EditInspector: A Benchmark for Evaluation of Text-Guided Image Edits

Ron Yosef¹, Moran Yanuka², Yonatan Bitton³, Dani Lischinski¹ ¹The Hebrew University of Jerusalem, ²Tel Aviv University, ³Google Research,

Abstract

Text-guided image editing, fueled by recent advancements in generative AI, is becoming increasingly widespread. This trend highlights the need for a comprehensive framework to verify text-guided edits and assess their quality. To address this need, we introduce EditInspector, a novel benchmark for evaluation of text-guided image edits, based on human annotations collected using an extensive template for edit verification¹. We leverage EditInspector to evaluate the performance of state-of-the-art (SoTA) vision and language models in assessing edits across various dimensions, including accuracy, artifact detection, visual quality, seamless integration with the image scene, adherence to common sense, and the ability to describe editinduced changes. Our findings indicate that current models struggle to evaluate edits comprehensively and frequently hallucinate when describing the changes. To address these challenges, we propose two novel methods that outperform SoTA models in both artifact detection and difference caption generation.

1 Introduction

The ability to create and modify images is vital in fields such as social media, marketing, and graphic design. Recent advancements in generative AI have greatly democratized this ability. In particular, natural language enables high-quality, customized visual content creation with minimal effort.

Text-guided editing models require a source image and instruction (Kawar et al., 2023; Zhang et al., 2022; Brooks et al., 2023; Wu et al., 2023b; Zhang et al., 2024b), sometimes allowing multiturn editing (Sheynin et al., 2023; He et al., 2024; Wu et al., 2023a; Cui et al., 2023). For more precise spatial control a user might provide the source image, a mask, and a text prompt specifying changes for the masked area (Avrahami et al., 2022; Nichol et al., 2022; Couairon et al., 2022; Wang et al., 2023; Zhang et al., 2024a). Extensive human evaluations showed that mask-based text-guided editing produces superior results compared to mask-free editing (Wang et al., 2023; Zhang et al., 2024a).

Despite these advancements, evaluating the quality and accuracy of edits remains challenging, as demonstrated in Figure 1. Current methods often focus on whether the edited object matches the requested attributes (Wang et al., 2023) or use ranking scores for accuracy (Zhang et al., 2024a). However, they overlook pain points such as unintended artifacts, misalignment with user expectation, visual quality, and adherence to common sense. For example, in Figure 2, the edit changes teardrops to stars as instructed, but unintentionally adds a line and alters the wall's appearance.

To address these challenges, we propose EditInspector, a comprehensive benchmark for *assessing evaluators of text-guided image edits* (Section 2). EditInspector examines edits across five dimensions: (1) whether the edit accurately follows the instructions and aligns with user expectations; (2) introduction of unintended artifacts; (3) technical quality (low resolution, blur, etc.); (4) the accuracy of a description of the main difference; and (5) the accuracy of a detailed listing of the differences between the original and the edited images.

We begin by creating a human evaluation framework, shown in Figure 2, that assesses edits based on the dimensions outlined above (Section 2.1). Using this framework, we collected human annotations as edit inspectors through crowdsourcing, evaluating 783 edits from the MagicBrush (Zhang et al., 2024a) test set of 1,053 edits, to introduce the EditInspector benchmark (Section 2.2).

We then evaluate state-of-the-art vision and language models (VLMs) as edit inspectors on the EditInspector benchmark, comparing their performance with human annotations, as shown in Figure 1. The results show that all models perform

¹https://editinspector.github.io/



Figure 1: The assessments for the edit "Let the floor be made of wood" vary across different models, with 2–3 models answering each question correctly. Gemini 1.5 failed to detect any differences between the images, while GPT-40 successfully identified only the main difference. See Appendix A.7 for full-size prompts.

poorly across all tasks, with accuracy hovering around random chance (Section 3.3.1). Gemini-1.5 (Gemini Team, 2024) emerged as the top performer for the edit inspector questions, achieving 70.3% accuracy in the edit accuracy question. We evaluate models' ability to generate a summary of the main change and a detailed list of all differences as an upper-bound test of edit accuracy, artifact detection, and visual quality. In this task, GPT-40 achieved 39% accuracy in describing the main difference but detected only 12% of all differences, with only 40% aligning with human annotations, highlighting significant hallucinations. (Section 3.3.2).

We tackle the challenges of artifact detection and difference caption generation with two methods. First, we developed a zero-shot pipeline using Gemini as the visual backbone to generate instructiongrounded difference captions and metadata (Section 4.1). The pipeline analyzes image captions at three zoom levels around the edit area and outputs a difference caption, achieving 75% accuracy in describing the main difference, compared to 39% by the best SoTA model. Second, we introduced a novel artifact detection method that achieves 64% accuracy by analyzing object segmentation probabilities around the edited area (Section 4.2).

Finally, we introduce an end-to-end fine-tuned model that rivals much larger models, delivering competitive SoTA performance while reducing computational costs (Section 5). To train our model we use two augmentation methods to generate 31,059 training instances. The first method creates negative examples with objects closely resembling the original (Section 5.1), and the second reverses the edit direction, e.g., by changing an

"Add" edit to a "Remove" edit (Section 5.2).

In summary, our main contributions are: (1) A comprehensive framework for image edit evaluation, and the EditInspector benchmark, which we release for future work and future model assessment; (2) A thorough evaluation of SoTA VLMs as edit inspectors, showing that, across all aspects, none can effectively assess edits; (3) Two new methods outperforming SoTA models for artifact detection and difference caption generation; and, (4) An end-to-end fine-tuned model that rivals much larger models in performance.

2 EditInspector Dataset

Our goal is to develop a dataset and framework for image editing verification that offers a comprehensive evaluation of edits, addressing overlooked pain points like unintended artifacts, instruction inconsistencies, scene misalignment, and technical flaws. To achieve this, we introduced the human evaluation framework in Section 2.1 and annotated 783 MagicBrush edits using it to create our benchmark in Section 2.2. MagicBrush is a manually annotated dataset for instruction-guided, mask-provided image editing. The statistics and analysis of our benchmark are presented in Section 2.3.

2.1 Human Evaluation Framework

Our motive was to develop a comprehensive framework that evaluates multiple aspects of image editing. We tested and refined templates and questions using internal and crowdsourced feedback, resulting in the framework shown in Figure 2.

The evaluation begins with *Accuracy Level*, where annotators assess whether the edit follows



Figure 2: This is an example of our annotation user interface. The edit appears to be accurately executed but includes unexpected elements, such as differences in the door layers and a tilted star edge. There are mild artifacts, including a shadow behind the wall and a thick gray line beneath the star cutout. Clicking the tree icons opens decision trees that help annotators follow the evaluation guidelines (See Appendix A.16).

the instruction and meets user expectations. If it fully follows the instruction, annotators select *Accurate* or *Accurate*, *But Unexpected* if it deviates from expectations. For partial adherence to the instruction, they select *Inaccurate*, *Reflects Instruction*, and for no adherence, *Inaccurate*.

For any selection other than *Accurate*, annotators are asked to explain under *Contextual Consistency* how the edit failed to meet expectations or align with the instruction, image scene, or common sense. Under the *Technical Precision* question annotators comment on pixel-level details like resolution, blurriness, and smoothness.

For example, in Figure 2 a teardrop cutout was changed to a star-shaped hole, but all annotators marked it as "*Accurate, But Unexpected*" due to the tilted star edge and the unexpected material appearance, as seen in *Contextual Consistency* feedback.

Next, the *Artifacts* evaluation involves annotators identifying any unintended distortions or anomalies in the edit. Artifacts are classified into two levels: *Significant* or *Mild*, based on their severity. In the example in Figure 2, two *Mild* artifacts are present: an unintended shadow and an extra line beneath the star-shaped hole.

Finally, to collect a difference caption that describes all differences between the original and edited images as an upper-bound evaluation, we start with an automatically generated caption that describes the main difference (Section 4.1). Humans then review it, either accepting or correcting it and expand it to include additional differences if artifacts are present, as shown in Figure 2.

2.2 Human Annotation

We employed Amazon Mechanical Turk (AMT) to evaluate image edits using human annotators, as shown in Figure 2, with three annotations per edit. Quality annotators were selected through a paid qualification test, and multiple steps were taken to ensure the instructions were clear and accessible in the UI (See Appendices A.6 and A.17).

2.3 Human Evaluation Analysis

Majority vote label distribution is presented in Table 1. Despite the task's subjectivity, agreement averaged 80% to 86%, compared to random chance of 25% for Accuracy and 33% for Artifacts. Majority agreement hit 95% for Accuracy and 97% for Artifacts. Full agreement among all annotators was achieved for 42% to 57% of edits. In 85% of examples, the edit reflected the instruction ("Accurate" or "Accurate, But Unexpected"), while 38% of edits contained significant artifacts.

The edit types were distributed as follows: Add 35.8%, Change Attribute 21.6%, Remove 7.3%, and Replace 31.3%, based on the image cap-

Category	Statistics (%)				
Accuracy Level	Accurate: 8% Accurate Unexpected: 77%	Inaccurate: 6% Inaccurate Reflects: 4%			
Artifacts Level	Significant: 38%	Mild: 57% No Artifact: 2%			
Technical Precision	Yes: 31%	No: 69%			
Visual Consistency	Yes: 18%	No: 82%			
Diff Caption Accuracy	Yes: 60%	No: 40%			

Table 1: Distribution of majority vote labels across categories. In 85% of examples, the edit reflected the instruction ("Accurate" or "Accurate, But Unexpected"), while 38% of edits contained significant artifacts.

tion pipeline metadata in Section 4.1. Figure 3 shows the percentage of issues reported by annotators in the Contextual Consistency and Technical Precision feedback, with resolution and shape/proportion concerns being particularly prominent. See Appendix A.9 for a full overview.

3 Auto-Evaluation

Using the EditInspector benchmark, we evaluate the ability of SoTA VLMs to serve as edit inspectors. The evaluation consists of two components: the first assesses the models' ability to verify edit accuracy and alignment with user expectations, while the second serves as an upper-bound test, examining their ability to generate captions that describe the main differences and all differences, including unintended artifacts (Section 3.3.2).

