Reasoning to Edit: Hypothetical Instruction-Based Image Editing with Visual Reasoning

Qingdong He [*1] **Xueqin Chen** [*2] **Chaoyi Wang** [3] **Yanjie Pan** [4] **Xiaobin Hu** [5] **Zhenye Gan** [1] **Chengjie Wang** [1] **Xiangtai Li** [6] **Jiangning Zhang** [1] **Yabiao Wang** [7 1]

Abstract

Instruction-based image editing (IIE) has advanced rapidly with the success of diffusion models. However, existing efforts primarily focus on simple and explicit instructions to execute editing operations such as adding, deleting, moving, or swapping objects. They struggle to handle more complex implicit hypothetical instructions that require deeper reasoning to infer plausible visual changes and user intent. Additionally, current datasets provide limited support for training and evaluating reasoning-aware editing capabilities. Architecturally, these methods also lack mechanisms for fine-grained detail extraction that support such reasoning. To address these limitations, we propose **Reason50K**, a large-scale dataset specifically curated for training and evaluating hypothetical instruction–reasoning image editing, along with **ReasonBrain**, a novel framework designed to reason over and execute implicit hypothetical instructions across diverse scenarios. Reason50K includes over 50K samples spanning four key reasoning scenarios: Physical, Temporal, Causal, and Story reasoning. ReasonBrain leverages a Multimodal Large Language Model (MLLM) for editing guidance generation and a diffusion model for image synthesis, incorporating a Fine-grained Reasoning Cue Extraction (FRCE) module to capture detailed visual and textual semantics essential for supporting instruction reasoning. To mitigate the semantic loss, we further introduce a Cross-Modal Enhancer (CME) that enables rich interactions between the fine-grained cues and MLLM-derived features. Extensive experiments demonstrate that ReasonBrain consistently outperforms state-of-the-art baselines on reasoning scenarios while exhibiting strong zero-shot generalization to conventional IIE tasks. Code is released.

[*]Equal contribution [1]Tencent Youtu Lab, China [2]Sichuan University, China [3]University of the Chinese Academy of Sciences, China [4]Fudan University, China [5]National University of Singapore, Singapore [6]Nanyang Technological University, Singapore [7]Zhejiang University, China. Correspondence to: Xueqin Chen <xqchen728@scu.edu.cn>, Yabiao Wang <yabiaowang@zju.edu.cn>.

Proceedings of the 43rd International Conference on Machine Learning, Seoul, South Korea. PMLR 306, 2026. Copyright 2026 by the author(s).

1. Introduction

The successful deployment of diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020) in text-to-image (T2I) generation (Ramesh et al., 2022; Rombach et al., 2022; Saharia et al., 2022; He et al., 2024) has significantly accelerated the development of *instruction-based image editing* (IIE) (Nguyen et al., 2024b). IIE focuses on performing precise and localized modifications on a source image in response to human commands, thereby enhancing the controllability and accessibility of visual manipulation (Fu et al., 2024). Prior works typically utilize the CLIP-based text encoder inherent in T2I diffusion models for instruction embedding (Brooks et al., 2023; Zhang et al., 2023; Hui et al., 2024; Zhao et al., 2024; Geng et al., 2024; Zhang et al., 2024), which is insufficient for comprehending complex instructions. To address this limitation, recent methods (Huang et al., 2024; Nguyen et al., 2024a; Wang et al., 2024b; Li et al., 2024; Zhou et al., 2025; Sun et al., 2025; Tian et al., 2025) have proposed substituting it with multimodal large language models (MLLMs) (Touvron et al., 2023; Liu et al., 2023), enabling richer cross-modal understanding and better alignment with user intent.

Despite the significant progress achieved by these efforts, the following limitations remain underexplored: (*L1*) **Overlooking the ambiguous nature of user intent.** Existing IIE methods are primarily designed to handle simple, explicit, and goal-directed instructions (e.g., "remove the dog"), which typically correspond to straightforward atomic editing operations such as adding, replacing, or deleting objects. However, users may not always hold a clear editing objective at the outset and instead express ambiguous intent through a hypothetical instruction (e.g., "What would happen if the ice cube was left in the sun?"). As illustrated in Fig. 1, current efforts struggle to interpret and act on such inputs. Success-

Figure 1. Current efforts fail to handle hypothetical instructions, producing incorrect results, while our ReasonBrain generates plausible, reasoning-aware edits.

fully addressing these cases requires models to go beyond surface-level edits based on atomic operations and instead perform deeper reasoning over real-world context, physical changes, and the causal or temporal implications of the instruction, to produce more complex and coherent edits that cannot be achieved by merely executing simple operations such as adding, deleting, or swapping objects. Moreover, currently there is no dataset specifically designed for hypothetical instruction-based editing that offers sufficient scale and scenario diversity (Huang et al., 2024; Zeng et al., 2025; Jin et al., 2024; Meng et al., 2024). (*L2*) **Insufficient reasoning ability.** While integrating MLLMs (Touvron et al., 2023; Liu et al., 2023) can enhance instruction comprehension and improve alignment with user commands to some extent, these models still lack dedicated mechanisms for deep reasoning over hypothetical instructions (Fig. 1–SmartEdit). We attribute this limitation to the prevailing paradigm's reliance on coarse-grained features extracted directly from the input image and instruction, which fail to capture the fine-grained semantic cues necessary for implicit reasoning or fully exploit the MLLM's embedded world knowledge (Wang et al., 2024a). For instance (Fig. 1(a)), reasoning over such an instruction requires jointly interpreting both visual and textual cues. Visually, elements like the cube's sharp edges, surface gloss, and surrounding environment reveal its physical state and context. Textually, features such as conditional phrasing, object references, and temporal expressions together suggest a melting process. By integrating these cues, the model should infer the implied physical transformation and determine coherent editing operations, rather than simply applying known atomic edits, to generate a plausible edited image.

To address these limitations, we first formalize the problem as a novel **Hypothetical Instruction-Reasoning Image Editing (HI-IE)** task, and propose a unified solution: a large-scale hypothetical instruction-based dataset, **Reason50K**, and a tailored reasoning-aware framework, **ReasonBrain**. Reason50K comprises diverse hypothetical instructions spanning four reasoning categories: physical, temporal, causal, and story-based reasoning, totaling 51,039 samples. Each sample consists of a source image, a hypothetical instruction, and a corresponding target image that reflects the intended edit. ReasonBrain is a hybrid framework that jointly and interactively performs reasoning and editing, thereby overcoming the need for multi-round refinement of textual instructions (Fu et al., 2024). This unified design mitigates reasoning uncertainty and reduces inference time. Specifically, ReasonBrain consists of an MLLM and a diffusion model, augmented with two specialized modules for visual guidance, reasoning, and semantic enrichment, i.e, the Fine-grained Reasoning Cue Extraction (FRCE) module and the Cross-Modal Enhancer (CME). The FRCE module extracts detailed reasoning cues through two branches: the visual reasoning branch captures both local and global visual semantics to model spatial relationships and object-level interactions, while the textual reasoning branch identifies key object references and contextual intent from hypothetical instructions, enriched with relevant visual context. These fine-grained features, combined with multi-scale image tokens and textual embeddings, are input to the MLLM alongside learnable tokens to implicitly generate reasoning-aware visual guidance. Finally, the CME enhances these signals via semantic complementarity across modalities, producing well-aligned, semantically rich guidance for diffusion-based image editing. In sum, our contributions are as follows:

- We systematically extend traditional Instruction-Based Image Editing (IIE) to **Hypothetical Instruction-Reasoning Image Editing (HI-IE)**. This task involves implicit, ambiguous, and hypothetical editing instructions that demand deeper reasoning over contextual cues, physical dynamics, and user intent.
- We curate a new large-scale dataset, **Reason50K**, specifically designed to support the reasoning of hy-

pothetical instruction in image editing. It contains 51,039 triplets of source image, hypothetical instruction, and target image, covering four distinct reasoning categories: physical, temporal, causal, and story-based reasoning.

- We propose **ReasonBrain**, a novel image editing framework that combines a MLLM with fine-grained reasoning cues extraction and a cross-modal enhancer. Together, these components endow the model with implicit cross-modal reasoning capabilities, enabling it to infer plausible, knowledge-grounded transformations and produce semantically coherent guidance for diffusion-based image editing under complex hypothetical scenarios.
- We conduct extensive experiments on both Reason50K and widely used benchmark datasets, demonstrating the effectiveness and generalization ability of ReasonBrain across reasoning-intensive and standard understanding-based editing scenarios.

Conflict of Interest Disclosure. The authors declare that they have no financial conflicts of interest related to this work.

2. Related Work

Instruction-based Image Editing (IIE). IIE (Brooks et al., 2023) aims to train generative models to manipulate a given image based on user-provided instructions. A milestone in this field is InstructPix2Pix (IP2P) (Brooks et al., 2023), which is the first to incorporate natural language instructions into image editing by fine-tuning a text-to-image (T2I) diffusion model on paired image-instruction datasets. Subsequent works have built upon IP2P by introducing novel curated datasets to enhance real-world editing performance and generate high-quality outputs (Zhang et al., 2023; Hui et al., 2024; Zhao et al., 2024; Liu et al., 2024; Yu et al., 2025). Others have focused on improving instruction-output alignment by incorporating advanced techniques such as reward learning (Zhang et al., 2024; Bai et al., 2024) and multi-task training (Sheynin et al., 2024). Recently, researchers have integrated multimodal large language models (MLLMs) (Touvron et al., 2023; Liu et al., 2023) into existing image editing paradigms to enhance the model's ability to understand complex instructions. We refer to these approaches as MLLM-enhanced methods (Fu et al., 2024; Huang et al., 2024; Li et al., 2024; Wang et al., 2024b; Zhou et al., 2025; Tian et al., 2025). For instance, MGIE (Fu et al., 2024) leverages MLLMs to generate expressive instructions and provide explicit guidance, thereby enhancing editing performance. SmartEdit (Huang et al., 2024) further introduces a bidirectional interaction module that facilitates mutual understanding between the MLLM output and the input image. Despite recent advances, most existing efforts still rely on direct and explicit instructions (e.g., "Remove the ice"), limiting their ability to handle hypothetical instructions (e.g., "What would happen if the ice cube melted?") that require deeper reasoning. While MLLMs offer general world knowledge, current frameworks lack mechanisms to extract and utilize fine-grained reasoning details. Our **ReasonBrain** builds upon the MLLM-enhanced paradigm by introducing dedicated reasoning-aware modules for generating precise visual guidance, with an emphasis on accurately inferring the implicit intent and real-world context embedded in hypothetical instructions.

