

FunEditor: Achieving Complex Image Edits via Function Aggregation with Diffusion Models

Mohammadreza Samadi¹, Fred X. Han¹, Mohammad Salameh¹,
Hao Wu³, Fengyu Sun³, Chunhua Zhou³, Di Niu²

¹Huawei Technologies Canada,

²Dept. ECE, University of Alberta,

³Huawei Kirin Solution, China

{mohammadreza.samadi1, fred.xuefei.han1, mohammad.salameh, wuhao135, zhouchunhua}@huawei.com,
sunfengyu@hisilicon.com, dniu@ualberta.ca

Abstract

Diffusion models have demonstrated outstanding performance in generative tasks, making them ideal candidates for image editing. Recent studies highlight their ability to apply desired edits effectively by following textual instructions, yet with two key challenges remaining. First, these models struggle to apply multiple edits simultaneously, resulting in computational inefficiencies due to their reliance on sequential processing. Second, relying on textual prompts to determine the editing region can lead to unintended alterations to the image. We introduce FunEditor, an efficient diffusion model designed to learn atomic editing functions and perform complex edits by aggregating simpler functions. This approach enables complex editing tasks, such as object movement, by aggregating multiple functions and applying them simultaneously to specific areas. Our experiments demonstrate that FunEditor significantly outperforms recent inference-time optimization methods and fine-tuned models, either quantitatively across various metrics or through visual comparisons or both, on complex tasks like object movement and object pasting. In the meantime, with only 4 steps of inference, FunEditor achieves 5–24 times inference speedups over existing popular methods.


Project — <https://mhmdsmidi.github.io/funeditor>

Introduction

Diffusion models (DM) like DALLE-3 and Adobe Firefly have significantly advanced the field of image editing, marking a paradigm shift in how digital content is generated and refined. These models have revolutionized the creative process by allowing users to produce high-quality, photorealistic images from simple textual descriptions, a capability known as text-to-image (T2I) generation. Beyond generating images, diffusion models have also excelled in following text-based instructions to perform specific editing tasks, referred to as instruction-based editing. This approach makes image manipulation more intuitive and accessible by allowing users to guide the editing process with natural language commands (Brooks, Holynski, and Efros 2023; Sheynin et al. 2024; Geng et al. 2024). One prominent example of an instruction-based DM is InstructPix2Pix

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

(a) Basic Functions



(b) Function Aggregation

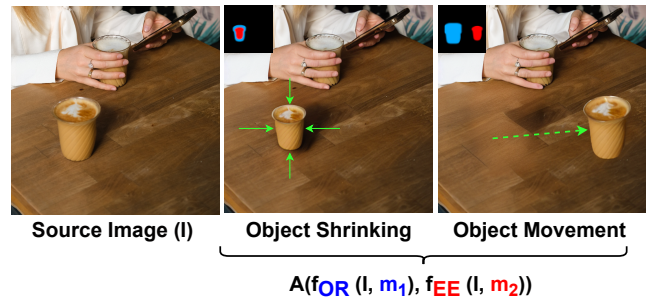


Figure 1: (a) Results of applying *basic functions*—Edge Enhancement (f_{EE} in the center) and Object Removal (f_{OR} on the right)—using their respective masks on the input image Best viewed when enlarged. (b) Demonstration of *function aggregation* using the proposed method. By simultaneously applying Object Removal and Edge Enhancement on different masks, complex edits such as object shrinking (middle) and object movement (right) can be achieved. A represents the operation of function aggregation.

(IP2P) (Brooks, Holynski, and Efros 2023), which can edit images using textual instructions. EmuEdit (Sheynin et al. 2024) further improves the generalizability of IP2P by training on more diversified tasks. They found that training on multi-task datasets can enhance the models’ editing capabilities when handling challenging instructions.

In recent years, the scope of instruction-based editing has expanded to include more complex tasks such as object movement (Wang et al. 2024; Mou et al. 2024a), subject composition (Yang et al. 2022), and virtual try-on (Chen