3.1 Models

We evaluate GPT-4, GPT-4o, GPT-4-turbo (OpenAI, 2024), Gemini-Pro-Vision (Gemini Team, 2023), Gemini-Pro-1.5 (Gemini Team, 2024), Qwen2.5-VL (Bai et al., 2025) and InternVL3 (Zhu et al., 2025) on all tasks using their latest versions as of August 2024 (Section A.12). We prioritized prompts that best conveyed user instructions and improved overall performance (See Appendix A.7).

3.2 Auto-Evaluation Setup

Edit Inspector Questions. Preliminary experiments revealed that models struggled to handle multiple categories, especially in detecting mild artifacts. To enhance clarity and relevance, we simplified the categorization by replacing multiple-



Figure 3: Frequency of issues identified by human annotators in the Contextual Consistency and Technical Precision textual feedback. Shape/Proportion concerns being particularly prominent.

choice questions with binary questions. For the accuracy question, both "Accurate" and "Accurate But Unexpected" were grouped under "Accurate," while in the artifacts question, only "Significant Artifacts" were counted as artifacts.

Difference Caption Generation. Traditional caption metrics (BLEU, METEOR, ROUGE, CIDEr) rely on N-gram overlaps but fail to distinguish edited objects, penalize stylistic variations, ignore edit sequences, and miss semantic misalignments. As shown in Table 2, these limitations lead to misleadingly high scores for incorrect captions. Section A.1 provides further examples and analysis.

To address these limitations, we propose two novel evaluation metrics tailored for **all differences caption** comparisons: Model Precision (MP) and Hallucination Rate (HR). MP is the percentage of human-annotated differences matching modeldetected ones, while HR is the percentage of modeldetected differences that do not correspond to any human-annotated differences.

We calculate these metrics by generating Difference Triplets (DTs) with the source object, target object, and action type for each change in the model and human captions. The two resulting sets of DTs are then used to compute MP and HR. A match between two DTs is determined if the edit action types are identical, and the source and target objects are similar, as evaluated by GPT-40. The

Example	Metrics
Ground Truth Caption: The main difference is	
the first image has a blue vase, and the second	MP : 0
image has a brown vase.	BL: 0.68
Generated Caption: The main difference is the	RO: 0.81
first image has a squirrel, and the second image	ME: 0.78
does not.	
Ground Truth Caption: A brown squirrel was	

added to the image.	MP : 1
Generated Caption: The difference between the	BL: 0.55
two images is that the first image has a blue vase	. RO: 0.60
The second image has a blue vase and a squirrel	ME: 0.57
next to it.	

Ground Truth Caption: In the first image, the **MP**: 0 tree was removed, and new flowerbed was added. BL: 0.73 **Generated Caption:** In the first image, the RO: 0.79 flowerbed was removed, and new tree was added. ME: 0.76

Table 2: Comparison of traditional linguistic metrics (BLEU, ROUGE, METEOR) against our proposed evaluation metric (MP). The first example shows high scores despite missing the edited object. The second penalizes correct but longer captions. The third fails to detect reversed edits, while our metric captures these issues.

similarity check between source and target objects is relaxed, allowing matches for objects with different attributes. A stricter check would have caused models to fail completely.

In addition, we introduced MP_{soft} and HR_{soft} , which count DT matches also in case of a reversed source and target object match, offering a more comprehensive analysis of model performance. See Section A.2 for mathematical formulations of the metrics, and Section A.4 for an intuitive example.

We evaluate the model's **main difference caption** by comparing it to the main difference extracted from the human-provided difference caption, which describes all of the edit's differences. GPT-4 is used to assess whether the main modelidentified difference matches the human one. Extracting the main difference is not complex, as the main change is typically mentioned first.

3.3 Auto-Evaluation Results

3.3.1 Edit Inspector Questions Results The results for the Yes/No questions are presented in Table 3. GPT-40 achieved the highest score on most questions except 'Edit Accuracy' and 'Difference Caption Accuracy', where Gemini-1.5 scored the highest. Below, we summarize our main observations from these results.

Struggling with Inaccurate Edits and Artifact Classification. Detection of inaccurate edits was challenging, with most models correctly classifying only 0-30%, except Qwen2.5-VL(39%) and GPT-40 (47%). All models mistakenly predicted edits as visually consistent, with precision scores between 0-22.3%. Differentiating artifacts from non-artifacts was also challenging. While GPT-40 had the highest accuracy (65.7%) it missed many artifacts with low recall (52.7%). All models frequently misclassified non-artifacts (18-30%), with Gemini misclassifying 72%.

Assessing the accuracy of inconsistent edits is challenging. There is a strong conditional dependency between the edit accuracy and contextual consistency questions. A discrepancy up to 40% was observed in the accuracy question when edits lacked contextual consistency. Conversely, models had difficulty with the contextual consistency question in accurate edits, with a 23% drop in performance. This dependency was also present (up to 12%) between the caption accuracy and contextual consistency questions.

Remove edits are challenging for models. While Gemini 1.5, GPT-4, GPT-4-turbo, and InternVL3 struggled with 'Remove' edits, showing accuracy gaps up to 65% in edit accuracy, GPT-4o excelled with 91% accuracy, making it the only model to handle these edits well.

Alongside Yes/No questions, we assessed models' feedback on Contextual Consistency and Technical Precision, finding it misaligned with human feedback in most cases (see Appendix A.8).

3.3.2 Difference Caption Generation Results

Main Differences Captions: Table 3 shows the percentage of instances where the model-identified main difference matched the human-reported one, with GPT-40 leading at 39% accuracy. Across all models, performance improved by up to 98% with accurate edits, a trend also seen in generating all differences captions. 'Remove' edits had the lowest performance, with accuracy dropping by up to 50% compared to the best-performing 'Replace' edits.

All Differences Captions: Table 3 shows that GPT-40 and Qwen2.5-VL achieves the highest Model Precision (MP) at 12%, while Qwen2.5-VL demonstrates the lowest Hallucination Rate (HR) at 58%, along with notable improvements in soft metrics, suggesting confusion between source and target objects. Overall performance remains suboptimal, as model predictions often misalign with human annotations. On average, models describe 1-3.1 differences per image, whereas human annotators identified six differences on average. This

	Gemini	Gemini-1.5	GPT-4	GPT-40	GPT-4 Turbo	Qwen2.5 VL	InternVL3	LLaVA	LLaVA (Supervised)
	Edi	t Inspectors Q	uestions						
Accuracy	49.9%	70.3%	67.3%	67.8%	66.9%	67.7%	70.2%	58.9%	67.2%
Contextual Consistency	50.4%	51.1%	50.4%	55.7%	48.2%	49.5%	49.2%	52.0%	-
Technical Precision	50.1%	46.3%	53.7%	55%	49.3%	48.4%	46.7%	50.1%	-
Artifacts	49.4%	58.5%	50.7%	65.7%	52.8%	50%	49.8%	47.6%	51.7%
Difference Caption Acc	53.9%	66.3%	63.9%	64.3%	64%	58.2%	58.2%	50.0%	54.5%
	Differe	nces Caption	Generatio	n					
Main Difference	31%	31%	27%	39%	24%	38%	26%	8%	10%
MP	-	8%	8%	12%	8%	12%	7%	-	-
MP _{soft}	-	9%	10%	14%	9%	14%	10%	-	-
HR	-	67%	78%	60%	75%	58%	87%	-	-
HR _{soft}	-	65%	75%	56%	72%	52%	83%	-	-
Avg. Diff	-	1	2.5	1.8	1.5	1.5	3.1	-	-
No Diffs	-	24%	0.7%	0.3%	6%	0.7%	3.4%	-	-

Table 3: Combined performance on Edit Inspectors questions, and the Difference Caption Generation task. GPT-40 model demonstrates the best performance in Edit Inspectors questions, achieving the highest or third-highest scores across all questions. Qwen2.5-VL achieves the highest precision in predicting differences, with the lowest hallucination rate. Avg. Diff indicates the average number of differences detected per edit, while No Diffs represents the percentage of edits where no differences were predicted. Human annotators identified an average of 6 differences per edit. The main difference row reports the percentage of predicted main difference captions correctly describing the main difference. The LLaVA (Supervised) column presents the performance of the finetuned model; see Section 5.3 for further analysis.

gap highlights models' difficulty in capturing subtle differences and their tendency to overlook details or introduce hallucinated changes.

Additionally, we observed that models tend to hallucinate less where the edits are accurately performed, leading to a 29% improvement in HR and a threefold increase in MP across all models.

Models vary significantly in predicting no differences between images. For example, Gemini-1.5 predicts no differences in 24% of the examples, compared to only 0.3% for GPT-40. Gemini-1.5's higher rate of "no difference" predictions lowers its HR but causes it to identify fewer differences than GPT-40, which detects 80% more differences while keeping a lower HR. When the edit is contextually consistent, most models predict no differences up to 3 times more often, suggesting they are more sensitive to semantic flaws than pixel-level ones.

Models struggle with Remove edits while excelling in Add edits. All models perform best on Add edits and worst on Remove edits, with Model Precision (MP) differing by up to 2.7x. The Hallucination Rate (HR) for Remove edits is significantly worse, increased up to 50% compared to Add edits.

Models are sensitive to scene complexity (i.e.,

the number of objects). Figure 6 in the Appendix shows that as the number of objects increases, all models exhibit declining precision and rising hallucination rates. GPT-4 and GPT-4-turbo, in particular, struggle more with complex scenes, showing sharp increases in hallucinations. While Gemini-1.5 and GPT-40 also degrade, their decline is less steep. This trend was not observed in the Edit Inspector questions (Yes/No questions).

4 New Methods

To tackle the challenges models face in generating accurate difference captions and detecting unintended artifacts, we developed a zeroshot pipeline for producing detailed, instructiongrounded captions (Section 4.1) and an artifact detection method using segmentation model probabilities (Section 4.2). Our methods are competitive with the best models, and in the main difference generation task outperform them by 36% margin.

4.1 Difference Caption Pipeline

Our pipeline generates detailed, instructiongrounded difference captions and rich metadata by selecting image captions of the edited object area



Figure 4: Example of our pipeline generating an instruction-grounded difference caption with rich metadata. Edit images are split into three zoom levels, with Gemini extracting and prioritizing captions to generate the metadata.

that align with the edit instructions. It achieves 75% accuracy in describing the main edit, surpassing GPT-40's 39% accuracy.

