

FlowCycle: Target-aware Image Editing via Cycle-consistent Constraints

Yanghao Wang, Zhen Wang, and Long Chen*

The Hong Kong University of Science and Technology
 ywangtg@connect.ust.hk, zhenwang@ust.hk, longchen@ust.hk
<https://github.com/HKUST-LongGroup/FlowCycle>




Fig. 1: Comparisons between target-agnostic and target-aware corruption manner.

Abstract. Recent pre-trained text-to-image flow models have enabled remarkable progress in text-based image editing. Mainstream approaches adopt a **corruption-then-restoration** paradigm, where the source image is first corrupted into an editable “intermediate state” and then restored to the target image under the prompt guidance. However, current methods construct this intermediate state in a **target-agnostic** manner, *i.e.*, they mainly focus on realizing source image reconstruction while neglecting the semantic gaps towards the specific editing target. This design inherently results in limited editability or inconsistency when the desired modifications substantially deviate from the source. In this paper, we argue that the intermediate state should be **target-aware**, *i.e.*, selectively corrupting editing-relevant contents while preserving editing-irrelevant ones. Thus, we propose **FlowCycle**, an inversion-free and flow-based editing framework that parameterizes corruption with learnable noises and optimizes them through a cycle-consistent process. By iteratively editing the source to the target and recovering back to the source

* Long Chen is the corresponding author.

with dual consistency constraints, FlowCycle learns to produce a target-aware intermediate state, enabling faithful modifications while preserving source consistency. For efficiency, we further accelerate the optimization by dynamically adjusting the sampling steps. Extensive ablations demonstrated that FlowCycle achieves superior editing performance.

Keywords: Diffusion Model · Target-aware Editing · Cycle Consistency




Fig. 2: Corruption-then-restoration paradigm for existing flow-based image editing methods. Existing methods obtain target-agnostic intermediate states for subsequent editing. In contrast, FlowCycle obtains target-aware intermediate states via cycle-consistency optimization, which is both editable and source-preserving.

1 Introduction

Pre-trained text-to-image flow models [8, 20] have recently proven strong generative capacity, which has gathered growing interest in adapting them for text-based editing [19, 30, 42]. Unlike standard text-to-image generation, editing requires not only producing new content but also preserving certain source details. Specifically, given a source image and its description, the key challenge for text-based editing is to modify the image according to a target prompt while maintaining non-target content as the source, *i.e.*, *source consistency*. To achieve this goal, most existing flow-based editing approaches [16, 19, 30, 42] closely follow well-established diffusion-based methods, since both flow-based and diffusion-based [12, 32] models share a similar generative principle, *i.e.*, they corrupt the clean image with progressive noise and then learn to reverse this process to generate new samples. Consequently, existing flow-based editing methods generally adopt a common **corruption-then-restoration** paradigm. As shown in Figure 2(a), this paradigm consists of two main stages:

1. **Corruption.** The source image is first transformed into an *intermediate state* (*i.e.*, the corrupted image), which allows flexible editing while maintaining the source consistency required to restore the source image. As shown in Figure 2(b), the representative approaches include *inversion-based* methods [7, 46] that perform ODE inversion [32] and subsequent improvements [17, 30, 36]

to enhance the source reconstruction accuracy. As well as *inversion-free* methods [16, 19, 24] that directly sample noise and add it into the source image to improve the efficiency of the corruption process.

2. **Restoration.** Starting from the intermediate state, the pre-trained flow model restores the image under the guidance of the target prompt. To further enhance the preservation of source consistency, many methods inject source conditions (*cf.*, Figure 2(c)) during the restoration stage [17, 30, 36, 46].

In that case, an *ideal intermediate state* of this paradigm should strike a balance between editability and preservation of source consistency. As shown in Figure 1(b), given a specific source-target prompt pair that aims to edit the cat (*e.g.*, ...**cat**... $\rightarrow$...**tiger**...), the corruption should be applied to the editing-relevant contents (*e.g.*, **cat**) to ensure faithful modifications, while the source information of editing-irrelevant contents (*e.g.*, **fish tank**) should be preserved to maintain reliable source consistency. However, existing methods all generally acquire the intermediate state in a **target-agnostic** manner, *i.e.*, they emphasize the reconstruction of the entire source image without considering the specific editing targets, relying solely on the prior knowledge (*e.g.*, source prompt and pre-trained model) to achieve this balance. As shown in Figure 1(a), such intermediate states reveal an approximate reconstruction of the source image, preserving unbiased global layout structures and semantic details. When restored with different target prompts focusing on distinct contents, this source-aware but target-agnostic design often leads to unsatisfactory results, especially when the editing target significantly deviates from the source image or the prior knowledge. For instance, as shown in the Figure 1(a), when the target is to change the shape of the bubble (*e.g.*, ...**bubbles**... $\rightarrow$...**heart bubbles**...), the target-agnostic intermediate state not only fails to reliably preserve editing-irrelevant contents (*e.g.*, **cat** and **goldfish**) but also conflict with editing-relevant contents (*e.g.*, **bubbles**), making them difficult to modify as desired. Meanwhile, even injecting additional source conditions during restoration cannot resolve the problem, since the intermediate state itself remains misaligned with the editing target.