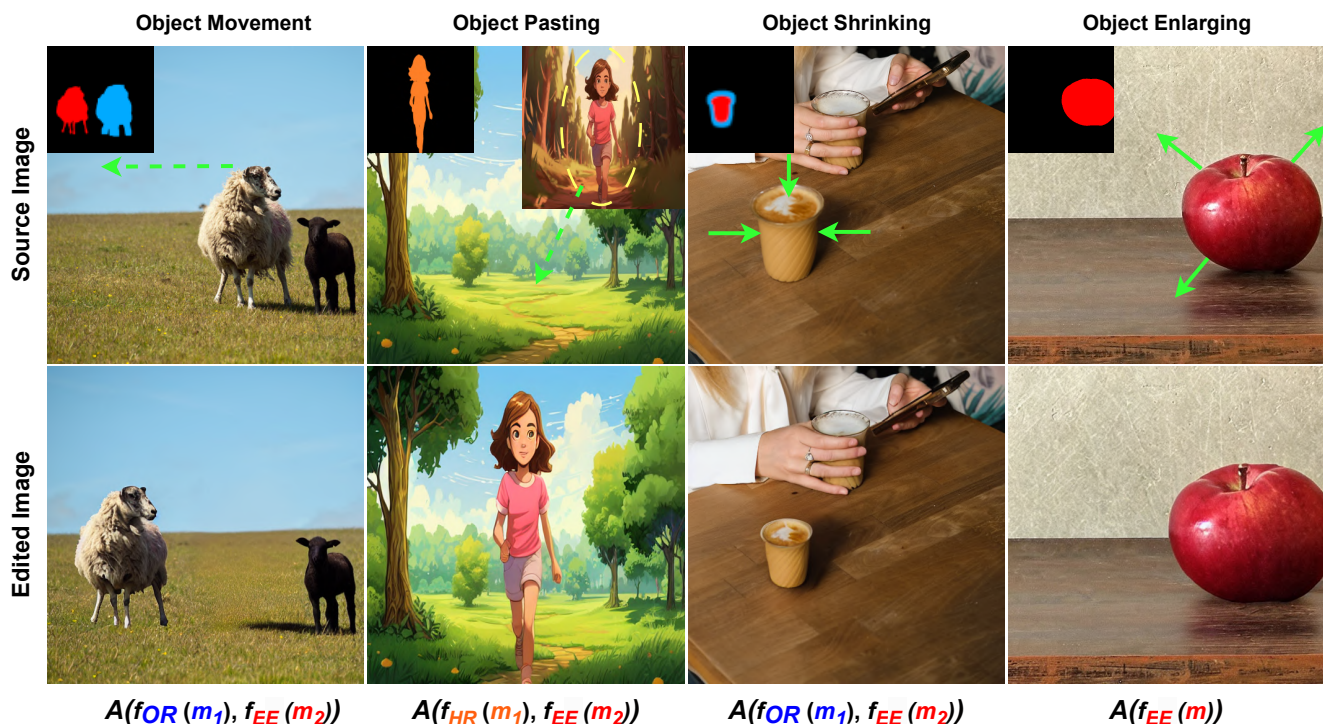


Figure 2: Our approach is capable of composing multiple editing functions and applying them simultaneously. This enables it to perform complex edit functions such as object movement, object resizing, and object pasting in **4 steps**. f_{OR} , f_{EE} , and f_{HR} refer to object removal, edge enhancement, and harmonization functions, respectively. Each function is applied only to the specified mask region. To save space, source image I is omitted from the function arguments.

et al. 2024b). In real-world scenarios, a single edit may involve multiple types of modifications to an image, resulting in a complex edit. For instance, the object movement task involves a combination of simpler editing functions: the model needs to fill in the missing region at the original object’s location with a natural background, while enhancement functions are often required for the moved object to ensure consistency with its new location. Recent studies (Guo and Lin 2024) reveal that instruction-based DMs, such as IP2P, struggle to faithfully follow editing instructions when multiple edits are required on an image, highlighting potential areas for improvement. Moreover, both IP2P and EmuEdit are limited to sequential editing, where each subsequent edit is performed on the results of the previous one. This approach increases the number of denoising steps during inference and is prone to error accumulation. To address this challenge, developing methods that can efficiently handle complex editing tasks is essential.

Previous works have attempted to tackle complex tasks either through inference-time optimizations (Mou et al. 2024b,a) or supervised fine-tuning (Wang et al. 2024; Winter et al. 2024). DiffEditor (Mou et al. 2024a) and DragonDiffusion (Mou et al. 2024b) rely on inference-time optimizations using energy-based loss functions, which involve an additional outer optimization loop and multiple calls to the UNet during inference, leading to high latency. Supervised fine-tuning approaches, such as AnyDoor (Chen et al. 2024b),

address this issue by directly training a diffusion model capable of complex edits. However, collecting training data for all possible complex tasks is challenging, often requiring significant cost and human effort.

In this work, we propose FunEditor, an efficient diffusion model that performs complex edits through the aggregation of simpler functions. Under this formulation, FunEditor learns to effectively integrate multiple localized edit functions and apply them simultaneously. This is achieved by defining learnable task tokens during training and combining them during inference. For localization, each simple task requires a binary mask as input and performs cross-attention modification (Guo and Lin 2024; Simsar et al. 2023). Compared to prior editing models, FunEditor offers three unique advantages. **First**, during training, it only requires simple and atomic editing tasks such as object removal, harmonization, edge enhancement, etc., to learn new task tokens. During inference, by passing a sequence of learned task tokens into diffusion model, FunEditor can apply multiple functions simultaneously to accomplish a complex edit, requiring no energy guidance during inference. Figure 1 highlights the output of FunEditor for object movement and object resizing by aggregating object removal and edge enhancement. **Second**, FunEditor is data-efficient as we only require training on each atomic task individually, for which training data is either abundant or easy to collect, making it a superior, low-cost option when compared to prior training-

based approach like AnyDoor (Chen et al. 2024b). **Third**, FunEditor prioritizes localized editing where each simple edit function only applies to a specified region in the input image, leading to improved consistency. **Lastly**, as FunEditor is compatible with pre-trained few-step diffusion models such as LCM (Luo et al. 2023), these atomic functions can be combined with pixel-space operation like copy and pasting to perform high-quality simultaneous edits in only 4 steps without the need of subject regeneration or extra optimization steps during inference as indicated in Figure 2.

