

Move and Act: Enhanced Object Manipulation and Background Integrity for Image Editing

Pengfei Jiang¹, Mingbao Lin², Fei Chao^{1*}

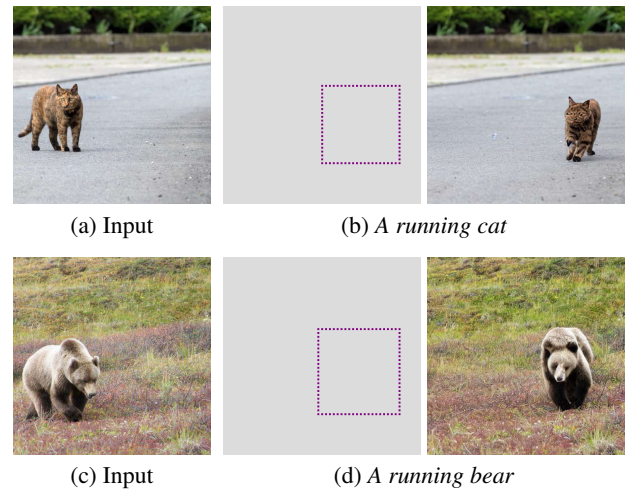
¹Key Laboratory of Multimedia Trusted Perception and Efficient Computing,
Ministry of Education of China, Xiamen University, China

²Skywork AI Singapore

pengfeijiang@stu.xmu.edu.cn, linmb001@outlook.com, fchao@xmu.edu.cn

Abstract

Current methods commonly utilize three-branch structures of inversion, reconstruction, and editing, to tackle consistent image editing task. However, these methods lack control over the generation position of the edited object and have issues with background preservation. To overcome these limitations, we propose a tuning-free method with only two branches: inversion and editing. This approach allows users to simultaneously edit the object’s action and control the generation position of the edited object. Additionally, it achieves improved background preservation. Specifically, we transfer the edited object information to the target area and repair or preserve the background of other areas during the inversion process at a specific time step. In the editing stage, we use the image features in self-attention to query the key and value of the corresponding time step in the inversion to achieve consistent image editing. Impressive image editing results and quantitative evaluation demonstrate the effectiveness of our method.



Introduction

Recent advancements in text-to-image generation using diffusion models (Nichol et al. 2022; Saharia et al. 2022; Ramesh et al. 2022; Rombach et al. 2022) have significantly improved image generation capabilities. By capitalizing on extensive training data, these models (Saharia et al. 2022; Ramesh et al. 2022) excel in generating remarkable images based on user prompts. Compared to generative adversarial networks (GANs), diffusion models provide superior quality and diversity in image generation. Consequently, numerous studies (Hertz et al. 2023; Tumanyan et al. 2023; Parmar et al. 2023; Dong et al. 2023; Cao et al. 2023; Huang et al. 2023) have begun investigating diffusion models for image editing tasks. Predominantly, these methods involve text-conditioned editing, where users supply text input to modify images. Although techniques like (Hertz et al. 2023; Tumanyan et al. 2023; Parmar et al. 2023; Dong et al. 2023) can adeptly edit image style, object appearance, and object category based on text input, they do not alter scene layout or modify object actions in the image. Nonetheless, there are instances where users desire to edit an object’s action

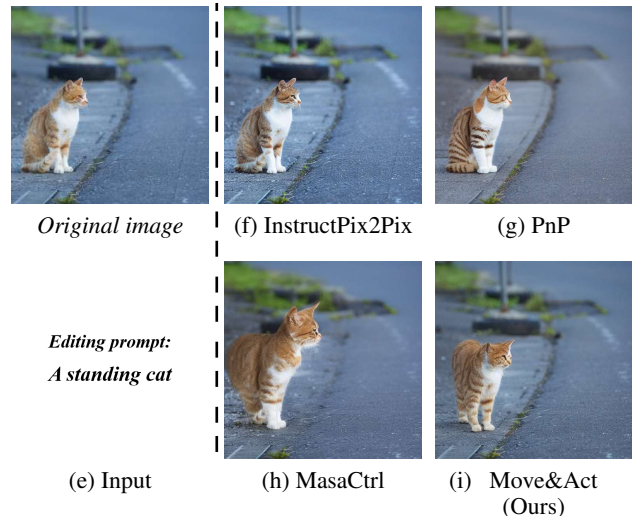


Figure 1: (a) - (d): our method allows users to control where to place the edited object. (e) - (i): altering only action, our method captures the “standing” pose and maintains consistent backgrounds with the input image.

*Corresponding author

while maintaining its appearance and background. This specific editing process is called consistent image editing.