The pipeline extracts image captions at three zoom levels around the edit in-painting mask for both the source and target images. We then select the captions that best match the edit instructions, measured by the number of shared nouns or their synonyms using WordNet (Fellbaum, 1998). Using these grounded captions and the edit instruction, we employ a one-shot prompt with GPT-4 (OpenAI, 2024) to generate a detailed difference caption with metadata, as shown in Figure 4.

We found this method most effective for generating a main difference caption. Other methods, such as asking object-specific questions or requesting long image descriptions, often resulted in significant hallucinations and incorrect or biased descriptions. This issue persisted with different visual backbones, such as GPT-4 (OpenAI, 2024), LLAVA 1.5 (Liu et al., 2024), etc. Integrating human instructions with edited area descriptions allow information sharing as seen in Figures 16, 22.

4.2 Artifact Detection

We developed two artifact detection methods using the extracted metadata from our pipeline. The first method uses the Detic model (Zhou et al., 2022) to analyze the segmentation probability of each object intersected by the edit mask. A drop of the probability score by more than 4% as a result of the edit is considered an artifact.

The second method identifies elements that intersect with the mask area, have disappeared from the image, and do not overlap with the edited object's bounding box. This often occurs when the mask is large, but the edited object is small.

Combined, our methods achieve 64% balanced accuracy in detecting "Significant" artifacts, competitive only with GPT-40 scoring 65.7%. Figure 5 shows the first artifact detection method. If the small car intersecting with the inpainting area had been unintentionally removed, it would illustrate the second method.

An oracle that combines the optimal predictions from GPT-40 and our artifact detection method reaches a score of 86.8% with 100% precision. **This indicates that our artifact method and GPT-40 correctly classify different sets of examples.**

5 Model Supervision

We introduce an end-to-end fine-tuned LLaVA (Language-Vision Alignment) model that rivals much larger models in performance. It offers edit evaluation abilities equivalent to SoTA models while significantly reducing computational costs, providing an efficient solution for AI-generated image edit evaluation.

We trained the model using the MagicBrush train set consists of 8,808 edits. A balanced set of 5,422 edits was used for artifact detection. For edit accuracy and caption generation, the entire set was used, and 31,059 training instances were produced using the two augmentation methods described below. Further details are provided in Appendix A.11.

5.1 Negative Edit Augmentation

This method generates negative edits by selecting a deceptive target object and producing corresponding metadata, including instructions and difference captions. In Figure 7, a similarly sized scene object



Figure 5: The first method for detecting artifacts using the Detic model for the edit "turn the stop sign to a lollipop". Comparing Detic probabilities for objects intersecting the turquoise in-painting mask between the pre-edit (left) and post-edit (right) images reveals two artifacts, the truck and small car, whose probability drops exceeds our threshold.

(an umbrella) was chosen as the deceptive target, and new metadata was generated using GPT-3.5 with few-shot prompting. For Add and Replace edits, the deceptive object is a visually similar absent object, like a cactus instead of a potted plant. For Change Attribute edits, attributes are modified, like altering a coat's color from blue to red.

5.2 Reverse Edit Augmentation

This augmentation focuses on reversing the edit using few-shot prompts with GPT-3.5. Add edits are changed to Remove edits, Replace edits involve switching the source and target objects, and Change Attribute edits reverse the attribute modification. For example, in Figure 7, the edit "Remove one potted plant" is reversed to "Add one potted plant." Applied on top of the negative augmentation, this process expands the dataset fourfold, providing comprehensive training data for our model.

5.3 Supervision Results

Our model demonstrates competitive performance against SoTA VLMs. As shown in Table 3, it outperforms Gemini, GPT-4, and GPT-4-turbo in artifact detection, with only Gemini-1.5 (58.5%) and GPT-40 (65.7%) performing better. For Edit Accuracy, it achieves 67.2%, surpassing Gemini (49.9%) and GPT-4 Turbo. It also maintains competitive performance in the Difference Caption Accuracy (54.5%), surpassing Gemini model (53.9%). These results validate our augmentation methods and highlight the value of our training data.

6 Additional Editing Methods

To keep up with the latest image editing models and provide more robust evaluations of models as Edit Inspectors, we used 100 edits from the MagicBrush test set to generate mask-guided image edits using UltraEdit (Zhao et al., 2024) and Imagen3 (Imagen-Team-Google et al., 2024).

We annotated these edits using the same methods described in Section 2.2, with the distribution shown in Table 5. Comparing the human labels, we found that MagicBrush edits achieved the highest overall accuracy (85.3%). Imagen3 had the highest "Visual Consistency" (37.7%), "Technical Precision" (61.6%), and the fewest edits without artifacts (10.10%). UltraEdit showed the lowest "Visual Consistency" (10.67%) and lowest accuracy rate (31%), highlighting variation in edit quality across models. These findings show that the latest image editing models exhibit notable weaknesses.

Tables 6, 7, and 8 show model performance as Edit Inspectors. We observe a consistent decline in performance on the Edit Inspectors Questions for both Imagen3 and UltraEdit edits. GPT-40 remains relatively strong in identifying main differences, but its performance, as well as that of other models, drops across core quality dimensions such as accuracy, technical precision, etc. In contrast, performance on Difference Caption Generation remains comparatively stable. These results further support our observation that current models struggle to evaluate edits comprehensively and frequently hallucinate when describing changes.

7 Related Work

Recent advances in text-guided image editing enable modifications via natural language (Sheynin et al., 2023; He et al., 2024; Wu et al., 2023a; Cui et al., 2023), with some models supporting multiturn refinement (Cui et al., 2023). Others use spatial masks for precise, localized edits (Avrahami et al., 2022; Nichol et al., 2022; Wang et al., 2023), which offer better control than text-only methods (Wang et al., 2023; Zhang et al., 2024a).

Edit quality is often measured using pixel-level similarity (L1/L2 norms) and CLIP-based cosine similarity (Radford et al., 2021). However, these metrics poorly align with human judgment (Basu et al., 2023), offering only quantitative scores without qualitative insights.

Image editing benchmarks like EditBench (Wang et al., 2023) and EditVal (Basu et al., 2023) assess editing models through automatic and human evaluations, focusing on instruction adherence and object or scene preservation. In contrast, our work evaluates models as edit inspectors on overlooked edit aspects such as scene integration, pixellevel issues, and artifact detection. We also introduce the category "Accurate, But Unexpected" to capture technically correct edits that deviate from user expectations and collect textual feedback and detailed difference captions to provide deeper insights into edit quality.

8 Conclusion

In this work, we introduce EditInspector, a public benchmark for assessing evaluators of text-guided image edits across several dimensions: accuracy, artifact detection, visual quality, seamless integration with the image scene, adherence to common sense, and the ability to describe edit-induced changes. Using the EditInspector benchmark, we show that state-of-the-art vision and language models perform poorly as edit inspectors and cannot effectively assess edits. To address these limitations, we propose two novel methods and fine-tune an edit inspector that outperforms these models. We hope that our benchmark and proposed methods will drive advancements in edit evaluation and inspire further research in this domain.

9 Future Work

Future work can refine difference caption generation and explore new approaches to address existing model limitations. Additionally, several directions could further expand the benchmark's coverage and evaluation capabilities:

- Incorporating complex multi-object and multi-operation edits: Expanding the benchmark to edits involving multiple objects and operations, such as simultaneously adding new objects while removing others, would enable a broader assessment of model generalization.
- **Supporting multi-turn edit evaluation:** Extending the benchmark to include sequential edits would allow evaluation of models' ability to maintain visual and semantic consistency across multiple editing steps.

10 Limitations

Our benchmark is based exclusively on the MagicBrush dataset for evaluating edits, which, while covering diverse scenarios, is limited to natural images and mask-guided edits. Recent studies have shown promising results with free-text methods (Sheynin et al., 2023) and growing interest in editing of synthetic images. Additionally, the distribution of edit types in the test set reflects the natural distribution of human edits from the MagicBrush dataset, as determined by a human study. While this mirrors real-world editing trends, it may not equally represent all edit types. These limitations highlight distinct research directions that could be explored independently of our current work.

References

- Omri Avrahami, Dani Lischinski, and Ohad Fried. 2022. Blended diffusion for text-driven editing of natural images. In 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. 2025. Qwen2.5-vl technical report. *Preprint*, arXiv:2502.13923.
- Samyadeep Basu, Mehrdad Saberi, Shweta Bhardwaj, Atoosa Malemir Chegini, Daniela Massiceti, Maziar Sanjabi, Shell Xu Hu, and Soheil Feizi. 2023. Editval: Benchmarking diffusion based text-guided image editing methods. *Preprint*, arXiv:2310.02426.