To evaluate the effectiveness of FunEditor, we conducted extensive evaluations on object movement and object pasting as complex editing tasks using the COCOEE and ReS datasets. We compared FunEditor against a wide range of recently proposed editing baselines, including both inference-time optimization methods (Mou et al. 2024b) and fine-tuned models (Chen et al. 2024b). The results demonstrate that FunEditor significantly outperforms all baselines across various evaluation metrics, such as IQA and object-background consistency. Notably, FunEditor performs complex edits in just 4 steps, making it 5–24× faster than the baselines in the object movement task. In addition to quantitative results, visual comparisons in Figures 5 and 6 further highlight FunEditor’s superiority in complex editing tasks.

Related Work

Image Editing with Diffusion Model. The success of text-to-image (T2I) diffusion models in generating high-quality images has led to their widespread adoption in image editing (Epstein et al. 2023; Wang et al. 2024; Mou et al. 2024b,a; Alaluf et al. 2023; Hertz et al. 2023; Jeong et al. 2024; Patashnik et al. 2023). Training-free editing methods, such as Prompt-to-Prompt (Hertz et al. 2022), leverage these models by swapping cross-attention maps (CA-maps) of specific tokens during inference, enabling real-time editing without the need for fine-tuning. Plug-and-Play (Tumanyan et al. 2023) extends this approach by introducing spatial features and attention maps to guide the generation process and localize changes using masks. These methods exploit the generative capabilities of T2I diffusion models to apply edits and precisely target specific areas. Another line of work defines the editing function during inference using an energy function to update the noisy latent space (Mou et al. 2024a; Epstein et al. 2023; Shi et al. 2023; Mou et al. 2024b). This approach guides the editing direction and samples noise to align the latent space with the editing goal, as demonstrated in DragonDiffusion (Mou et al. 2024b) and DiffEditor (Shi et al. 2023). However, these techniques often require additional steps, such as optimization and image inversion using methods like DDIM (Dhariwal and Nichol 2021), which can reduce efficiency.

Training-based Editing Model. As the editing functions become more complex, the T2I diffusion model may show inconsistencies and poor performance. This is primarily because the model has not seen similar examples during training. Training-based approaches address this by using large editing datasets (Zhang et al. 2024a,b; Hui et al. 2024), which include pairs of source and edited images along with

the corresponding edit instructions, to fine-tune a T2I diffusion model specifically for editing tasks. InstructPix2Pix (IP2P) (Brooks, Holynski, and Efros 2023) leverages GPT-3 (Brown et al. 2020) and Prompt-to-Prompt (Hertz et al. 2022) to generate synthetic editing datasets. To enable the source image to be passed to the UNet as an additional input, they added a new channel to the input convolutional layer, allowing the model to preserve details in unintended areas. Building on the success of InstructPix2Pix (IP2P), subsequent works, such as MagicBrush (Zhang et al. 2024a), have manually annotated images to create higher-quality datasets. By utilizing the same training procedure to fine-tune their editing models, these approaches have demonstrated improved performance (Zhang et al. 2024b; Hui et al. 2024). However, the main shortcoming of IP2P-based editing models is the occurrence of unintended changes to other parts of the source image (Simsar et al. 2023; Guo and Lin 2024). To address this issue, FoI (Guo and Lin 2024) introduced an inference technique that modulates attention and masks the sampled noise for the text-conditioned output, thereby limiting changes only to the specific area of the edit.

Multi-task Diffusion Model. Unifying multiple tasks within a single model has been an intriguing research direction in the literature. InstructDiffusion (Geng et al. 2024) investigates the potential benefits of joint training across various tasks, including keypoint detection, segmentation, image enhancement, and editing. Their results suggest that multi-task training enhances the model’s performance across these diverse tasks. Similarly, EmuEdit (Sheynin et al. 2024) introduces a unified diffusion model capable of performing multiple editing and computer vision tasks, such as object removal, global editing, and object detection. EmuEdit uses task embeddings as conditions to switch between tasks, effectively allowing the model to adapt its behavior to the specific task at hand. While these approaches show impressive results on individual tasks, they are limited by their inability to apply two tasks simultaneously, requiring multiple user interactions to achieve complex edits. In this work, we address this limitation by bridging the gap between multi-task training and the ability to perform multiple editing tasks in a single operation, enabling efficient complex editing.