Imagic (Kawar et al. 2023) adopts a complex prompt fine-tuning process for text-based consistent image editing. MasaCtrl (Cao et al. 2023) offers a tuning-free approach, transforming self-attention into mutual self-attention. These methods edit an object’s action without allowing control over the object’s generated position. DragonDiffusion (Mou et al. 2024) facilitates object displacement within the image, yet it restricts the editing of the relocated object’s action. In practical applications, users may require simultaneous editing of an object’s action and position. As illustrated in the upper part of Figure 1, we allow users to input editing text to adjust the object’s action, while also providing a bounding box to dictate the object’s generated position.

We also maintain better background preservation. By designating the object’s generation position as its original location, our method effectively reduces to MasaCtrl’s editing function, solely modifying the object’s action. In the lower part of Figure 1, the backgrounds generated by InstructPix2Pix (Brooks, Holynski, and Efros 2023), PnP (Tumanyan et al. 2023) and MasaCtrl (Cao et al. 2023) exhibit a certain loss in detail, whereas our method produces a background more closely resembling the input image. Also, the cat’s action in our results aligns more accurately with the editing prompt and the surrounding environment.

We introduce a training-free Move&Act method, for consistent image editing, empowering users to simultaneously modify an object’s action and its generated position within an image. As depicted in Figure 2, our method transfers object information from the original area to a user-defined target bounding box, effectively creating the edited object in the desired location. Since cross-attention maps linked to an object token encompass spatial position information of the object in the image, we propose the transfer of high-response areas from the original region to the target region.

To achieve this, we maximize attention scores in the target area while minimizing them in non-target regions. However, this strategy alone is inadequate for a comprehensive object information transfer, as larger attention scores may linger in the original area. To address this, we utilize the surrounding background to inpaint the object’s initial location, employing an efficient feature extractor to obtain features from the original and surrounding background regions, and calculating the L_1 loss between them. This inpainting process not only harmonizes the background in the original and adjacent areas but also effectively eliminates residual object information, facilitating a complete transfer to the target region. Moreover, we incorporate a background preservation loss to guarantee consistency before and after editing. The object information transfer procedure is finalized during the inversion stage, while the subsequent reverse stage ensures consistency in the object’s appearance and image background by leveraging the image feature in self-attention to query the key and value at the corresponding inversion stage timestep, as opposed to the image feature itself.

We have performed extensive experiments to substantiate our efficacy. Both qualitative and quantitative outcomes demonstrate that our Move&Act proficiently edits the ob-

ject’s action in line with the editing prompt and accurately generates the edited object within the target bounding box.

Related Work

Text-to-Image Generation

Recent research utilizing diffusion models have made significant strides in the realm of image generation. Notably, (Dhariwal and Nichol 2021) has outperformed the majority of generative adversarial networks (GANs) in terms of generation metrics, indicating a substantial progression in the field. The use of diffusion models in the generation of images from text has produced remarkable outcomes.

To align generated images and prompts, Glide (Nichol et al. 2022) integrated diffusion models with guidance techniques. Large-scale diffusion models have been created (Saharia et al. 2022; Ramesh et al. 2022), exhibiting even more remarkable performance in generating images from text. However, the substantial size of these models renders them unsuitable for average users with limited computing capabilities. (Rombach et al. 2022) proposed the latent diffusion model (LDM), which employs an autoencoder to map images from the image space to the latent space, thereby reducing the computational demands of diffusion models. Furthermore, LDM incorporates a cross-attention layer within its structure, achieving outstanding results in the text-to-image task. Our work builds on LDM, specifically utilizing a pre-trained LDM called stable diffusion (Rombach et al. 2022).

Image Editing

Text Driven Image Editing. Recently, diffusion model-based image editing has achieved unparalleled results. For example, P2P (Hertz et al. 2023) utilizes cross-attention control to edit images from one prompt to another, while PnP (Tumanyan et al. 2023) introduces a plug-and-play method for text-driven image translation. Pix2Pix-Zero (Parmar et al. 2023) utilizes regularized forward DDIM and cross-attention guidance to achieve text-driven image translation. PTI (Dong et al. 2023) edits particular objects by optimizing the condition embedding during the image reconstruction process. This updated condition embedding is then utilized to finalize the image editing. NTI (Mokady et al. 2023) discovers that optimizing null-text Embedding can accurately reconstruct the original image while avoiding fine-tuning the model, thereby preserving the structure of image and enabling high-fidelity editing of real images.

Consistent Image Editing. Consistent image editing (Kawar et al. 2023; Cao et al. 2023) seeks to modify a single object in an image while leaving the rest of the image intact. DragGAN (Pan et al. 2023) was the first to introduce drag-style image editing to achieve this consistency. However, the quality and generalizability of DragGAN’s editing are constrained by its dependence on GANs. Capitalizing on the success of diffusion models, (Mou et al. 2024; Shi et al. 2023) have employed diffusion models to execute drag-style image editing, resulting in promising outcomes.

