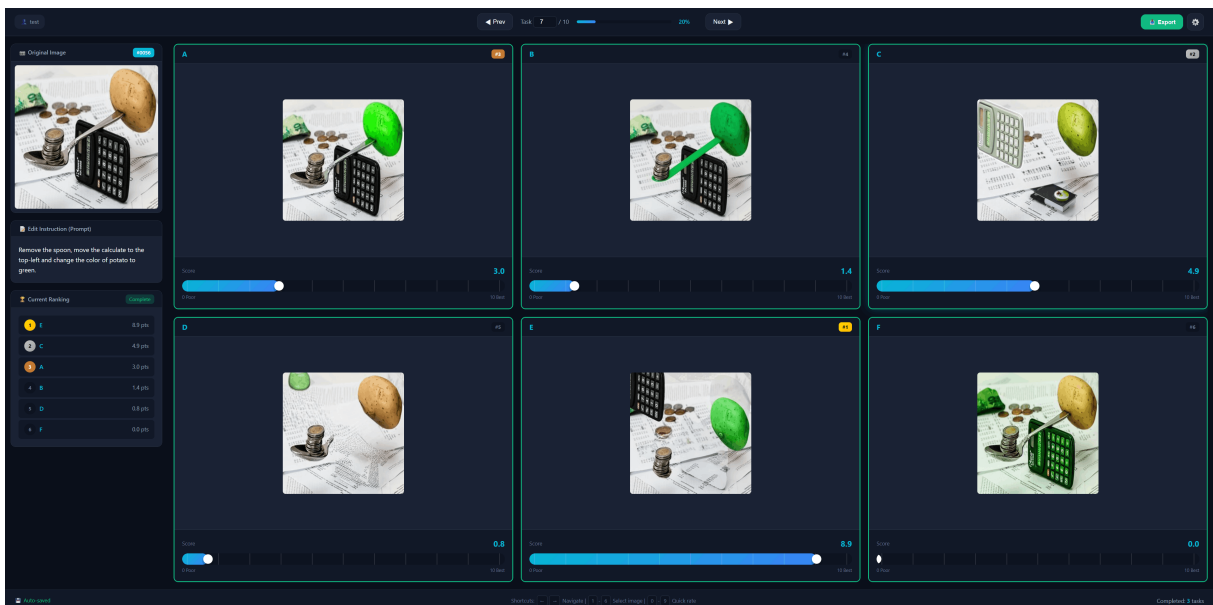


Figure 10: Example of human evaluation system.