Reasoning-aware Datasets for Image Editing. Only a few works have explored reasoning-aware datasets for image editing (Huang et al., 2024; Zeng et al., 2025; Jin et al., 2024). ReasonEdit (Huang et al., 2024) is designed primarily for evaluating the reasoning capabilities of image editing models and contains only a small set of textual samples, making it insufficient for training. Moreover, it focuses on object-level reasoning based on explicit instructions, while overlooking reasoning about editing operations themselves. EditWorld (Zeng et al., 2025) aims to inject physical dynamics simulation capabilities into models across both realistic and virtual scenarios, but does not emphasize instruction-level reasoning. ReasonPix2Pix (Jin et al., 2024) introduces indirect instructions in a descriptive, goal-oriented style but lacks the depth and diversity needed for more advanced hypothetical reasoning. Recently, a benchmark ComplexEdit (Yang et al., 2025) is introduced to assess the sequential editing capabilities of IIE models handling instructions across varying complexity. However, it still focuses on explicit instructions composed of predefined atomic operations, rather than explicitly evaluating the model's ability to reason about the underlying user intent. In contrast to these prior efforts, we curate **Reason50K**, a large-scale dataset specifically tailored for hypothetical instruction reasoning, enabling models to understand and execute complex edits grounded in physical, temporal, causal, and story scenarios.

3. Methodology

The goal of our work is to perform image editing by reasoning from a user-provided *hypothetical instruction*–an implicit, often ambiguous prompt that requires the model to infer the intended transformation through deeper reasoning about real-world context, physical dynamics, or potential outcomes (e.g., "What would happen if the sun were setting?"). We refer to this task as **Hypothetical Instruction-Reasoning Image Editing (HI-IE)**. To this end, we propose a unified solution comprising a large-scale dataset, **Reason50K**, specifically curated to support the challenging task of hypothetical HI-IE, and a reasoning-aware image editing framework, **ReasonBrain**.

3.1. Reason50K for Reasoning Injecting

To support our proposed HI-IE task, we construct a novel large-scale dataset Reason50K specifically curated to inject *hypothetical instruction-reasoning ability* into image editing models. Reason50K contains 51,039 samples, each consisting of an input image, a corresponding hypothetical instruction, and a target edited image. Unlike existing datasets that primarily feature explicit or goal-oriented prompts (e.g., "Remove the animal in the mirror" or "Make it snowy"), our instructions are implicit, open-ended, and reasoning-driven (e.g., "What would happen if the ice cube were left at room temperature?" or "What would this bouquet look like if it were split into two separate rose bouquets?"). This shift introduces a significantly higher level of abstraction and demands real-world understanding. Reason50K is constructed using an inverse-style strategy, where the source image is generated from the target. Specifically, we leverage GPT (Achiam et al., 2023) to produce hypothetical instructions based on user-provided prompts, and employ a diffusion model to generate multiple candidate source images. A hybrid scoring and evaluation scheme is then applied to select the most appropriate image, forming the final image pairs. Due to space limitations, the full data curation process is detailed in App. C. Moreover, Reason50K is organized around four distinct reasoning scenarios: **Physical**, **Temporal**, **Causal**, and **Story** Reasoning, each illustrated with representative examples and instructions in Fig. 2.

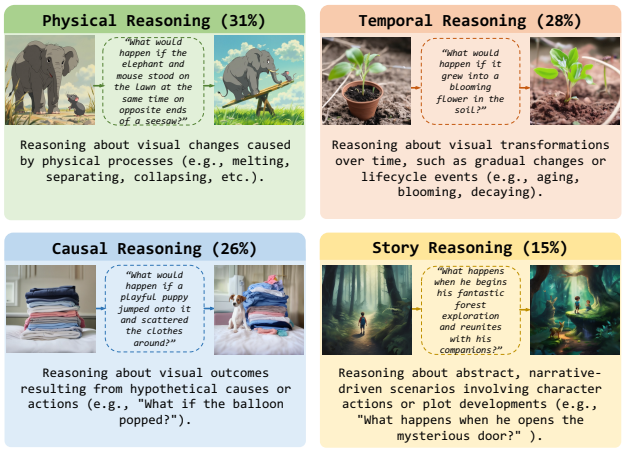

Figure 2. Reasoning scenarios in Reason50K. The percentages in parentheses indicate the proportion of each category. The text below each sample shows an instruction of the corresponding reasoning type.

Reason50K vs. Existing Datasets. Tab. 1 summarizes the key differences between our dataset and existing ones. In comparison: 1) *ReasonPix2Pix* (Jin et al., 2024) enhances instruction reasoning by using LLMs to rewrite explicit editing commands into goal-oriented implicit instructions. However, the dataset primarily consists of descriptive instructions and lacks more complex, abstract hypothetical instructions that require deeper contextual and causal reasoning. In addition, it lacks systematic categorization. 2) *EditWorld* (Zeng et al., 2025) focuses on simulating world dynamics across both real and virtual scenarios. Although it includes some hypothetical instructions, its primary objective is to inject physical and temporal simulation capabilities into image editing models, rather than equipping them with implicit reasoning skills for understanding hypothetical instructions. Additionally, the overall scale of the EditWorld dataset is significantly smaller than ours. 3) *ReasonEdit* (Huang et al., 2024) is a small-scale dataset designed primarily for evaluation. It focuses on object-level reasoning from explicit instructions and lacks both diversity and editing operation reasoning. Most importantly, it is not sufficient for model training. 4) *Complex-Edit* (Yang et al., 2025) is a recent benchmark also curated for evaluation, with a particular focus on varying levels of instruction complexity. However, it mainly uses a chain-of-thought-like paradigm to construct the dataset and assess sequential editing, rather than explicit instruction reasoning. In contrast, our **Reason50K** is the first to provide *systematic, large-scale support* for training and evaluating *hypothetical instruction reasoning* across *diverse scenarios*. Each instruction requires deep reasoning grounded in contextual understanding and world knowledge—spanning physical, causal, temporal, and story-based scenarios—to guide image editing. This goes beyond surface-level transformations or explicit object manipulation. By incorporating carefully crafted hypothetical instructions, the dataset significantly promotes deeper semantic reasoning within image editing models.

Table 1. Comparison of different datasets. Note that since ReasonPix2Pix (Jin et al., 2024) is not publicly available, we reference sample cases from their paper and exclude it from further experimental comparison (see Fig. 5).

Datasets	ReasonPix2Pix	EditWorld	ReasonEdit	Complex-Edit	Reason50K (Ours)
Reasoning Instruction	✓	✗	✗	✗	✓
Automatic Generated	✓	✓	✓	✓	✓
Open Domain	✗	✓	✗	✓	✓
#Edits	40,212	8,674	219	1,062	51,039
Source					
Instruction	*A colorful insect has landed*	*Shifted her gaze to the left.*	*Please remove the empty plate.*	*Replace the palm trees with cheery blossom trees* ... *Add a car to the road beneath the bridge to complete the setting.*	*What would happen if the cat stood in front of a mirror?*
Target					

It is noted that our dataset Reason50K is constructed from synthetic data, as it is difficult to obtain image pairs from videos that support hypothetical editing with clear reasoning structures. In addition, there is currently no standardized benchmark or evaluation protocol for the frame-by-frame assessment of reasoning-based edits. Nevertheless, the generation process of Reason50K is carefully designed to ensure both semantic consistency and high visual quality, thereby enabling the synthetic dataset to effectively support generalizable research.

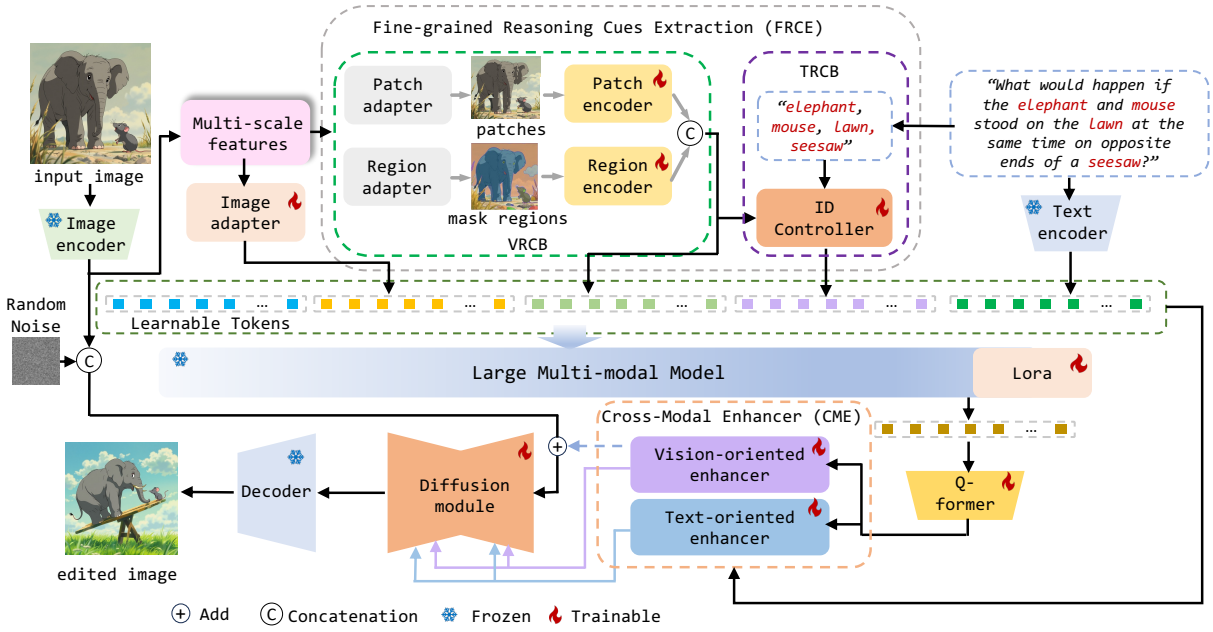

Figure 3. The overall framework of ReasonBrain. Given an input image I and a hypothetical instruction H, ReasonBrain first encodes them into multi-scale visual features and textual tokens using the image encoder $\mathcal{E}_I(\cdot)$ and text encoder $\mathcal{E}_T(\cdot)$, respectively. These features are then passed into the FRCE module to learn detailed reasoning cues. Subsequently, all learned features are fed into the MLLM to generate visual guidance, which is further transformed via a QFormer to align with the diffusion model's latent space. Finally, the resulting visual guidance interacts with the previously extracted fine-grained cues through a CME module to enhance semantic representation, which is then used to condition the diffusion model for final image generation.

3.2. ReasonBrain

ReasonBrain comprises an MLLM for visual guidance, reasoning, and a diffusion model responsible for conditional image generation. To support knowledge reasoning, we incorporate a Fine-Grained Reasoning Cues Extraction (FRCE) module. In addition, we introduce a Cross-Modal Enhancer (CME) to further enrich semantic representations through modality-specific refinement. The overall framework of ReasonBrain is illustrated in Fig. 3.

Fine-grained Reasoning Cues Extraction. Existing efforts utilize $\mathcal{E}_I(I)$ and $\mathcal{E}_T(H)$ directly for visual guidance generation, overlooking fine-grained cues critical for implicit reasoning, such as local object attributes, subtle spatial relationships, and context-dependent semantics between image regions and instructions, etc. To address this limitation, we introduce a Fine-Grained Reasoning Cues Extraction (FRCE) module, which comprises two specialized branches designed to capture fine-grained visual and textual reasoning cues, respectively.