- Tim Brooks, Aleksander Holynski, and Alexei A. Efros. 2023. Instructpix2pix: Learning to follow image editing instructions. *Preprint*, arXiv:2211.09800.
- Guillaume Couairon, Jakob Verbeek, Holger Schwenk, and Matthieu Cord. 2022. Diffedit: Diffusionbased semantic image editing with mask guidance. *Preprint*, arXiv:2210.11427.
- Xing Cui, Zekun Li, Peipei Li, Yibo Hu, Hailin Shi, and Zhaofeng He. 2023. Chatedit: Towards multi-turn interactive facial image editing via dialogue. *Preprint*, arXiv:2303.11108.
- Christiane Fellbaum. 1998. *WordNet: An Electronic Lexical Database*. Bradford Books.
- Gemini Team. 2023. Gemini: A family of highly capable multimodal models. *Preprint*, arXiv:2312.11805.
- Gemini Team. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *Preprint*, arXiv:2403.05530.
- Yingqing He, Zhaoyang Liu, Jingye Chen, Zeyue Tian, Hongyu Liu, Xiaowei Chi, Runtao Liu, Ruibin Yuan, Yazhou Xing, Wenhai Wang, Jifeng Dai, Yong Zhang, Wei Xue, Qifeng Liu, Yike Guo, and Qifeng Chen. 2024. Llms meet multimodal generation and editing: A survey. *Preprint*, arXiv:2405.19334.
- Imagen-Team-Google, :, Jason Baldridge, Jakob Bauer, Mukul Bhutani, Nicole Brichtova, Andrew Bunner, Lluis Castrejon, Kelvin Chan, Yichang Chen, Sander Dieleman, Yuqing Du, Zach Eaton-Rosen, Hongliang Fei, Nando de Freitas, Yilin Gao, Evgeny Gladchenko, Sergio Gómez Colmenarejo, Mandy Guo, Alex Haig, Will Hawkins, Hexiang Hu, Huilian Huang, Tobenna Peter Igwe, Christos Kaplanis, Siavash Khodadadeh, Yelin Kim, Ksenia Konyushkova, Karol Langner, Eric Lau, Rory Lawton, Shixin Luo, Soňa Mokrá, Henna Nandwani, Yasumasa Onoe, Aäron van den Oord, Zarana Parekh, Jordi Pont-Tuset, Hang Qi, Rui Qian, Deepak Ramachandran, Poorva Rane, Abdullah Rashwan, Ali Razavi, Robert Riachi, Hansa Srinivasan, Srivatsan Srinivasan, Robin Strudel, Benigno Uria, Oliver Wang, Su Wang, Austin Waters, Chris Wolff, Auriel Wright, Zhisheng Xiao, Hao Xiong, Keyang Xu, Marc van Zee, Junlin Zhang, Katie Zhang, Wenlei Zhou, Konrad Zolna, Ola Aboubakar, Canfer Akbulut, Oscar Akerlund, Isabela Albuquerque, Nina Anderson, Marco Andreetto, Lora Aroyo, Ben Bariach, David Barker, Sherry Ben, Dana Berman, Courtney Biles, Irina Blok, Pankil Botadra, Jenny Brennan, Karla Brown, John Buckley, Rudy Bunel, Elie Bursztein, Christina Butterfield, Ben Caine, Viral Carpenter, Norman Casagrande, Ming-Wei Chang, Solomon Chang, Shamik Chaudhuri, Tony Chen, John Choi, Dmitry Churbanau, Nathan Clement, Matan Cohen, Forrester Cole, Mikhail Dektiarev, Vincent Du, Praneet Dutta, Tom Eccles, Ndidi Elue, Ashley Feden, Shlomi Fruchter, Frankie Garcia, Roopal Garg, Weina Ge, Ahmed Ghazy, Bryant Gipson, Andrew Goodman, Dawid Górny, Sven

Gowal, Khyatti Gupta, Yoni Halpern, Yena Han, Susan Hao, Jamie Hayes, Jonathan Heek, Amir Hertz, Ed Hirst, Emiel Hoogeboom, Tingbo Hou, Heidi Howard, Mohamed Ibrahim, Dirichi Ike-Njoku, Joana Iljazi, Vlad Ionescu, William Isaac, Reena Jana, Gemma Jennings, Donovon Jenson, Xuhui Jia, Kerry Jones, Xiaoen Ju, Ivana Kajic, Christos Kaplanis, Burcu Karagol Ayan, Jacob Kelly, Suraj Kothawade, Christina Kouridi, Ira Ktena, Jolanda Kumakaw, Dana Kurniawan, Dmitry Lagun, Lily Lavitas, Jason Lee, Tao Li, Marco Liang, Maggie Li-Calis, Yuchi Liu, Javier Lopez Alberca, Matthieu Kim Lorrain, Peggy Lu, Kristian Lum, Yukun Ma, Chase Malik, John Mellor, Thomas Mensink, Inbar Mosseri, Tom Murray, Aida Nematzadeh, Paul Nicholas, Signe Nørly, João Gabriel Oliveira, Guillermo Ortiz-Jimenez, Michela Paganini, Tom Le Paine, Roni Paiss, Alicia Parrish, Anne Peckham, Vikas Peswani, Igor Petrovski, Tobias Pfaff, Alex Pirozhenko, Ryan Poplin, Utsav Prabhu, Yuan Qi, Matthew Rahtz, Cyrus Rashtchian, Charvi Rastogi, Amit Raul, Ali Razavi, Sylvestre-Alvise Rebuffi, Susanna Ricco, Felix Riedel, Dirk Robinson, Pankaj Rohatgi, Bill Rosgen, Sarah Rumbley, Moonkyung Ryu, Anthony Salgado, Tim Salimans, Sahil Singla, Florian Schroff, Candice Schumann, Tanmay Shah, Eleni Shaw, Gregory Shaw, Brendan Shillingford, Kaushik Shivakumar, Dennis Shtatnov, Zach Singer, Evgeny Sluzhaev, Valerii Sokolov, Thibault Sottiaux, Florian Stimberg, Brad Stone, David Stutz, Yu-Chuan Su, Eric Tabellion, Shuai Tang, David Tao, Kurt Thomas, Gregory Thornton, Andeep Toor, Cristian Udrescu, Aayush Upadhyay, Cristina Vasconcelos, Alex Vasiloff, Andrey Voynov, Amanda Walker, Luyu Wang, Miaosen Wang, Simon Wang, Stanley Wang, Qifei Wang, Yuxiao Wang, Ágoston Weisz, Olivia Wiles, Chenxia Wu, Xingyu Federico Xu, Andrew Xue, Jianbo Yang, Luo Yu, Mete Yurtoglu, Ali Zand, Han Zhang, Jiageng Zhang, Catherine Zhao, Adilet Zhaxybay, Miao Zhou, Shengqi Zhu, Zhenkai Zhu, Dawn Bloxwich, Mahyar Bordbar, Luis C. Cobo, Eli Collins, Shengyang Dai, Tulsee Doshi, Anca Dragan, Douglas Eck, Demis Hassabis, Sissie Hsiao, Tom Hume, Koray Kavukcuoglu, Helen King, Jack Krawczyk, Yeqing Li, Kathy Meier-Hellstern, Andras Orban, Yury Pinsky, Amar Subramanya, Oriol Vinyals, Ting Yu, and Yori Zwols. 2024. Imagen 3. Preprint, arXiv:2408.07009.

- Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. 2023. Imagic: Text-based real image editing with diffusion models. *Preprint*, arXiv:2210.09276.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024. Improved baselines with visual instruction tuning. *Preprint*, arXiv:2310.03744.
- Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. 2022. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *Preprint*, arXiv:2112.10741.

- OpenAI. 2024. GPT-4 technical report. *Preprint*, arXiv:2303.08774.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. *Preprint*, arXiv:2103.00020.
- Shelly Sheynin, Adam Polyak, Uriel Singer, Yuval Kirstain, Amit Zohar, Oron Ashual, Devi Parikh, and Yaniv Taigman. 2023. Emu edit: Precise image editing via recognition and generation tasks. *Preprint*, arXiv:2311.10089.
- Su Wang, Chitwan Saharia, Ceslee Montgomery, Jordi Pont-Tuset, Shai Noy, Stefano Pellegrini, Yasumasa Onoe, Sarah Laszlo, David J. Fleet, Radu Soricut, Jason Baldridge, Mohammad Norouzi, Peter Anderson, and William Chan. 2023. Imagen editor and editbench: Advancing and evaluating text-guided image inpainting. *Preprint*, arXiv:2212.06909.
- Chenfei Wu, Shengming Yin, Weizhen Qi, Xiaodong Wang, Zecheng Tang, and Nan Duan. 2023a. Visual chatgpt: Talking, drawing and editing with visual foundation models. *Preprint*, arXiv:2303.04671.
- Tsung-Han Wu, Long Lian, Joseph E. Gonzalez, Boyi Li, and Trevor Darrell. 2023b. Selfcorrecting llm-controlled diffusion models. *Preprint*, arXiv:2311.16090.
- Kai Zhang, Lingbo Mo, Wenhu Chen, Huan Sun, and Yu Su. 2024a. Magicbrush: A manually annotated dataset for instruction-guided image editing. *Preprint*, arXiv:2306.10012.
- Shu Zhang, Xinyi Yang, Yihao Feng, Can Qin, Chia-Chih Chen, Ning Yu, Zeyuan Chen, Huan Wang, Silvio Savarese, Stefano Ermon, Caiming Xiong, and Ran Xu. 2024b. Hive: Harnessing human feedback for instructional visual editing. *Preprint*, arXiv:2303.09618.
- Zhixing Zhang, Ligong Han, Arnab Ghosh, Dimitris Metaxas, and Jian Ren. 2022. Sine: Single image editing with text-to-image diffusion models. *Preprint*, arXiv:2212.04489.
- Haozhe Zhao, Xiaojian Ma, Liang Chen, Shuzheng Si, Rujie Wu, Kaikai An, Peiyu Yu, Minjia Zhang, Qing Li, and Baobao Chang. 2024. Ultraedit: Instructionbased fine-grained image editing at scale. *Preprint*, arXiv:2407.05282.
- Xingyi Zhou, Rohit Girdhar, Armand Joulin, Philipp Krähenbühl, and Ishan Misra. 2022. Detecting twenty-thousand classes using image-level supervision. In *ECCV*.
- Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, Zhangwei Gao, Erfei Cui, Xuehui Wang, Yue Cao, Yangzhou Liu, Xingguang Wei,

Hongjie Zhang, Haomin Wang, Weiye Xu, Hao Li, Jiahao Wang, Nianchen Deng, Songze Li, Yinan He, Tan Jiang, Jiapeng Luo, Yi Wang, Conghui He, Botian Shi, Xingcheng Zhang, Wenqi Shao, Junjun He, Yingtong Xiong, Wenwen Qu, Peng Sun, Penglong Jiao, Han Lv, Lijun Wu, Kaipeng Zhang, Huipeng Deng, Jiaye Ge, Kai Chen, Limin Wang, Min Dou, Lewei Lu, Xizhou Zhu, Tong Lu, Dahua Lin, Yu Qiao, Jifeng Dai, and Wenhai Wang. 2025. Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. *Preprint*, arXiv:2504.10479.

A Appendix

A.1 Common Caption Comparison Metrics

Common metrics for comparing image captions, such as BLEU, METEOR, and ROUGE, rely on N-gram overlaps between generated and reference texts. However, they fall short of our core requirement to ensure accurate alignment between the edited objects and actions described in the captions. As shown in Table 4, while these metrics suggest that GPT-4 generates captions most similar to the ground truth, in practice, it is the least accurate model, exhibiting the highest hallucination rate and the largest number of average changes detected. Below, we provide a brief explanation of these metrics, followed by several scenarios illustrating their limitations in effectively evaluating difference captions.