In this paper, we argue that the *ideal intermediate state* should be constructed in a **target-aware** manner, selectively corrupting editing-relevant contents while preserving editing-irrelevant ones based on specific editing targets to enable faithful and effective editing. To this end, we introduce **FlowCycle**, a novel inversion-free and flow-based framework that faithfully incorporates target awareness, enabling precise and reliable text-based image editing. Notably, such a target-aware state exhibits a natural symmetry between the source and target image: the editing-relevant regions should be interchangeable, while the editing-irrelevant regions should remain invariant across both. Motivated by this, we construct a *cycle* between the source and target image, allowing the model to exploit this inherent symmetry and learn a target-aware corruption.

Specifically, instead of relying on fixed inversion or random noise sampling, our method parameterizes the corruption process with learnable noises through a cycled framework. As shown in Figure 2(d), it consists of three main steps. 1) *Source to target* (*i.e.*, the red path of Figure 2(d)): the source image is first

corrupted with a learnable source noise and then restored under the guidance of the target prompt, producing an initial target image that reflects the desired edit. 2) *Target to source* (*i.e.*, the blue path of Figure 2(d)): the initial target image is also corrupted with another learnable target noise and then restored with the source prompt, yielding a recovery image that traces back to the source. 3) *Cycle-consistent Optimization*: we impose two constraints to optimize the learnable noises — (i) align the recovery image to the source image, and (ii) align the corrupted source image with the corrupted target image. This joint regularization stabilizes the optimization and enables target-aware corruption. By iterating this process, learned noises gradually produce a target-aware intermediate state (*cf.*, Figure 1(b)), where the corrupted source image and the corrupted target image share an ideal role, *i.e.*, corrupt the editing-relevant content while maintaining the editing-irrelevant parts based on a specific editing target.

Although this optimization process introduces additional computational overhead due to iterative generation, it can be naturally accelerated by leveraging the sampling flexibility of flow models. Accordingly, we further propose a progressive sampling strategy that dynamically adjusts the number of sampling steps across optimization rounds to improve efficiency while maintaining editing quality.

In summary, we made three contributions in this paper:

- We identify a fundamental limitation in prevalent corruption-then-restoration editing methods: the intermediate state is acquired in a *target-agnostic* manner, which often leads to unsatisfactory results. Moreover, we argue that the target-aware paradigm is of significant necessity.
- To achieve target-aware editing, we propose the first *target-aware* corruption strategy and introduce FlowCycle, an inversion-free, flow-based framework that adaptively focuses corruption on editing-relevant content while preserving irrelevant content. Meanwhile, we introduce an efficient acceleration strategy to reduce the computational overhead of the optimization process.
- Extensive experiments across benchmarks and diffusion backbones have empirically shown the robust effectiveness of FlowCycle over existing methods.

2 Related Work

Existing text-based image editing methods largely converge on the *corruption-then-restoration* paradigm, improving editing performance at either the corruption or the restoration stage. While a few alternative approaches, such as unified diffusion models [9, 35, 40, 41] or MLLM-driven models [2, 6, 37], also contributed to the field, this paper focuses primarily on methods within this framework due to its broad adoption and foundational role.

Inversion-based Corruption. Methods in this category construct the intermediate state from the source image via ODE inversion [7, 32, 46], which allows the preservation of structural and semantic details for reconstruction (*i.e.*, restore with an empty prompt). However, the iterative inversion is prone to error accumulation [13, 27], which degrades the fidelity of source information and weakens

the reliability of subsequent editing. Thus, recent works focus on improving inversion accuracy [5, 17, 30, 36], enabling more faithful reconstruction of the source image and providing a stronger foundation for high-quality editing.

Inversion-free Corruption. To alleviate the high computational overhead caused by ODE inversion, inversion-free methods [14, 16, 19, 24] bypass the inversion step by directly sampling random noise to the source image to construct the intermediate state. Compared with inversion-based strategies, these methods achieve significantly faster corruption. However, they remain highly sensitive to noise level [26] and hyperparameter choices [16, 19]. Excessive corruption risks erasing key source information, while insufficient perturbation limits editability. Thus, the central challenge lies in striking the right balance between editability and faithful preservation of source consistency.

Injection-based Restoration. Solely relying on the intermediate state for restoration often results in partial information loss and inconsistency in the edited image [13]. To address this, a number of methods propose to explicitly inject source information throughout the restoration process, ensuring stronger alignment between the edited image and the source image. This can be achieved through additional model inputs, *e.g.*, optimized prompt [27] and latent [13, 39]. Intermediate model features [17, 36, 44] derived from the source image, *e.g.*, attention maps [10, 34]. And enhanced model outputs [30, 38, 46] such as predicted noise or velocity. Nevertheless, without a properly constructed and target-aware intermediate state, such injection strategies still offer limited improvements.

Cycle Consistency in Generation. Well-recognized cycle consistency [47] has also been explored, particularly in image-to-image translation tasks, where consistency regularization is enforced through translation cycles between source and target domains. Such methods typically require images from two specific domains for training or optimization [21, 31, 43]. While recent work has attempted to adapt cycle consistency to image editing [3], their “cycle” still focuses solely on source inversion and reconstruction. In contrast, our work leverages cycle consistency to enable a novel *target-aware* corruption during text-based editing, where the cycle is established between source and target to explicitly align editing-relevant and irrelevant contents.