Imagic (Kawar et al. 2023) necessitates optimizing both textual embedding and diffusion model. It linearly interpo-

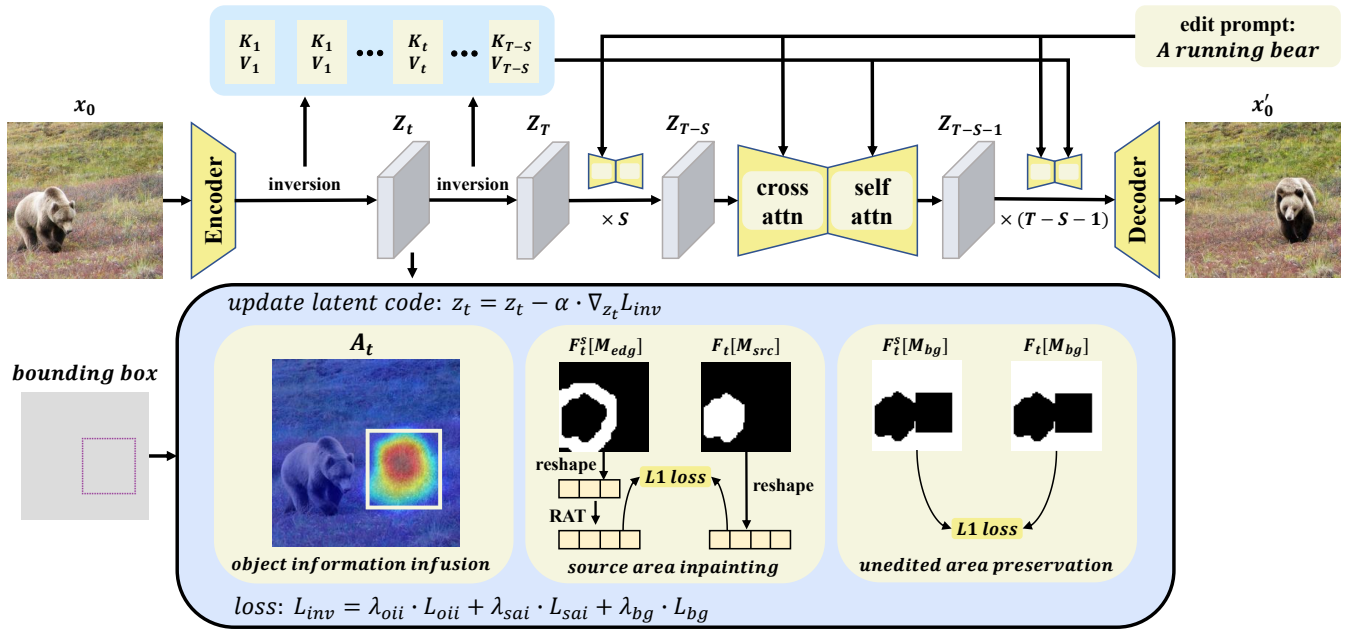


Figure 2: Move&Act pipeline. In the inversion stage, we update latent code and transfer object information to the target area while repairing the source background. In the reverse stage, we use self-attention image features for consistent image editing.

lates the target text embedding and the optimized embedding, and then inputs the fine-tuned model to generate the edited image. MasaCtrl (Cao et al. 2023) employs mutual self-attention which involves querying image content from source images using a self-attention mechanism to achieve consistent image editing. These methods are restricted to modifying the actions of objects and lack control over the generated position. In contrast, ours consists of an inversion branch and an edit branch, allowing for simultaneous editing of the object’s action and control over its generated position.

Methodology

The pipeline of our Move&Act is depicted in Figure 2, where we are given an image and an object to be edited. Our objective is to generate the edited object within the designated target bounding box while ensuring that it adheres to the editing prompt. In the subsequent sections, we refer to the user-provided bounding box area as the target area, the region where the object originally resides as the source area, and the complement of the union of both target and source areas as the unedited area (or background area).

To accomplish this editing task, we employ a two-step process. Firstly, we transfer the object information from the source area to the target area during a specific step in the inversion stage. Secondly, we modify the object’s action throughout the reverse stage, resulting in a seamlessly edited image with a sophisticated and coherent appearance.

Object Information Transfer

To facilitate the generation of the edited object within the target area, we infuse the object information into this designated area by updating the latent code during the 35-*th* time

step of the inversion process. Following the relocation of the object to the target area, it is essential to seamlessly inpaint the source area. In the ensuing sections, we will delve deeper into the mechanics of object information infusion and the intricacies of source area inpainting.