- BLEU: Computes the number of matches in unigrams, bigrams, trigrams, and 4-grams between generated and reference text. Includes a brevity penalty to discourage shorter outputs.
- ROUGE: ROUGE-1 calculates the F1 score for unigrams. ROUGE-2 calculates the F1 score for bigrams.
- METEOR: Incorporates features such as stemming, synonym matching, and paraphrase recognition. Computes the unigram F1 score.
- CIDEr: Measures the similarity between generated and reference captions using TF-IDF weighted n-grams (unigrams to 4-grams). Emphasizes consensus between generated captions and multiple human references while penalizing overuse of common n-grams.

Although these metrics are widely used in image captioning, they have severe limitations when evaluating difference captions for image edits.

	Gemini-1.5	GPT-4	GPT-40	GPT-4 Turbo		
Differences Caption Generation						
Main Difference	31%	27%	39%	24%		
MP	8%	8%	12%	8%		
HR	67%	78%	60%	75%		
METEOR	0.11	0.22	0.19	0.19		
ROUGE-1	0.15	0.36	0.29	0.30		
ROUGE-2	0.04	0.09	0.08	0.07		
BLEU	0.01	0.02	0.03	0.02		

Table 4: Comparison of models on the Difference Caption Generation task. GPT-4 achieves the best results on METEOR, ROUGE-1, and ROUGE-2 metrics, while GPT-40 ranks highest in BLEU.

Miss Weighting the Edited Objects and Actions. These metrics struggle to differentiate between critical objects and less significant words in the context of difference captions. For instance, consider the ground truth caption: "The main difference between the two images is the first image has a blue vase and the second image a brown vase." If the generated caption states, "The main difference between the two images is the first image has a squirrel and the second image does not," linguistic metrics might still assign relatively high scores (e.g., BLEU: 0.68, ROUGE-1 Recall: 0.81, METEOR: 0.78) due to superficial word overlaps. However, these scores fail to reflect the semantic misalignment between the captions. In contrast, our proposed metric assigns a score of 0, accurately reflecting the discrepancy in the identified edited object and action.

Accounting for Unchanged Objects, Varying Length, and Stylistic Differences. Conventional metrics often penalize captions that include mentions of unchanged objects, vary in length, or differ stylistically, even when accurately describing the detected changes. For instance, consider the generated caption: "The difference between the two images is that the first image has a blue vase. The second image has a blue vase and a squirrel next to it." Our metric would assign this caption a perfect score of 1, as it correctly identifies the key difference (the addition of the squirrel) in alignment with the ground truth caption: "A brown squirrel was added to the image." In contrast, linguistic metrics would score close to 0 due to the inclusion of details about the unchanged "blue vase" and penalties for variations in length and phrasing. This demonstrates the robustness of our metric in handling linguistic variability while focusing on the accuracy of detected changes.

Capturing the Order of Edits. The above mentioned metrics overlook the importance of edit sequence order. For instance, consider the ground truth captions: "In the first image, the tree was removed, and a new flowerbed was added" and the generated caption "In the first image, the flowerbed was removed, and a new tree was added." Although both captions involve the same objects (tree and flowerbed) and actions (added and removed), the sequence of edits conveys entirely different meanings. The n-gram based metrics would assign high scores to these captions because they mention the same words (objects and actions), regardless of their order, failing to penalize semantic misalignment. In contrast, our metric explicitly evaluates the edit sequence order, ensuring that generated captions accurately reflect the correct sequence of changes.

A.2 Mathematical Explanation of Metrics

We evaluate model performance on **all differences captions** using two metrics: **Model Precision** (**MP**) and **Hallucination Rate** (**HR**). These are computed based on Difference Triplets (DTs), defined as:

DT = (source object, target object, action type),

where *source object* is the original object affected by the edit, *target object* is the resulting object of the edit, and *action type* is the type of edit (e.g., "add," "remove," "replace"). **Model Precision (MP):** Measures the percentage of humanannotated DTs (\mathcal{H}) matched by model-detected DTs (\mathcal{M}):

$$\mathrm{MP} = \frac{|\mathcal{H} \cap \mathcal{M}|}{|\mathcal{H}|} \times 100,$$

where $|\mathcal{H} \cap \mathcal{M}|$ is the number of matched DTs, and $|\mathcal{H}|$ is the total human-annotated DTs. **Hallucination Rate (HR):** Measures the percentage of model-detected DTs (\mathcal{M}) not matching any humanannotated DTs (\mathcal{H}):

$$\mathrm{HR} = \frac{|\mathcal{M} \setminus \mathcal{H}|}{|\mathcal{M}|} \times 100,$$

where $|\mathcal{M} \setminus \mathcal{H}|$ is the number of hallucinated DTs, and $|\mathcal{M}|$ is the total model-detected DTs. **Soft Metrics:** MP_{soft} and HR_{soft} allow matches when source and target objects in DTs are reversed:

$$\begin{split} \mathbf{MP}_{\text{soft}} &= \frac{|\mathcal{H}_{\text{soft}} \cap \mathcal{M}|}{|\mathcal{H}|} \times 100, \\ \mathbf{HR}_{\text{soft}} &= \frac{|\mathcal{M} \setminus \mathcal{H}_{\text{soft}}|}{|\mathcal{M}|} \times 100. \end{split}$$

Matching Criteria: A DT match requires identical *action type* and similar *source/target objects* (assessed by GPT-4). Relaxed matching (\mathcal{H}_{soft}) accounts for reversed source and target objects.

A.3 Additional Editing Methods

UltraEdit Annotation Agreement. Accuracy stood out with a high complete majority rate (78%) and the highest average agreement rate (92.67%). For the more detailed **accuracy levels**, majority vote was reached in 81% of cases, though the complete majority rate dropped to 53%, with an average agreement rate of 95% among those with a majority vote.

In the **artifacts** category, annotators reached a 64% complete majority rate, with an average agreement rate of 88%. The more granular **artifact levels** followed a similar pattern, with 98% majority vote, 52% complete majority rate, and 82.67% average agreement.

Technical Precision showed a complete majority in only 36% of cases, with a moderate average agreement rate of 78.67%. Finally, **Visual Consistency** achieved a strong 91% complete majority rate, and a 73% average agreement rate, reflecting consistent annotator judgment in this category.

Imagen3 Annotation Agreement. Accuracy achieved a complete majority rate of 83.84% and the highest average agreement rate (94.61%) across all categories. For the more detailed **accuracy levels**, majority vote was reached in 95.96% of cases, though the complete majority rate dropped to 36.36%, with an average agreement rate of 76.09%.

In the **artifacts** category, annotators reached a 58.59% majority vote and a 64% complete majority rate, with an average agreement rate of 86.20%.

The more granular **artifact levels** showed a 94.95% majority vote, 46.46% complete majority rate, and 78.79% average agreement.

Technical Precision had a complete majority rate of 44. 44% and an average agreement rate of 81.48%. **Visual Consistency** achieved a 45.45% complete majority rate and 81.82% average agreement, indicating relatively stable annotator consensus in this category.

MagicBrush Annotation Agreement. Accuracy levels in MagicBrush exhibited strong annotator alignment, with a 95% majority vote rate and an average agreement rate of 80.67%, though the complete majority rate was moderate at 52%.

For the **artifacts** levels, annotators reached a 99% majority vote, with an average agreement rate of 83% and a complete majority rate of 51%, closely mirroring previous patterns observed in Imagen3.

Technical Precision (annotated as good quality) showed a complete majority in 41% of cases, with an average agreement rate of 80.33%.

Finally, **Visual Consistency** achieved a 57% complete majority rate and the highest average agreement rate in MagicBrush at 85.67%, suggesting relatively consistent annotator judgments in this category.

A.4 Metrics Example

We calculate the MP and HR metrics using Figure 1 GPT-40 and the human-annotated difference caption. The ground truth lists the following human-annotated differences (\mathcal{H}) :

(carpet floor, wooden floor, Replace),

(None, door, Add),

(fridge bottom, extended fridge bottom, Change),

(yellow box, extended yellow box, Change),

(yellow box text, None, Remove),

(text, image, Replace)

GPT-40 detects only one difference:

 $\mathcal{M} = \{(\text{carpet floor}, \text{wooden floor}, \text{Replace})\}.$

Model Precision (MP): Model Precision (MP) measures the percentage of human-annotated DTs (\mathcal{H}) matched by model-detected DTs (\mathcal{M}) :

$$\mathbf{MP} = \frac{|\mathcal{H} \cap \mathcal{M}|}{|\mathcal{H}|} \times 100.$$

Category	Imagen3 (%)	UltraEdit (%)	MagicBrush (%)
Accuracy Level	Accurate: 15.82 Accurate Unexpected: 46.46 Inaccurate: 18.86 Inaccurate Reflects: 18.86	Accurate: 6.3 Accurate Unexpected: 46 Inaccurate: 31 Inaccurate Reflects: 16.7	Accurate: 16 Accurate Unexpected: 69.3 Inaccurate: 7 Inaccurate Reflects: 7.7
Artifacts Level	Significant: 36.7 Mild: 53.2 No Artifact: 10.1	Significant: 29.3 Mild: 65 No Artifact: 5.7	Significant: 36.7 Mild: 58 No Artifact: 5.3
Technical Precision	Yes: 61.6 No: 38.3	Yes: 37.6 No: 62.3	Yes: 39 No: 61
Visual Consistency	Yes: 37.7 No: 62.3	Yes: 10.7 No: 89.3	Yes: 22 No: 78

Table 5: Distribution of annotation values across categories for Imagen3, UltraEdit, and MagicBrush. The table summarizes the percentage breakdown for each evaluation category. Notably, MagicBrush had the highest percentage of overall Accurate edits, while Imagen3 showed the strongest performance in Technical Precision.