(1) Visual Reasoning Cues Branch (VRCB). This branch aims to extract fine-grained visual cues from both local and global perspectives. Specifically, the *local perspective* focuses on capturing object parts, textures, and spatial patterns, which are essential for implicit reasoning tasks involving fine object distinctions and subtle appearance changes. Inspired by the MAE framework (He et al., 2022), we first divide the visual features $\mathcal{E}_I(I)$ into patches using a patch adapter $\mathcal{P}(\cdot)$, and then apply a patch-level feature extractor $E_{\mathcal{P}}(\cdot)$ to obtain localized visual representations. This process is formalized as: $R_{local} = E_{\mathcal{P}}(\mathcal{P}(\mathcal{E}_I(I)))$. In contrast, the *global perspective* captures inter-object relationships and contextual information across the entire scene, which are beneficial for reasoning tasks that require an understanding of object interactions, event causality, and broader scene dynamics. We first employ an instance segmentation model (e.g., SAM (Kirillov et al., 2023; Ravi et al., 2024)) to segment objects from the background in the image. Subsequently, a region-level feature extractor $E_{\mathcal{R}}(\cdot)$ is trained to learn holistic semantic representations for each segmented instance. The entire process is formalized as:

$$R_{global} = E_{\mathcal{R}}(\text{SAM}(\mathcal{E}_I(I))). \tag{1}$$

After this dual-level operation, we concatenate R_{local} and R_{global} to form the final visual reasoning features R_V used for subsequent processing.

(2) Textual Reasoning Cues Branch (TRCB). This branch aims to extract the key object referenced in H, serving as a bridge between linguistic intent and visual reasoning. Specifically, we first employ GPT (Achiam et al., 2023) to extract the referenced object from the instruction and serialize it into a structured object token O. We then introduce an *ID Controller* to enhance the model's ability to perform object-grounded reasoning by facilitating interaction between the object token and the visual reasoning features R_V. This module not only enriches the object to-

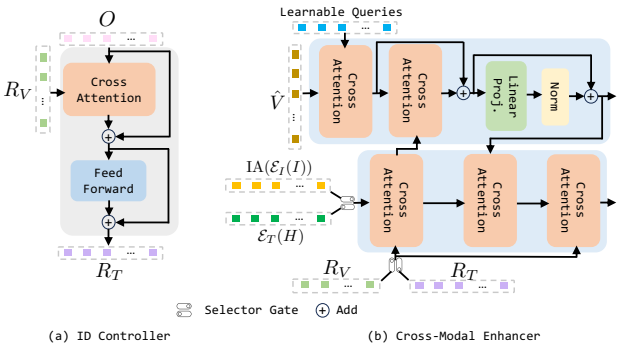

Figure 4. Network design of the (a) ID Controller and (b) Cross-Modal Enhancer.

ken with visual context, enabling the model to reason about the object beyond its textual description, but also aligns the linguistic reference with its corresponding visual entity, helping to resolve ambiguities and prevent semantic drift during generation. The architecture of the ID Controller is illustrated in Fig. 4(a) and is implemented using a cross-attention layer (Lin et al., 2022) followed by a feed-forward network, formally defined as:

$$R_T = \text{FF}(\text{Cross-Atten}(R_V, O)). \tag{2}$$

Subsequently, following prior works (Huang et al., 2024; Zhou et al., 2025), we concatenate r additional learnable tokens $\mathcal{Q} = \{[\text{IMG}_1], \ldots, [\text{IMG}_r]\}$ with the extracted features and feed them into the MLLM for guidance generation. This operation transforms the implicit reasoning process into a token prediction task, where the MLLM learns to generate the embeddings of these r tokens, which serve as visual editing guidance. The process is formalized as:

$$\mathcal{T} = [\text{IA}(\mathcal{E}_I(I)), R_V, R_T, \mathcal{E}_T(H), \mathcal{Q}], V = \text{MLLM}(\mathcal{T}; \theta). \tag{3}$$

Here, $[\cdot, \cdot]$ denotes concatenation, and $\text{IA}(\cdot)$ is a trainable image adapter that maps visual features into the MLLM's latent space. The output $V \in \mathbb{R}^{r \times d}$ represents the hidden embeddings of the r learnable tokens. In addition, we employ a QFormer (Li et al., 2023) to align the feature space between the MLLM and the diffusion model, defined as $\hat{V} = \text{QFormer}(V)$.

Cross-Modal Enhancer. To compensate for the potential loss of visual and textual details in $\hat{V}$, we introduce a Cross-Modal Enhancer (CME). The CME consists of a *visual-oriented enhancer* and a *textual-oriented enhancer*, both implemented using the same bidirectional interaction mechanism. Specifically, each enhancer comprises five hybrid cross-attention blocks, followed by a linear projection layer and a normalization layer (see Fig. 4(b)). These modules reuse intermediate features generated both before and after the MLLM layer, enabling fine-grained semantic enhancement within each modality. Here, we illustrate the process using the visual-oriented enhancer. First, a set of learnable tokens Q[1] interacts with $\hat{V}$ (as key and value) via a cross-attention block to produce F_1. In parallel, $\mathcal{E}_I(I)$ (query) attends to R_V in a second cross-attention block, yielding F_2. Next, F_1 (query) and F_2 (key and value) are fused via a third block, followed by a residual connection and normalization to obtain $\bar{V}$. This refined representation then serves as key and value in a fourth block, interacting with F_2 (query). The output is further refined through a final cross-attention with R_V, producing the enhanced visual representation $\bar{R}_V$. Similarly, the textual-oriented enhancer is implemented by replacing $\mathcal{E}_I(I)$ and R_V with $\mathcal{E}_T(H)$ and R_T, respectively. Finally, the CME module outputs four enhanced features, which are passed to the diffusion model together with $\mathcal{E}_I(I)$ and a noisy latent z for final image generation. Due to space limitations, the training objectives of ReasonBrain are provided in App. D.

[1]The newly introduced symbols are used only in this section to illustrate the visual-oriented enhancer.

4. Experiments

4.1. Quantitative Results

Performance Comparison on Reasoning Scenarios. The experimental results on Reason50K are summarized in Tab. 2, covering four distinct types of reasoning scenarios. ReasonBrain consistently outperforms all SOTA baselines across almost all metrics, demonstrating superior ability to infer hypothetical instructions and produce logically accurate edits aligned with world knowledge. While UltraEdit and PixWizard achieve competitive performance, their advantage mainly stems from exposure to massive datasets with diverse instructions rather than true reasoning capabilities. Earlier MLLM-enhanced methods (MGIE and SmartEdit) still underperform, suggesting that simply integrating MLLMs or fine-tuning on reasoning datasets is insufficient for HI-IE. Latest works, such as MIGE and SmartFreeEdit, improve semantic consistency and achieve substantially better performance, they still lag behind our ReasonBrain due to the absence of fine-grained reasoning cues. Although unified foundation models (Bagel–GPT Image 1.5) that are trained on large-scale corpora and equipped with more parameters, they still do not achieve satisfactory performance. This indicates that scaling laws alone cannot bridge the reasoning gap without dedicated architectural modules and training objectives. Additionally, we select five representative methods for evaluation on ReasonEdit, EditWorld, and Complex-Edit. As shown in Fig. 5, ReasonBrain, trained solely on our Reason50K, generalizes effectively to novel reasoning scenarios and achieves the best overall performance across all datasets. We also report the identity preservation metric (IP) (Yang et al., 2025), on which our ReasonBrain achieves the highest score, indicating that it more effectively retains core elements from the source image while following the instruction. These demonstrate that our model trained on Reason50K is actually performing reasoning and not just learning a shortcut.

Table 2. Results on the Reason50K dataset for ReasonBrain and selected baselines. ↓ indicates that lower values are better, while ↑ indicates that higher values are better. The best results are highlighted in **bold**, and the second-best results are underlined.

Method	Physical Reasoning			Temporal Reasoning			Causal Reasoning			Story Reasoning			Total			
	CLIP↑	MLLM↑	Ins-Align↑	CLIP↑	MLLM↑	Ins-Align↑	CLIP↑	MLLM↑	Ins-Align↑	CLIP↑	MLLM↑	Ins-Align↑	CLIP↑	MLLM↑	Ins-Align↑	IP↑
InstructPix2Pix (Brooks et al., 2023)	0.083	0.825	0.211	0.207	0.846	0.678	0.153	0.785	0.191	0.196	0.628	0.225	0.160	0.771	0.326	5.13
MagicBrush (Zhang et al., 2023)	0.102	0.844	0.335	0.228	0.877	0.725	0.162	0.802	0.344	0.209	0.635	0.359	0.175	0.790	0.441	7.52
MGIE (Fu et al., 2024)	0.098	0.802	0.288	0.213	0.832	0.685	0.155	0.761	0.328	0.201	0.622	0.322	0.167	0.754	0.406	7.63
SmartEdit (Huang et al., 2024)	0.118	0.849	0.602	0.226	0.881	0.779	0.165	0.823	0.385	0.211	0.655	0.361	0.180	0.802	0.532	8.02
UltraEdit (Zhao et al., 2024)	0.156	0.861	0.584	0.231	0.922	0.826	0.193	0.869	0.482	0.209	0.669	0.458	0.197	0.830	0.588	5.68
PixWizard (Lin et al., 2025)	0.161	0.881	0.481	0.234	0.951	0.833	0.132	0.863	0.393	0.172	0.697	0.389	0.175	0.848	0.524	8.55
MIGE (Tian et al., 2025)	0.159	0.853	0.612	0.253	0.917	0.881	0.178	0.851	0.465	0.216	0.673	0.447	0.202	0.826	0.612	8.64
SmartFreeEdit (Sun et al., 2025)	0.169	0.869	0.659	**0.268**	0.934	**0.949**	0.189	0.866	0.501	0.229	0.685	0.482	0.214	0.841	0.659	8.81
Bagel (Deng et al., 2025)	0.136	0.804	0.482	0.216	0.861	0.694	0.152	0.804	0.366	0.184	0.636	0.352	0.172	0.781	0.482	7.85
Bagel-Thinking (Deng et al., 2025)	0.146	0.819	0.514	0.232	0.876	0.740	0.163	0.819	0.391	0.198	0.648	0.376	0.185	0.795	0.514	8.08
Flux Kontext (Labs et al., 2025)	0.152	0.836	0.537	0.241	0.893	0.773	0.169	0.835	0.409	0.206	0.662	0.393	0.192	0.812	0.537	8.22
Qwen-Image-Edit (Wu et al., 2025)	0.139	0.805	0.495	0.221	0.862	0.713	0.155	0.805	0.376	0.189	0.637	0.362	0.176	0.782	0.495	7.92
Qwen-Image-Edit (Finetuned) (Wu et al., 2025)	0.159	0.851	0.568	0.252	0.915	0.818	0.178	0.849	0.432	0.215	0.672	0.415	0.201	0.824	0.568	8.38
ChatGPT-4o	0.144	0.819	0.524	0.228	0.876	0.755	0.161	0.819	0.398	0.195	0.648	0.383	0.182	0.795	0.524	8.15
GPT Image 1.5	0.162	0.858	0.612	0.257	0.922	0.881	0.181	0.856	0.465	0.220	0.677	0.447	0.205	0.831	0.612	8.67
ReasonBrain (ours)	**0.186**	**0.902**	**0.846**	0.267	**0.977**	0.894	**0.282**	**0.891**	**0.858**	**0.301**	**0.736**	**0.798**	**0.259**	**0.877**	**0.847**	**9.72**

Table 3. Results on Emu Edit and MagicBrush test set for ReasonBrain and selected baselines.