Object Information Infusion. Our observation reveals that the attention map, associated with each object token within the cross-attention mechanism of stable diffusion, encompasses spatial position information about the object within the image. This observation aligns with the findings of previous studies (Chefer et al. 2023; Xie et al. 2023). Leveraging this insight, several studies (Xie et al. 2023; Xue et al. 2023; Chen, Laina, and Vedaldi 2024; Phung, Ge, and Huang 2023) have utilized this mechanism to manipulate the arrangement of objects within images. In a similar vein to (Chefer et al. 2023; Xie et al. 2023), we utilize a 16×16 resolution attention map to update the latent code, as this particular resolution attention map is rich in semantic information. Consequently, to facilitate the relocation of the object to the target area, we ensure that the attention map, corresponding to the edited object, achieves the highest attention score within the confines of the target area:

$$L_{in} = 1 - \frac{1}{k} \text{top}(A_t[M_{tgt}], k), \quad (1)$$

where M_{tgt} symbolizes the mask of the target area, and A_t represents the cross attention map that corresponds to the t -*th* time step of the edited object token, with $t = 35$. The $\text{top}(\cdot, k)$ operation is employed to extract the largest k attention scores within the target area.

Conversely, we minimize the attention score attributed to areas outside the designated target area, ensuring a precise




Figure 3: We use cross attention map to locate the source area of the object, and use dilation operation to obtain the edge area around the source region.

and controlled editing process:

$$L_{out} = \frac{1}{N} \sum A_t [1 - M_{tgt}]. \quad (2)$$

The final loss for object information infusion becomes:

$$L_{oii} = L_{in} + L_{out}. \quad (3)$$

Source Area Inpainting. To perform inpainting on the source area, we must first accurately locate it. Given that the high-response region of the cross-attention map corresponding to the object token represents the object’s spatial position within the image, we employ a threshold to isolate this high-response area, which serves as the object’s source area, as illustrated in Figure 3. To circumvent conflicts arising from maximizing the attention score of the target area, we inpaint the actual source area by subtracting the union of the source and target areas from the original source area.

Intuitively, the background of the source area should exhibit similarities to the surrounding background. To achieve this, we first conduct a dilation operation on the mask corresponding to the source area, thereby obtaining the region surrounding the source area, as demonstrated in Figure 3. We denote the expanded surrounding area as the edge area.

In order to harmonize the background of the source area with that of the edge area, we extract features from both areas of the latent code and subsequently compute the L_1 loss between the features of these regions. DIFT (Tang et al. 2023) discovered that the UNet of the diffusion model is an efficient feature extractor, thus we utilize UNet to extract features from the latent code. The decoder of the UNet denoiser comprises four blocks of varying scales. According to DIFT, the second layer contains more semantic information, while the third layer encompasses more geometric information. In our method, we opt for the third layer.

Since the dimensions of the edge area and the source area may not necessarily be equal, we repeat and truncate the features of the edge area, to align features in both regions. The loss for source area inpainting becomes:

$$L_{sai} = \|F_t[M_{src}] - RAT(F_t^s[M_{edg}])\|_1, \quad (4)$$

where M_{src} and M_{edg} represent the source area and edge area. F_t^s denotes the feature of the original latent code, and F_t symbolizes the feature of the updated latent code, with $t = 35$. $RAT(\cdot)$ is a repeat and truncate operation.

Apart from its impact on the source area, L_{sai} also serves a potentially crucial role. Employing solely L_{oii} may not




Figure 4: Efficacy of source area inpainting loss L_{sai} . With/Without L_{sai} , the object fully/partly transfers to target area.

facilitate the complete transfer of object information to the target area, even though L_{oii} minimizes the attention scores of the non-target area. In Figure 4, when only L_{oii} is utilized, remnants of the object information persist in the source area, which could hinder the generation of the edited object within the target area. By concurrently employing L_{sai} and L_{oii} , L_{sai} will completely erase the object in the source area.

Background Preservation

Upon relocating the edited object from the source area to the target region, we observe that the infusion of object information and the subsequent inpainting of source area may inadvertently cause damage to the image’s background. To address this issue and maintain the integrity of the unedited areas, we propose a background-preserving loss function.

Utilizing the rich image structure and texture information contained within the intermediate features of the UNet denoiser, we employ UNet to extract features from the latent code. Subsequently, we compute the L_1 loss for the unedited regions between the original and updated latent codes:

$$L_{bg} = \|F_t[M_{bg}] - F_t^s[M_{bg}]\|_1, \quad (5)$$

where M_{bg} represents the unedited area. As illustrated in Figure 5, the absence of L_{bg} may result in the loss of some background information in the edited image. However, by introducing L_{bg} , we ensure consistency between the pre- and post-editing background areas. Simultaneously, the object information in the unedited region is suppressed, guaranteeing that the edited object adheres more closely to the user-specified bounding box constraints.

The ultimate loss function utilized for infusing object information and preserving background integrity at the 35-th time step of inversion process is:

$$L_{inv} = \lambda_{oii} \cdot L_{oii} + \lambda_{sai} \cdot L_{sai} + \lambda_{bg} \cdot L_{bg}, \quad (6)$$

where λ_{oii} , λ_{sai} , and λ_{bg} serve as balancing parameters. By applying L_{inv} , latent code z_t is updated with a step of α :

$$z_t = z_t - \alpha \cdot \nabla_{z_t} L_{inv}. \quad (7)$$

As updating step increases, the step size α undergoes linear decay. This approach ensures a smooth and consistent image editing process, preserving the background and seamlessly integrating the edited object into the target area.