3 Method

3.1 Preliminaries: Flow Matching (FM)

Flow Matching (FM) models [22, 23] are trained to fit a velocity field $u_t(x_t)$ on time $t \in [0, 1]$ such that they can transport data from one distribution (*e.g.*, normal distribution π_1) to another distribution (*e.g.*, images distribution π_0). Specifically, given a data point $x_1 \sim \pi_1$, the transport path can be simulated by solving an ordinary differential equation (ODE):

$$dx_t = u_t(x_t)dt, \quad (1)$$

which converts x_1 into $x_0 \sim \pi_0$. To match the underlying velocity field $u_t(x_t)$, a velocity prediction network [4] $v_\theta(x_t, t)$ was trained by minimizing the such loss:

$$\mathcal{L}_{FM} = \mathbb{E}_{x_t, t} \|v_\theta(x_t, t) - u_t(x_t)\|^2. \quad (2)$$

However, the ground-truth velocity field $u_t(x_t)$ is not accessible. Alternatively, prior studies [22, 23] propose to optimize the approximated conditional FM loss:

$$\mathcal{L}_{CFM} = \mathbb{E}_{x_0, p_t(x_t|x_0), t} \|v_\theta(x_t, t) - u_t(x_t|x_0)\|^2, \quad (3)$$

where $p_t(x_t|x_0)$ is usually defined as the optimal transport [25] probability path, *i.e.*, $p_t(x_t|x_0) = \mathcal{N}(x_t; (1-t)x_0, t^2I)$ while I is the identity matrix. Since π_1 is standard normal distribution, we can sample $x_1 \sim \pi_1$ and get:

$$x_t = (1-t)x_0 + tx_1 \Rightarrow u_t(x_t|x_0) = \frac{dx_t}{dt} = x_1 - x_0. \quad (4)$$

By substituting $u_t(x_t|x_0) = x_1 - x_0$ into Eq. (3), get final training objective:

$$\mathcal{L}_{CFM} = \mathbb{E}_{x_0 \sim \pi_0, x_1 \sim \pi_1, t} \|v_\theta(x_t, t) - (x_1 - x_0)\|^2. \quad (5)$$

In the inference stage, we can start from a randomly sampled Gaussian noise $x_1 \sim \pi_1$ and then integrate to generate a clean image: $x_0 = x_1 - \int_1^0 v_\theta(x_t, t) dt$.

3.2 Corruption-then-Restoration Paradigm

Given a source image x_0^{src} , source prompt c_{src} , and target prompt c_{tar} , the paradigm has two stages:

1) Corruption: Convert the source image x_0^{src} into an intermediate state x_t^{src} :

$$x_t^{src} = \psi(x_0^{src}, t, c_{src}), \quad t \in (0, 1], \quad (6)$$

where ψ is a corruption function such as directly adding random noise (inversion-free) or calculating an inversion (inversion-based). The principle of both corruptions is obtaining an intermediate state x_t^{src} while maintaining the information of the source image x_0^{src} . In that case, starting from x_t^{src} can benefit *source consistency*, *i.e.*, the editing result can maintain the editing-irrelevant part. It is worth noting that current corruption is target-agnostic, *i.e.*, the noisy image x_t^{src} is fixed for different editing targets. In contrast, we aim to find an ideal target-aware x_t^{src} that can specifically destroy the editing-relevant part while maintaining the editing-irrelevant part.

2) Restoration: After getting x_t^{src} , the most straightforward approach is to take it as the starting point of restoration directly (*i.e.*, $x_t^{tar} = x_t^{src}$) and denoising it under the guidance of the target prompt c_{tar} . However, only relying on the starting point that contains the source image’s information is not enough for *source consistency*. Thus, existing methods inject a drift item Δ_t (containing the source image’s information) during the denoising:

$$dx_t^{tar} = \xi(v_\theta, x_t^{tar}, t, c_{tar}, \Delta_t) dt, \quad (7)$$

where ξ represents different injection functions. For instance, adding Δ_t to the middle attention maps, predicted velocity v_θ , or input x_t^{tar} . The principle is




Fig. 3: Framework of FlowCycle. We rectify the corruption process by optimizing two learnable noises (*i.e.*, ϵ_{src} and ϵ_{tar}) under the cycle-consistent constraints of whole editing. After optimization, a target-aware intermediate state x_t^{src} can be obtained by adding ϵ_{src} to the source image x_0^{src} . Starting the restoration from this optimized x_t^{src} , the quality of the editing result x_0^{tar} can be significantly improved.

enhancing the *source consistency* of the final editing result by injecting a source-relevant drift item Δ_t during denoising. This drift item can rectify the denoising path and help the final editing result maintain more editing-irrelevant parts.

Our motivation is to find a smart target-aware intermediate state x_t^{src} as the start point of restoration. In that case, even if there is no source condition injection (*i.e.*, Δ_t), we can still gain a good editing result since the source information is preserved in the intermediate state quite well. In the following Section 3.3, we propose one optimization-based solution to find this ideal x_t^{src} .