Method	Emu Edit Test set					MagicBrush Test Set				
	$CLIP_{dir}$↑	$CLIP_{im}$↑	$CLIP_{out}$↑	L1↓	DINO↑	$CLIP_{dir}$↑	$CLIP_{im}$↑	$CLIP_{out}$↑	L1↓	DINO↑
InstructPix2Pix (Brooks et al., 2023)	0.078	0.834	0.219	0.121	0.762	0.115	0.837	0.245	0.093	0.767
MagicBrush (Zhang et al., 2023)	0.090	0.838	0.222	0.100	0.776	0.123	0.883	0.261	0.058	0.871
MGIE (Fu et al., 2024)	0.083	0.746	0.231	0.163	0.594	0.116	0.745	0.251	0.162	0.577
SmartEdit (Huang et al., 2024)	0.092	0.858	0.274	0.119	0.771	0.119	0.895	0.262	0.094	0.820
Emu Edit (Sheynin et al., 2024)	0.109	0.859	0.231	0.094	0.819	0.135	0.897	0.261	0.052	0.879
UltraEdit (Zhao et al., 2024)	0.107	0.844	0.283	0.071	0.793	-	0.868	-	0.088	0.792
PixWizard (Lin et al., 2025)	0.104	0.845	0.248	0.069	0.798	0.124	0.884	0.265	0.063	0.876
MIGE (Tian et al., 2025)	0.112	0.862	0.255	0.082	0.825	0.131	0.902	0.268	0.056	0.882
SmartFreeEdit (Sun et al., 2025)	0.115	0.868	0.261	0.078	0.832	0.133	0.905	0.270	0.054	0.887
Bagel (Deng et al., 2025)	0.101	0.842	0.240	0.075	0.790	0.121	0.882	0.260	0.065	0.868
Bagel-Thinking (Deng et al., 2025)	0.103	0.846	0.243	0.073	0.794	0.124	0.885	0.262	0.062	0.872
Flux Kontext (Labs et al., 2025)	0.105	0.850	0.246	0.072	0.797	0.126	0.887	0.263	0.060	0.875
Qwen-Image-Edit (Wu et al., 2025)	0.111	0.860	0.253	0.068	0.822	0.130	0.890	0.264	0.057	0.881
ChatGPT-4o	0.102	0.843	0.242	0.074	0.788	0.122	0.880	0.259	0.064	0.865
GPT Image 1.5	0.110	0.861	0.254	0.067	0.806	0.132	0.892	0.266	0.055	0.885
ReasonBrain (ours)	**0.126**	**0.923**	**0.302**	**0.051**	**0.898**	**0.139**	**0.928**	**0.281**	**0.049**	**0.893**

Figure 5. Results on ReasonEdit, EditWorld, and Reason50K for ReasonBrain and selected SOTA methods, highlighting performance under various reasoning datasets.

Generalization Comparison on Understanding Scenarios. As shown in Tab. 3, ReasonBrain achieves the best overall performance among all SOTA methods, demonstrating strong zero-shot generalization to new datasets. This result also highlights that, although trained solely on hypothetical instructions, ReasonBrain retains a robust understanding of clear, goal-directed editing commands. Furthermore, the performance of those IIE-trained methods suggests that scaling with large and diverse datasets can further enhance effectiveness on conventional instruction-based image editing tasks. In contrast, unified models not specifically optimized for IIE generally achieve inferior performance.

4.2. Qualitative Results

As shown in Fig. 6, we visualize editing results of ReasonBrain and selected SOTA methods across four different reasoning scenarios. Our method demonstrates superior capability in executing hypothetical editing instructions by accurately reasoning about user intent, affected objects, and their plausible state transitions, while also maintaining stability in non-edited regions (original scene/object identity (ID)). For instance, in the first row (physical reasoning), our method successfully generates a physically plausible and contextually coherent scene–depicting the elephant and mouse standing on opposite ends of a seesaw, where the heavier elephant naturally tilts the seesaw downward. Additionally, our ReasonBrain preserves the original lawn, seesaw structure, and object proportions. This demonstrates that our model can not only reason about relative weight, spatial arrangement, and the physical dynamics implied by the instruction, but also maintain a certain degree of identity consistency. In contrast, InstructPix2Pix fails to capture the core intent, producing an irrelevant result. Other methods may include the correct objects mentioned in the instruction–such as the elephant, mouse, and seesaw–but fall short in modeling their physical interactions, resulting in unrealistic or semantically inconsistent compositions. Notably, ChatGPT-4o produces an implausible outcome in which the mouse outweighs the elephant, contradicting basic physical intuition. We attribute this failure to hallucinations introduced by the language model (Huang et al., 2025).

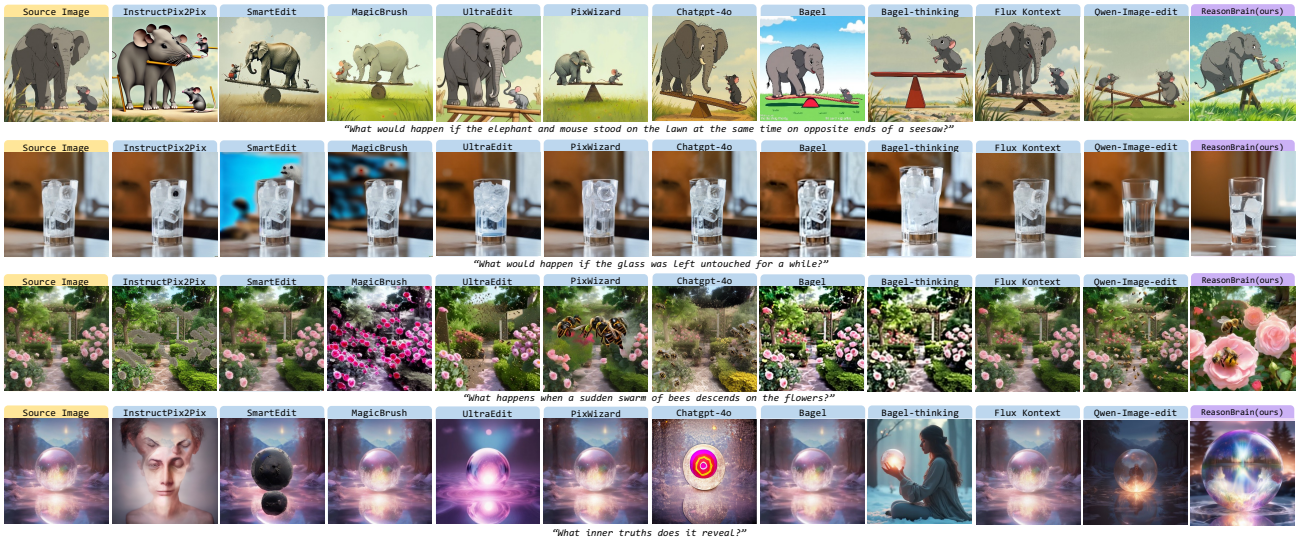

Figure 6. Qualitative comparison on Reason50K between ReasonBrain and selected SOTA methods. ReasonBrain demonstrates a strong ability to reason over implicit hypothetical instructions and produce semantically plausible edits grounded in world knowledge.

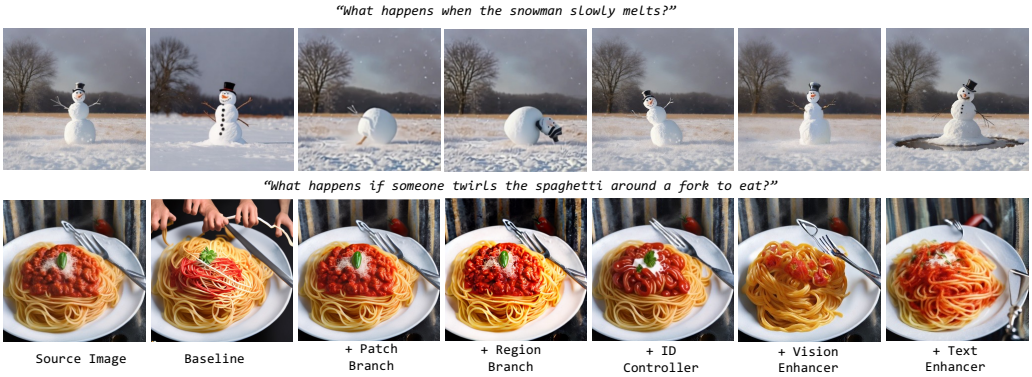

Figure 7. Qualitative comparison of ablation variants in ReasonBrain.

Moreover, although ReasonBrain may introduce slight scene adjustments compared to the source image, these changes are strictly bounded by the hypothetical instruction. They serve as logical supplements to the core modification rather than arbitrary alterations, and non-target regions (e.g., walls, lawns) remain visually identical to the original. For example, in the third row (casual reasoning), ReasonBrain adjusts the camera distance and local contrast to make the bee–flower interaction visible. This is a necessary scene–level refinement to truly convey the hypothetical event. If one were to rigidly forbid any background adaptation, the generated bees would be barely visible or visually unnatural. In contrast, baselines (e.g., InstructPix2Pix cannot model the environmental interplay between the bees and the flowers) preserve the overall scene appearance but fail to express the required interaction, illustrating that high appearance consistency does not imply correct reasoning. ReasonBrain strikes a balance between identity preservation and instruction expressiveness. Core ID elements (e.g., flower clusters, lawn layout, and non-target objects) remain intact, while subtle background refinements (such as camera distance, mild lighting, or contrast adjustments) are applied only when necessary to ensure the hypothetical change is visually coherent and semantically meaningful.

A similar pattern appears in the last row (story reasoning), where the instruction requires uncovering subtle, implicit transformations of the subject (e.g., hidden textures or symbolic details). Baselines such as InstructPix2Pix produce irrelevant outputs lacking reasoning, while SmartEdit fails to highlight implicit cues due to overly rigid background preservation. ReasonBrain, by contrast, achieves a balanced result in which the core identity (subject shape, global layout, and non-target regions) is retained, and only minimal, reasoning-driven background adjustments (e.g., localized light enhancement or slight contrast tuning) are applied to make the "inner truths" visually perceptible. These refinements are not excessive but necessary. Without them, the instruction-induced changes would remain imperceptible, undermining the purpose of implicit hypothetical editing. More qualitative results can be found in Fig. A3.

4.3. Ablation and Analysis

Effectiveness of Each Component. To evaluate the contribution of each component in ReasonBrain, we conduct

ablation experiments by progressively integrating key modules. As shown in Tab. 4 and Fig. 7, the fine-grained visual features help the model capture subtle visual cues essential for reasoning. The ID Controller significantly improves performance by preserving object identity during cross-modal alignment, which is critical for accurate instruction reasoning. Additionally, the CME module enhances overall generation quality by reinforcing modality-specific semantics and providing more detailed guidance. Furthermore, we also observe that removing any individual component leads to a performance drop across all three metrics, indicating that each module contributes meaningfully to the overall effectiveness of the proposed framework.