Figure 5: Editing prompt: A standing cat. Efficacy of background preservation loss L_{bg} . Without L_{bg} , background details like meadow become distorted resembling MasaCtrl. With L_{bg} , the background remains well-preserved.

Content Editing Consistency

Upon obtaining z_T through the inversion process, we employ it as the initial point for the reverse denoising procedure. To ensure that the edited object generates the same action as indicated by the editing prompt, we integrate the editing prompt into the reverse process using the cross attention mechanism of UNet. Besides, to preserve the consistency of the image content prior to and following the editing process, we replace the key and value within the self attention mechanism with their counterparts from the corresponding time step in the inversion stage, commencing from the 7-th time step during the reverse procedure.

In contrast to the approach presented in (Cao et al. 2023), we query the key and value during the inversion stage, as we believe that the key and value from this stage contain richer information about the original image’s object appearance and background compared to using the key and value from the reconstruction stage as in (Cao et al. 2023). Furthermore, utilizing the key and value from the inversion stage diminishes the image reconstruction cost. Consequently, our optimization process transforms the original three branches of inversion, reconstruction, and editing (Cao et al. 2023) into two branches: inversion and editing. The exchange of key and value in the optimized operation can be expressed:

$$f = \begin{cases} \text{attention}(Q, K_{inv}, V_{inv}), & t > S \text{ and } l \in L, \\ \text{attention}(Q, K, V), & \text{otherwise,} \end{cases} \quad (8)$$

where K_{inv} and V_{inv} denote the key and value of the corresponding time step in the inversion stage. When the time step exceeds $S = 7$ and the self attention layer is part of the predefined set L , we alter the self attention key and value to K_{inv} and V_{inv} . By enabling image features to access and interact with the key and value of the input image, we facilitate a uniform and congruent image editing process, ultimately yielding a cohesive and polished result.

Experiments

Implementation Details

In the experimental section, we employ the stable diffusion (Rombach et al. 2022) using checkpoint v1.4. We utilize a 50-step DDIM inversion process, and the classifier-free guidance is set at 7.5. During the 35-th step of inversion, we update the latent code. The hyperparameter values for λ_{oi} ,

λ_{sai} , and λ_{bg} in Eq. (6) are set to 0.5, 0.25, and 0.25, respectively. In order to update the latent code in accordance with Eq. (7), we establish the step size α at 150 and designate the number of iterations as 50. From the 7-th step of the reverse process onward, we substitute the key and value within the self attention mechanism as per Eq (8). All experiments are conducted on a single 3090 GPU, well demonstrating the democratization of our introduced method in this paper.

Evaluation Dataset

We manually annotate 200 condition-image pairs from the web, the collected data focuses on non-rigid objects editing. Each image is linked to a specific condition, which includes the inversion prompt, editing prompt, and bounding box. Inversion prompts follow a common template: “A photo of <object>”, e.g., “A photo of cat”. The editing prompt template is “A <action> <object>”, e.g., “A running cat”. The bounding box designates the object’s generation location.

Qualitative Result

The image editing results of our method are vividly depicted in Figure 6. Our approach exhibits notable strengths in three key areas: (1) the edited object accurately adheres to the instructed action; (2) the edited object precisely aligns with the designated position box; (3) the background is seamlessly integrated into the edited image. As such, our method significantly enhances the manipulation of objects and the preservation of background integrity during the image editing process. This ultimately results in a more refined and visually appealing output, demonstrating the efficacy of our proposed technique.

We configure the generated position of the edited object to match the object’s original location, for comparing with existing methods. As shown in Figure 1 and Figure 7, our method exhibits superior background preservation, and the actions of the edited objects align more closely with the user-provided prompts. This comparison highlights our method in achieving enhanced consistency.

Quantitative Result

We calculate the following metrics to compare with previous methods: (1) CLIP Score (Radford et al. 2021), which assesses the association between text and images. A higher CLIP Score value signifies a greater correlation between text-image pairs. (2) YOLO Score (Li et al. 2021), which evaluates the accuracy of the generated position of the edited object using the AP_{50} metric. To detect the bounding box and predict the category of the edited object in the generated image, we employ YOLOv8 (Reis et al. 2023).

As illustrated in Table 1, our Move&Act achieves a higher CLIP Score, indicating that the editing results produced by our method exhibit greater consistency with the user-provided prompts. We also report our AP_{50} score. Since the methods we compared do not possess the capability to control the generated position of the edited object, the AP_{50} values for these methods are left blank.