Method

In this section, we explain how complex localized editing tasks could be formulated into a problem of function aggregation. This formulation allow us to achieve simultaneous edits during inference in few-steps. Then, we showcase how FunEditor solve this problem by a carefully designed training framework. Figure 3 provides an overview of our proposed pipeline for both training and inference.

Editing as Function Aggregation

Let f be a localized editing function to a certain region of the source image I , indicated by a binary mask m . $f(I, m)$ denotes the edited image. We then define an aggregation function F as:

$$F = \mathcal{A}(f_1(I, m_1), f_2(I, m_2), \dots, f_n(I, m_n)). \quad (1)$$

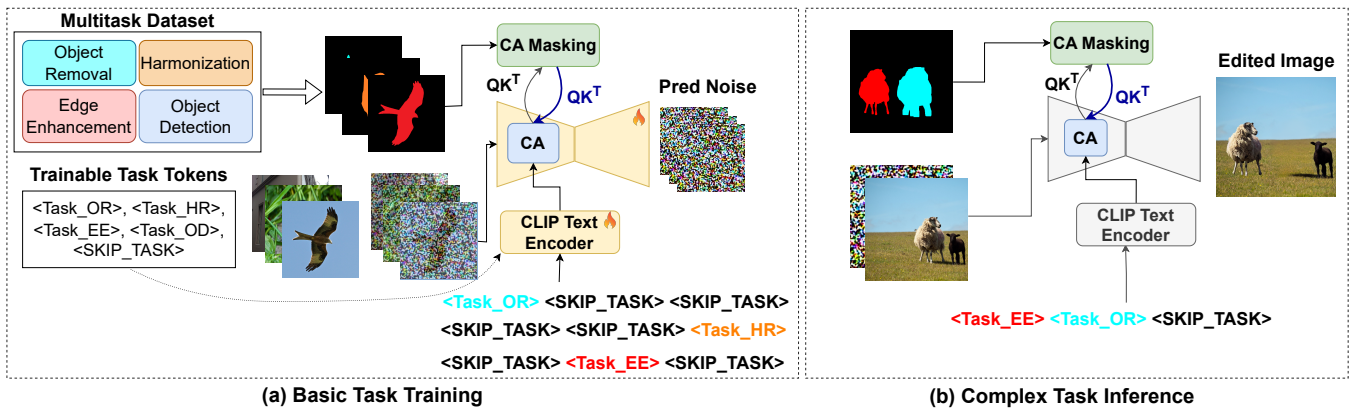


Figure 3: Overview of our proposed training and inference pipeline. During the basic task training phase (a) the diffusion model learns to perform various simple tasks based on the provided task tokens and masks. During inference (b), we could implement complex edit functions by combining multiple task masks and tokens.

In essence, F represents aggregating different f and applying them simultaneously to multiple regions of the source image in one shot. The main advantage of such formulation is that it enables the divide-and-conquer of complex tasks (see Figure 1). Given that it is challenging to collect sufficient high-quality training data to directly learn F , our goal is to learn each f individually during training and aggregate their effects during inference. In FunEditor, each basic function f_i is represented by a trainable task token \mathbf{T}_i . For example, providing the token $\langle \text{Task_OR} \rangle$ to the UNet as input would enable the task of object removal, as indicated by Figure 3(a). Although EmuEdit also define task embedding vectors that are of similar purposes, the main difference is that we directly add the task tokens to the vocabulary of the prompt tokenizer. This design offers two advantages compared to EmuEdit. First, it allows FunEditor to freely combining tasks by passing the corresponding task tokens as prompts to the UNet. To disable tasks, we introduce another special token named $\langle \text{SKIP_TASK} \rangle$, which ensures a consistent prompt length during training and inference, improving the overall generalizability. Second, task tokens pave the way towards localized, simultaneous editing during inference. To achieve this, we simply set the desired task tokens in the input prompt and keep the remaining ones as $\langle \text{SKIP_TASK} \rangle$. Similar to instruction-based editing (Simsar et al. 2023; Guo and Lin 2024), our goal is to apply each editing function f_i to its corresponding mask m_i to maintain control over the editing region and prevent leakage into other areas.

Since the UNet processes the task tokens via cross-attentions (CA), it is natural to inject localization information via the same medium. Therefore, we employ a cross-attention (CA) masking technique to ensure each task f_i does not affect the unmasked area $1 - m_i$. During both training and inference, we apply this CA masking technique to all CA layers. For each CA layer, we interpolate the mask into its resolution and set the values of QK^T of \mathbf{T}_i token corresponding to the unmasked area $1 - m_i$ to its minimum value.