3.3 FlowCycle

Overview. As shown in Figure 3(a), our FlowCycle contains three steps, including: 1) *Source to Target*: We first add a parameterized learnable noise ϵ_{src} (initialized as a Gaussian noise) to the source image x_0^{src} and get the intermediate state x_t^{src} . After getting x_t^{src} , we leverage the pretrained flow matching model to denoise it conditioned on the target prompt c_{tar} . The denoised result is the target image x_0^{tar} . 2) *Target to Source*: After getting the initialized target image x_0^{tar} , we add another parameterized learnable noise ϵ_{tar} (initialized as a Gaussian noise) to it and then get the noisy image z_t^{tar} . After that, we use the flow matching model conditioned on the source prompt c_{src} to denoise it into a recovery image z_0^{src} . 3) *Cycle-consistent Optimization*: After (1) and (2), we calculate Mean Squared Error (MSE) loss between x_t^{src} and z_t^{tar} as well as the MSE loss between x_0^{src} and z_0^{src} . Then we optimize the parameterized ϵ_{src} and ϵ_{tar} to decrease the loss. By repeating the three steps, ϵ_{src} and ϵ_{tar} are iteratively

optimized. In the inference stage (*cf.*, Figure 3(b)), the optimized ϵ_{src} is used for getting the target-aware intermediate state x_t^{src} . Then we can directly denoise from this intermediate state under the guidance of the target prompt without any other source condition injection and gain a satisfactory editing result.

Source to Target. We first randomly sample a noise ϵ_{src} from the standard normal distribution and parameterize it as a learnable vector. Then, we follow the forward process of FM to add ϵ_{src} to the source image x_0^{src} :

$$x_t^{src} = (1 - t) * x_0^{src} + t * \epsilon_{src} \quad t \in (0, 1), \quad (8)$$

where x_t^{src} is the intermediate state obtained by linearly interpolating the source image and noise. The interpolation ratio is controlled by t . We take x_t^{src} as the starting point of the following restoration process, *i.e.*, $x_t^{tar} = x_t^{src}$. By inputting x_t^{tar} into the pretrained text-to-image FM model (*i.e.*, the velocity prediction network v_θ) under the guidance of target prompt c_{tar} , we can get the predicted velocity $v_\theta(x_t^{tar}, t, c_{tar})$, which can transport x_t^{tar} to x_0^{tar} :

$$x_{t-\Delta}^{tar} = x_t^{tar} - v_\theta(x_t^{tar}, t, c_{tar}) * \Delta \Rightarrow x_0^{tar}, \quad (9)$$

where Δ is the step size. After the above process, we can get a naive target editing image x_0^{tar} (*e.g.*, SDEdit [26] directly takes it as the final editing result). However, it turns out x_0^{tar} suffers from the low *source consistency* because the random corruption indiscriminately destroyed the information of the source image.

Target to Source. To further improve x_0^{tar} , we try to recover the source image from x_0^{tar} . Considering its symmetrical nature, a better recovery result indicates that x_0^{tar} is of higher *source consistency*. Specifically, we parameterize another noise ϵ_{tar} (initialized from standard Gaussian noise) and add it to x_0^{tar} :

$$z_t^{tar} = (1 - t) * x_0^{tar} + t * \epsilon_{tar} \quad t \in (0, 1), \quad (10)$$

where z_t^{tar} is the noised version (the intermediate state of the reverse editing process) of the generated target image x_0^{tar} . We take z_t^{tar} as the starting point of the following restoration process, *i.e.*, $z_t^{src} = z_t^{tar}$ and denoise it under the guidance of source prompt c_{src} :

$$z_{t-\Delta}^{src} = z_t^{src} - v_\theta(z_t^{src}, t, c_{src}) * \Delta \Rightarrow z_0^{src}, \quad (11)$$

where z_0^{src} is the recovery of source image. The better recovery (*i.e.*, z_0^{src} is closer to x_0^{src}) indicates that the generated target image x_0^{tar} maintains editing-irrelevant information of the source image better. After the above two steps, we can conduct the following cycle-consistent optimization.

Cycle-consistent Optimization. As shown in Figure 3(a), the recovery image z_0^{src} is quite different from the source image x_0^{src} because the two intermediate

states (*i.e.*, x_t^{src} and z_t^{tar}) of the whole cycle were indiscriminately destroyed and the global source information was lost. To get the ideal target-aware intermediate state, we constrain the whole cycle to satisfy two cycle-consistent objectives:

1) The two intermediate states x_t^{src} and z_t^{tar} should share the aligned semantics as corrupted images, *i.e.*, destroy the editing-relevant parts of the source image (*e.g.*, **white dog** and **cat**) while maintaining the editing-irrelevant parts (*e.g.*, **background** and **car**). To achieve that, we use an alignment constraint to force x_t^{src} and z_t^{tar} to be close in the latent space:

$$\mathcal{L}_{align} = \|x_t^{src} - z_t^{tar}\|_2^2. \quad (12)$$

2) Since x_t^{src} and z_t^{tar} share the same semantics, starting from z_t^{tar} , the recovery result (*i.e.*, z_0^{src}) should be highly similar to the source image x_0^{src} because the ideal z_t^{tar} can maintain the editing-irrelevant parts. Thus, we take a recovery constraint to encourage z_t^{tar} to be a good start point of recovering to the source image x_0^{src} :

$$\mathcal{L}_{rec} = \|z_0^{src} - x_0^{src}\|_2^2. \quad (13)$$

To decrease the above two losses, we leverage the property that $\epsilon_{src}, \epsilon_{tar}$ can effectively rectify x_t^{src} and z_t^{tar} based on Eq. (8) and Eq. (10) and further rectify the whole cycle. Thus, we can jointly optimize the two learnable noises to decrease the $\mathcal{L}_{align}$ and $\mathcal{L}_{rec}$:

$$\epsilon_{src}, \epsilon_{tar} = \underset{\epsilon_{src}, \epsilon_{tar}}{\operatorname{argmin}} (\mathcal{L}_{rec} + \lambda \mathcal{L}_{align}), \quad (14)$$

where λ is a weight hyperparameter to balance the two loss items. Through the above cycle-consistent optimization, $\epsilon_{src}, \epsilon_{tar}$ can selectively destroy the information according to whether it is editing-relevant. Then the intermediate state x_t^{src} can be rectified towards the ideal role to be both editable and editing-irrelevant preserved. During inference, we can directly denoise the target-aware intermediate state x_t^{src} to get an improved editing result (*cf.*, Figure 3(b)).