Table 4. Results of the ablation study on each component of ReasonBrain. P, R, ID, V, and T denote the patch branch, region branch, ID controller, vision enhancer, and text enhancer, respectively. × and ✓ indicate without/with the component.

P	R	ID	V	T	CLIP↑	MLLM↑	Ins-Align↑
×	×	×	×	×	0.163	0.752	0.388
✓	×	×	×	×	0.187	0.786	0.466
✓	✓	×	×	×	0.206	0.802	0.529
✓	✓	✓	×	×	0.239	0.833	0.758
✓	✓	✓	✓	×	0.251	0.865	0.822
✓	✓	✓	×	✓	0.248	0.845	0.785
✓	✓	×	✓	✓	0.231	0.838	0.776
✓	×	✓	✓	✓	0.246	0.847	0.788
×	✓	✓	✓	✓	0.240	0.842	0.781
✓	✓	✓	✓	✓	**0.259**	**0.877**	**0.847**

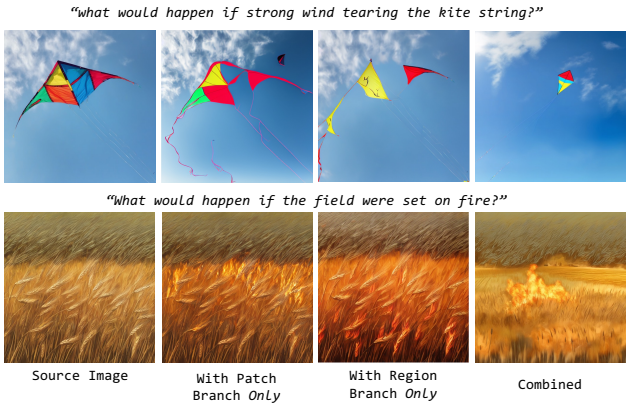

Figure 8. Qualitative comparison of patch and region branches in VRCB.

Impact of Each Visual Branch in VRCB. As shown in Fig. 8, we visualize the individual and combined contributions of the patch and region branches in our VRCB. The patch branch captures local details such as texture variations and appearance changes, which are useful for modeling subtle transformations (e.g., wind effects on kite fabric or shifting crop colors). In contrast, the region branch focuses on global semantic structures and object-level understanding, such as the overall layout of the kite or the boundary of the burning field. When used independently, each branch captures only partial aspects of the intended edit. The patch branch may introduce local distortions without preserving object integrity, while the region branch may lack detailed variation. Their combination enables complementary integration of local precision and global semantics, resulting in more coherent, semantically accurate, and visually plausible edits under hypothetical instructions.

Impact of MLLM backbone. The default MLLM backbone used in this work is LLaVA-v1.1-7B, and we replace it with various scale and architecture MLLMs and report the result in Tab. 5. The results show that although stronger MLLM backbones naturally lead to better performance, the major gains come from our proposed FRCE and CME modules, as these two decoupled modules address the key bottlenecks of MLLMs in HI-IE task, i.e., lack of fine-grained reasoning cues and insufficient cross-modal interaction. More quantitative analysis can be found in App. F and App. G.

Table 5. Ablation results with different MLLM backbones and model configurations.

MLLM Backbone	Model Configuration	CLIP↑	MLLM↑	Ins-Align↑
LLaVA-v1.1-7B (Default)	Vanilla	0.163	0.752	0.388
	+ FRCE + CME	**0.259**	**0.877**	**0.847**
LLaVA-NeXT-7B	Vanilla	0.182	0.786	0.425
	+ FRCE + CME	**0.271**	**0.889**	**0.862**
LLaVA-NeXT-13B	Vanilla	0.195	0.802	0.458
	+ FRCE + CME	**0.283**	**0.901**	**0.875**
Qwen-VL-7B	Vanilla	0.176	0.779	0.412
	+ FRCE + CME	**0.265**	**0.882**	**0.855**
Qwen-VL-2.5-7B	Vanilla	0.188	0.791	0.437
	+ FRCE + CME	**0.276**	**0.892**	**0.868**
Qwen-VL-3-7B	Vanilla	0.197	0.810	0.452
	+ FRCE + CME	**0.285**	**0.904**	**0.871**
MiniCPM-V-7B	Vanilla	0.171	0.772	0.403
	+ FRCE + CME	**0.262**	**0.879**	**0.851**

5. Conclusion

We extend IIE to a hypothetical instruction reasoning setting and propose a unified solution from both dataset and model perspectives. Specifically, we curate Reason50K, a large-scale dataset of 51,039 samples specifically designed to support hypothetical instruction reasoning across four diverse categories: physical, temporal, causal, and story reasoning. Simultaneously, we introduce ReasonBrain, a novel framework that enhances instruction reasoning by combining a MLLM and a fine-grained feature extraction module. We further integrate a cross-modal enhancer to enrich the semantics of the guidance used for diffusion-based editing. Extensive experiments on both reasoning-intensive and understanding scenarios, demonstrate that ReasonBrain exhibits strong reasoning capabilities as well as robust generalization performance. In addition, our Reason50K can facilitates broader advancements in reasoning-aware image generation, providing an extensible resource for future research in this emerging direction.

Acknowledgment

Qingdong He and Xueqin Chen contributed equally to the methodology and writing of this work. This work was partially funded by the Fundamental Research Funds for the Central Universities, Sichuan University, China (No. 1082204112M89).

Impact Statement

This work aims to advance the Image Editing field by introducing a novel task, dataset, and methodology. We plan to make the dataset and associated code publicly available for research. We release our code under an open-source license with explicit stipulations to mitigate the risk of misuse.

References

Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

AI, S. Stable diffusion 2.1 base. https://huggingface.co/stabilityai/stable-diffusion-2-1-base, 2022. Accessed: March 5, 2025.

Aminabadi, R. Y., Rajbhandari, S., Awan, A. A., Li, C., Li, D., Zheng, E., Ruwase, O., Smith, S., Zhang, M., Rasley, J., et al. Deepspeed-inference: enabling efficient inference of transformer models at unprecedented scale. In *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–15. IEEE, 2022.

Bai, J., Chow, W., Yang, L., Li, X., Li, J., Zhang, H., and Yan, S. Humanedit: A high-quality human-rewarded dataset for instruction-based image editing. *arXiv preprint arXiv:2412.04280*, 2024.

Brooks, T., Holynski, A., and Efros, A. A. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 18392–18402, 2023.

CompVis. Stable diffusion v1.4. https://huggingface.co/CompVis/stable-diffusion-v1-4, 2022. Accessed: March 5, 2025.

Deng, C., Zhu, D., Li, K., Gou, C., Li, F., Wang, Z., Zhong, S., Yu, W., Nie, X., Song, Z., et al. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*, 2025.

Fu, T.-J., Hu, W., Du, X., Wang, W. Y., Yang, Y., and Gan, Z. Guiding Instruction-based Image Editing via Multimodal Large Language Models. In *International Conference on Learning Representations (ICLR)*, 2024.

Geng, Z., Yang, B., Hang, T., Li, C., Gu, S., Zhang, T., Bao, J., Zhang, Z., Li, H., Hu, H., et al. Instructdiffusion: A generalist modeling interface for vision tasks. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pp. 12709–12720, 2024.

He, K., Chen, X., Xie, S., Li, Y., Dollár, P., and Girshick, R. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16000–16009, 2022.

He, Q., Peng, J., Xu, P., Jiang, B., Hu, X., Luo, D., Liu, Y., Wang, Y., Wang, C., Li, X., et al. Dynamiccontrol: Adaptive condition selection for improved text-to-image generation. *arXiv preprint arXiv:2412.03255*, 2024.

He, S., Ming, A., Li, Y., Sun, J., Zheng, S., and Ma, H. Thinking image color aesthetics assessment: Models, datasets and benchmarks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 21838–21847, 2023.

Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

Hore, A. and Ziou, D. Image quality metrics: Psnr vs. ssim. In *2010 20th international conference on pattern recognition*, pp. 2366–2369. IEEE, 2010.

Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.

Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B., and Liu, T. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Trans. Inf. Syst.*, 43(2), January 2025. ISSN 1046-8188. doi: 10.1145/3703155.

Huang, Y., Xie, L., Wang, X., Yuan, Z., Cun, X., Ge, Y., Zhou, J., Dong, C., Huang, R., Zhang, R., et al. 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*, pp. 8362–8371, 2024.

Hui, M., Yang, S., Zhao, B., Shi, Y., Wang, H., Wang, P., Zhou, Y., and Xie, C. Hq-edit: A high-quality dataset for instruction-based image editing. *arXiv preprint arXiv:2404.09990*, 2024.

Huo, F., Xu, W., Zhang, Z., Wang, H., Chen, Z., and Zhao, P. Self-introspective decoding: Alleviating hallucinations for large vision-language models. In *The Thirteenth International Conference on Learning Representations*, 2025.

Jin, Y., Ling, P., Dong, X., Zhang, P., Wang, J., and Lin, D. Reasonpix2pix: instruction reasoning dataset for advanced image editing. *arXiv preprint arXiv:2405.11190*, 2024.

Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W.-Y., Dollár, P., and Girshick, R. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4015–4026, 2023.

Labs, B. F. Flux. https://github.com/black-forest-labs/flux, 2024.

Labs, B. F., Batifol, S., Blattmann, A., Boesel, F., Consul, S., Diagne, C., Dockhorn, T., English, J., English, Z., Esser, P., et al. Flux. 1 kontext: Flow matching for in-context image generation and editing in latent space. *arXiv preprint arXiv:2506.15742*, 2025.

Li, J., Li, D., Savarese, S., and Hoi, S. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pp. 19730–19742. PMLR, 2023.

Li, Y., Bian, Y., Ju, X., Zhang, Z., Shan, Y., Zou, Y., and Xu, Q. Brushedit: All-in-one image inpainting and editing. *arXiv preprint arXiv:2412.10316*, 2024.

Lin, H., Cheng, X., Wu, X., and Shen, D. Cat: Cross attention in vision transformer. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*, pp. 1–6, 2022.

Lin, W., Wei, X., Zhang, R., Zhuo, L., Zhao, S., Huang, S., Xie, J., Peng, G., and Li, H. Pixwizard: Versatile image-to-image visual assistant with open-language instructions. In *International Conference on Learning Representations*, pp. 38680–38711, 2025.

Liu, C., Li, X., and Ding, H. Referring image editing: Object-level image editing via referring expressions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13128–13138, 2024.

Liu, H., Li, C., Wu, Q., and Lee, Y. J. Visual instruction tuning. In *Advances in Neural Information Processing Systems*, volume 36, pp. 34892–34916. Curran Associates, Inc., 2023.

Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.

Meng, F., Shao, W., Luo, L., Wang, Y., Chen, Y., Lu, Q., Yang, Y., Yang, T., Zhang, K., Qiao, Y., et al. Phybench: A physical commonsense benchmark for evaluating text-to-image models. *arXiv preprint arXiv:2406.11802*, 2024.