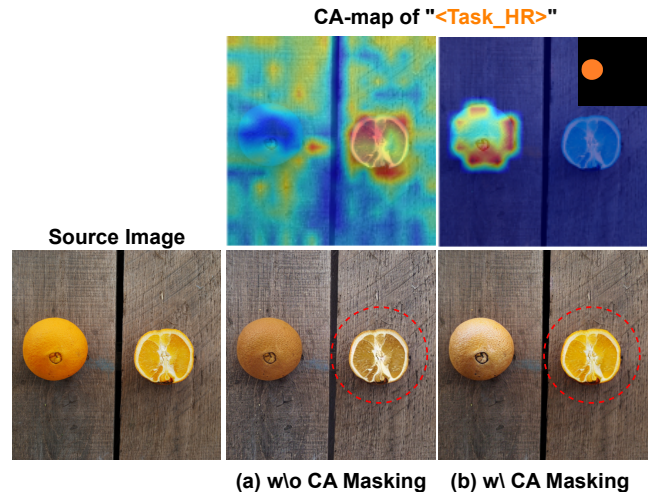


Figure 4: Harmonization without cross-attention masking affects the entire image (a). While with masking, edits are confined to the masked region, (b), preventing changes to unmasked areas. Mask is indicated by the top right mini-figure. The prompt is “ $\langle \text{Task_HR} \rangle \langle \text{SKIP_TASK} \rangle \langle \text{SKIP_TASK} \rangle$ ”

Mathematically represented as:

$$QK^T[(1 - m_i), \mathbf{T}_i] = \min(QK^T). \quad (2)$$

Figure 4 highlights the importance of CA masking for localized editing. Notably without CA masking the model could attend to other regions of the source image and generates undesired changes outside of the masked region.

FunEditor Training Strategy

The training procedure for FunEditor builds upon the IP2P architecture. In IP2P, a source image I_{src} , a target image I_{trg} , and a text instruction T are provided, with noise ϵ added to I_{trg} to create a noisy input z_t . The UNet ϵ_θ is



Figure 5: Qualitative comparison between our approach and baseline methods for object repositioning within an image, demonstrating the superior performance of our method. To move an object, FunEditor composes object removal and edge enhancement functions.




Figure 6: Visual comparison between our method and baseline methods for object pasting from a reference image into a target image. FunEditor applies harmonization and edge enhancement functions to seamlessly paste an object into another image.

trained to predict ϵ using the following objective:

$$\min_{\theta} \mathbb{E}_{y, \epsilon, t} [\|\epsilon - \epsilon_{\theta}(z_t, t, I_{src}, I_{trg}, T)\|_2^2], \quad (3)$$

where $y = (I_{src}, I_{trg}, T)$ represents a sample from the IP2P editing dataset. In FunEditor, instead of using a text instruction T , we incorporate new task tokens into the CLIP text tokenizer (Radford et al. 2021) to represent the desired editing function. Additionally, we use masks corresponding to each editing function within the cross-attention. Thus, our training sample takes the form $(I_{src}, \mathbf{T}, m, I_{trg})$, where I_{src} is the source image, \mathbf{T} is the task token, m is the mask, and I_{trg} is the target image.

The idea of task embeddings provide the foundation for enabling simultaneous edits in FunEditor. EmuEdit defines learnable task embedding vectors, which are provided to the UNet as task-selectors via cross-attention interaction and the timestep embedding layers. In contrast, we incorporate the task embeddings via cross-attention to have more control

on their spatial effects on the editing regions. The following equation defines our objective for training FunEditor, where we learn an embedding vector for each task token that represents a basic edit:

$$\min_{\theta, v_1, \dots, v_k} \mathbb{E}_{g, \epsilon, t} [\|\epsilon - \epsilon_{\theta}(z_t, t, I_{src}, m, \mathbf{T})\|_2^2]. \quad (4)$$

Experimentally we find that setting the number of task tokens in the prompt equal to the number of supported simple tasks generates the best results. During training we randomly activate one task token and keep the remaining tokens as $\langle \text{SKIP-TASK} \rangle$. Since the task token ordering is not important, we also randomly shuffle the task tokens in every training step to prevent the model from memorizing their positions.

Experiments

In this section, we evaluate the effectiveness of our proposed method by evaluating it on two complex image editing

Method	#Steps	NFEs	Avg. Latency	std
SelfGuidance	50	100	11	0.6
DragonDiffusion	50	160	23	0.7
DiffEditor	50	176	24	0.4
Paint-by-Example	50	50	5	0.6
AnyDoor	50	50	12	0.4
FunEditor (ours)	4	4	1	0.4

Table 1: Efficiency is compared in terms of the number of function evaluations (NFEs) and latency (seconds). Latency is measured as the average wall-clock time for editing one image over 10 runs on a single Nvidia V100 (32GB) GPU.

tasks: Object Movement (F_{OM}) and Object Pasting (F_{OP}). Building on the formulation introduced in the previous section, we define each complex function as an aggregation of the following basic functions:

$$F_{OM}(I, M_{src}, M_{trg}) = \mathcal{A}(f_{OR}(I, M_{src}), f_{EE}(I, M_{trg})),$$

$$F_{OP}(I, M_{trg}) = \mathcal{A}(f_{HR}(I, M_{trg}), f_{EE}(I, M_{trg})).$$

In these formulations, f_{OR} , f_{EE} , and f_{HR} represent the basic functions for object removal, edge enhancement, and harmonization, respectively. The masks M_{src} and M_{trg} correspond to the source object and the target location within the image.