3.4 Optimization Acceleration

Our primary motivation is to obtain a target-aware intermediate state that preserves editing-irrelevant regions while making editing-relevant parts mutable. As the first work in this direction, FlowCycle offers a straightforward and effective solution. Although this optimization-based framework inevitably introduces higher computational cost than optimization-free alternatives, its efficiency can be naturally improved by exploiting the inherent sampling flexibility of flow models. Specifically, the dominant computational overhead arises from generating two images (target image x_0^{tar} and recovery image z_0^{src}) within each optimization round, resulting in a large number of function evaluations (NFEs). Therefore, accelerating the sampling trajectory during generation provides a direct way to reduce the overall cost. This can be achieved through dynamic step scheduling, advanced ODE solvers, or early-stopping strategies, all of which are compatible with our framework.




Fig. 4: Quantitative comparisons. The closer to the bottom right corner, the better the balance between Structure Distance (source consistency) and CLIP Score (semantic alignment). Connected markers represent different hyperparameters (see Appendix). Ours can achieve the best balance across datasets and diffusion models.

In this paper, we pioneer a simple yet effective acceleration strategy by dynamically adjusting the sampling steps during optimization. The key observation is that in early optimization rounds, the learnable noise is far from the ideal state, thus we can use fewer sampling steps. As the optimization proceeds, more fine-grained updates are required, motivating the use of increased sampling steps in later rounds. Therefore, we design a monotonically increasing sampling schedule to decide the sampling steps for each optimization round. For example, the scheduler can be:

$$f(n) = \text{rounding}(\sin(\frac{\pi n}{2N}) * M), \quad (15)$$

where N is the overall optimization rounds, n is the current optimization round index, and M is the largest sampling steps. Moreover, this scheduler can be flexibly replaced by other monotonically increasing functions, such as linear or exponential schedules, to further enhance the acceleration effect.

4 Experiments

Implementation Details. We used SD-3-medium [8] from HuggingFace as the pretrained text-to-image flow model. The overall timestep number is set to 50 and the corruption step $t = 33$ (*cf.*, Eq. (8)). We used the Adam optimizer [18] with $1e-1$ learning rate for 100 rounds of noise optimization. The loss weight λ was set to 0.2. More details are in the Appendix and the codes.

Benchmarks. We evaluated our method on two benchmarks. 1) PIE-Bench [15]: a text-based image editing benchmark. It contains 700 (source image, source prompt, target prompt) tuples, which cover nine editing types (*e.g.*, “add object” and “change style”). Each image is 512^2 resolution. Besides, there is usually




Fig. 5: Qualitative comparisons on 512^2 resolution PIE-Bench (left half) and 1024^2 resolution DIV2K (right half). Compared with various baselines across different editing types (e.g., object or style change), FlowCycle can significantly maintain the editing-irrelevant parts of the source image while adhering to the editing target. Meanwhile, these two datasets cover different resolutions and editing complexity (e.g., DIV2K contains multi-object edit cases), indicating our robust effectiveness.

only a single editing target in each prompt. 2) DIV2K [1] subset: containing over 70 real images with around 280 (source image, source prompt, target prompt) tuples. Each image is 1024^2 resolution. Meanwhile, this dataset has longer target prompts that may contain multiple editing targets, which is more challenging. For the evaluation metrics, we reported the *CLIP Score* [11] (for semantic alignment evaluation) and *Structure Distance* [33] (for source consistency evaluation). More details about benchmarks are in the Appendix.

4.1 Comparison with State-of-the-Art Methods

Baselines. We compared our method with five state-of-the-art flow-based editing methods that were implemented with SD-3-medium: 1) *ODE inversion* (*ODE-Inv*), which gets the inversion by using the Euler Solver and denoises with the target prompt. 2) *SDEdit* [26], which adds random noise and then denoises with the target prompt. 3) *iRFDS* [45], which is a SDS-based editing method. 4) *FlowEdit* [19] and 5) *FlowAlign* [16] are two inversion-free editing methods. The results and hyperparameters for the reimplementations of these methods follow previous works. Moreover, since the target-aware editing manner is not limited




Fig. 6: Optimization rounds ablation. From left to right, the quality of the editing is gradually improved (especially for source consistency) across optimization rounds.

to Flow-based models, our method can also be compatible with other diffusion backbones like DDPM-based models. Thus, we also compared our method with Prompt-to-Prompt [10] (*P2P*) and pix2pix-zero [28] (*P2P Zero*) based on SD-1.5 [29]. We also compared with the editing model Qwen-Image-Edit-2509 (*Qwen-Image*) [40]. More details are in the Appendix.