	GPT-4	GPT-4o	GPT-4 Turbo	Qwen2.5 VL	InternVL3	LLaVA	LLaVA (Supervised)
	Ed	it Inspecto	ors Questi	ons			
Accuracy	63%	51.8%	57.5%	62.1%	54.4%	52.6%	54.3%
Contextual Consistency	48.5%	41.6%	37.1%	58.2%	61.9%	-	-
Technical Precision	53.4%	46.4%	48.5%	43.8%	46.2%	-	-
Artifacts	42.9%	56.3%	51%	50%	41%	52.7%	55.6%
Difference Caption Acc	55.6%	55.6%	57.5%	54.4%	44.4%	50%	48.1%
	Differ	ences Cap	tion Gene	ration			
Main Difference	36%	44%	36%	8%	22%	4%	9%
MP	7%	9%	7%	7%	8%	-	-
MPsoft	9%	11%	8%	9%	12%	-	-
HR	82%	74%	80%	84%	90%	-	-
HR _{soft}	80%	68%	77%	80%	86%	-	-
Avg. Diff No Diffs	2.5 0.8%	1.9 0%	1.5 4.5%	2.5 0.8%	3.4 3.2%	-	-

Table 6: Models performance on UltraEdit edits across models Edit Inspectors questions and Difference Caption Generation. The first section reports binary question accuracy on core evaluation criteria (Accuracy, Contextual Consistency, Technical Precision, and Artifacts). The second section presents difference caption metrics: percentage of predicted main difference captions correctly describing the main difference, hallucination rates (HR), and average number of predicted differences. *Avg. Diff* indicates the mean number of differences per edit, and *No Diffs* reports the percentage of edits with no predicted differences.

	GPT-4	GPT-40	GPT-4 Turbo	Qwen2.5 VL	InternVL3	LLaVA	LLaVA (Supervised)
	Ed	it Inspecto	rs Questi	ons			
Accuracy	55.3%	49.9%	49.2%	52.9%	51.5%	54.4%	58.8%
Contextual Consistency	47.7%	49.5%	55.6%	42.1%	48.9%	-	-
Technical Precision	51.4%	50.9%	49.1%	47.5%	54.5%	-	-
Artifacts	47.1%	56.5%	49.5%	50.0%	48.6%	43.6%	48.8%
Difference Caption Acc	54.2%	58.2%	55.9%	51.8%	50.8%	50%	53.9%
	Differ	ences Cap	tion Gene	ration			
Main Difference	29%	44%	27%	29%	13%	4%	11%
MP	9%	11%	8%	10%	9%	-	-
MPsoft	11%	12%	9%	12%	12%	-	-
HR	78%	73%	82%	73%	92%	-	-
HR _{soft}	73%	70%	80%	67%	88%	-	-
Avg. Diff	2.5	1.9	1.5	1.5	3.2	-	-
No Diffs	0.8%	0%	4.5%	1.2%	4%	-	-

Table 7: Models performance on Imagen3 edits Edit Inspectors questions and Difference Caption Generation. The first section reports binary question accuracy on core evaluation criteria (Accuracy, Contextual Consistency, Technical Precision, and Artifacts). The second section presents difference caption metrics: percentage of predicted main difference captions correctly describing the main difference, hallucination rates (HR), and average number of predicted differences. *Avg. Diff* indicates the mean number of differences per edit, and *No Diffs* reports the percentage of edits with no predicted differences.

	GPT-4	GPT-40	GPT-4 Turbo	Qwen2.5 VL	InternVL3	LLaVA	LLaVA (Supervised)
	Ed	it Inspecto	ors Questi	ons			
Accuracy	73.7%	71%	60%	64%	63.6%	58.5%	62.3%
Contextual Consistency	48.2%	56.7%	59.7%	38.4%	45.6%	-	-
Technical Precision	49.2%	47.6%	49.8%	47.5%	42.6%	-	-
Artifacts	51.6%	60.6%	54.1%	50%	52.5%	44.8%	56.2%
Difference Caption Acc	63.2%	60.7%	61.1%	59.8%	58.4%	50%	46.5%
	Differ	ences Cap	tion Gene	ration			
Main Difference	34%	50%	34%	37%	25%	5%	12%
MP	8%	11%	7%	11%	8%	-	-
MPsoft	8%	12%	8%	1%	11%	-	-
HR	75%	61%	78%	63%	88%	-	-
HR _{soft}	74%	57%	75%	59%	83%	-	-
Avg. Diff	2.5	1.9	1.5	1.5	3.4	-	_
No Diffs	0.8%	0%	4.5%	0.6%	2.3%	-	-

Table 8: Models performance on MagicBrush edits Edit Inspectors questions and Difference Caption Generation. The first section reports binary question accuracy on core evaluation criteria (Accuracy, Contextual Consistency, Technical Precision, and Artifacts). The second section presents difference caption metrics: percentage of predicted main difference captions correctly describing the main difference, hallucination rates (HR), and average number of predicted differences. *Avg. Diff* indicates the mean number of differences per edit, and *No Diffs* reports the percentage of edits with no predicted differences.

The only match between \mathcal{H} and \mathcal{M} is:

(carpet floor, wooden floor, Replace)

Therefore:

$$|\mathcal{H} \cap \mathcal{M}| = 1, \quad |\mathcal{H}| = 6,$$

$$\mathbf{MP} = \frac{1}{6} \times 100 \approx 16.67\%.$$

Hallucination Rate (HR): Hallucination Rate (HR) measures the percentage of model-detected DTs (\mathcal{M}) that do not match any human-annotated DTs (\mathcal{H}):

$$\mathrm{HR} = \frac{|\mathcal{M} \setminus \mathcal{H}|}{|\mathcal{M}|} \times 100$$

Here, all model-detected DTs match humanannotated DTs, so:

$$\mathcal{M} \setminus \mathcal{H} = \emptyset, \quad |\mathcal{M}| = 1,$$
$$\mathrm{HR} = \frac{0}{1} \times 100 = 0\%.$$
$$\mathrm{MP} = 16.67\%, \quad \mathrm{HR} = 0\%.$$

A.5 Detailed Description of New Methods

Caption Pipeline We introduce a structured pipeline to generate difference captions that describe the specific visual changes between the source and edited images. This pipeline combines both image-level and region-level vision-language descriptions with a prompt-based large language model (LLM) querying strategy to produce rich and accurate textual metadata for each edit.

The process begins by extracting bounding boxes from the user-provided edit mask. If a single bounding box is found, we crop the source and target images in two ways: a tight crop and a padded version, where the box is either doubled in size or expanded to cover at least 15% of the image dimensions. If multiple bounding boxes are present or if cropping fails, we fall back to using the full images.

We then apply a vision-language model (Gemini) to generate descriptions of the cropped and full image regions. For each image, we obtain textual descriptions of the masked region, the full image, and the padded area (if applicable). These serve as candidate visual contexts for the next stage of the pipeline.

To ensure that the most relevant image region is used for caption generation, we apply a noun-based grounding mechanism to select the best image description for each image. We extract nouns from the instruction and compare them—using both exact matches and synonym overlap—to nouns in each candidate region description. If the default crop lacks sufficient alignment, we select the region (tight crop, padded crop, or full image) with the highest degree of noun-level overlap. When no region aligns directly, we fall back to the one that shares object mentions with at least one other candidate. This step ensures that the final prompt is grounded in a region of the image that is semantically aligned with the user instruction.

A predefined prompt template is then populated with the instruction and the selected image descriptions. This prompt is passed to GPT-4, which returns a structured response containing multiple fields: the predicted action type, short and extensive difference captions, a revised instruction, source and target object names, and a brief explanation of the edit.

To ensure reliability, the pipeline includes fallback mechanisms if the model returns an invalid action type (e.g., "None"), reattempting generation with alternative image regions. Only examples with valid, well-formed responses are retained.

Artifact Detection - Edit Mask Intersection Our segmentation-based artifact detection method analyzes object presence and score variation around the edited region to identify potential unintended artifacts. For each object detected by the Detic model in the source and target images, we compute the intersection between its segmentation mask and the user-provided edit mask. We use the object's binary mask contours rather than bounding boxes, allowing for precise spatial comparisons. To improve robustness, we perform these comparisons at two different resolutions: the source image size and the edited image size, since mask-based scores can fluctuate slightly (by a few percentage points) when resized.

To reduce noise, we exclude very small regions (objects with less than 2.4% intersection with the mask) as well as objects that are almost entirely contained within the mask (above 97% intersection), as these are typically intentional edit targets or too minor to evaluate reliably. In addition, we focus specifically on objects that are partially affected by the edit those whose segmentation masks intersect with the edit region by more than 0% but less than 40%. This range helps isolate objects that may have been unintentionally damaged during the edit process, as opposed to ones that were fully changed.

For each object class, we retain all detected instances that meet these intersection criteria. To avoid expensive pixel-level comparisons between non-overlapping objects, we first check whether their bounding boxes intersect. This provides a fast pre-filter, since if the bounding boxes do not overlap, their masks cannot intersect either. Only objects with overlapping bounding boxes are further compared across the source and edited images to measure changes in detection confidence. An object is flagged as containing an artifact if its confidence score drops by more than 4% after the edit, provided it is not fully masked or too small.

Artifact Detection - Unintended Object Additions or Removals In addition to evaluating score-based degradations near the edit mask, we introduce a second method designed to detect unintended additions or removals of secondary objects within the mask area. This method specifically applies to Add and Remove edits, where the risk of unintentionally modifying unrelated parts of the scene is higher. Unlike the first method, which looks at changes in detection confidence, this approach focuses on complete object disappearance or appearance that may not be aligned with the instruction.

The process begins by identifying the main edit object bounding boxes from the source and target images using object detection. These bounding boxes are used to define the intended region of change.

Next, we detect all other objects present in the source and edited images, and classify them as secondary objects. We then filter out any objects whose class labels are semantically similar to the main object. This includes direct matches, shared noun forms, and synonym relationships. This filtering step ensures that we focus only on truly unrelated objects that are not supposed to be changed.

From this filtered set of secondary objects, we further isolate those that intersect with the mask area but do not spatially overlap with the main object's bounding boxes. This spatial condition helps distinguish unintended changes from those that are part of the intended edit.