We compare our method to baseline approaches in terms of quality and efficiency. The following subsections outline the experimental setup, datasets, evaluation metrics, and results, highlighting our method’s performance against baseline techniques.

Experimental Settings

Implementation Details We fine-tuned our model using 4 Nvidia V100 (32GB) GPUs, with the IP2P UNet (Brooks, Holyński, and Efros 2023) serving as our diffusion backbone. The AdamW optimizer (Loshchilov and Hutter 2017) was employed with a learning rate of 5×10^{-5} and a batch size of 4. During inference, we distilled the trained model into 4 steps using SD1.5 LCM-LoRA from HuggingFace¹. For the baselines, we used SDv1.5 (Rombach et al. 2022) from HuggingFace², adhering to all default inference hyperparameters. As a preprocessing step for object removal, we dilated the mask with a kernel size of 20 to prevent any pixel leakage into the painted region.

Datasets The COCOEE dataset, compiled by Yang et al. (2022), features 3,500 images manually selected from the MSCOCO (Microsoft Common Objects in Context) (Lin et al. 2014) validation set. A human operator used the Segment Anything model (Kirillov et al. 2023) to extract segments and assign diff vectors, resulting in a benchmark of 100 images with corresponding masks and diff vectors for object movement tasks. Wang et al. (2024) also released the ReS dataset, comprising 100 pairs of real-world images that present challenging cases of object movement.

¹<https://huggingface.co/latent-consistency/lcm-lora-sdv1-5>

²<https://huggingface.co/runwayml/stable-diffusion-v1-5>

Metrics In our evaluations, we utilize four groups of metrics: overall image quality, object consistency, background consistency, and semantic consistency. For Image Quality Assessment (IQA), we assess the perceptual quality of the images using TOPIQ (Chen et al. 2024a), MUSIQ (Ke et al. 2021), and LIQE (Zhang et al. 2023). Object Consistency is measured by comparing the similarity of the moved object to its original counterpart in the source image, using the LPIPS (Zhang et al. 2018) and PSNR metrics. Background Consistency is evaluated by assessing the similarity between the background in the edited image and the original image, using the same metrics as for Object Consistency. Finally, Semantic Consistency is determined by evaluating the similarity between the semantics of the source and edited images through CLIP image encoder, which measures image-to-image similarity based on CLIP embeddings (Radford et al. 2021).

Baselines We evaluate our approach against both training-based and training-free methods on complex editing tasks. Given a limited number of open-source training-based model specifically for complex tasks (Winter et al. 2024; Wang et al. 2024), we include training-free approaches to comprehensively assess the quantitative performance of our method. Our baselines for the complex editing tasks are as follow:

- **Paint-by-Example** (Yang et al. 2022): This study explores a training-based approach for reference-guided image editing, training a model to fill a designated mask area with a reference object.
- **AnyDoor** (Chen et al. 2024b): A training-based method that enables teleportation of an object into a user-specified location and shape, leveraging object identity information to accurately reconstruct the object in the intended form.
- **SelfGuidance** (Epstein et al. 2023): A training-free method that utilizes energy functions on attention maps to optimize the noisy latent representation during inference, allowing for edits such as object movement.
- **DragonDiffusion** (Mou et al. 2024b): A training-free approach that employs an energy function with two components: one for applying changes within the target mask and another for preserving the unmasked area.
- **DiffEditor** (Mou et al. 2024a): Similar to DragonDiffusion, this method enhances the quality of edits by optimizing an energy function, incorporating score-based gradient guidance and time interval strategies into the diffusion noise sampling process.

Experimental Results

Efficiency Table 1 demonstrates that FunEditor, using only 4 steps, requires 172 fewer NFEs and is 23 seconds faster than the state-of-the-art method for object movement, DiffEditor (Mou et al. 2024a), on the COCOEE dataset. SelfGuidance, DragonDiffusion, and DiffEditor, rely on optimizing energy functions to guide complex edits during inference. These method requires multiple calls to the UNet,

Method	IQA			Object Consistency		Background Consistency		Semantic Consistency
	TOPIQ \uparrow	MUSIQ \uparrow	LIQE \uparrow	LPIPS \downarrow	PSNR \uparrow	LPIPS \downarrow	PSNR \uparrow	CLIP-I2I \uparrow
SelfGuidance	0.554	65.91	3.90	0.083	22.77	0.259	17.86	0.897
DragonDiffusion	0.571	68.87	4.27	0.034	28.59	0.098	23.99	0.965
DiffEditor	0.579	69.09	4.27	0.036	28.49	0.094	24.23	0.967
Paint-by-Example	0.595	69.57	4.31	0.066	23.12	0.089	25.23	0.937
AnyDoor	0.558	68.37	4.12	0.053	25.43	0.122	23.55	0.954
FunEditor (ours)	0.611	70.39	4.36	0.017	34.27	0.066	26.54	0.969

Table 2: Quantitative evaluation of our approach compared to the baselines on the object movement task using the COCOEE dataset. To move an object, FunEditor encompasses object removal and edge enhancement functions.