Qualitative Evaluation. We gave qualitative comparisons in Figure 5. We have three observations: 1) Our method produced high-quality editing results (high source consistency and semantic alignment) on cases from various editing types. This indicates that our method is generally effective for extensive editing requirements. 2) Highlightly, our method can significantly maintain the editing-irrelevant parts. This advantage exactly verifies that the target-aware corruption can smartly hold the editing-irrelevant source content while making the editing-relevant contents editable. 3) Across two different datasets (different resolution and prompt complexity), FlowCycle demonstrates consistent superiority, indicating our robust effectiveness. More qualitative results are in the Appendix.

Quantitative Evaluation. As shown in Figure 4, by only optimizing noises under the cycle constraints and getting a target-aware intermediate state, our method can maintain a remarkable source consistency (*i.e.*, the lower *Structure Distance*). The source consistency is a significant outcome benefited from the cycle consistency constraints. At the same time, our method can also achieve a good semantic alignment (*i.e.*, the higher *CLIP score*) to the target prompt. Overall, our method can achieve a good balance between source consistency and semantic alignment (closest to the bottom right corner). As a highlight, on the PIE-Bench, our method based on SD-1.5 can get a very close performance to the strong baseline FlowEdit based on SD-3. This further demonstrates our effectiveness. We provide complete quantitative results in the Appendix.

4.2 Ablation Study

Optimization Rounds. To analyze how optimization with cycle constraints influences the editing performance, we provided visualized examples in Figure 6. Each example contained the editing results of different optimization rounds (0,20,60,100). Before the optimization (*i.e.*, adding a random noise), it already achieves a good target semantic alignment. For example, the **cat** has already appeared to replace the **rabbit** (0 round of the first row), and the **girl’s short hair** has also appeared to replace the **long hair** (0 round of the second row). However, the editing-irrelevant parts (*e.g.*, **background**, **candies**, and **girl**) are quite different from the source images. As the optimization proceeds, the source consistency has gradually improved, while the target semantics remain well-aligned. This indicates that FlowCycle can maintain the editing-irrelevant parts while accomplishing edit the target parts. Quantitative results in the Appendix.

Intermediate State Transfer. To further verify our main claim (*i.e.*, the target-aware corruption can make the editing-relevant part editable while maintaining the editing-irrelevant part), we conducted an intermediate state transfer experiment. As shown in Figure 7, given a source image and its source prompt “...white dog...”. We first got three different intermediate states:

- 1) *Random State*: It was acquired by adding a random noise to the source image.
- 2) *Transfer State*: We optimized a noise under the “...white dog...” (source) - “...husky...” (target) prompt pair using our method. Then we acquired the transfer state by adding this noise to the source image.
- 3) *Optimized States*: We optimized three noises under the subsequent corresponding editing prompt pairs, respectively. Then we acquired three different optimized states by adding these noises to the source image, respectively.

After getting the above states, we edited the source image by following three target prompts, *i.e.*, “...gold retriever...”, “...lion...”, and “...lion with crown...”. As shown in Figure 7, we have two observations: 1) The *Transfer State* obviously can lead to better editing results than *Random States*. This demonstrates there is good transferable ability of target-aware intermediate states across similar editing patterns (*e.g.*, change the white dog to another animal). 2) Although the *Transfer State* gets good performance when transferring between cases with similar patterns, it still can not outperform the *Optimized States*. Because the *Optimized States* have more accurate target-aware corruption and source maintaining capacity, leading to the best results among them. This experiment verifies the claim and significance of target-aware editing, and it also provides a protocol to reuse the optimized intermediate states.

Editing Deviation Magnitudes. As shown in the first row of Figure 7, from left to right, the editing target deviation from the source image is increasing (from “white dog → golden retriever” to “white dog → lion with crown”). We can see that the background preservation is also becoming worse along with the increasing deviation. This demonstrates the limitation of the current target-aware manner, *i.e.*, not flexible to varied editing deviations. In contrast, our target-aware states (the third row) can adapt better to this type of large-deviation editing (with consistent background preservation). This indicates that target-aware




Fig. 7: Transfer ability of intermediate state. It shows that the intermediate state (e.g., obtained from “white dog” into “husky dog” pair) can work well for other targets that have a similar edit pattern (e.g., “white dog” into “golden retriever”).

Table 1: Computational overhead comparisons between our accelerated method and other methods on PIE-Bench with SD3. Results are based on a single A800-80GB GPU.

	Method	Distance ↓	CLIP ↑	Time (s)	GPU (GB)
Optimization-free	FlowAlign	0.028	25.28	4	16
	Qwen-Image-Edit	0.054	25.20	92	60
	iRFDS	0.069	25.12	132	34
Optimization-based	Ours	0.013	25.48	247	16
	Ours-accelerate-weak	0.014	25.41	160	16
	Ours-accelerate-mid	0.014	25.34	126	16
	Ours-accelerate-strong	0.017	25.31	90	16

states can be adjusted according to the specific editing target, thus making it more flexible for varied editing deviation magnitudes.

4.3 Computational Overhead Analysis

We gave computational overhead (including editing time per image and required GPU memory) comparisons between the current optimization-free methods, the optimization-based methods, and our FlowCycle.