Nguyen, T.-T., Nguyen, D.-A., Tran, A., and Pham, C. Flexedit: Flexible and controllable diffusion-based object-centric image editing. *arXiv preprint arXiv:2403.18605*, 2024a.

Nguyen, T. T., Ren, Z., Pham, T., Huynh, T. T., Nguyen, P. L., Yin, H., and Nguyen, Q. V. H. Instruction-guided editing controls for images and multimedia: A survey in llm era. *arXiv preprint arXiv:2411.09955*, 2024b.

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PmLR, 2021.

Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., and Chen, M. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.

Ravi, N., Gabeur, V., Hu, Y.-T., Hu, R., Ryali, C., Ma, T., Khedr, H., Rädle, R., Rolland, C., Gustafson, L., et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024.

Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10684–10695, 2022.

Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E. L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35: 36479–36494, 2022.

Sheynin, S., Polyak, A., Singer, U., Kirstain, Y., Zohar, A., Ashual, O., Parikh, D., and Taigman, Y. Emu edit: Precise image editing via recognition and generation tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8871–8879, 2024.

SimianLuo. Lcm dreamshaper v7. https://huggingface.co/SimianLuo/LCM_Dreamshaper_v7, 2024. Accessed: March 5, 2025.

Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., and Ganguli, S. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pp. 2256–2265. pmlr, 2015.

Sun, Q., Luo, J., Zhang, D., and Li, X. Smartfreeedit: Mask-free spatial-aware image editing with complex instruction understanding. In *Proceedings of the 33rd ACM International Conference on Multimedia*, MM '25, pp. 8244–8252, 2025. doi: 10.1145/3746027.3755395.

Tang, F., Huang, Z., Liu, C., Sun, Q., Yang, H., and Lim, S.-N. Intervening anchor token: Decoding strategy in alleviating hallucinations for MLLMs. In *The Thirteenth*

International Conference on Learning Representations, 2025.

Tian, X., Li, W., Xu, B., Yuan, Y., Wang, Y., and Shen, H. Mige: Mutually enhanced multimodal instruction-based image generation and editing. In *Proceedings of the 33rd ACM International Conference on Multimedia*, MM '25, pp. 10622–10631, 2025. doi: 10.1145/3746027.3755811.

Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

Wang, Y., Chen, W., Han, X., Lin, X., Zhao, H., Liu, Y., Zhai, B., Yuan, J., You, Q., and Yang, H. Exploring the reasoning abilities of multimodal large language models (mllms): A comprehensive survey on emerging trends in multimodal reasoning. *arXiv preprint arXiv:2401.06805*, 2024a.

Wang, Y., Yang, S., Zhao, B., Zhang, L., Liu, Q., Zhou, Y., and Xie, C. Gpt-image-edit-1.5m: A million-scale, gpt-generated image dataset, 2025. URL https://arxiv.org/abs/2507.21033.

Wang, Z., Li, A., Li, Z., and Liu, X. Genartist: Multimodal llm as an agent for unified image generation and editing. *Advances in Neural Information Processing Systems*, 37: 128374–128395, 2024b.

Wu, C., Li, J., Zhou, J., Lin, J., Gao, K., Yan, K., Yin, S.-m., Bai, S., Xu, X., Chen, Y., et al. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*, 2025.

Yang, S., Hui, M., Zhao, B., Zhou, Y., Ruiz, N., and Xie, C. Complex-edit: Cot-like instruction generation for complexity-controllable image editing benchmark. *arXiv preprint arXiv:2504.13143*, 2025.

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

Zeng, B., Yang, L., Liu, J., Xu, M., Zhang, Y., Wan, P., Zhang, W., and Yan, S. Editworld: Simulating world dynamics for instruction-following image editing. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pp. 12674–12681, 2025.

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

Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 586–595, 2018.

Zhang, S., Yang, X., Feng, Y., Qin, C., Chen, C.-C., Yu, N., Chen, Z., Wang, H., Savarese, S., Ermon, S., et al. Hive: Harnessing human feedback for instructional visual editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9026–9036, 2024.

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

Zhou, J., Li, J., Xu, Z., Li, H., Cheng, Y., Hong, F.-T., Lin, Q., Lu, Q., and Liang, X. Fireedit: Fine-grained instruction-based image editing via region-aware vision language model. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 13093–13103, 2025.

Appendix

A. Reproducibility Statement

We have already elaborated on all the models or algorithms proposed, experimental configurations, and benchmarks used in the experiments in the main body or appendix of this paper. Furthermore, we declare that the entire code used in this work will be released after acceptance.

B. The Use of Large Language Models

We use large language models solely for polishing our writing, and we have conducted a careful check, taking full responsibility for all content in this work.

C. Dataset Generation

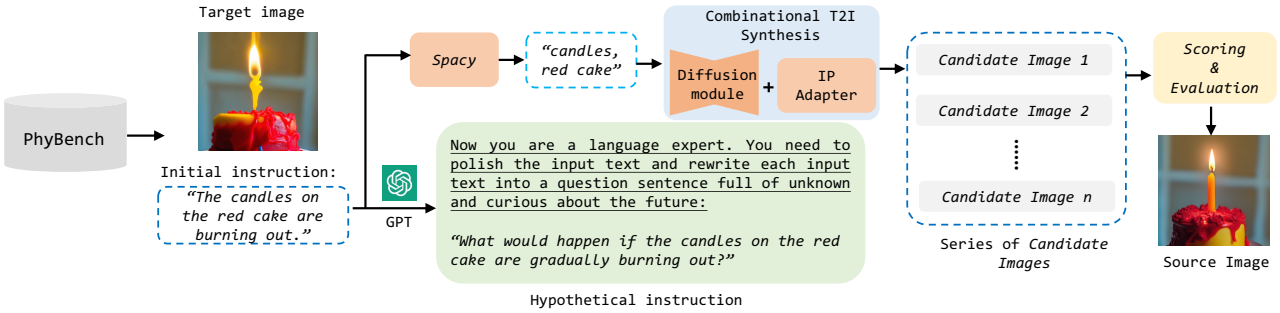

Figure A1. Data generation process.

The construction of our data consists of two parts. The first part (over 90% of the entire dataset) is generated following the pipeline illustrated in Fig. A1, which adopts an inverse generation strategy–deriving the source image from the target. Specifically, we first adopt the same procedure as PhyBench (Meng et al., 2024) to generate target images along with their initial instructions. Each initial instruction is then rewritten into a hypothetical form using prompt-based rewriting with GPT (Achiam et al., 2023). In parallel, we use SpaCy [2] to perform named entity recognition (NER) on the initial instruction to extract candidate objects for source image generation. These candidates are passed to a diffusion model equipped with an IP-Adapter to synthesize multiple image variants. Each candidate image is subsequently evaluated by GPT, and the top-N images are selected based on a combination of GPT scores and perceptual quality metrics (He et al., 2023; Hore & Ziou, 2010; Zhang et al., 2018). Each selected image, along with the corresponding hypothetical instruction and target image, constitutes a sample in our dataset. The second part (less than 10%) consists of story-type samples derived from EditWorld (Zeng et al., 2025). We apply the same instruction rewriting, scoring, and filtering process to refine and select high-quality samples for inclusion.

D. Training Objectives of ReasonBrain

The training of ReasonBrain comprises two components: fine-tuning the MLLM and optimizing the diffusion model. Specifically, for fine-tuning the MLLM, we freeze most of its parameters and apply Low-Rank Adaptation (LoRA) (Hu et al., 2022) for efficient adaptation. The objective is defined as:

$$\mathcal{L}_{MLLM} = -\sum_{i=1}^{r} \log p_{\{\theta \cup \theta_{LoRA}\}} \left([\mathrm{IMG}_i] \mid \mathrm{IA}(\mathcal{E}_I(I)), R_V, R_T, \mathcal{E}_T(H), [\mathrm{IMG}_1], \ldots, [\mathrm{IMG}_{i-1}] \right), \tag{A1}$$

where θ_{LoRA} denotes the trainable parameters introduced by LoRA. This loss minimizes the negative log-likelihood of predicting each learnable token $[\mathrm{IMG}_i]$ conditioned on the fine-grained features and previously predicted tokens. For image generation, we adopt a latent diffusion objective:

$$\mathcal{L}_{DM} = \mathbb{E}_{\mathcal{E}_I(\hat{I}), \mathcal{E}_I(I), H, \epsilon, t} \left[\left\| \epsilon - \epsilon_\delta \left(t, [z_t, \mathcal{E}_I(I)] + \bar{R}_{visual} + \bar{R}_{text}, [\bar{e}_{visual}, \bar{e}_{text}] \right) \right\|_2^2 \right], \tag{A2}$$

[2] https://spacy.io/

where $\epsilon \sim \mathcal{N}(0, 1)$ is the sampled noise and z_t is the noisy latent at timestep t. $\epsilon_\delta(\cdot)$ denotes the denoising network trained to predict the added noise based on the timestep and the provided visual reasoning guidance. The overall training objective is defined as the sum of the MLLM and diffusion losses: $\mathcal{L} = \mathcal{L}_{MLLM} + \mathcal{L}_{DM}$. Moreover, to mitigate hallucination risks, we adopted a context- and instruction-aware token selection strategy inspired by SID (Huo et al., 2025), combined with the dynamic token propagation mechanism from TAME (Tang et al., 2025) in the training process of MLLM.

E. Setups: Datasets, Metrics, and Details

Datasets: We use Reason50K for both training and evaluation. Specifically, for each reasoning category, 400 samples are randomly selected for validation, while the remaining samples are used for training. In addition, we assess the reasoning capability of ReasonBrain on two external benchmarks: *ReasonEdit* (Huang et al., 2024), *EditWorld* (Zeng et al., 2025), and *Complex-Edit* (Yang et al., 2025). To evaluate generalization on conventional understanding scenarios, we further test on the *MagicBrush Test Set* (Zhang et al., 2023) and the *Emu Edit Test Set* (Sheynin et al., 2024).

Metrics: To evaluate performance under reasoning scenarios, we adopt three metrics: *CLIP Score* (Radford et al., 2021), *MLLM Score* (Zeng et al., 2025), and *Instruction Alignment (Ins-Align)* (Huang et al., 2024). Here, CLIP Score measures the semantic similarity between the edited image and the expected output text using CLIP's image-text embeddings. MLLM Score employs an MLLM to assess instruction-following performance. Following (Zeng et al., 2025), we provide the input description, editing instruction, and output description along with the edited image to Video-LLaVA. The prompt is defined as: "The input description: [object Object], the editing instruction: [object Object], and the output description: [object Object]. Please evaluate if the given edited image has been successfully edited. If yes, return 1; if not, return 0." The final score is the average of model judgments across all samples. In addition, Ins-Align Score evaluates how well the edited image aligns with the given instruction. We follow (Yang et al., 2025) to calculate IP, which better illustrates identity preservation, i.e., whether elements of the input image that should remain unchanged are indeed preserved. Following (Huang et al., 2024), ten human annotators independently rated the outputs on the Reason50K dataset, and we report the average alignment score. For understanding scenarios, we adopt *L1 distance*, *CLIP image similarity*, *DINO similarity*, *CLIP text-image similarity*, and *CLIP text-image direction similarity* as evaluation metrics.