Method	IQA			Object Consistency		Background Consistency		Semantic Consistency
	TOPIQ \uparrow	MUSIQ \uparrow	LIQE \uparrow	LPIPS \downarrow	PSNR \uparrow	LPIPS \downarrow	PSNR \uparrow	CLIP-I2I \uparrow
SelfGuidance	0.586	69.41	3.61	0.064	24.21	0.273	17.92	0.869
DragonDiffusion	0.690	74.95	4.72	0.030	29.68	0.083	25.38	0.934
DiffEditor	0.691	74.94	4.73	0.032	29.59	0.083	25.44	0.933
Paint-by-Example	0.717	74.94	4.75	0.052	24.20	0.076	26.27	0.907
AnyDoor	0.679	73.68	4.51	0.045	25.89	0.106	23.48	0.923
FunEditor (ours)	0.702	75.00	4.75	0.015	35.61	0.063	26.48	0.942

Table 3: Quantitative evaluation of our approach compared to the baselines on the object movement task was conducted using the ReS dataset. To move an object, FunEditor encompasses object removal and edge enhancement functions.

as it must perform both inpainting and blending of the object with its surrounding environment, leading to a more time-consuming process. Comparing to training-based approaches, our model is 12 times faster than the training-based baseline AnyDoor (Chen et al. 2024b) and 5 times faster than Paint-by-Example (Yang et al. 2022). These methods regenerate the object to fit the target position demanding a higher number of steps to achieve the desired result. In contrast, FunEditor achieves superior results in fewer steps by employing simpler enhancement functions, such as edge enhancement, and pixel operations, without the need for fully regenerating the object. Instead, these functions refine the composition and harmonization of the pasted object in pixel space, enabling us to achieve complex edits more efficiently. It should be noted that the studies that used synthetic data generation for object movement (Winter et al. 2024; Wang et al. 2024) did not open-source their models, therefore we could not use them as baselines to assess their latency in a multi-turn fashion.

Image Quality We evaluate the performance of FunEditor through quantitative and qualitative experiments on object movement and object pasting. Tables 2 and 3 presents our quantitative comparison with the baselines on the object movement task on COCOEE and ReS datasets. Our model achieves higher quality across almost all IQA metrics compared to the baselines, improving the accuracy of object movement while preserving the details of both the object and the background. By aggregating atomic functions, FunEditor successfully removes objects, enhancing semantic consistency and overall image quality compared to optimization-based approaches that rely on energy functions. Additionally, since we use the original object in the pixel space and apply enhancement functions such as edge

enhancement and harmonization, we can faithfully move or paste the object without altering its appearance, unlike the baselines that regenerate the object and modify its appearance.

Figure 5 and 6 provide qualitative comparisons between our method and the baselines for object movement and object pasting. As shown in Figure 5, the baselines face two primary issues: unsuccessful object removal and unsuccessful object pasting. In contrast, our method seamlessly repositions and blends the object into the destination while effectively removing the original object. Supporting the quantitative results in Table 2, FunEditor faithfully preserves the background, applying changes only to the intended editing area. For example, SelfGuidance and DiffEditor alter the shape of the floor lamp in the first image. In the case of object pasting, as shown in Figure 6, regenerating the object can lead to issues such as unfaithful or incomplete regeneration like the first example.

Conclusion

In this paper, we introduced FunEditor, an efficient editing approach that utilizes function aggregation to address complex editing tasks such as object movement. FunEditor offers several key advantages: it requires only publicly available datasets for atomic tasks, enables the simultaneous execution of multiple functions to achieve complex edits, and is compatible with few-step diffusion models. Our approach not only outperforms baselines specifically designed for complex tasks like object movement and object pasting in terms of quality but also offers superior efficiency. Additionally, we evaluated FunEditor across a wide range of numerical metrics and visual quality, comparing it against both training-based and training-free baselines.