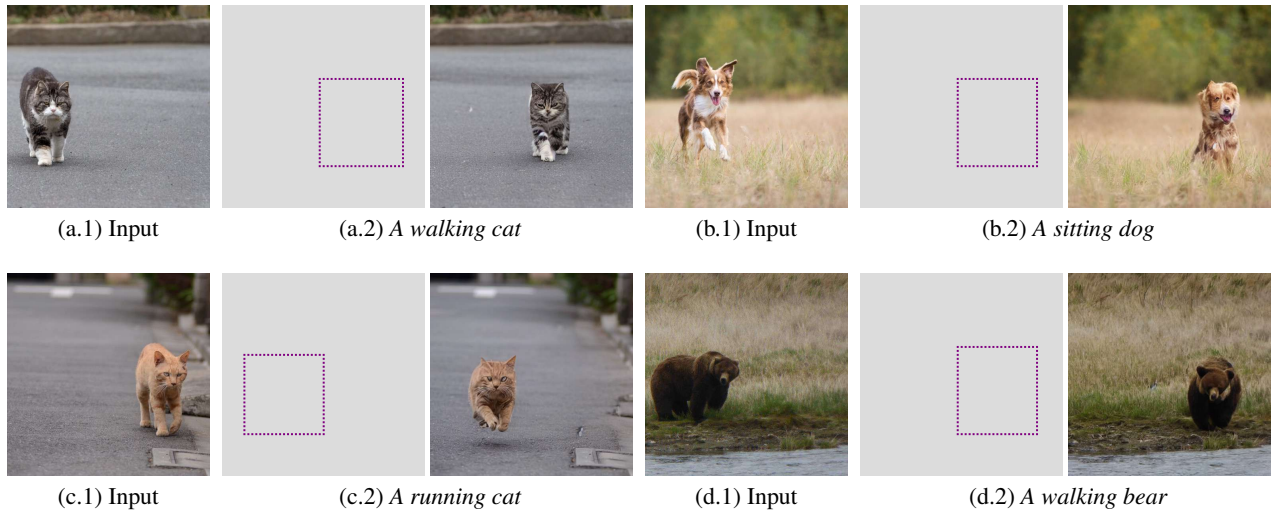


Figure 6: Visualization of our proposed Move&Act for consistent image editing.

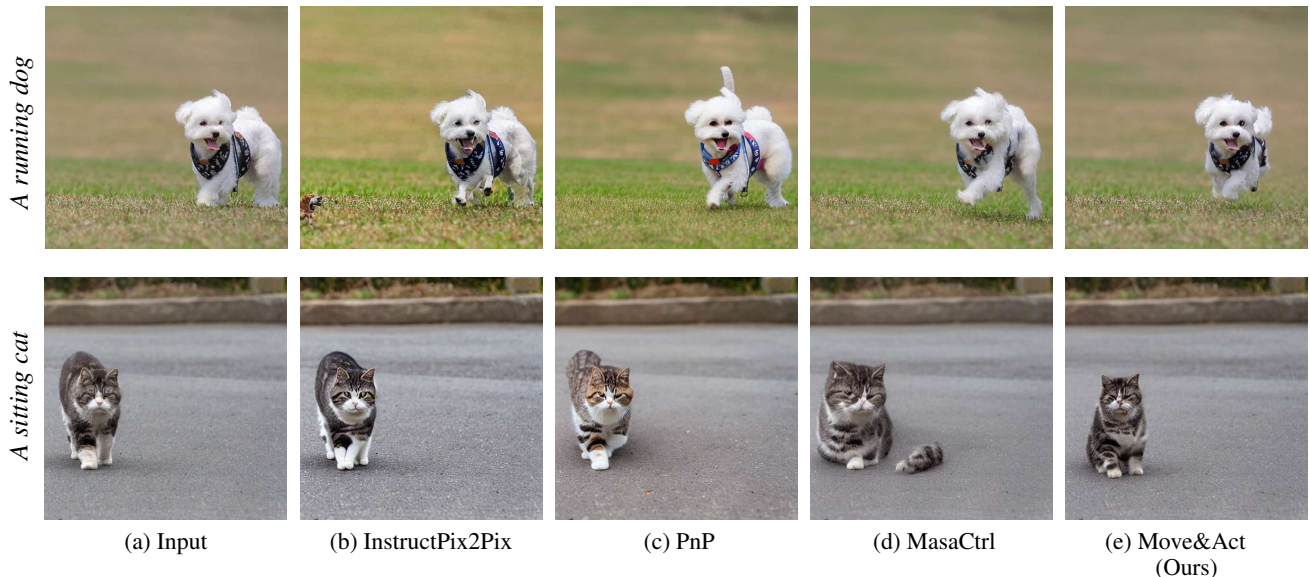


Figure 7: Visual comparison of different methods for modifying only actions.