Implementation Details: During training, we adopt the pre-trained LLaVAv1.1-7B (Liu et al., 2023) (default setting) and QFormer (Li et al., 2023) and employ DeepSpeed (Aminabadi et al., 2022) Zero-2 to perform LoRA (Hu et al., 2022) fine-tuning, with rank and alpha of 8 and 16, respectively. Following (Huang et al., 2024; Fu et al., 2024), we expand the original LLM vocabulary with 32 new tokens, and the QFormer is composed of 6 transformer layers and 77 learnable query tokens. For the base editing model, we implement it with Flux (Labs, 2024) using FLUX.1-dev, which consists of 12B parameters. Models for other qualitative results are implemented using SD series (CompVis, 2022; SimianLuo, 2024; Rombach et al., 2022; AI, 2022) with their original codebase. Our model is trained with a batch size of 16, and we use AdamW (Loshchilov & Hutter, 2019) optimizer to train the model with the weight decay as 1e-2 and the learning rate as 1e-3. All the experiments are conducted on 16 H20 GPUs. For a fair comparison, all baseline models are fine-tuned on the same training set used by ReasonBrain.

F. More Discussion of the Framework Design

Functional group studies: To clarify the pivotal components and avoid fragmented analysis, we reorganize the ablations into coarse-grained functional groups, focusing on the core reasoning–editing interaction. As shown in Tab. A1 and visualized in Fig. A2, the Baseline Group (MLLM + diffusion only) produces only minimal or unrelated changes, achieving the lowest Ins-Align Score of **0.388**. This demonstrates that simply combining an MLLM with a diffusion model is insufficient for understanding and executing implicit hypothetical instructions. The FRCE Core Group (Method IDs 2–3) forms the foundation of reasoning–editing interaction. Without the ID Controller (Method 2), the model produces semantically inconsistent outcomes—such as generating an additional intact egg while also showing partial cracking—indicating a loss of object identity and clear semantic drift. With the ID Controller added (Method 3, full FRCE Core), the model preserves the correct object identity and generates a more coherent "dropped-egg" outcome, where the same egg is broken in a physically plausible way. This raises the Ins-Align Score to 0.758, a 43.3% improvement, confirming that the ID Controller is crucial for binding reasoning cues to the correct object and preventing unintended identity changes. Finally, the CME Enhancement Group (Method 4) achieves the strongest overall performance (CLIP: **0.259**, MLLM: **0.877**, Ins-Align: **0.847**). As seen in the rightmost result of Fig. A2, CME produces a clean, physically plausible "dropped and splattered" egg consistent with the instruction. CME refines the cross-modal alignment between FRCE-derived cues and diffusion-based editing features, reducing mismatches between the intended reasoning and the resulting edits. It acts as a targeted enhancer rather than a

standalone reasoning module. In summary, our framework's effectiveness relies on two pivotal components: (1) the FRCE Core, with the ID Controller as the key mechanism for linking reasoning to identity-preserving edits; and (2) the CME module, which further enforces cross-modal consistency and elevates the overall reasoning quality.

Table A1. Quantitative results for the functional groups of ReasonBrain.

Functional Group	Method ID	with Patch Branch	with Region Branch	with ID Controller	with Vision Enchancer	with Text Enchancer	CLIP Score↑	MLLM Score↑	Ins-Align Score↑
Baseline Group	1	×	×	×	×	×	0.163	0.752	0.388
FRCE Core Group	2	✓	✓	×	×	×	0.206	0.802	0.529
FRCE Core Group + ID	3	✓	✓	✓	×	×	0.239	0.833	0.758
CME Enhancement	4	✓	✓	✓	✓	✓	**0.259**	**0.877**	**0.847**

Figure A2. Qualitative results for the functional component groups of ReasonBrain.

Table A2. Ablation study on different design choices.

Exp ID	Plain	Simple CA	CME	CLIP↑	MLLM ↑	Ins-Align ↑
1	✓			0.239	0.833	0.758
2		✓		0.241	0.841	0.766
Ours			✓	**0.259**	**0.877**	**0.847**

Effectiveness of CME module: To validate the effectiveness of the bidirectional information interaction in our proposed CME module, we conducted two comparative experiments. In *Exp 1*, we remove the CME module entirely and directly feed the feature output from QFormer into the diffusion model. This ablation study is designed to evaluate the effectiveness of the information interaction introduced by the CME module. In *Exp 2*, we aim to assess the necessity of bidirectional information interaction. Specifically, we retain only the cross-attention block on the image feature branch, discarding all other components of the CME module. As a result, the textual features from QFormer are applied to the image features in a unidirectional manner. The results are presented in Tab. A2, demonstrating the marginal gains brought by each component and highlighting the overall benefit of the complete CME design.

Bidirectional interaction validation: We conduct additional experiments to further validate the design of the *bidirectional interaction* between reasoning and editing within a unified framework. Specifically, we evaluate the FRCE module and the CME module independently. The FRCE module extracts fine-grained reasoning cues (e.g., physical object weight, temporal changes), injects them into the MLLM's visual-guidance generation, and binds them to the diffusion model via the QFormer. In contrast, the CME module enables mutual alignment between MLLM-generated reasoning tokens and diffusion-based editing features. The corresponding results are provided in Tab. A3 and Tab. A4. Our results show that removing the FRCE–diffusion interaction leads to a 12.3% drop in Ins-Align Score and an 8.7% decrease in visual plausibility (human evaluation). Using CME without bidirectional interaction further reduces the Ins-Align Score by 4.1%. These findings demonstrate the necessity of both modules and highlight the importance of their bidirectional integration.

G. Other Quantitative Results

Computational costs & lightweight alternative: Tab. A5 shows the inference time of our ReasonBrain alongside all baselines. We find that our model demonstrates strong reasoning capabilities while maintaining an inference time comparable to that of existing methods. To further accelerate the model, we explored a lightweight variant, **ReasonBrain-3B**, by adopting a smaller pretrained MLLM. The model settings and performance of ReasonBrain and **ReasonBrain-3B** are

Table A3. Performance comparison: with vs. w/o FRCE–Diffusion interaction

Method	CLIP Score (↑)	MLLM Score (↑)	Ins-Align Score (↑)
ReasonBrain (Full, FRCE + Diffusion Interaction)	0.259	0.877	0.847
ReasonBrain (w/o FRCE-Diffusion Binding, MLLM Only)	0.228	0.801	0.743
Performance Drop	-12.0%	-8.7%	-12.3%

Table A4. Performance comparison: bidirectional vs. unidirectional CME

CME Mode	CLIP Score (↑)	MLLM Score (↑)	Ins-Align Score (↑)
Bidirectional	0.259	0.877	0.847
Unidirectional (Reasoning→Editing only)	0.248	0.842	0.806

listed in Tab. A6. We observe that, compared to ReasonBrain-7B, ReasonBrain-3B is significantly faster but exhibits only a slight performance drop. Nevertheless, it still outperforms all baselines reported in Tab. 2. This indicates that even in instruction-based reasoning scenarios, a lightweight MLLM can not only accelerate inference but also effectively identify the editing target and execute precise edits, supported by its strong reasoning ability and rich world knowledge.

Table A5. Comparison of inference times.

Method	Inference Time (s)
InstructPix2Pix	26
MagicBrush	28
MGIE	37
SmartEdit	33
UltraEdit	28
PixWizard	35
Ours	**32**

Table A6. Settings and Performance of ReasonBrain and its lightweight variant.

Attribute	ReasonBrain-7B	ReasonBrain-3B
MLLM Backbone	LLaVA-7B	LLaVA-3B
Diffusion Backbone	SD 2.1	SD 1.5
FRCE Params	82M	41M
CME Params	45M	22M
Trainable Params	800M	320M
CLIP ↑	0.238	0.259
MLLM ↑	0.852	0.877
Ins-Align ↑	0.822	0.847
Inference Time (s)	24	32

Table A7. Performance under different editing steps.

Steps	CLIP ↑	MLLM ↑	Ins-Align ↑
1-step	0.259	0.877	0.847
2-steps	0.257	0.878	0.848
3-steps	0.260	0.877	0.847

Table A8. Comparison of one-to-many editing settings.

Setting	CLIP ↑	MLLM ↑	Ins-Align ↑	IP ↑
one-to-one	0.259	0.877	0.847	9.72
one-to-two	0.152	0.680	0.315	4.25
one-to-three	0.088	0.258	0.204	1.66

Different editing settings: *(1) Multi-step editing*. We extend the original one-step experiments to include 2-step and 3-step editing scenarios (multi-hop reasoning chains, e.g., "ice melting → water evaporating → damp ground forming"). The results (shown in Tab. A7) indicate that the model maintains robust performance across multiple editing steps, demonstrating the effectiveness of ReasonBrain in handling sequential multi-step reasoning-based editing. We attribute this to the ID Controller module, which operates independently at each editing step, enabling effective processing of chained instructions without causing significant feature drift. *(2) Mapping settings*. We conduct an additional ablation study by extending our original *one-to-one mapping* to *one-to-two* and *one-to-three* mappings. The results are shown in Tab. A8, we find that incorporating multiple targets (i.e., one-to-two/one-to-three mappings) led to significant performance degradation across all metrics. These settings introduced training instability due to conflicting optimization signals from multiple targets for the same input. In particular, Ins-Align and Identity Preservation dropped significantly, indicating that using multiple targets for the same input makes it difficult to retain key elements from the source image–further emphasizing the importance of maintaining a one-to-one structure in our task.

Human evaluation: We randomly select 50 images corresponding to four distinct reasoning scenarios. For each image, we obtain the results of InstructPix2Pix, MagicBrush, MGIE, SmartEdit, UltraEdit and PixWizard. Then, randomly shuffle the order of these method results. For each set of images, we ask 30 participants to independently select the three best pictures. The first one is the best picture corresponding to the text prompt (i.e., Instruct Alignment), and the second one is the picture with the highest visual quality (i.e., Image Quality). The result is shown in Tab. A9. We found that 75.71% of participants believed ReasonBrain better reflected the correct reasoning behind the instructions, and 63.55% preferred the results generated by ReasonBrain. In contrast, all baseline methods received less than 10% on both evaluations.

Influences of data generation: The key steps in our data generation process (see App. C and Fig. A1 for details) are:

Table A9. Human evaluation results.

Method	Instruct-Alignment (%)	Image Quality (%)
InstructPix2Pix	2.62	2.14
MagicBrush	3.11	5.65
MGIE	3.10	6.68
SmartEdit	3.55	6.85
UltraEdit	6.09	6.11
PixWizard	5.82	9.02
Ours	**75.71**	**63.55**

(a) Utilizing PhyBench to acquire a target image and its initial prompt;

(b) Extracting candidate objects from the prompt (e.g., object identity, action, attributes);

(c) Using these candidate objects to guide a T2I model to generate candidate source images;

(d) Selecting the best candidate based on a combination of evaluation metrics;

(e) Rewriting the initial prompt into a hypothetical instruction format.