References

- Alaluf, Y.; Garibi, D.; Patashnik, O.; Averbuch-Elor, H.; and Cohen-Or, D. 2023. Cross-image attention for zero-shot appearance transfer. *arXiv preprint arXiv:2311.03335*.
- Brooks, T.; Holynski, A.; and Efros, A. A. 2023. Instruct-pix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18392–18402.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Chen, C.; Mo, J.; Hou, J.; Wu, H.; Liao, L.; Sun, W.; Yan, Q.; and Lin, W. 2024a. Topiq: A top-down approach from semantics to distortions for image quality assessment. *IEEE Transactions on Image Processing*.
- Chen, X.; Huang, L.; Liu, Y.; Shen, Y.; Zhao, D.; and Zhao, H. 2024b. Anydoor: Zero-shot object-level image customization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6593–6602.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- Epstein, D.; Jabri, A.; Poole, B.; Efros, A.; and Holynski, A. 2023. Diffusion self-guidance for controllable image generation. *Advances in Neural Information Processing Systems*, 36: 16222–16239.
- Geng, Z.; Yang, B.; Hang, T.; Li, C.; Gu, S.; Zhang, T.; Bao, J.; Zhang, Z.; Li, H.; Hu, H.; et al. 2024. Instructdiffusion: A generalist modeling interface for vision tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12709–12720.
- Guo, Q.; and Lin, T. 2024. Focus on your instruction: Fine-grained and multi-instruction image editing by attention modulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6986–6996.
- Hertz, A.; Mokady, R.; Tenenbaum, J.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2022. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*.
- Hertz, A.; Voynov, A.; Fruchter, S.; and Cohen-Or, D. 2023. Style aligned image generation via shared attention. *arXiv preprint arXiv:2312.02133*.
- Hui, M.; Yang, S.; Zhao, B.; Shi, Y.; Wang, H.; Wang, P.; Zhou, Y.; and Xie, C. 2024. HQ-Edit: A High-Quality Dataset for Instruction-based Image Editing. *arXiv preprint arXiv:2404.09990*.
- Jeong, J.; Kim, J.; Choi, Y.; Lee, G.; and Uh, Y. 2024. Visual Style Prompting with Swapping Self-Attention. *arXiv preprint arXiv:2402.12974*.
- Ke, J.; Wang, Q.; Wang, Y.; Milanfar, P.; and Yang, F. 2021. Musiq: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, 5148–5157.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, 740–755. Springer.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Luo, S.; Tan, Y.; Huang, L.; Li, J.; and Zhao, H. 2023. Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378*.
- Mou, C.; Wang, X.; Song, J.; Shan, Y.; and Zhang, J. 2024a. DiffEditor: Boosting Accuracy and Flexibility on Diffusion-based Image Editing. *arXiv preprint arXiv:2402.02583*.
- Mou, C.; Wang, X.; Song, J.; Shan, Y.; and Zhang, J. 2024b. DragonDiffusion: Enabling Drag-style Manipulation on Diffusion Models. In *The Twelfth International Conference on Learning Representations*.
- Patashnik, O.; Garibi, D.; Azuri, I.; Averbuch-Elor, H.; and Cohen-Or, D. 2023. Localizing object-level shape variations with text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 23051–23061.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis With Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10684–10695.
- Sheynin, S.; Polyak, A.; Singer, U.; Kirstain, Y.; Zohar, A.; Ashual, O.; Parikh, D.; and Taigman, Y. 2024. Emu edit: Precise image editing via recognition and generation tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8871–8879.
- Shi, Y.; Xue, C.; Pan, J.; Zhang, W.; Tan, V. Y.; and Bai, S. 2023. Dragdiffusion: Harnessing diffusion models for interactive point-based image editing. *arXiv preprint arXiv:2306.14435*.
- Simsar, E.; Tonioni, A.; Xian, Y.; Hofmann, T.; and Tombari, F. 2023. Lime: Localized image editing via attention regularization in diffusion models. *arXiv preprint arXiv:2312.09256*.
- Tumanyan, N.; Geyer, M.; Bagon, S.; and Dekel, T. 2023. Plug-and-play diffusion features for text-driven image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1921–1930.

Wang, Y.; Cao, C.; Dong, Q.; Li, Y.; and Fu, Y. 2024. Repositioning the Subject within Image. *arXiv preprint arXiv:2401.16861*.

Winter, D.; Cohen, M.; Fruchter, S.; Pritch, Y.; Rav-Acha, A.; and Hoshen, Y. 2024. ObjectDrop: Bootstrapping Counterfactuals for Photorealistic Object Removal and Insertion. *arXiv preprint arXiv:2403.18818*.

Yang, B.; Gu, S.; Zhang, B.; Zhang, T.; Chen, X.; Sun, X.; Chen, D.; and Wen, F. 2022. Paint by Example: Exemplar-based Image Editing with Diffusion Models. *arXiv preprint arXiv:2211.13227*.

Zhang, K.; Mo, L.; Chen, W.; Sun, H.; and Su, Y. 2024a. Magicbrush: A manually annotated dataset for instruction-guided image editing. *Advances in Neural Information Processing Systems*, 36.

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

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

Zhang, W.; Zhai, G.; Wei, Y.; Yang, X.; and Ma, K. 2023. Blind image quality assessment via vision-language correspondence: A multitask learning perspective. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14071–14081.