According to Section 3.4, we can accelerate our optimization process for faster editing. To this end, we gave the results of three different acceleration versions. 1) *Ours-accelerate-weak*: the $f(n) = \text{rounding}(\sin \frac{\pi n}{2N} * M)$ and it can accelerate around 1.5 times over original FlowCycle. 2) *Ours-accelerate-mid*: the $f(n) = \text{rounding}(\frac{nM}{N})$, and it can accelerate around two times over the original

FlowCycle. 3) *Ours-accelerate-strong*: the $f(n) = \text{rounding}((1 - \cos \frac{\pi n}{2N}) * M)$, and it can accelerate around three times over the original FlowCycle.

As shown in Table 1, we can have three observations: 1) The acceleration operation can significantly reduce the editing time (*e.g.*, the time overhead of the “strong” version is only around 1/3 of the no-acceleration version) with a little performance degradation. 2) Compared with the SOTA optimization-free method *FlowAlign*, our method can gain better editing quality at the cost of a higher time cost. 3) Compared with the large editing model based method *Qwen-Image* and optimization-based *iRFDS*, *Ours-accelerate-strong* can outperform them on both editing quality and cost.

5 Conclusion

In this paper, we revealed the limitations of the target-agnostic intermediate state in the existing corruption-then-restoration editing paradigm and proposed to acquire target-aware intermediate states by optimizing two noises with cycle-consistent constraints. To satisfy consistency, the editing result is improved to maintain the editing-irrelevant contents of the source image for better source consistency. We conducted extensive experiments to show the effectiveness and significance of the target-aware manner.

Acknowledgment. This work was supported by Hong Kong SAR Research Grants Council (RGC) General Research Fund (16219025), National Natural Science Foundation of China (NSFC) Young Scientists Fund Category B (62522216), Hong Kong SAR Research Grants Council (RGC) Early Career Scheme (26208924), and National Natural Science Foundation of China (NSFC) Young Scientists Fund Category C (62402408).

References

1. Agustsson, E., Timofte, R.: Ntire 2017 challenge on single image super-resolution: Dataset and study. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 126–135 (2017) 11
2. AI, I., Gong, B., Zou, C., Zheng, C., Zhou, C., Yan, C., Jin, C., Shen, C., Zheng, D., Wang, F., et al.: Ming-omni: A unified multimodal model for perception and generation. arXiv preprint arXiv:2506.09344 (2025) 4
3. Beletskii, I., Kuznetsov, A., Alanov, A.: Inverse-and-edit: Effective and fast image editing by cycle consistency models. arXiv preprint arXiv:2506.19103 (2025) 5
4. Chen, R.T., Rubanova, Y., Bettencourt, J., Duvenaud, D.K.: Neural ordinary differential equations. Advances in neural information processing systems 31 (2018) 5
5. Dao, T., Wang, Z., Pham, K.T., Chen, L.: Steerflow: Steering rectified flows for faithful inversion-based image editing. arXiv preprint arXiv:2604.01715 (2026) 5
6. Deng, C., Zhu, D., Li, K., Gou, C., Li, F., Wang, Z., Zhong, S., Yu, W., Nie, X., Song, Z., et al.: Emerging properties in unified multimodal pretraining. arXiv preprint arXiv:2505.14683 (2025) 4

7. Deng, Y., He, X., Mei, C., Wang, P., Tang, F.: Fireflow: Fast inversion of rectified flow for image semantic editing. arXiv preprint arXiv:2412.07517 (2024) [2](#), [4](#)
8. Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., et al.: Scaling rectified flow transformers for high-resolution image synthesis. In: Forty-first international conference on machine learning (2024) [2](#), [10](#)
9. Fu, T.J., Qian, Y., Chen, C., Hu, W., Gan, Z., Yang, Y.: Univg: A generalist diffusion model for unified image generation and editing. arXiv preprint arXiv:2503.12652 (2025) [4](#)
10. Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., Cohen-Or, D.: Prompt-to-prompt image editing with cross attention control. arXiv preprint arXiv:2208.01626 (2022) [5](#), [12](#)
11. Hessel, J., Holtzman, A., Forbes, M., Le Bras, R., Choi, Y.: Clipscore: A reference-free evaluation metric for image captioning. In: Proceedings of the 2021 conference on empirical methods in natural language processing. pp. 7514–7528 (2021) [11](#)
12. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in neural information processing systems **33**, 6840–6851 (2020) [2](#)
13. Huberman-Spiegelglas, I., Kulikov, V., Michaeli, T.: An edit friendly ddpm noise space: Inversion and manipulations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12469–12478 (2024) [4](#), [5](#)
14. Jiang, Y., Wang, Z., Wang, Y., Yu, J., Zhuang, Y., Xiao, J., Chen, L.: Flowdc: Flow-based decoupling-decay for complex image editing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 25757–25766 (2026) [5](#)
15. Ju, X., Zeng, A., Bian, Y., Liu, S., Xu, Q.: Pnp inversion: Boosting diffusion-based editing with 3 lines of code. International Conference on Learning Representations (ICLR) (2024) [10](#)
16. Kim, J., Hong, Y., Ye, J.C.: Flowalign: Trajectory-regularized, inversion-free flow-based image editing. arXiv preprint arXiv:2505.23145 (2025) [2](#), [3](#), [5](#), [11](#)
17. Kim, J., Park, J., Song, Y., Kwak, N., Rhee, W.: Reflex: Text-guided editing of real images in rectified flow via mid-step feature extraction and attention adaptation. arXiv preprint arXiv:2507.01496 (2025) [2](#), [3](#), [5](#)
18. Kingma, D.P.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014) [10](#)
19. Kulikov, V., Kleiner, M., Huberman-Spiegelglas, I., Michaeli, T.: Flowedit: Inversion-free text-based editing using pre-trained flow models. arXiv preprint arXiv:2412.08629 (2024) [2](#), [3](#), [5](#), [11](#)
20. Labs, B.F.: Flux. <https://github.com/black-forest-labs/flux> (2024) [2](#)
21. Lee, H., Kang, M., Han, B.: Text-driven image editing using cycle-consistency-driven metric learning (2023) [5](#)
22. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. arXiv preprint arXiv:2210.02747 (2022) [5](#), [6](#)
23. Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. arXiv preprint arXiv:2209.03003 (2022) [5](#), [6](#)
24. Mao, J., Wang, K., Xiang, Y., Chen, K.: Tweezedit: Consistent and efficient image editing with path regularization. arXiv preprint arXiv:2508.10498 (2025) [3](#), [5](#)
25. McCann, R.J.: A convexity principle for interacting gases. Advances in mathematics **128**(1), 153–179 (1997) [6](#)
26. Meng, C., Song, Y., Song, J., Wu, J., Zhu, J.Y., Ermon, S.: Sdedit: Image synthesis and editing with stochastic differential equations. arXiv preprint arXiv:2108.01073 (2021) [5](#), [8](#), [11](#)