Method	CLIP-Score	AP_{50}
InstructPix2Pix	0.2966 ± 0.0013	-
PnP	0.3035 ± 0.0009	-
MasaCtrl	0.3066 ± 0.0008	-
Move&Act (Ours)	0.3098 ± 0.0011	33.70 ± 0.74

Table 1: A comparison of quantitative results.

Ablation Study

Time Step to Update z_t During Inversion. To determine the optimal time step for updating the latent code during inversion branch in Eq. (6), we experiment with the 25-*th*, 35-

th, and 45-*th* time steps. As illustrated in Figure 8, setting the time step to 25 results in incomplete object information transfer, causing the object not to appear in the target area. Conversely, a time step of 45 allows for object information transfer but yields unsatisfactory visual results. The attention map visualization reveals that updating the latent code at the 35-*th* time step achieves the most effective object information transfer.

Iterations of Latent Updating. We perform experiments to find the suitable number of iterations to converge Eq. (6), as shown in Figure 9. With 25 iterations, the edited image struggles to maintain the original background effectively, and the cat’s action appears distorted. Conversely, setting the iteration number to a higher value of 75 allows for better

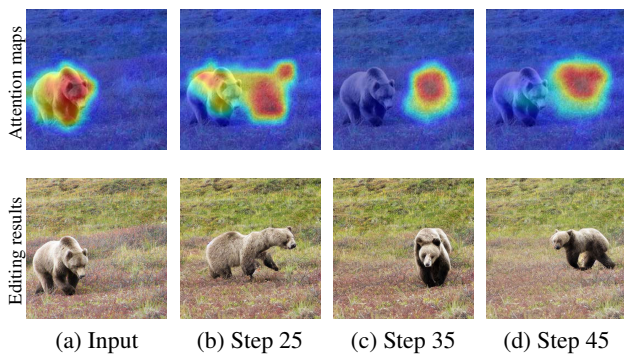


Figure 8: Editing results of different time step settings to update z_t during inversion.



Figure 9: Editing results of different iteration settings.

preservation of the original background; however, the cat’s action does not align well with the user’s provided prompt.

Taking these observations into account, we opt for an iteration number of 50 steps. This setting strikes a balance between preserving the background and accurately editing the object’s action, resulting in a more visually appealing and coherent edited image that adheres to the editing prompt.

Starting Time Step for Editing Consistency. In Eq. (8), our approach utilizes image features in self-attention to query the key and value at the corresponding time step during inversion, starting from time step S in the reverse process. As depicted in Figure 10, we experimented with different values of S . Setting S to a larger time step results in the edited object losing its original appearance, with alterations in the dog’s decorations, hair, and the image background. When S is set to 7, the editing outcome aligns more closely with the user’s prompt.

Self-attention Layer Setting. We only replace the key and value in the specified self attention layer in Eq. (8). In order to test which self-attention layers in UNet play a major role in restoring the object appearance and the image background, we implement an experiment and test three replacement schemes, which were to replace all self-attention key and value in UNet, replace only self-attention key and value of the UNet encoder and replace only self-attention key and value of the UNet decoder. As shown in Figure 11, decoder self-attention plays a major role in restoring the object appearance and the image background, while encoder self-attention plays a minor role. Thus, we only replace the key and value of the UNet decoder’s self attention.

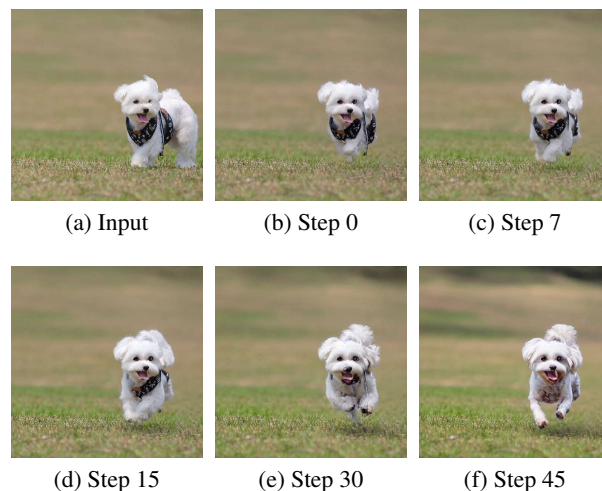


Figure 10: Time step for editing consistency.




Figure 11: Results of different self-attention layer settings.

Limitation and Future Work

Complex image input is a challenge for our method, particularly in effectively restoring the source area’s background when it is surrounded by complex background. Additionally, the kernel size choice for the dilation operator is pivotal in our approach. A smaller kernel size results in an edge area with deficient features, leading to suboptimal outcomes. Conversely, an excessively large kernel size produces a broad edge area, encompassing background regions too distant from the source area to be relevant for restoration. To refine our method, future research will concentrate on investigating source area inpainting loss with an enhanced capacity for background recovery, thereby ensuring more precise and consistent results across various editing scenarios.