Finally, we compare the presence of these secondary objects across the source and edited images. For Remove actions, if a secondary object is present in the source image but missing in the edited image, we flag it as an unintended removal. For Add actions, if a new secondary object appears in the edited image that was not in the original, we flag it as an unintended addition. If at least one such change is detected, we mark the edit as containing a secondary artifact.

This method captures a different failure mode than the first: it identifies whole-object additions or losses in the masked area that are not part of the intended instruction. Full implementation details, including semantic class filtering, mask-intersection checks, and bounding box exclusion, are available in our released code.

A.6 Additional Annotation Information

Each image edit was annotated by three annotators, with annotations conducted in batches of 27-54 edits. Annotators were paid at a rate of \$0.70 per sample, resulting in an average hourly wage of \$18.

To ensure the quality of annotations, we implemented a qualification test to select quality annotators. We provided detailed instructions, including decision trees that visually guide the answering process. These decision trees were accessible via the user interface ("tree icon"), allowing annotators to follow the guidelines while annotating image edits.

Additionally, a settings window was available, enabling annotators to customize the UI, including font size, width, and padding, to suit their personal preferences (See Appendix A.16).

To assess annotation consistency, we computed Fleiss' Kappa for each question. The accuracy question showed moderate agreement ($\kappa > 0.41$), suggesting annotators generally aligned on whether an edit was accurate. The artifact question, artifact severity level, accuracy level, and difference caption accuracy all exhibited fair agreement (κ between 0.21 and 0.40), indicating moderate consistency across annotators. In contrast, the visual consistency and technical precision questions exhibited only slight agreement (κ between 0.01 and 0.20), highlighting greater subjectivity or ambiguity in how these aspects were interpreted.

A.7 Tasks Prompts

Model performance varied greatly with different prompts, suggesting that models may struggle to fully understand the task. We selected prompts that conveyed the user instructions and improved the overall performance.

• Difference Caption Accuracy Task (Yes/No)

You are provided with before and after images of an image edit for the edit instruction "{}". Does the difference caption "{}" describe the difference between the two images (Answer only Yes/No)?

• Visual Consistency Task (Yes/No) You are provided with before and after images of an image edit for the edit instruction "{}". Is the edited object or its area (in remove/replace actions) consistent with the edit instruction and the image scene in terms of shape, size, brightness, shadows, texture, color, etc. (Answer only Yes/No)?

• Is Accurate Task

You are provided with before and after images of an image edit for the edit instruction "{}". Was the edit instruction "{}" accurately executed and does it reflect the intended change (Answer only Yes/No)?

• Artifacts Task

You are provided with before and after images of an image edit for the edit instruction "{}". Are there any artifacts or alterations in the image not intended to be affected by the edit "{}" (Answer only Yes/No)?

• Technical Precision Task (Yes/No)

You are provided with before and after images of an image edit for the edit instruction "{}". Does the edited object or its area (in remove/replace actions) maintain the image resolution, exhibit blur, show any smoothness, etc. (Answer only Yes/No)?

- Generate all differences caption You are provided with before and after images of an image edit. Please describe all the differences between these two images. Focus only on the differences; do not include any irrelevant information. Ignore any style differences between the images, such as changes in artistic style, color grading, or filters.
- Generate main differences caption Please describe the main difference between the two images.

A.8 Textual Feedback

We compared the predicted feedback from the models with human annotations by using a zero-shot prompt with GPT-40 that determines whether two pieces of feedback share any common points (yielding a simple Yes or No). The models' feedback matched human feedback only in a very small percentage of cases. The contextual consistency feedback shared common points with human feedback in 7%-28% of cases, while technical precision feedback did so in 4%-51% of instances.

A.9 Categories of Feedback Issues

- **Shape/Proportion**: Captures distortions in the shape, size, or proportions of objects.
 - *Keywords*: shape, proportion, size, distorted, too big, too small
 - *Example*: "The bird has an odd shape and is also yellow."
- **Blur/Fuzziness**: Deals with visual issues related to blurred or unclear edges, lack of sharpness, and fuzziness.
 - *Keywords*: blurry, fuzzy, smudged, blurred edges, not clear
 - *Example*: "The cat's fur is smoothened and texture is changed."
- **Texture**: Focuses on objects with unrealistic or unnatural textures, often described as too smooth or grainy.
 - *Keywords*: texture, smooth, grainy, patchy, unnatural
 - *Example*: "The building texture is unnatural."
- **Lighting/Brightness**: Involves issues where shadows are inconsistent or missing, or where lighting is overexposed or underexposed.
 - *Keywords*: shadows, lighting, brightness, overexposed, underexposed
 - *Example*: "The white bright part on the pan gives it an unrealistic look."
- **Color**: Captures cases where colors are oversaturated, under-saturated, or do not align with the scene.
 - *Keywords*: color, too bright, saturated, unnatural color
 - *Example*: "The fox is bright and inconsistent with the rest of the image."

- Unreal/Artificial Look: Describes objects that appear cartoonish, toy-like, or overly artificial, failing to blend with the rest of the scene.
 - *Keywords*: cartoon, toy, artificial, fake, graphical
 - *Example*: "The helicopter's texture resembles a toy."
- **Placement**: Refers to objects that are misaligned or incorrectly oriented in the scene.
 - *Keywords*: placement, misaligned, incorrect angle, orientation
 - *Example*: "The curtain is hanging in the air instead of the bar."
- Missing/Extra Objects: Captures cases where objects are unexpectedly added or removed, causing inconsistencies.
 - *Keywords*: missing, removed, added, extra, inconsistent
 - *Example*: "The man's face was removed and replaced by a mask."
- Edges: Focuses on issues related to sharp, uneven, or poorly blended edges.
 - Keywords: edges, sharp, uneven, jagged
 - *Example*: "The edges of the pizza are not even."
- **Resolution**: Refers to cases where the visual clarity or quality of the image is degraded, often appearing pixelated or with visual noise.
 - *Keywords*: resolution, clarity, pixelated, low quality
 - *Example*: "The image of the bird looks pixelated and low in resolution."

A.10 Analysis Methodology

Our categorization process followed these steps:

- Examining the Workers' Feedback: We reviewed detailed textual feedback from workers who evaluated the instruction-based edits. Each piece of feedback was carefully analyzed to identify recurring issues.
- 2. **Identifying Categories**: We identified common themes in the feedback and organized them into meaningful categories representing distinct visual and technical issues.

- 3. Extracting Keywords for Categories: For each category, we identified specific keywords and phrases that workers frequently used to describe the issues. These keywords were used to group similar feedback together.
- 4. Generating Statistics: We quantified the frequency of each category across the entire dataset to understand which types of issues were most prevalent. This analysis provided insights to guide future improvements in the edits.

A.11 Supervision Details

The model was fine-tuned for 1 epoch using AdamW with a 2×10^{-4} learning rate. Since it accepts a single image input, we concatenated the before-and-after images.

A.12 Model Versions

• GPT Models

- GPT-40 (Released: 2024-08-06)
- GPT-4 Turbo (Released: 2024-04-09)
- GPT-4 (Version: 0613)

• Gemini Models

- Gemini 1.5 Pro (001)
- Gemini 1.0 Pro (001)

• Image Editing Models

- imagen-3.0-capability-001 via Vertex AI API, guidance scale set to 12
- UltraEdit (Original publicly available version)

A.13 Additional Experiments

Hallucination Rates as a Function of the Number of Objects Figure 6 presents the precision and hallucination rates as a function of the number of objects in the edited images. There is a performance drop in all models as the number of objects in the images increases, highlighting a trend where more complex scenes contribute to higher hallucination rates and lower precision.

Out-of-Distribution Evaluation of the EditInspector Model To evaluate the outof-distribution (OOD) performance of our fine-tuned model, we conducted an additional experiment using a balanced set of 120 samples from the Image Editing Request (IER) dataset. The IER dataset consists of human-authored edits,



Figure 6: Comparison of model precision and hallucination rates as a function of the number of objects in the edited images. The performance of all models decreases as the number of objects in the images increases, highlighting a trend where more complex scenes contribute to higher hallucination rates and lower precision.

originally collected from Reddit, and created using tools like Photoshop. Each edit is paired with a free-form language instruction, covering a wide range of real-world edits across both local and global scenarios.

We selected 30 high-quality, localized edits that are content-safe and appropriate for public release. Each edit was then reversed by swapping the source and target images and inverting the instruction, resulting in 60 examples. This approach allowed us to both increase the number of test cases and evaluate the model's ability to handle edit directionality and instruction symmetry.

For each edit (original and reversed), we also created a manually crafted distractor instruction and difference caption, introducing plausible but incorrect variations. For example, if the original instruction removed an elephant, the distractor might reference removing a nearby person. This setup challenges the model to distinguish accurate vs. misleading edits in realistic and diverse scenarios.

We tested both the base LLaVA model and our fine-tuned model on two key tasks: (1) Edit Accuracy, and (2) Difference Caption Accuracy.

Edit Accuracy Task: The fine-tuned model achieved a balanced accuracy of 59.2%, compared to 54.2% for the base model.

Notably, the base model showed a strong bias toward answering "Yes" on nearly all examples (111 out of 120), resulting in a true negative rate of only 3.3%. In contrast, the fine-tuned model exhibited a more balanced response, predicting "Yes" for 63 out of 120 edits, with a significantly higher true negative rate of 56.6%.

Difference Caption Accuracy Task: The base model again defaulted to "Yes" for all edits, yield-ing a 0% true negative rate.

The fine-tuned model provided more calibrated responses, splitting answers approximately evenly and achieving a true negative rate of 56.7%.

These results demonstrate that our fine-tuned model generalizes better to unfamiliar editing distributions. It shows a stronger grasp of task semantics, provides more accurate judgments in OOD scenarios, and avoids the base model's tendency to over-predict positive responses.

A.14 Augmentation methods

A.15 Licenses

All use of scientific artifacts is consistent with their intended use. This work focuses on evaluating existing models in the English language using images from the MagicBrush dataset and does not introduce new models, generate new images, or employ technologies that could pose ethical, societal, or safety risks. We collected anonymous human annotations using Amazon Mechanical Turk crowdsourcing platform. The images are used in accordance with the MagicBrush license, and the evaluation code and dataset are released under the CC-BY-4.0 license.