Finally, each data sample is composed of a selected candidate image, the rewritten instruction, and the original target image. We argue that the potential sources of influence on the final performance include:

- **Noise in Source Image Generation (Step a):** Some target images may contain distortions or semantically irrelevant content.
- **Missing Object Elements (Step b):** Key candidate objects may be omitted during prompt decomposition.
- **Lack of Constraints (Step c):** The T2I generation process does not explicitly enforce identity or background preservation when generating candidate source images.
- **Limited Evaluation Dimensions (Step d):** The scoring and selection criteria may overlook fine-grained visual details such as object consistency or scene coherence.

To this end, we conduct a preliminary ablation study to assess how these factors affect model performance. The results are summarized in Tab. A10. We find that when the initial step produces low-quality target images (error image), the model's performance drops significantly. Similarly, in the source image generation step, missing object elements (missing ID text) or ignoring ID correlations (w/o ID adapter) negatively impact performance. Finally, during the selection phase, using only a single metric (i.e., single selection) results in a considerable performance gap, highlighting the importance of applying a comprehensive set of evaluation metrics when constructing the final dataset.

Table A10. Impact of the data generation process on performance.

Setting	CLIP ↑	MLLM ↑	Ins-Align ↑
Error Image	0.104	0.425	0.226
Missing ID Text	0.125	0.465	0.233
w/o ID Adapter	0.087	0.312	0.158
Single Selection	0.172	0.705	0.505
Ours	**0.259**	**0.877**	**0.847**

Instruction conversion study: First, we would like to emphasize that simply converting a hypothetical query into a declarative instruction cannot fulfill the HI-IE requirement. For example, an LLM may convert the query "What would happen if the ice cube was left in the sun?" into "Replace the ice cube with melted water." However, this rephrased instruction omits important details, such as the melting progression of the ice cube, the spread of water stains, lighting changes caused by transparent water, and the overall physical consistency of the scene. To demonstrate our claim, we conduct a zero-shot comparison in which an LLM first converts the hypothetical instruction into an explicit one, and an IIE model then performs the editing. As shown in Table A11, even with GPT-4o for instruction conversion, the best-performing GPT-4o+Qwen-ImageEdit achieves an Ins-Align score of only 0.512, which is about 39.5% lower than ReasonBrain's score of 0.847.

Table A11. Comparison with simple instruction-conversion pipelines.

Pipeline	CLIP↑	MLLM↑	Ins-Align↑
Llama-3-8B-Instruct Conversion			
+ InstructPix2Pix (Fine-tuned on Reason50K)	0.172	0.765	0.385
+ Flux-Kontext (Zero-shot)	0.185	0.782	0.402
+ Qwen-ImageEdit (Zero-shot)	0.191	0.795	0.421
GPT-3.5-Turbo Conversion			
+ InstructPix2Pix (Fine-tuned on Reason50K)	0.188	0.791	0.432
+ Flux-Kontext (Zero-shot)	0.197	0.804	0.451
+ Qwen-ImageEdit (Zero-shot)	0.203	0.817	0.476
GPT-4o Conversion			
+ InstructPix2Pix (Fine-tuned on Reason50K)	0.195	0.808	0.457
+ Flux-Kontext (Zero-shot)	0.208	0.825	0.489
+ Qwen-ImageEdit (Zero-shot)	0.215	0.838	0.512
ReasonBrain (Ours)	**0.259**	**0.877**	**0.847**

H. More Qualitative Results

To further demonstrate the editing performance of ReasonBrain compared to SOTA methods, we present additional qualitative results in Fig. A3. It is evident that ReasonBrain consistently outperforms other SOTA approaches, producing more coherent and visually plausible edits. These results highlight ReasonBrain's ability to reason effectively from hypothetical instructions and generate outputs that closely align with the transformations implied by the underlying reasoning.

I. Limitations

Limitations on Real-world Data Collection: The construction of our dataset Reason50K follows the common practice in the image editing community (e.g., EmuEdit (Sheynin et al., 2024) and GPT-Image-Edit (Wang et al., 2025)), relying primarily on synthetic data. While synthetic datasets provide semantic clarity and high visual quality, they may not fully capture the complexities, artifacts, and temporal dynamics inherent in real-world video data. Collecting high-quality datasets from real-world video sources remains particularly challenging due to the scarcity of suitable image pairs with explicit reasoning logic and the absence of standardized benchmarks or evaluation protocols for reasoning-based edits. Moreover, prior efforts such as EditWorld (Zeng et al., 2025) show that frames extracted from videos often suffer from low resolution and poor aesthetics, making them suboptimal for training high-fidelity models. Thus, addressing this gap by curating reasoning-aligned, high-quality video benchmarks constitutes an important direction for our future work.

Style & Identity Preservation: In Fig. 6, we find that in some cases the model may introduce stronger-than-desired global appearance changes (e.g., savanna vs. grassland). However, such cases are limited, and in most qualitative results the overall style remains consistent. Additionally, we emphasize that some minor scene changes (e.g., slight camera-distance shifts or local lighting/contrast adjustments) are introduced only to make the inferred visual changes perceptible and strictly constrained by the hypothetical instruction, rather than to arbitrarily regenerate the scene. For example, in the case "What happens when a sudden swarm of bees descends on the flowers?", our model slightly adjusts the camera distance to make the bees visible, while preserving the core scene content, including the flower clusters, lawn layout, background trees, and overall lighting style. Across all case studies, non-edited regions such as walls, lawns, and background objects remain visually consistent with the source image, and only regions related to the hypothetical change are modified. To alleviate this limitation, we will explore stronger style constraints or more disentangled control to improve fine-grained fidelity in our future work.

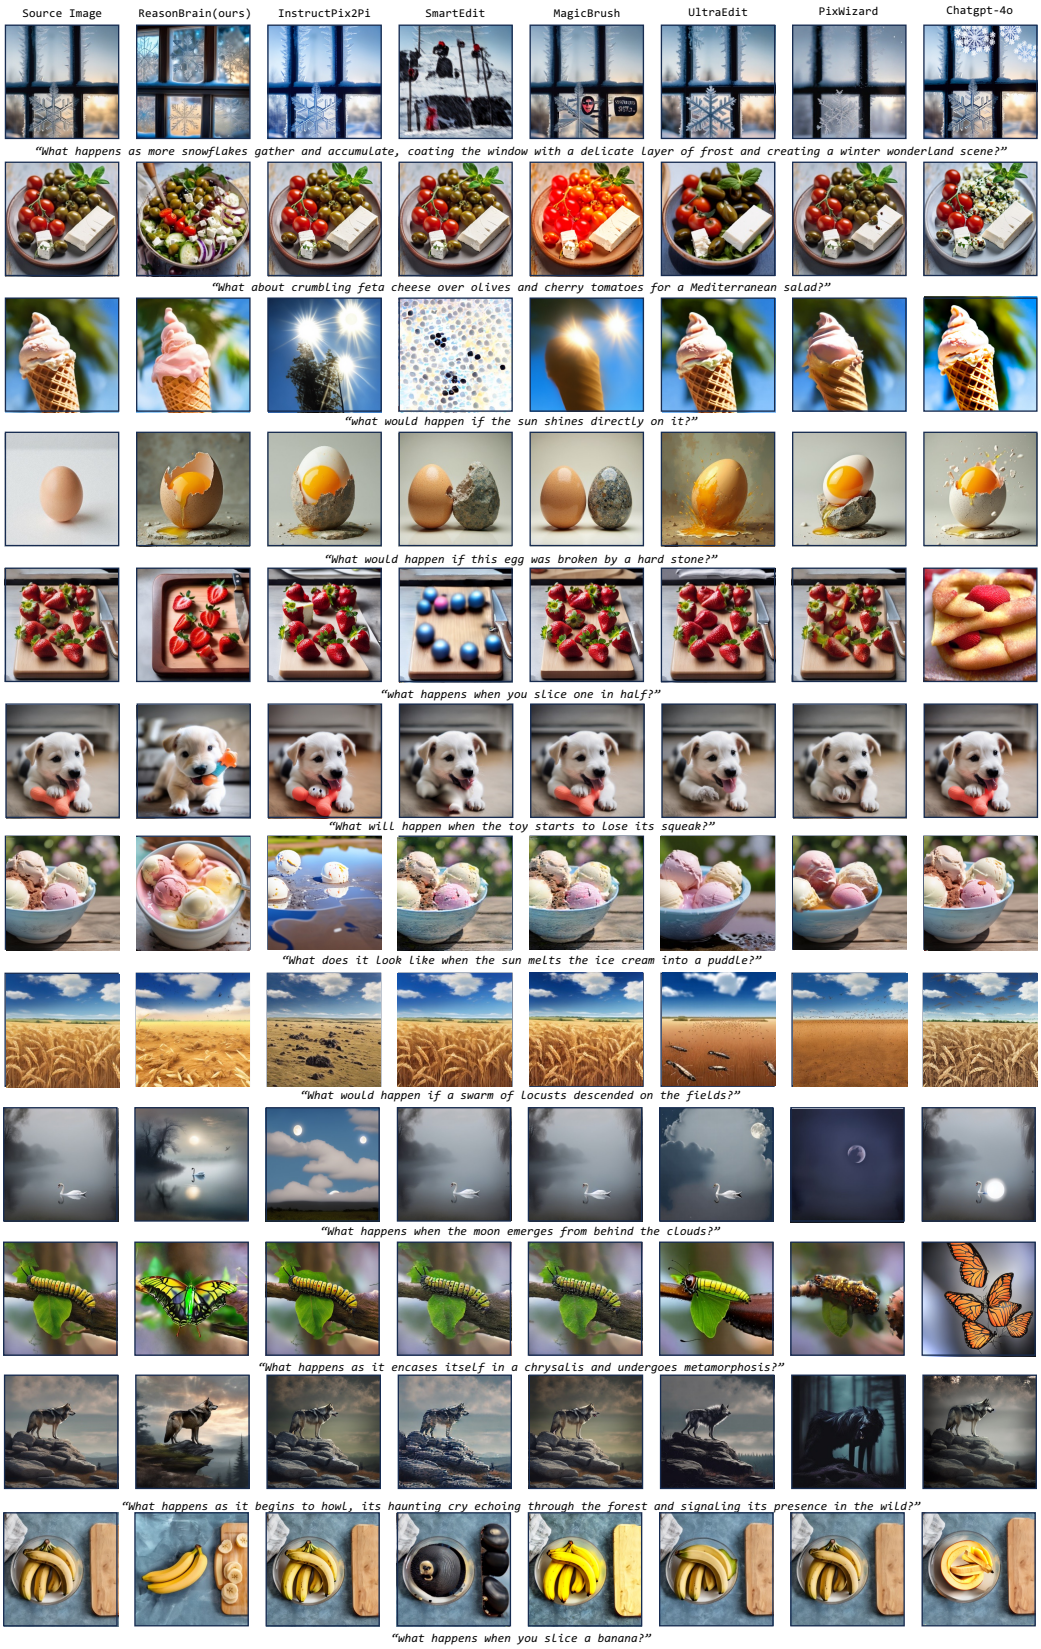

Figure A3. Qualitative comparison on Reason50K between ReasonBrain and selected SOTA methods.