Conclusion

We have introduced a tuning-free approach for achieving consistent image editing, enabling users to modify an object’s action and control its generated position, while ensuring superior background preservation. Our method transfers object information from the source area to the target area during the inversion stage and employs the features of the UNet denoiser to restore the source area’s background, preserving the unedited area’s background. This innovative approach facilitates more precise control over image objects and enhances background consistency before and after editing, causing a seamless and visually appealing outcome.

Acknowledgements

This work was supported by National Science and Technology Major Project (No. 2022ZD0118201), the National Science Fund for Distinguished Young Scholars (No.62025603), the National Natural Science Foundation of China (No. U21B2037, No. U22B2051, No. U23A20383, No. U21A20472, No. 62176222, No. 62176223, No. 62176226, No. 62072386, No. 62072387, No. 62072389, No. 62002305 and No. 62272401), and the Natural Science Foundation of Fujian Province of China (No. 2021J06003, No.2022J06001).

References

- 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.
- Cao, M.; Wang, X.; Qi, Z.; Shan, Y.; Qie, X.; and Zheng, Y. 2023. MasaCtrl: Tuning-Free Mutual Self-Attention Control for Consistent Image Synthesis and Editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 22503–22513.
- Chefer, H.; Alaluf, Y.; Vinker, Y.; Wolf, L.; and Cohen-Or, D. 2023. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. *ACM Transactions on Graphics*, 1–10.
- Chen, M.; Laina, I.; and Vedaldi, A. 2024. Training-free layout control with cross-attention guidance. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 5343–5353.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. In *Advances in Neural Information Processing Systems*, 8780–8794.
- Dong, W.; Xue, S.; Duan, X.; and Han, S. 2023. Prompt Tuning Inversion for Text-Driven Image Editing Using Diffusion Models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7396–7406.
- Hertz, A.; Mokady, R.; Tenenbaum, J.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2023. Prompt-to-prompt image editing with cross attention control. In *International Conference on Learning Representations*.
- Huang, J.; Liu, Y.; Qin, J.; and Chen, S. 2023. Kv inversion: Kv embeddings learning for text-conditioned real image action editing. In *Chinese Conference on Pattern Recognition and Computer Vision*, 172–184.
- Kawar, B.; Zada, S.; Lang, O.; Tov, O.; Chang, H.; Dekel, T.; Mosseri, I.; and Irani, M. 2023. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6007–6017.
- Li, Z.; Wu, J.; Koh, I.; Tang, Y.; and Sun, L. 2021. Image synthesis from layout with locality-aware mask adaption. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13819–13828.
- Mokady, R.; Hertz, A.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2023. Null-text inversion for editing real images using guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6038–6047.
- Mou, C.; Wang, X.; Song, J.; Shan, Y.; and Zhang, J. 2024. Dragondiffusion: Enabling drag-style manipulation on diffusion models. In *International Conference on Learning Representations*.
- Nichol, A.; Dhariwal, P.; Ramesh, A.; Shyam, P.; Mishkin, P.; McGrew, B.; Sutskever, I.; and Chen, M. 2022. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In *International Conference on Machine Learning*, 16784–16804.
- Pan, X.; Tewari, A.; Leimkühler, T.; Liu, L.; Meka, A.; and Theobalt, C. 2023. Drag your gan: Interactive point-based manipulation on the generative image manifold. In *ACM SIGGRAPH 2023 Conference Proceedings*, 1–11.
- Parmar, G.; Kumar Singh, K.; Zhang, R.; Li, Y.; Lu, J.; and Zhu, J.-Y. 2023. Zero-shot image-to-image translation. In *ACM SIGGRAPH 2023 Conference Proceedings*, 1–11.
- Phung, Q.; Ge, S.; and Huang, J.-B. 2023. Grounded Text-to-Image Synthesis with Attention Refocusing. *arXiv preprint arXiv:2306.05427*.
- 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.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 3.
- Reis, D.; Kupec, J.; Hong, J.; and Daoudi, A. 2023. Real-Time Flying Object Detection with YOLOv8. *arXiv preprint arXiv:2305.09972*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10684–10695.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems*, 36479–36494.
- 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*.
- Tang, L.; Jia, M.; Wang, Q.; Phoo, C. P.; and Hariharan, B. 2023. Emergent correspondence from image diffusion. In *Advances in Neural Information Processing Systems*.
- 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.

Xie, J.; Li, Y.; Huang, Y.; Liu, H.; Zhang, W.; Zheng, Y.; and Shou, M. Z. 2023. Boxdiff: Text-to-image synthesis with training-free box-constrained diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7452–7461.

Xue, H.; Huang, Z.; Sun, Q.; Song, L.; and Zhang, W. 2023. Freestyle Layout-to-Image Synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14256–14266.