A.16 Annotation UI

A.17 Annotation Examples

🔸 Original Edit ► Original edit details 👳 Remove one potted plant One potted plant was removed Augmentation - Similar size object 🤖 Remove one umbrella One umbrella was removed 🖛 Reverse Edit 🖛 Reversed edit details 👮 Add one potted plant One potted plant was added

Figure 7: Illustration of our augmentation methods for a remove edit. The pre-edit image (left) shows a potted plant, while the post-edit image (right) depicts the scene with the plant removed. In the first augmentation method, the instruction and difference caption is modified by replacing the "potted plant" with an object of similar size (umbrella). In the second augmentation, we reverse the edit by switching the order of the images, changing the instruction and difference caption from "remove potted plant" to "add potted plant," and introducing a negative instruction for a visually similar object (e.g., cactus plant), which is absent in the post-edit image.

Remove one cactus plant One cactus plant was removed...

Figure 8: The accuracy scheme tree that was provided to annotators to guide the answering process.

Figure 9: The contextual consistency scheme tree that was provided to annotators to guide the answering process.

Figure 10: The technical precision scheme tree that was provided to annotators to guide the answering process.

Figure 11: The artifacts scheme tree that was provided to annotators to guide the answering process.

Figure 12: The difference caption instructions provided to annotators to guide the answering process.

_				
Inn	Settings Menu	×	Accurate But Unexpected	Accurate
re there	Slider Width	780	ficant ⊖ Mild ⊖ No	۲
oes the	Slider Font Size:	20	ks on the shelf were	changed to be le
ound." శ	Questions Font Size:	20	Yes 🔾 No 📀	
	Form Padding:	15		
	Form Width:	858		
	Save & Refresh Reset &	Refresh		

Figure 13: The setting menu for customizing the form font size, width etc.

Figure 14: Example of image edit verification sample - before image (Add a wild pig).

Edit Instruction: "Add How accurately did the	a wild pig." e edit executed?		
Innacurate	Innacurate, Reflects	Accurate But	Accurate
Did the edit maintain o Shadows, Color, Style	Instruction ontextual consistency (Shap , Texture, etc.)?	Unexpected e, Size, Brightness,	🔿 Yes 💿 No 👔
Please explain why the edited ob The pig is disproportiona body and legs lack natur a different color tone and	ect isn't consistent with the edit instruction tely large compared with the per al texture, giving it an unrealistic I texture, which disrupts the natu	, image scene, and commonsense ople, but its legs are dispr look. The wooden floor w aral flow of the floor.	e oportionately thin. The pig vhere the pig stands on has
Did the edit demonstra smoothness)?	ate high technical precision (resolution, blur, and	🔿 Yes 💿 No 🤅
Please explain why the technical The pig appears low reso	precision of the edit is not high plution, with smooth legs and ba	ck. The edges of the legs	look blurred.
Are there any artifact	s or alterations? O Signific	ant 💿 Mild 🔾 No 🤇	
Does the difference ca	ption "A brown and black w	wild pig was added to	the wooden room with
the bicycle." accurate	ely describe the difference?	🔘 Yes 🔾 No 🔮	
Please modify the difference capt A brown and black wild p modified, blurred and sm pant is cut off and size is the bicycle tires are little height and width are incr blurred, white rods behind	ion to accurately describe the difference (<i>i</i> ig was added to the wooden roc oothened. The left edge of the p reduced, the brown pant is little distorted. The table behind is dif eased, legs are smoothened an d the table are removed and alt	uddress the artifacts of the edit). m with the bicycle. The fl pants on the right side of t cut off and the edges are storted; legs are removed d blurred, the front edge of ered, the area under the t	oor around the pig is he pig is modified; the blue b blurry. The right edge of and added, some legs of the table is distorted and able is distorted and altered

Figure 15: Example of image edit verification sample - after image (Add a wild pig).

and a white dot is added.

rrent Image: After Im

Figure 16: Example of image edit verification sample - before image (Cake on the plate).

Current Image: Afte	er Image		
Edit Instruction: "Ca	an it be a cake on the p	late?"	
How accurately did	the edit executed?		
, , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , , ,			
		•	
Innacurate	Innacurate, Reflects Instruction	Accurate But Unexpected	Accurate
Did the edit maintai Brightness, Shadov	n contextual consistency vs, Color, Style, Texture,	/ (Shape, Size, etc.)?	● Yes ○ No
Did the edit demoneration and smoothness)?	strate high technical pred	cision (resolution, blur,	⊖ Yes ⊚ No (
Please explain why the technical p The cake and affected ar sandwich that was origina	precision of the edit is not high ea are much blurrier than the ite ally on the plate. The edges of th	ms replaced and doesn't mate	the focus of the
Are there any artifa	cts or alterations?	Significant 💿 Mild 🔾	No 🌲
Does the difference	caption "The sandwich	on the white plate w	as replaced wit
piece of dark brov	n layered cake." accur	ately describe the differ	rence?
⊙ Yes ○ No 📀			
Please modify the difference capti The sandwich on the whi frosting on top and betwee white paper on the table	on to accurately describe the difference (A te plate was replaced with a pier en each layer. An upside down widened and extended up to the	ddress the artifacts of the edit). Se of dark brown layered cake fork appeared under the perso plate.	. The cake has gold n's thumb, and the

Figure 17: Example of image edit verification sample - after image (Cake on the plate).

Current Image: Before Image

white paper on the table widened and extended up to the plate.

Edit Instruction: "delete the table" How accurately did the edit executed? Innacurate Innacurate, Reflects Accurate But Accurate Unexpected Instruction Did the edit maintain contextual consistency (Shape, Size, 🔾 Yes 💿 No 🌲 Brightness, Shadows, Color, Style, Texture, etc.)? Please explain why the edited object isn't consistent with the edit instruction, image scene, and comm The shadow from one of the table legs is still there and somehow extended. Did the edit demonstrate high technical precision (resolution, blur, O Yes ○ No ▲ and smoothness)? Are there any artifacts or alterations? O Significant
Mild O No Does the difference caption "The white wicker coffee table with a glass top was removed from the image." accurately describe the difference? () Yes () No (?) ence caption to accurately describe the difference (Address the artifacts of the edit The white wicker coffee table with a glass top was removed from the image. The texture of the bottom of the curtain got altered. The texture and design pattern of the floor beneath the table got altered and now it has parallel lines instead of boxes. The texture of the couch behind it got altered along with the metal frame beside it.

Figure 18: Example of image edit verification sample - before image (Delete the table).

Current Image: After Image Edit Instruction: "delete the table" How accurately did the edit executed? Innacurate, Reflects Accurate But Accurate Innacurate Instruction Unexpected Did the edit maintain contextual consistency (Shape, Size, 🔾 Yes 💿 No 🌲 Brightness, Shadows, Color, Style, Texture, etc.)? The shadow from one of the table legs is still there and somehow extended. Did the edit demonstrate high technical precision (resolution, blur, 🔵 Yes 🔿 No 🌘 and smoothness)? Are there any artifacts or alterations? O Significant
Mild O No Does the difference caption "The white wicker coffee table with a glass top was removed from the image." accurately describe the difference? () Yes () No (?) The white wicker coffee table with a glass top was removed from the image. The texture of the bottom of the curtain got altered. The texture and design pattern of the floor beneath the table got altered and now it has parallel lines instead of boxes. The texture of the couch behind it got altered along with the metal frame beside it.

Figure 19: Example of image edit verification sample - after image (Delete the table).

Figure 20: Example of image edit verification sample - before image (Empty the table).

Figure 21: Example of image edit verification sample - after image (Empty the table).

Current Image: Before Image Edit Instruction: "let the man cut a pineapple" How accurately did the edit executed?

		•	
Innacurate	Innacurate, Reflects Instruction	Accurate But Unexpected	Accurate
Did the edit maintai Brightness, Shadow	🔾 Yes 💿 No 🚺		
Please explain why the edited obje The pineapple piece look	ect isn't consistent with the edit instruction s to be missing texture details.	image scene, and commonsense	
Did the edit demons and smoothness)?	strate high technical pre	cision (resolution, blur,	🔿 Yes 💿 No 🌘
Please explain why the technical p The pineapple is blurred a	recision of the edit is not high and of lower quality and resoluti	on.	***
Are there any artifa	cts or alterations?	Significant 🔾 Mild 🔾	No 🌲
Does the difference	caption "The beef roas	st that the man was sl	icing was
replaced with a pir	neapple." accurately de	scribe the difference?	⊙ Yes ◯ No 🤇
Please modify the difference capti The beef roast that the m replaced with a knife, and pattern of the stains on th	on to accurately describe the difference (A an was slicing was replaced wit a light reflection has been rem e chopboard changed. The har	ddress the artifacts of the edit). h a pineapple. The two-pronge oved from the glass lid near th dle of the chopboard got narro	ed fork has been he chopping board. The bwed.

Current Image: After Image Edit Instruction: "let the man cut a pineapple" How accurately did the edit executed?

		<u> </u>	
Innacurate	Innacurate, Reflects Instruction	Accurate But Unexpected	Accurate
Did the edit maintai Brightness, Shadov	🔾 Yes 💿 No 🌲		
Please explain why the edited obj The pineapple piece look	ect isn't consistent with the edit instruction, is to be missing texture details.	image scene, and commonsense	4
Did the edit demonand smoothness)?	strate high technical prec	cision (resolution, bl	ur, 🔿 Yes 💿 No 🌲
Please explain why the technical The pineapple is blurred	precision of the edit is not high and of lower quality and resolution	on.	
Are there any artifa	icts or alterations?	Significant O Mild	○ No 🌲
Does the difference	caption "The beef roas	t that the man was	s slicing was
replaced with a pi	neapple." accurately des	scribe the difference	∋? ⊙Yes ⊖No 🕜
Please modify the difference capt The beef roast that the m replaced with a knife, and pattern of the stains on th	on to accurately describe the difference (A an was slicing was replaced with d a light reflection has been remo ne chopboard changed. The han	ddress the artifacts of the edit). h a pineapple. The two-pr oved from the glass lid ne dle of the chopboard got i	onged fork has been ar the chopping board. The narrowed.

Figure 23: Example of image edit verification sample - after image (Cut a pineapple).