27. Mokady, R., Hertz, A., Aberman, K., Pritch, Y., Cohen-Or, D.: Null-text inversion for editing real images using guided diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 6038–6047 (2023) [4](#), [5](#)
28. Parmar, G., Kumar Singh, K., Zhang, R., Li, Y., Lu, J., Zhu, J.Y.: Zero-shot image-to-image translation. In: ACM SIGGRAPH 2023 conference proceedings. pp. 1–11 (2023) [12](#)
29. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022) [12](#)
30. Rout, L., Chen, Y., Ruiz, N., Caramanis, C., Shakkottai, S., Chu, W.S.: Semantic image inversion and editing using rectified stochastic differential equations. arXiv preprint arXiv:2410.10792 (2024) [2](#), [3](#), [5](#)
31. Sasaki, H., Willcocks, C.G., Breckon, T.P.: Unit-ddpm: Unpaired image translation with denoising diffusion probabilistic models. arXiv preprint arXiv:2104.05358 (2021) [5](#)
32. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020) [2](#), [4](#)
33. Tumanyan, N., Bar-Tal, O., Bagon, S., Dekel, T.: Splicing vit features for semantic appearance transfer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10748–10757 (2022) [11](#)
34. Tumanyan, N., Geyer, M., Bagon, S., Dekel, T.: 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. pp. 1921–1930 (2023) [5](#)
35. Wang, B., Yue, Z., Zhang, F., Chen, S., Bi, L., Zhang, J., Song, X., Chan, K.Y., Pan, J., Wu, W., et al.: Selftok: Discrete visual tokens of autoregression, by diffusion, and for reasoning. arXiv preprint arXiv:2505.07538 (2025) [4](#)
36. Wang, J., Pu, J., Qi, Z., Guo, J., Ma, Y., Huang, N., Chen, Y., Li, X., Shan, Y.: Taming rectified flow for inversion and editing. arXiv preprint arXiv:2411.04746 (2024) [2](#), [3](#), [5](#)
37. Wang, Z., Li, L., Chen, L.: Recent advances in finetuning multimodal large language models. AI Magazine **46**(3), e70025 (2025) [4](#)
38. Wang, Z., Xu, Y., Xiao, J., Chen, L.: Flowmotion: Training-free flow guidance for video motion transfer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 38143–38153 (2026) [5](#)
39. Wu, C.H., De la Torre, F.: A latent space of stochastic diffusion models for zero-shot image editing and guidance. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7378–7387 (2023) [5](#)
40. Wu, C., Li, J., Zhou, J., Lin, J., Gao, K., Yan, K., Yin, S.m., Bai, S., Xu, X., Chen, Y., et al.: Qwen-image technical report. arXiv preprint arXiv:2508.02324 (2025) [4](#), [12](#)
41. Xiao, S., Wang, Y., Zhou, J., Yuan, H., Xing, X., Yan, R., Li, C., Wang, S., Huang, T., Liu, Z.: Omnigen: Unified image generation. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 13294–13304 (2025) [4](#)
42. Xu, S., Huang, Y., Pan, J., Ma, Z., Chai, J.: Inversion-free image editing with language-guided diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9452–9461 (2024) [2](#)
43. Xu, S., Ma, Z., Huang, Y., Lee, H., Chai, J.: Cyclenet: Rethinking cycle consistency in text-guided diffusion for image manipulation. Advances in Neural Information Processing Systems **36**, 10359–10384 (2023) [5](#)

- 44. Xu, Y., Wang, Z., Li, K., Xiao, J., Chen, L.: Freetuner: Any subject in any style with training-free diffusion. *ACM Transactions on Multimedia Computing, Communications and Applications* (2024) 5
- 45. Yang, X., Chen, C., Yang, X., Liu, F., Lin, G.: Text-to-image rectified flow as plug-and-play priors. *arXiv preprint arXiv:2406.03293* (2024) 11
- 46. Zhou, Z., Deng, Y., He, X., Dong, W., Tang, F.: Multi-turn consistent image editing. *arXiv preprint arXiv:2505.04320* (2025) 2, 3, 4, 5
- 47. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2223–2232 (2017) 5