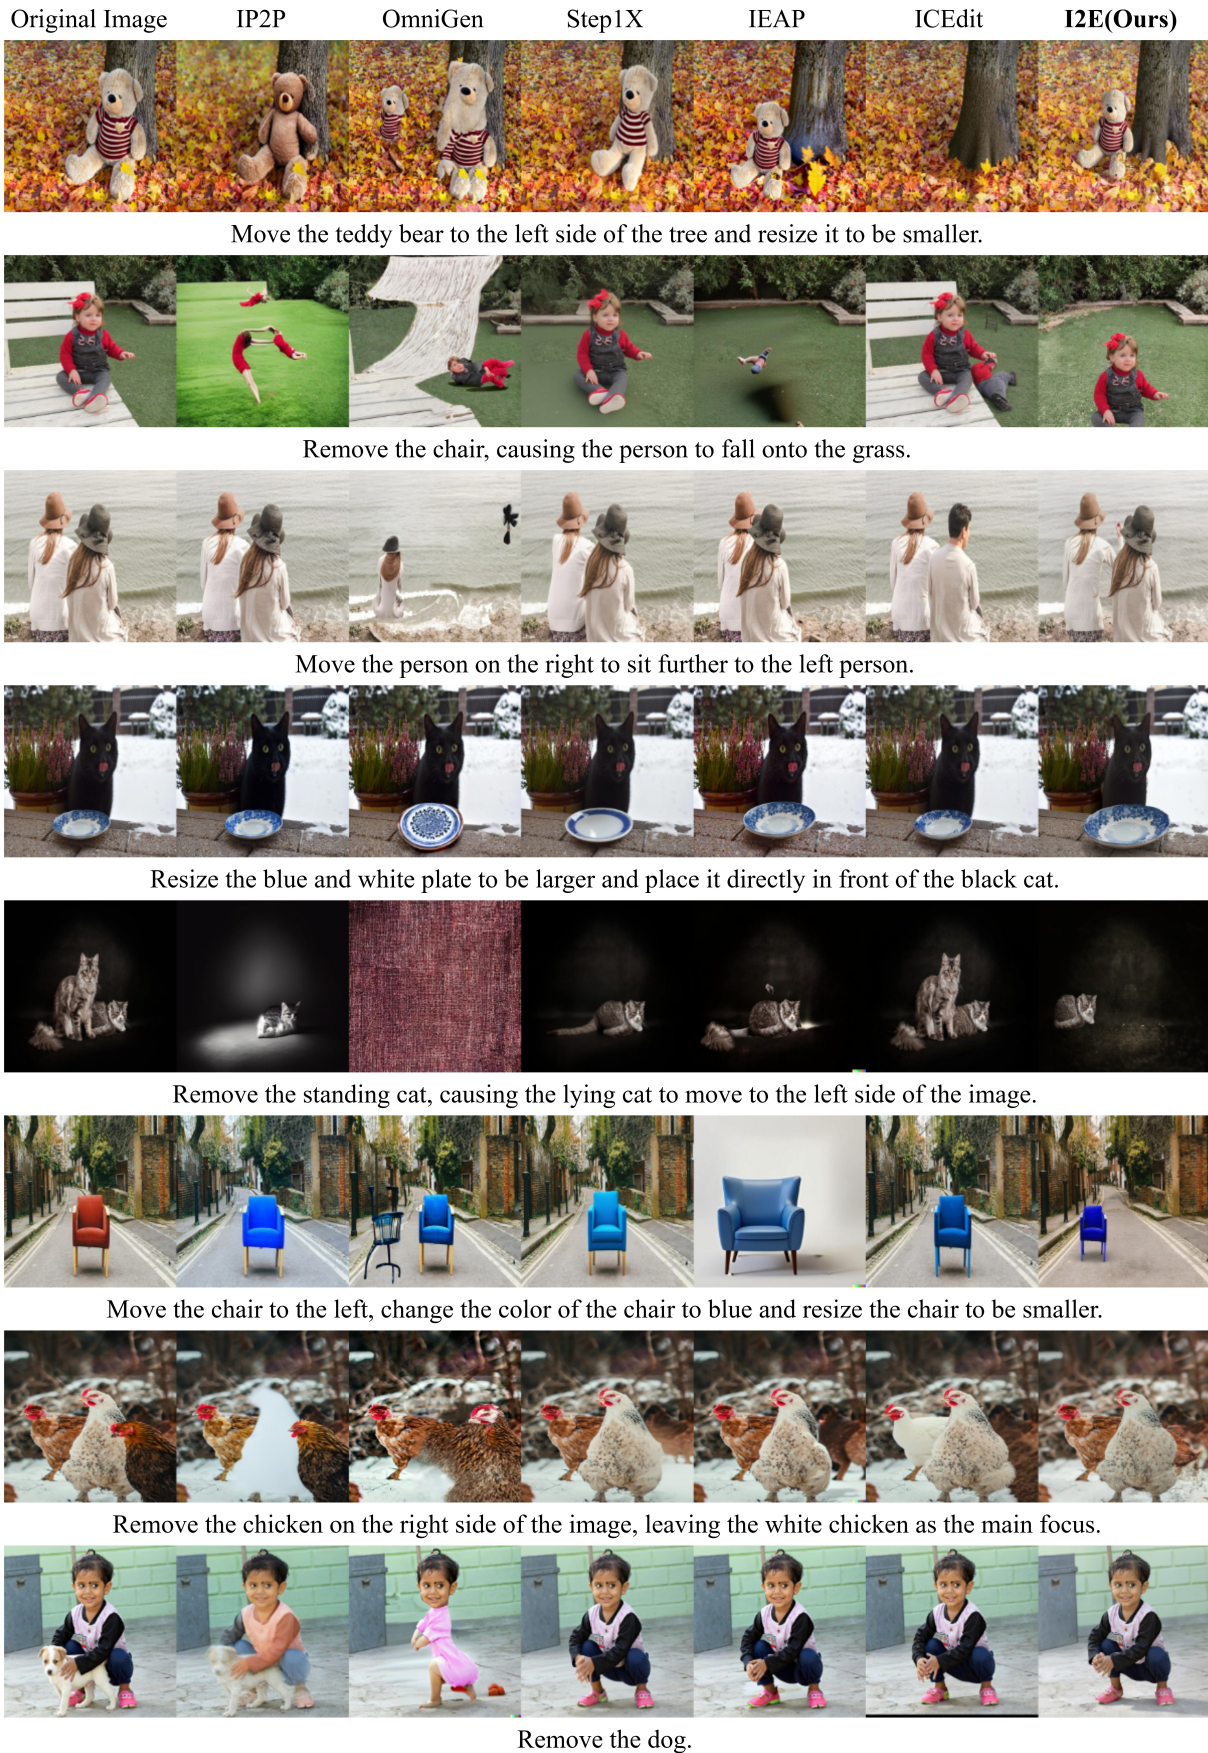


Figure 11: More qualitative results on our proposed I2E-Bench.



Figure 12: More qualitative results on our proposed I2E-Bench.




Remove the chair.




Remove the left brown horse.



Move the tiger to the right side of the image and then have the tiger cub follow it closely.



Edit the color of the dog in the middle to brown, remove the right dog.




Move the female duck from the right side to the center of the image.



Resize the cat near the grass to be larger and move it closer to the person.



Edit the color of the sofa to be dark blue, move the chair slightly to the right.



Move the seagull to the left of the starfish and resize the starfish to be slightly larger.

Figure 13: More qualitative results on our proposed I2E-Bench.

快手 Editing Effects of Real Product Images

				
<p>I2E </p>				
				
<p>Remove the label text on the top of the image</p>	<p>Swap the positions of the two objects below</p>	<p>Remove the watering device in the middle</p>	<p>Give a close-up of the camera and logo of the mobile phone below</p>	<p>Move the middle object to the top-right corner</p>

Figure 14: Editing Effects of Real Product Images from Kuaishou Mall.